Optimizing Noisy Complex Systems Liable to Failure

Davin Lunz

OPTIMISING NOISY COMPLEX SYSTEMS LIABLE TO FAILURE

DAVIN LUNZ*

Abstract. Inspired by complex systems in social and industrial contexts, we consider a family of coupled diffusion processes modeling system components, and an associated system objective. Each process is inherently noisy, driven by a controllable drift, and fails upon reaching a critical state. Interdependence is captured via the global objective and the governing dynamics (correlated noise, cascading failures). Analytical and numerical calculations reveal that the optimal strategies to steer such systems so as to maximise the objective are highly coupled, depending strongly on the state of the entire system. Strikingly, they exhibit a rich set of bifurcations, describing qualitatively different strategies throughout the parameter space.

Key words. Stochastic processes with reset, optimal control of PDEs, asymptotic analysis, discrete adjoint approach

AMS subject classifications. 35Q84, 49K20, 49N90, 90B25, 35B40

1. Introduction.

Complex systems encompass collective dynamics born out of the interactions between constituent components, and exhibit fascinating emergent features. The theory of complex systems — with broad application in power transmission, information cascade, disease outbreak, and biochemical processes to name a few — seeks to understand emergent behaviour as a function of the individual building blocks and their interdependence, with the aim of controlling such systems to guarantee robustness [54].

Network structure and dynamics are a prominent focus of complex systems theory [8, 15, 24, 49, 53]. In particular, cascading failures are widely studied as flow through a capacity-limited network. When a component fails the load redistribution can cause neighbouring components to fail, propagating an avalanche of failure [45, 46, 48, 51]. Characterising criticality by studying what network features allow this instability to manifest provides insight for failure-resilient design [1, 14, 32] and response strategies [47]. While previous work revealed the conditions for a single network to fail, less attention has been devoted to interacting networks, each liable to failure [10].

In this work, we address the question of interacting networks liable to failure from a slightly different perspective. We consider a network of interconnected, controllable diffusion processes. Each process represents the performance of a single network. Crucially, the diffusion processes undergo 'reset' upon reaching a critical state, modeling a single network failure. By modeling the single network by a diffusion process, we abstract away the intra-network failure mechanism, and instead focus on properties of the inter-network system, including the case of inter-network cascading failures.

Failures are costly but may not be catastrophic in the long run. This motivates us to consider a long-term objective function that rewards high performance over long time horizons while also accounting for costly failure. This proves a vexing challenge for systems where high performance comes hand-in-hand with high risk of failure, begging the question of how to best balance the risks and rewards. Our aim is to determine the optimal control drift (that which maximises the objective) subject to the complex system dynamics, thereby balancing the inherent trade-off

*Inria Saclay – Île de France, 91120 Palaiseau, France (davin.lunz@inria.fr)
École Polytechnique, CMAP 91128 Palaiseau, France
Institut Pasteur, 75015 Paris, France.

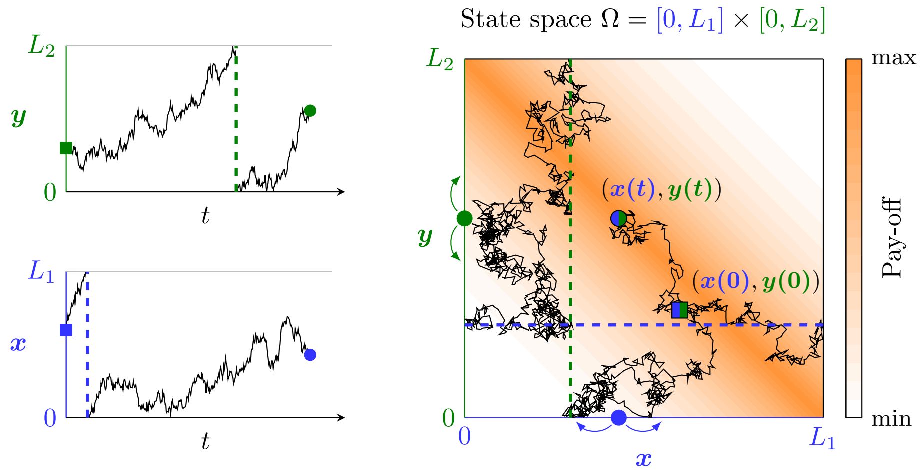

Fig. 1: Sample paths of two processes subject to reset, $x(t)$ and $y(t)$, each evolving on a domain $[0, L_i]$, resetting to zero upon hitting L_i (dashed lines). Note that the reset events act to set the state from L_i to 0, therefore, dashed lines show state changes from top to bottom in the one-dimensional plots on the left, where the corresponding coordinate jumps to zero in the two-dimensional state space $(x(t), y(t))$ on the right. The initial state is marked with a square, and the state at some later time t is marked with a circle. The processes may be *coupled* via their governing dynamics or by a pay-off function (here we depict a coupling pay-off function of $x + y$).

while accounting for the complexity of process interdependence.

The elementary processes in our model are one-dimensional diffusion processes subject to reset. Each process state evolves via a tunable drift under the influence of noise. A reset law governs jumps back to an origin. Aptly, resetting stochastic processes are versatile models, finding application from stochastic search algorithms [17, 18, 22, 25, 26, 40], through biological systems of DNA polymerase binding [7, 16, 50] and active matter [20, 35, 52], to ecological populations [5, 9, 36, 37, 38], communications queues [11, 33, 34], social [42] and mechanical systems in the study of reliability [12, 13, 41]. Resetting holds the process away from equilibrium [19], making it a useful model for non-equilibrium components of a complex system. Refs. [20, 21, 43] provide more extensive references.

We couple the individual processes via the global objective function (e.g. in power grids, where aggregate supply must match demand and an excess of either supply or demand is undesirable, the ideal output of any single power station depends on the other stations; see Figure 1). In the first instance we assume that each process observes the entire system state, and later we consider partial observability. Furthermore, we allow noise sources to be correlated among different processes, capturing inter-network noise covariance (e.g. fluctuations in wind conditions affecting wind turbines may be correlated depending on location).

The article is structured as follows. In section 2 we formulate the optimal control problem directly in terms of the underlying stochastic process. In section 3 we reformulate the problem in terms of its associated law by appealing to renewal theory. In section 4 we outline the numerical discretisation and the adjoint optimisation schemes we employ to compute solutions to the reformulated problem. The adjoint setting motivates an accompanying continuum adjoint analysis. In section 5 we illustrate concrete solutions, showcasing some of the bifurcations solutions undergo as problem parameters vary. Finally, in section 6 we discuss our findings and provide an outlook for future work.

2. Problem statement. Consider a family of d diffusion processes with state $\boldsymbol{x}(t) = (x_1(t), \dots, x_d(t))$ at time t, evolving within the domain $\Omega = [0, L_1] \times \cdots \times [0, L_d] \subset \mathbb{R}^d$ governed by the Itô SDE

$$\mathrm{d}\boldsymbol{x}(t) = \boldsymbol{v}(\boldsymbol{x}(t))\,\mathrm{d}t + \boldsymbol{\sigma}(\boldsymbol{x}(t))\,\mathrm{d}\boldsymbol{W}(t), \qquad \boldsymbol{x}(0) = \boldsymbol{x}_0, \tag{2.1}$$

for a drift coefficient $\boldsymbol{v} : \Omega \to \mathbb{R}^d$, a diffusion coefficient $\boldsymbol{\sigma} : \Omega \to \mathbb{R}^{d\times m}$, where $\boldsymbol{W}(t) \in \mathbb{R}^m$ is a standard m-dimensional Wiener process. The initial condition $\boldsymbol{x}_0 \in \Omega$ may be chosen from a distribution over the state space $P^0(\boldsymbol{x})$ for $\boldsymbol{x} \in \Omega$ or a distribution over the 0-boundary $\Phi^0(\boldsymbol{x})$ for $\boldsymbol{x} \in \partial\Omega_0$, or a combination of the two, such that

$$\int_\Omega P^0(\boldsymbol{x})\,\mathrm{d}\boldsymbol{x} + \int_{\partial\Omega_0} \Phi^0(\boldsymbol{x})\,\mathrm{d}\boldsymbol{x} = 1. \tag{2.2}$$

The domain Ω is enclosed by $\partial\Omega$ comprising the hyperplanes $x_i = 0$ and $x_i = L_i$ for each $i = 1, \dots, d$. We distinguish between these boundaries, denoting $\partial\Omega = \partial\Omega_0 \cup \partial\Omega_L$ for

$$\partial\Omega_0 = \{x_i = 0 \text{ for any } i\}, \qquad \partial\Omega_L = \{x_i = L_i \text{ for any } i\},$$

which we call the 0-boundary and L-boundary, respectively. The intersection $\partial\Omega_0 \cap \partial\Omega_L$ is nonempty but is of zero measure, and so while $\partial\Omega_0$ and $\partial\Omega_L$ are not strictly distinct, we nevertheless treat them as such. For points on the boundary we denote by $\boldsymbol{x}^c$ the point opposite $\boldsymbol{x}$: if the ith coordinate of $\boldsymbol{x}$ is on the boundary, $x_i \in \{0, L_i\}$, then $\boldsymbol{x}^c$ is the vector sharing all coordinates $x^c_j = x_j$ for all $j \neq i$ and $x^c_i = L_i - x_i$. While this is not well-defined for points $\boldsymbol{x} \in \partial\Omega_0 \cap \partial\Omega_L$, these are neglected since the intersection is of zero measure.

Upon hitting an L-boundary, $\boldsymbol{x} \in \partial\Omega_L$, a process has reached criticality and 'resets', that is, it is transported to the 0-boundary, $\boldsymbol{x}^c \in \partial\Omega_0$. This accounts exclusively for the reset of the process that reached criticality. In section 3.1 we discuss cascading failures, where the failure of one component results in the knock-on failure of other components, captured by setting multiple components to zero upon hitting an L-boundary. A particle in the vicinity of a 0-boundary experiences a reflecting boundary (that is, the resetting jump is only one way).

We impose bounds on the drift $\boldsymbol{v}(\boldsymbol{x}) \in [u, U]^d$, to capture the fact that our control of the processes is limited. We consider long-time pay-off functions $R(\boldsymbol{v})$ of the form

$$R(\boldsymbol{v}) = \lim_{T\to\infty} \frac{1}{T}\left[\int_0^T f(\boldsymbol{x}(t))\,\mathrm{d}t - CN(T)\right], \tag{2.3}$$

which reward some regions of state space through f (that is, $f : \Omega \to \mathbb{R}$ is path-independent), and penalise resets via the cost C, with $N(T)$ counting the number of resets until time $t = T$.

We thus arrive at the problem formulation: we seek an optimal drift, maximising the objective

$$\max_{\boldsymbol{v}(\boldsymbol{x})\in[u,U]^d} R(\boldsymbol{v}), \tag{2.4}$$

subject to the dynamics (2.1) and the aforementioned reset rules.

Typically, the objective function to optimise requires an expectation to be taken over the stochastic process. However, since we are interested exclusively in the ergodic limit, there is no expectation in this formulation.

Our approach is to study the law of the process (2.1) subject to reset. We analyse the Fokker–Planck equation governing the law in the large-time limit, and leverage renewal theory to represent (2.4) as a PDE-constrained optimisation problem. We then pursue both analytical and numerical results that shed light on the solution structure.

3. The law and its analysis. Consider the law of the stochastic process (2.1) subject to reset, which is the probability density $P(\boldsymbol{x}, t)$ of the particle being at location $\boldsymbol{x}$ at time t, and is governed by the Fokker–Planck equation

$$\frac{\partial P}{\partial t}(\boldsymbol{x}, t) + \nabla \cdot \boldsymbol{\Phi}(\boldsymbol{x}, t) = 0, \tag{3.1a}$$

for $\boldsymbol{x} \in \mathring{\Omega}$ and $t > 0$, where the probability flux $\boldsymbol{\Phi}$ is given by

$$\boldsymbol{\Phi}(\boldsymbol{x}, t) = \boldsymbol{v}(\boldsymbol{x})P(\boldsymbol{x}, t) - \nabla \cdot (\boldsymbol{D}(\boldsymbol{x}, \boldsymbol{v}(\boldsymbol{x}))P(\boldsymbol{x}, t)), \tag{3.1b}$$

for the diffusivity tensor $\boldsymbol{D} = \boldsymbol{\sigma}\boldsymbol{\sigma}^\top/2$. The reset mechanism and reflecting boundaries are represented by the boundary conditions:

$$P(\boldsymbol{x}, t) = 0, \qquad \boldsymbol{x} \in \partial\Omega_L, \tag{3.1c}$$

$$\boldsymbol{\Phi}(\boldsymbol{x}, t) \cdot \boldsymbol{n}(\boldsymbol{x}) = -\Phi^0(\boldsymbol{x})\delta(t) + \boldsymbol{\Phi}(\boldsymbol{x}^c, t) \cdot \boldsymbol{n}(\boldsymbol{x}), \qquad \boldsymbol{x} \in \partial\Omega_0, \tag{3.1d}$$

where $\boldsymbol{n}(\boldsymbol{x})$ is the outward-facing normal at $\boldsymbol{x}$ in both terms. The first term on the right-hand side of (3.1d) captures an initial injection of probability mass density Φ^0 along $\partial\Omega_0$. The minus sign is because the normal is outward-facing, thus Φ^0 is positive. We specify an initial probability density P^0 in the domain interior $\mathring{\Omega}$ via

$$P(\boldsymbol{x}, 0) = P^0(\boldsymbol{x}), \qquad \boldsymbol{x} \in \mathring{\Omega}. \tag{3.1e}$$

We also impose the unit-mass requirement (2.2) for the initial probability within the domain and injected into it.

The system is ergodic when the diffusivity is non-degenerate [2, 3, 6, 28, 29, 39], therefore, the law will converge to a unique stationary distribution. Since we are interested in long time horizons during which systems recover from failure, as reflected in the long-term pay-offs (2.3), we may neglect the initial transient (and the initial conditions) and focus exclusively on the steady-state solution of (3.1). With this restriction to the large-time limit in mind, we now apply the renewal approach.

We introduce the simpler first-passage process corresponding to (3.1) without the resetting mechanism. We denote the corresponding density $p(\boldsymbol{x}, t)$ and flux $\boldsymbol{\phi}(\boldsymbol{x}, t)$, satisfying

$$\frac{\partial p}{\partial t}(\boldsymbol{x}, t) + \nabla \cdot \boldsymbol{\phi}(\boldsymbol{x}, t) = 0, \tag{3.2a}$$

for $\boldsymbol{x} \in \mathring{\Omega}$ and $t > 0$, where the flux is given by

$$\boldsymbol{\phi}(\boldsymbol{x}, t) = \boldsymbol{v}(\boldsymbol{x})p(\boldsymbol{x}, t) - \nabla \cdot (\boldsymbol{D}(\boldsymbol{x}, \boldsymbol{v}(\boldsymbol{x}))p(\boldsymbol{x}, t)). \tag{3.2b}$$

We impose the absorbing and reflecting boundary conditions

$$p(\boldsymbol{x}, t) = 0, \qquad \boldsymbol{x} \in \partial\Omega_L, \tag{3.2c}$$

$$\boldsymbol{\phi}(\boldsymbol{x}, t) \cdot \boldsymbol{n}(\boldsymbol{x}) = -\hat{\Phi}^0(\boldsymbol{x})\delta(t), \qquad \boldsymbol{x} \in \partial\Omega_0, \tag{3.2d}$$

with the initial injection contribution, and impose the initial conditions

$$p(\boldsymbol{x}, 0) = \hat{P}^0(\boldsymbol{x}), \qquad \boldsymbol{x} \in \mathring{\Omega}. \tag{3.2e}$$

The only difference between (3.1) and (3.2) is the flux boundary condition on $\partial\Omega_0$.

We introduce $Q(t)$ the probability density of the first-passage time to the L-boundary being t, which is given by the probability flux at the L-boundary

$$Q(t) = \int_{\partial\Omega_L} \boldsymbol{\phi}(\boldsymbol{x}, t) \cdot \boldsymbol{n}(\boldsymbol{x}) \, \mathrm{d}\boldsymbol{x}. \tag{3.3}$$

We may now express the large-time probability density P recursively, with the help of p and Q via

$$P(\boldsymbol{x}, t) = p(\boldsymbol{x}, t) + \int_0^t Q(\tau) P(\boldsymbol{x}, t - \tau) \, \mathrm{d}\tau, \tag{3.4}$$

representing the probability of a particle to be located at $\boldsymbol{x}$ if it is yet to hit an L-boundary, or, if it first hit the boundary at a previous time and arrived at $\boldsymbol{x}$ after reset. We introduce $N(t)$ the average number of resets up until time t. This may also be expressed recursively with the aid of Q via

$$N(t) = \int_0^t Q(\tau)(1 + N(t - \tau)) \, \mathrm{d}\tau, \tag{3.5}$$

representing the density of the first reset occurring at any intermediate time $0 \leq \tau \leq t$ followed by (one more than) the average number of resets occurring in the remaining time until t.

The renewal-theory approach, where the resetting process is expressible recursively, relies on the assumption that, after each reset, the process is described by p in (3.2) until the subsequent reset. This observation allows us to determine the initial distribution $\hat{P}^0$ and initial injection density $\hat{\Phi}^0$ as follows. Since the process resets to the 0-boundary but not the domain interior $\mathring{\Omega}$ it follows that p has no initial distribution within the domain, $\hat{P}^0 = 0$. Similarly, from the periodic boundary conditions (3.1d) we see that the initial injection density $\hat{\Phi}^0$ must coincide with the steady-state normal flux at the L-boundary. In other words, the underlying assumption of the renewal approach is satisfied for a zero initial distribution and the stationary flux exiting the L-boundary injected at the 0-boundary. Of course, the law (3.1) of the original process (2.1) may have an arbitrary initial condition P^0 and injection density Φ^0 imposed; these restrictions are only required for p in order that the renewal theory be applicable in the large-time limit. Despite having deduced the form of $\hat{\Phi}^0$, we retain this term as is for the moment, since we do not know the steady-state normal flux at the L-boundary.

A Laplace transform, which we denote by a tilde, diagonalises the convolutions in the recursive forms (3.4) and (3.5) to give $\widetilde{P}$ and $\widetilde{N}$ in terms of $\widetilde{p}$ and $\widetilde{Q}$, namely

$$\widetilde{P}(\boldsymbol{x}, s) = \frac{\widetilde{p}(\boldsymbol{x}, s)}{1 - \widetilde{Q}(s)}, \qquad \widetilde{N}(s) = \frac{\widetilde{Q}(s)}{s(1 - \widetilde{Q}(s))}, \tag{3.6}$$

with s being the Laplace frequency. We see from (3.2b) and (3.3) that $\widetilde{Q}$ is a function only of $\widetilde{p}$. Therefore, it follows from (3.6) that $\widetilde{P}$ and $\widetilde{N}$ are also functions only of $\widetilde{p}$. We thus seek to solve (3.2) in Laplace space, namely

$$s\widetilde{p}(\boldsymbol{x}, s) + \nabla \cdot \widetilde{\boldsymbol{\phi}}(\boldsymbol{x}, s) = 0, \tag{3.7a}$$

for

$$\widetilde{\boldsymbol{\phi}}(\boldsymbol{x}, s) = \boldsymbol{v}(\boldsymbol{x})\widetilde{p}(\boldsymbol{x}, s) - \nabla \cdot (\boldsymbol{D}(\boldsymbol{x}, \boldsymbol{v}(\boldsymbol{x}))\widetilde{p}(\boldsymbol{x}, s)), \tag{3.7b}$$

subject to the boundary conditions

$$\widetilde{p}(\boldsymbol{x}, s) = 0, \qquad \boldsymbol{x} \in \partial\Omega_L, \tag{3.7c}$$

$$\widetilde{\boldsymbol{\phi}}(\boldsymbol{x}, s) \cdot \boldsymbol{n}(\boldsymbol{x}) = -\hat{\Phi}^0(\boldsymbol{x}), \qquad \boldsymbol{x} \in \partial\Omega_0. \tag{3.7d}$$

Solving the transformed (3.7) analytically in the general case appears to be intractable. Our aim is to solve in the large-time asymptotic limit $s \to 0$. We pose the expansion

$$\widetilde{p}(\boldsymbol{x}, s) \sim \widetilde{p}_0(\boldsymbol{x}) + s\widetilde{p}_1(\boldsymbol{x}) + \cdots, \tag{3.8}$$

and adopt an analogous notation for the expansion of the corresponding flux terms $\widetilde{\boldsymbol{\phi}}$. The leading-order equation takes the form,

$$\nabla \cdot \widetilde{\boldsymbol{\phi}}_0(\boldsymbol{x}) = 0, \tag{3.9a}$$

for $\boldsymbol{x} \in \mathring{\Omega}$, subject to the boundary conditions

$$\widetilde{p}_0(\boldsymbol{x}) = 0, \qquad \boldsymbol{x} \in \partial\Omega_L, \tag{3.9b}$$

$$\widetilde{\boldsymbol{\phi}}_0(\boldsymbol{x}) \cdot \boldsymbol{n}(\boldsymbol{x}) = -\hat{\Phi}^0(\boldsymbol{x}), \qquad \boldsymbol{x} \in \partial\Omega_0. \tag{3.9c}$$

The first-order equation takes the form

$$\nabla \cdot \widetilde{\boldsymbol{\phi}}_1(\boldsymbol{x}) = -\widetilde{p}_0(\boldsymbol{x}), \tag{3.10a}$$

for $\boldsymbol{x} \in \mathring{\Omega}$, subject to the boundary conditions

$$\widetilde{p}_1(\boldsymbol{x}) = 0, \qquad \boldsymbol{x} \in \partial\Omega_L, \tag{3.10b}$$

$$\widetilde{\boldsymbol{\phi}}_1(\boldsymbol{x}) \cdot \boldsymbol{n}(\boldsymbol{x}) = 0, \qquad \boldsymbol{x} \in \partial\Omega_0. \tag{3.10c}$$

Using the divergence theorem, it follows from (2.2) and (3.9) that the leading-order normal flux through $\partial\Omega_L$ is given by

$$\begin{aligned}\int_{\partial\Omega_L} \widetilde{\boldsymbol{\phi}}_0(\boldsymbol{x}) \cdot \boldsymbol{n}(\boldsymbol{x})\,\mathrm{d}\boldsymbol{x} &= \int_{\Omega} \nabla \cdot \widetilde{\boldsymbol{\phi}}_0(\boldsymbol{x})\,\mathrm{d}\boldsymbol{x} - \int_{\partial\Omega_0} \widetilde{\boldsymbol{\phi}}_0(\boldsymbol{x}) \cdot \boldsymbol{n}(\boldsymbol{x})\,\mathrm{d}\boldsymbol{x} \\ &= \int_{\partial\Omega_0} \hat{\Phi}^0(\boldsymbol{x})\,\mathrm{d}\boldsymbol{x} \\ &= 1.\end{aligned} \tag{3.11}$$

Similarly at first order, it follows from (3.10) that

$$\begin{aligned}\int_{\partial\Omega_L} \widetilde{\boldsymbol{\phi}}_1(\boldsymbol{x}) \cdot \boldsymbol{n}(\boldsymbol{x})\,\mathrm{d}\boldsymbol{x} &= \int_{\Omega} \nabla \cdot \widetilde{\boldsymbol{\phi}}_1(\boldsymbol{x})\,\mathrm{d}\boldsymbol{x} - \int_{\partial\Omega_0} \widetilde{\boldsymbol{\phi}}_1(\boldsymbol{x}) \cdot \boldsymbol{n}(\boldsymbol{x})\,\mathrm{d}\boldsymbol{x} \\ &= -\int_{\Omega} \widetilde{p}_0(\boldsymbol{x})\,\mathrm{d}\boldsymbol{x}.\end{aligned} \tag{3.12}$$

Combining (3.3), (3.11), and (3.12), we find that, up to first-order,

$$\begin{aligned}\widetilde{Q}(s) &\sim \int_{\partial\Omega_L} \widetilde{\boldsymbol{\phi}}_0(\boldsymbol{x})\cdot\boldsymbol{n}(\boldsymbol{x})\,\mathrm{d}\boldsymbol{x} + s\int_{\partial\Omega_L} \widetilde{\boldsymbol{\phi}}_1(\boldsymbol{x})\cdot\boldsymbol{n}(\boldsymbol{x})\,\mathrm{d}\boldsymbol{x}\\ &= 1 - s\int_{\Omega} \widetilde{p}_0(\boldsymbol{x})\,\mathrm{d}\boldsymbol{x}.\end{aligned} \tag{3.13}$$

From (3.6) and (3.13), after inverting the Laplace transform, we deduce that the large-time distribution and the average number of resets may be expressed as

$$P(\boldsymbol{x}) \sim \frac{\widetilde{p}_0(\boldsymbol{x})}{\int_\Omega \widetilde{p}_0(\boldsymbol{z})\,\mathrm{d}\boldsymbol{z}}, \qquad \frac{N(T)}{T} \sim \frac{1}{\int_\Omega \widetilde{p}_0(\boldsymbol{x})\,\mathrm{d}\boldsymbol{x}}, \tag{3.14}$$

that is, P is a normalisation of the $\widetilde{p}_0$ density.

Equipped with the large-time distribution (3.14) and large-time reset count (3.14), we are able to find expressions for objective functions of the form (2.3), namely

$$R = \lim_{T\to\infty} \frac{1}{T}\left[\int_0^T f(\boldsymbol{x}(t))\,\mathrm{d}t - CN(T)\right] = \frac{\int_\Omega f(\boldsymbol{x})\widetilde{p}_0(\boldsymbol{x})\,\mathrm{d}\boldsymbol{x} - C}{\int_\Omega \widetilde{p}_0(\boldsymbol{x})\,\mathrm{d}\boldsymbol{x}}. \tag{3.15}$$

Crucially, the objective (3.15) is given exclusively in terms of $\widetilde{p}_0$, the solution of (3.9). Ultimately, the ergodicity of the process has allowed us to replace the time averages in (2.3) with ensemble averages in (3.15).

The ergodicity ensures that the stationary distribution P is independent of the initial conditions P^0 and Φ^0. Nevertheless, we find that $\widetilde{p}_0$ depends on $\hat{\Phi}^0$ in (3.9), which we retained since we did not know the stationary boundary-normal flux of P. This may now be addressed by noting from (3.14) that the steady-state P is proportional to $\widetilde{p}_0$. Therefore, imposing the periodic flux boundary condition (3.1d) on $\widetilde{p}_0$ itself guarantees the correct boundary-normal up to a multiplicative constant. The exact value of the constant is unimportant since $\widetilde{p}_0$ is normalised in (3.14), and thus it suffices to ensure it is nonzero. Ultimately, we arrive at the equation governing $\widetilde{p}_0$, namely

$$\nabla\cdot\widetilde{\boldsymbol{\phi}}_0(\boldsymbol{x}) = 0, \tag{3.16a}$$

for $\boldsymbol{x}\in\mathring{\Omega}$, subject to the boundary conditions

$$\widetilde{p}_0(\boldsymbol{x}) = 0, \qquad \boldsymbol{x}\in\partial\Omega_L, \tag{3.16b}$$

$$\widetilde{\boldsymbol{\phi}}_0(\boldsymbol{x})\cdot\boldsymbol{n}(\boldsymbol{x}) = \widetilde{\boldsymbol{\phi}}_0(\boldsymbol{x}^c)\cdot\boldsymbol{n}(\boldsymbol{x}), \qquad \boldsymbol{x}\in\partial\Omega_0, \tag{3.16c}$$

where $\boldsymbol{n}(\boldsymbol{x})$ is the outward-facing normal at $\boldsymbol{x}$ in both terms and the unit probability mass condition (2.2) takes the form

$$\int_{\partial\Omega_0} \widetilde{\boldsymbol{\phi}}_0(\boldsymbol{x})\cdot\boldsymbol{n}(\boldsymbol{x})\,\mathrm{d}\boldsymbol{x} = -1. \tag{3.16d}$$

It is worth noting that this subtlety involving the initial and boundary conditions of the p problem has not been previously addressed in the literature. This is because the previous studies have restricted the geometry to one-dimension [12, 41] (or imposed symmetry [13] with the same effect), whereas in this study the high dimensionality and lack of symmetry necessitates a more careful analysis. When we consider cascading failures in section 3.1, the particular case where all processes fail upon any reset guarantees that the process resets to a single point whereby this problem is again eliminated (but not the case of partial cascades, that is, if there is any process where, upon reaching criticality, not all processes reset).

3.1. Cascading failures. A cascading failure is any dependency in the system such that the failure of one component (say the ith component for some $i \in \{1, \ldots, d\}$) results not only in its own reset, but also the reset of at least one other component (say j for any $i \neq j \in \{1, \ldots, d\}$). Such dependencies may be asymmetrical. That is, the failure of process i might always cause the reset of process j, but not the converse: the failure of j need not cause the reset of i. A cascade may cause the reset of all other processes, but may also be partial; causing the reset of some other but not all other components.

Cascading failures are incorporated in the model by changing the boundary conditions imposed on the law to account for the fact that multiple processes may be reset when a single process reaches criticality. This raises a minor technical issue regarding the normal flux boundary condition. A cascading failure results in a particle reset via absorption at an L-boundary and injection at a 0-boundary where $x_i = x_j = 0$ for some $i \neq j$. The domain geometry at such a point is singular since the normal is not defined, making a normal flux boundary condition indeterminate. This was not previously a problem because, in the periodic case considered until now, the probability mass injected at the singular regions of the domain was of zero measure, however, this is no longer the case. To circumvent this issue, we may consider injection at a location vanishingly near this point. To preserve symmetry, we choose the internal point $x_i = x_j = \epsilon$ for some $\epsilon \ll 1$ which is in $\mathring{\Omega}$. A boundary point could also be chosen at the cost of a small loss of symmetry. In practice, we choose ϵ to simply be one discrete grid unit.

The prior analysis remains analogous. For the sake of concreteness, we demonstrate the case where reaching any critical state results in a cascading failure of all components. The governing Fokker–Planck equation (3.1) takes the form

$$\frac{\partial P}{\partial t}(\boldsymbol{x}, t) + \nabla \cdot \boldsymbol{\Phi}(\boldsymbol{x}, t) = \delta(\boldsymbol{x} - \boldsymbol{\epsilon}) \int_{\partial \Omega_L} \boldsymbol{\Phi}(\boldsymbol{z}, t) \cdot \boldsymbol{n}(\boldsymbol{z}) \, \mathrm{d}\boldsymbol{z}, \tag{3.17a}$$

for $\boldsymbol{x} \in \mathring{\Omega}$ and $t > 0$, where $\boldsymbol{\epsilon}$ is the vector with ϵ in each entry, $\boldsymbol{n}(\boldsymbol{z})$ is the outward-facing normal at $\boldsymbol{z} \in \partial\Omega_L$, and the probability flux $\boldsymbol{\Phi}$ remains as in (3.1b). The reset mechanism and reflecting boundaries are represented by the boundary conditions:

$$P(\boldsymbol{x}, t) = 0, \qquad \boldsymbol{x} \in \partial\Omega_L, \tag{3.17b}$$

$$\boldsymbol{\Phi}(\boldsymbol{x}, t) \cdot \boldsymbol{n}(\boldsymbol{x}) = -\Phi^0(\boldsymbol{x})\delta(t), \qquad \boldsymbol{x} \in \partial\Omega_0, \tag{3.17c}$$

where Φ^0 is the probability mass density injected initially. The initial probability density P^0 in the domain interior $\mathring{\Omega}$ is specified

$$P(\boldsymbol{x}, 0) = P^0(\boldsymbol{x}), \qquad \boldsymbol{x} \in \mathring{\Omega}. \tag{3.17d}$$

We again require that the initial probability has unit mass, satisfying (2.2).

With the same objective (2.3), we may apply the renewal theory if the first-passage process p describes the law of P between consecutive resets. For P governed by (3.17), this condition amounts to there being no injection probability for p on the 0-boundary (since the reset occurs into the domain interior), and the initial distribution of p is to match the right-hand side of (3.17a) for the steady-state solution $\boldsymbol{\Phi}(\boldsymbol{x})$. As before, we do not know the total boundary flux of the steady-state solution, however, this is some nonzero constant that scales $\widetilde{p}_0$ but is normalised in P via (3.14) and thus may be taken as unity without loss of generality. Seeking the large-time asymptotic solution leads us analogously to the system governing the leading-order Laplace-transformed first-passage density $\widetilde{p}_0$, namely

$$\nabla \cdot \widetilde{\boldsymbol{\phi}}_0(\boldsymbol{x}) = \delta(\boldsymbol{x} - \boldsymbol{\epsilon}), \tag{3.18a}$$

for $\boldsymbol{x} \in \mathring{\Omega}$ and the usual flux form (3.7a), subject to the boundary conditions

$$\widetilde{p}_0(\boldsymbol{x}) = 0, \qquad \boldsymbol{x} \in \partial\Omega_L, \tag{3.18b}$$

$$\widetilde{\boldsymbol{\phi}}_0(\boldsymbol{x}) \cdot \boldsymbol{n}(\boldsymbol{x}) = 0, \qquad \boldsymbol{x} \in \partial\Omega_0. \tag{3.18c}$$

In this case, application of the divergence theorem proves that the inhomogeneous condition

$$\int_{\partial\Omega_L} \widetilde{\boldsymbol{\phi}}_0(\boldsymbol{x}) \cdot \boldsymbol{n}(\boldsymbol{x}) \, \mathrm{d}\boldsymbol{x} = 1, \tag{3.19}$$

is satisfied by the solution of (3.18) and need not to be imposed to ensure a non-zero solution. For partial cascading failures, where a reset causes only some other components to fail, the form becomes hybrid; mixing periodic normal flux boundary conditions with delta functions of magnitude equal to the integrated flux over the corresponding critical boundary, while imposing (3.19). We do not further discuss partial cascades in this work.

To recap, by analysing the law of the process, we have reformulated the original stochastic optimal control problem (2.4) as a PDE-constrained optimisation problem of maximising (3.15) subject to the independent-failure model (3.16). This formulation may be extended to cascading failures, such as in (3.17). We now proceed to detail a numerical approach to solving the PDE-constrained optimisation problems.

4. The adjoint perspective.

4.1. Numerical implementation. Our aim in this section is to develop a numerical scheme to solve the PDE-constrained optimisation problem of maximising the objective (3.15) subject to the independent-failure model (3.16) or the cascading-failure model (3.18).

The first step is to discretise the associated PDEs (3.16) and (3.18). In Appendix A we provide complete details of a consistent finite-difference scheme tailored for anisotropic diffusion. Here, we assemble these discrete schemes (including the discretisation of the differential operator, the associated boundary conditions, and any associated inhomogeneity constraints) in generic matrix form:

$$\boldsymbol{A}\boldsymbol{p} = \boldsymbol{b}, \tag{4.1}$$

where $\boldsymbol{A}$ is the finite-difference matrix operator, and $\boldsymbol{b}$ is the inhomogeneous right-hand side of the system. The vector $\boldsymbol{p}$ comprises the discretised solution and perhaps a Lagrange multiplier for an inhomogeneity constraint.

The next step is to discretise the objective function R in (3.15), which we write as

$$R \approx \frac{\left(\sum_{\boldsymbol{j}\in\Gamma} f(\boldsymbol{x}_{\boldsymbol{j}}) p_{\boldsymbol{j}}\right) H - C}{\left(\sum_{\boldsymbol{j}\in\Gamma} p_{\boldsymbol{j}}\right) H} =: \frac{(\boldsymbol{f}^\top \boldsymbol{p})H - C}{(\boldsymbol{1}^\top \boldsymbol{p})H}, \tag{4.2}$$

where H is the discretisation of $\mathrm{d}\boldsymbol{x}$ the differential element (that is, the product of the state spacing in each dimension, see (A.9)), $\boldsymbol{j} \in \Gamma$ indexes discrete points in the discretised domain, so $\boldsymbol{p} = (p_{\boldsymbol{j}})_{\boldsymbol{j}}$ is the discrete solution at states $\boldsymbol{x_j} \in \Omega$, $\mathbf{1}$ is a vector of ones with dimension matching $\boldsymbol{p}$, and $\boldsymbol{f}$ is the vector of values $f(\boldsymbol{x_j})$. In the case where $\boldsymbol{p}$ also holds a Lagrange multiplier, we assume that the corresponding entries in the $\boldsymbol{f}$ and $\mathbf{1}$ vectors are zero.

Note that the scheme (4.1) depends on the (discretised) control $\boldsymbol{v}$ via the operator $\boldsymbol{A} = \boldsymbol{A}(\boldsymbol{v})$. Thus we are equipped with a mapping from the discrete control $\boldsymbol{v}$, via the system (4.1) providing the solution $\boldsymbol{p}$, to the discretised objective function $R(\boldsymbol{v}) := R(\boldsymbol{p}(\boldsymbol{v}))$. The final step, of optimising $R(\boldsymbol{v})$ with respect to $\boldsymbol{v}$, may be achieved by several approaches. Gradient-free methods require only the aforementioned ingredients to optimise the objective, however, these typically require a large number of iterations which proves computationally demanding since the solve can be very costly with a highly resolved discretisation. For this reason, we choose the L-BFGS-B (limited-memory control-constrained BFGS) algorithm [44], a gradient-based quasi-Newton method, designed to handle large-scale optimisation problems. This approach provides significantly faster convergence by exploiting the gradient information, however, it requires the gradient of R with respect to the control $\boldsymbol{v}$: $\nabla_{\boldsymbol{v}} R$. Note that we use the terms 'drift' and 'control' interchangeably; the former describes its physical function while the latter describes its conceptual role in the optimisation setting.

The most straightforward approach to compute $\nabla_{\boldsymbol{v}} R$ is via finite-differences, however, this suffers from both numerical inaccuracy [4] and requires prohibitively many solves: one for each component of $\boldsymbol{v}$. This second limitation is particularly challenging as increasing the discretisation resolution increases the dimension of $\boldsymbol{v}$. To remedy both of these drawbacks, we calculate an exact expression for the derivative via the adjoint approach [23]. This requires the equivalent of a single extra solve, irrespective of the dimension of $\boldsymbol{v}$ (an improvement by a factor on the order of $\mathcal{O}(2500)$ for the simulations in this study). We now demonstrate how the adjoint approach is applied to our discrete scheme (4.1) to yield the gradient $\nabla_{\boldsymbol{v}} R$ for use in the optimisation algorithm.

First, we introduce the discrete Lagrangian

$$\mathcal{L} = \frac{\boldsymbol{f}^\top \boldsymbol{p} - C/H}{\mathbf{1}^\top \boldsymbol{p}} - \boldsymbol{\lambda}^\top \left(\boldsymbol{A}\boldsymbol{p} - \boldsymbol{b} \right), \tag{4.3}$$

where $\boldsymbol{\lambda}$ is an adjoint vector of dimension matching $\boldsymbol{p}$. The dependence of $\mathcal{L}$ on the drift $\boldsymbol{v}$ is explicit via $\boldsymbol{A}$ and implicit via the dependence of $\boldsymbol{p}$ on $\boldsymbol{v}$ (also stemming from $\boldsymbol{A}$) through (4.1). Therefore, by the chain rule, we see that

$$\nabla_{\boldsymbol{v}} \mathcal{L} = \frac{\partial \mathcal{L}}{\partial \boldsymbol{p}} \frac{\partial \boldsymbol{p}}{\partial \boldsymbol{v}} + \frac{\partial \mathcal{L}}{\partial \boldsymbol{v}} = \left[\frac{\boldsymbol{f}^\top}{\mathbf{1}^\top \boldsymbol{p}} - \frac{\boldsymbol{f}^\top \boldsymbol{p} - C/H}{\mathbf{1}^\top \boldsymbol{p}} \frac{\mathbf{1}^\top}{\mathbf{1}^\top \boldsymbol{p}} - \boldsymbol{\lambda}^\top \boldsymbol{A} \right] \frac{\partial \boldsymbol{p}}{\partial \boldsymbol{v}} - \boldsymbol{\lambda}^\top \frac{\partial \boldsymbol{A}}{\partial \boldsymbol{v}} \boldsymbol{p}. \tag{4.4}$$

Then, if $\boldsymbol{\lambda}$ satisfies the costate equation, $\partial \mathcal{L} / \partial \boldsymbol{p} = 0$, which takes the form

$$\boldsymbol{A}^\top \boldsymbol{\lambda} = \frac{\boldsymbol{f}}{\mathbf{1}^\top \boldsymbol{p}} - \frac{\boldsymbol{f}^\top \boldsymbol{p} - C/H}{\mathbf{1}^\top \boldsymbol{p}} \frac{\mathbf{1}}{\mathbf{1}^\top \boldsymbol{p}}, \tag{4.5}$$

we obtain the gradient of the reward, namely

$$\nabla_{\boldsymbol{v}} R = \nabla_{\boldsymbol{v}} \mathcal{L} = \frac{\partial \mathcal{L}}{\partial \boldsymbol{v}} = -\boldsymbol{\lambda}^\top \frac{\partial \boldsymbol{A}}{\partial \boldsymbol{v}} \boldsymbol{p}. \tag{4.6}$$

The first equality in (4.6) follows from the fact that the Lagrangian coincides with the reward for all choices of $\boldsymbol{\lambda}$, since $\boldsymbol{p}$ satisfies (4.1).

It is worth taking stock of the complete numerical optimisation routine. At each iteration i, we begin with the current control iterate $\boldsymbol{v}^i$ and solve the forward problem (4.1) to give $\boldsymbol{p}$, and thus R from (4.2). Then the adjoint system (4.5) is solved to give $\boldsymbol{\lambda}$. Note that the adjoint problem is of the same dimensionality as the forward problem, and thus of similar computational complexity. Finally, the gradient is calculated using the state and costate via (4.6). Thus, we have an efficient recipe for computing $R(\boldsymbol{v}^i)$ and $\nabla_{\boldsymbol{v}} R(\boldsymbol{v}^i)$ at each iteration. Given an initial guess $\boldsymbol{v}_0$, the control drift is iteratively improved by the L-BFGS-B quasi-Newton method based on the values $R(\boldsymbol{v}^i)$ and $\nabla_{\boldsymbol{v}} R(\boldsymbol{v}^i)$ until some convergence criterion is reached. Further details and an extended discussion on the numerical convergence is provided in Appendix B.

Before we delve into numerical results, we note that all of our simulations converged towards optimal drifts of type bang–bang: taking either their upper bound or lower bound but not intermediate values. This observation motivates the analysis in section 4.2, where we turn to the calculus of variations to argue that, under mild assumptions, this observation is expected. The adjoint analysis undertaken in the continuum setting complements the discrete approach outlined above, demonstrating that the discrete approach is consistent with the continuum variational calculus.

4.2. Bang–bang control. Numerical results suggest that, when $\boldsymbol{D}$ has no $\boldsymbol{v}$-dependence, the optimal controls are bang–bang, taking either their upper bound U or lower bound u but not intermediate values (recall from the problem statement (2.4) that we only consider bounded controls $\boldsymbol{v}(\boldsymbol{x}) \in [u, U]^d$). We offer a formal justification for this observation based on the calculus of variations. Consider an optimal drift $\boldsymbol{v}(\boldsymbol{x})$ with respect to some objective function $R(\boldsymbol{v})$ of the form (3.15), with reward function $f(\boldsymbol{x})$. Importantly, we assume that f is not constant in any open ball in $\Omega \subset \mathbb{R}^d$. By way of contradiction, we assume that the optimal drift $\boldsymbol{v}(\boldsymbol{x}) \in (u, U)^d$ in some open ball $\Lambda \subset \Omega$, that is, the optimal drift lies strictly within its bounds in some small region.

Due to the optimality of $\boldsymbol{v}$, we expect from the first-order optimality conditions that the first variation of R with respect to the control drift, δR, vanishes (for some set of admissible perturbations we will describe). To help calculate the first variation, we introduce the (continuum) Lagrangian

$$\mathcal{L} = R(\boldsymbol{v}) - \int_\Omega \lambda(\boldsymbol{x}) \left\{ \nabla \cdot [\boldsymbol{v}(\boldsymbol{x}) \widetilde{p}_0(\boldsymbol{x}) - \nabla \cdot (\boldsymbol{D}(\boldsymbol{x}) \widetilde{p}_0(\boldsymbol{x}))] \right\} \mathrm{d}\boldsymbol{x}, \tag{4.7}$$

where $\lambda(\boldsymbol{x})$ is an adjoint variable. Integrating by parts (and momentarily suppressing the dependence on $\boldsymbol{x}$, where clear, for brevity), we find that

$$\begin{aligned} \int_\Omega \lambda \left\{ \nabla \cdot [\boldsymbol{v} \widetilde{p}_0 - \nabla \cdot (\boldsymbol{D} \widetilde{p}_0)] \right\} \mathrm{d}\boldsymbol{x} &= \int_\Omega (-\nabla\lambda \cdot \boldsymbol{v} - \nabla\nabla\lambda : \boldsymbol{D}) \, \widetilde{p}_0 \, \mathrm{d}\boldsymbol{x} \\ &\quad + \int_{\partial\Omega_0} \nabla\lambda \cdot (\boldsymbol{D}\boldsymbol{n}) \widetilde{p}_0 \, \mathrm{d}\boldsymbol{x} + \int_{\partial\Omega_L} [\lambda(\boldsymbol{x}) - \lambda(\boldsymbol{x}^c)] \, \widetilde{\phi}_0 \cdot \boldsymbol{n} \, \mathrm{d}\boldsymbol{x}. \end{aligned} \tag{4.8}$$

It follows, by taking the first variation of $\mathcal{L}$ with respect to $\widetilde{p}_0$, that the costate equation is given by

$$\nabla\lambda(\boldsymbol{x}) \cdot \boldsymbol{v}(\boldsymbol{x}) + \nabla\nabla\lambda(\boldsymbol{x}) : \boldsymbol{D}(\boldsymbol{x}) = -\frac{f(\boldsymbol{x})}{\int_\Omega \widetilde{p}_0(\boldsymbol{z}) \, \mathrm{d}\boldsymbol{z}} - \frac{\int_\Omega f(\boldsymbol{z}) \widetilde{p}_0(\boldsymbol{z}) \, \mathrm{d}\boldsymbol{z} - C}{\left(\int_\Omega \widetilde{p}_0(\boldsymbol{z}) \, \mathrm{d}\boldsymbol{z}\right)^2}, \tag{4.9a}$$

subject to the boundary conditions

$$\nabla\lambda(\boldsymbol{x})\cdot(\boldsymbol{D}(\boldsymbol{x})\boldsymbol{n}) = 0, \qquad \boldsymbol{x}\in\partial\Omega_0, \tag{4.9b}$$

$$\lambda(\boldsymbol{x}) = \lambda(\boldsymbol{x}^c), \qquad \boldsymbol{x}\in\partial\Omega_L. \tag{4.9c}$$

We identify the right-hand side of (4.9) as the continuum analog of the right-hand side of (4.5). Armed with the costate that satisfies (4.9) we may now analyse the first-order optimality conditions: the first variation of the objective R with respect to the control, in all admissible perturbation directions $\boldsymbol{u}$, vanishes, namely

$$\delta R = \delta\mathcal{L} = \int_\Omega [\nabla\lambda(\boldsymbol{x})\cdot\boldsymbol{u}(\boldsymbol{x})]\,\widetilde{p}_0(\boldsymbol{x})\,\mathrm{d}\boldsymbol{x} = 0. \tag{4.10}$$

An admissible perturbation $\boldsymbol{u}$ is one with support in Λ such that, for small enough ϵ, the perturbed control is admissible: $\boldsymbol{v}(\boldsymbol{x}) + \epsilon\boldsymbol{u}(\boldsymbol{x}) \in [u, U]$ for all $\boldsymbol{x}\in\Lambda$. It suffices to take bounded functions supported on any closed subset of Λ. Then, since $\widetilde{p}_0$ is strictly positive on Λ (for non-degenerate diffusion the stationary distribution must be strictly positive on $\mathring{\Omega}$ since every location is reachable), by the fundamental lemma of the calculus of variations we deduce that $\nabla\lambda\equiv 0$ on Λ. Therefore, from (4.9a), we see that f must be constant on Λ, in contradiction to our non-stationarity assumption. The preceding argument is independent of the boundary conditions or inhomogeneities forcing $\widetilde{p}_0$, and therefore it remains valid in the case of cascading failures.

We have shown that the optimal drift must be bang–bang on the interior of the domain (assuming f is not stationary, similar to the condition in [41]). Note that when $\boldsymbol{D}$ depends on the drift $\boldsymbol{v}$ this result does not hold [41].

We now turn our attention to studying concrete numerical simulations.

5. A tale of two processes.

The optimal strategy, a function in $\Omega\to[u,U]^d$, may be easily visualised for symmetric systems (i.e. dynamics and objectives not distinguishing between different processes) where $d=2$. In this case, the symmetry ensures that the optimal drift satisfies $v_1(x,y) = v_2(y,x)$, and thus we may illustrate just $v_1:\Omega\to[u,U]$, say, for each set of parameters, which is simply a surface over the two-dimensional domain $\Omega\subset\mathbb{R}^2$. Symmetry also allows for significant computational efficiency via dimension reduction.

Remarkably, numerical calculations of the optimal drift turn out to be bang–bang controls, that is, $v_i(\boldsymbol{x})$ takes either its upper bound U or lower bound u, but not intermediate values. We demonstrated analytically in section 4.2 that this observation is a general characterisation of solutions for a large class of objective functions, noise couplings, and cascading-failure interdependence. This allows us to present optimal drift strategies $v_1:\Omega\to\{u,U\}$ as contour plots over the domain Ω with the contours delineating between just two function values, U and u, representing regions of maximum and minimum drift, which we shade and leave unshaded, respectively (further details in Appendix C). Since all contours on such a plot represent this single delineation, several such plots, for different parameter values, may be overlayed to capture how the contours (and the optimal drifts they illustrate) are affected by changing parameters. Unless specified otherwise, we use the forms

$$\begin{gathered} d=2,\quad L_1=L_2=U=1,\quad u=C=0,\\ \boldsymbol{\sigma}=\sqrt{2}\boldsymbol{I},\quad f(x,y) = -(1-x-y)^2, \end{gathered} \tag{5.1}$$

where $\boldsymbol{I}$ denotes the identity tensor and the associated diffusivity tensor is given by $\boldsymbol{D}=\boldsymbol{I}$. The reward function f in (5.1) encodes a pay-off for systems where the total

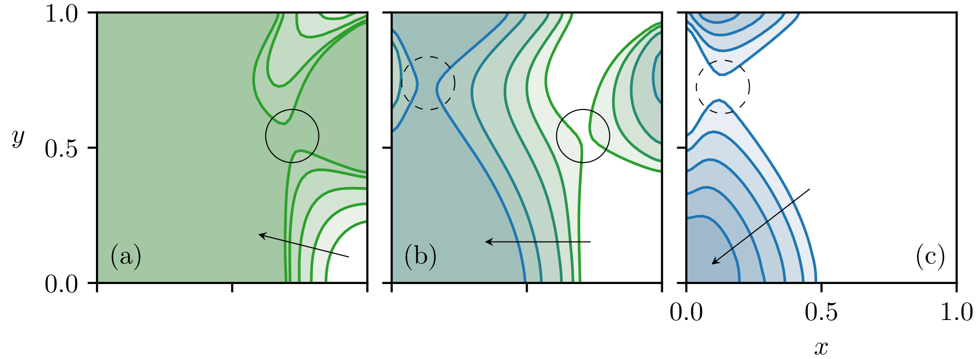

Fig. 2: Optimal drift $v_1(x, y)$ for various resetting costs: (a) $C \in \{-0.04, -0.03, -0.022, -0.016, -0.0135\}$, (b) $C \in \{-0.013, -.008, 0, 0.015, 0.029\}$, (c) $C \in \{0.032, 0.045, 0.065, 0.09, 0.12\}$. Maximal drift ($v_1 = U$) is shaded, minimum drift ($v_1 = u$) is unshaded. Arrows show directions of increasing C. All plots share the same (x, y) axes. Circles of matching style highlight the topology-changing bifurcations. Changes in colour are for visual aid but of no quantitative relevance.

'performance', that is, the sum of the states $x(t) + y(t)$, is near unity. We can think of this as a simple model rewarding a desired total system output, with both excess supply or deficient supply (i.e. excess demand) being undesirable. This induces a natural coupling between the processes.

First, we study the influence of the reset cost C. In Figure 2 we illustrate optimal strategies for various values of C, and observe how for negative reset costs (i.e. reset rewards) the state space is almost entirely shaded, as it is optimal to impose maximum drift except for corner regions (Figure 2a). The lower right is a region of high pay-off and resetting from $x = 1$ to $x = 0$ for small y results in jumping towards the low-pay-off origin. The ideal strategy is thus to drift only in the y direction (recall that $v_1(x, y) = v_2(y, x)$, therefore the y-drift is maximal in the lower right while the x-drift is minimal). This extends the duration in the high pay-off zone while driving the state to larger y such that a jump from $x = 1$ to $x = 0$ is further removed from the low pay-off zone. Similarly in the low-pay-off upper right, it is best to reset while avoiding proximity to $(x, y) = (1, 1)$ from which reset to the origin is likely.

As C is increased beyond some critical value near $C \approx -0.013$, the strategy bifurcates as the contours undergo a topological change: the regions of maximum drift are no longer connected (Figures 2a and 2b, solid circles). As C increases further, the region of maximum drift in the upper right vanishes, and at some critical value near $C \approx 0.03$ the remaining region of maximum drift on the left bifurcates, 'pinching off' into two (Figures 2b and 2c, dashed circles). Ultimately, for large values of C the optimal strategy is to impose minimum drift throughout most of the state space except for some corner regions (Figure 2c). In the lower left the pay-off is low and maximum drift still outweighs the increased risk of costly resets. In the upper left, it is better to drift in the x direction at the cost of slightly lower pay-off so that if a reset occurs the system is not near the origin.

The simplest (two-dimensional, symmetric) setting exemplifies how the deceptively simple-looking problem belies a rich set of highly non-trivial optimal strategies.

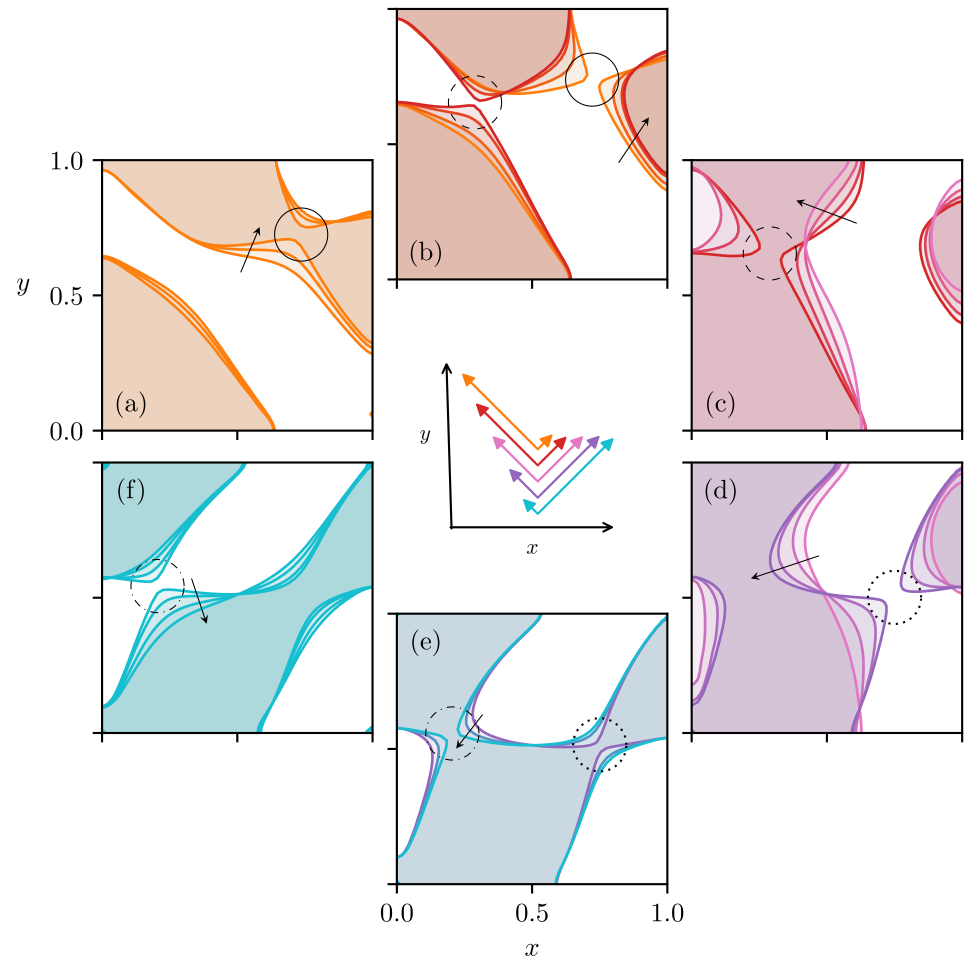

Fig. 3: Optimal drift $v_1(x, y)$ for the correlated noise (5.2) and various correlation parameters: (a) $s \in \{-0.2, -0.18, -0.165\}$, (b) $s \in \{-0.16, -0.14, -0.115, -0.105\}$, (c) $s \in \{-0.1, -0.07, -0.03, 0\}$, (d) $s \in \{0, 0.04, 0.09, 0.11\}$, (e) $s \in \{0.115, 0.125, 0.132\}$, (f) $s \in \{0.135, 0.145, 0.17, 0.2\}$. Maximal drift ($v_1 = U$) is shaded, minimum drift ($v_1 = u$) is unshaded. Black arrows show directions of increasing s. All plots share the same (x, y) axes. Circles of matching style highlight the topology-changing bifurcations. The axis in the figure centre shows the decomposition of the diffusion tensor (5.3) for $s \in \{-0.2, -0.05, 0, 0.05, 0.2\}$. Changes in colour are for visual aid but of no quantitative relevance.

Multiple bifurcations distinguish between different contour topologies, corresponding to the emergence of qualitatively different strategies.

Next, we study the influence of correlated noise via the one-dimensional family of diffusion coefficients:

$$\boldsymbol{\sigma}(s) = \sqrt{2}\begin{pmatrix} \sqrt{1-|s|} & \operatorname{sgn}(s)\sqrt{|s|} \\ \operatorname{sgn}(s)\sqrt{|s|} & \sqrt{1-|s|} \end{pmatrix}, \tag{5.2}$$

for $s \in (-1/2, 1/2)$. The sign of s allows us to model both positively and negatively correlated noise, with $s = 0$ reducing to the case of $\boldsymbol{D} = \boldsymbol{I}$. The magnitude of noise (the expected square-deviation due to fluctuations) is preserved in each process: $\sum_j \sigma_{ij}^2 = 2$ for each row i, irrespective of s. Therefore, the 'amount' of noise in the system is preserved while only the correlations vary. The associated diffusivity tensor $\boldsymbol{D}$ admits the spectral decomposition:

$$\begin{aligned}\boldsymbol{D} &= \begin{pmatrix} 1 & 2\,\mathrm{sgn}(s)\sqrt{|s|(1-|s|)} \\ 2\,\mathrm{sgn}(s)\sqrt{|s|(1-|s|)} & 1 \end{pmatrix} \\ &= \boldsymbol{Z}\begin{pmatrix} \lambda_+ & 0 \\ 0 & \lambda_- \end{pmatrix}\boldsymbol{Z}^\top, \quad \text{for} \quad \boldsymbol{Z} = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix},\end{aligned} \tag{5.3}$$

and $\lambda_\pm = 1 \pm 2\,\mathrm{sgn}(s)\sqrt{|s|(1-|s|)}$. We identify from $\boldsymbol{Z}$ that the factorised diffusion directions are independent of s. The parameter s merely fixes the diffusivity in each direction, under the constraints that $\lambda_\pm > 0$ and $\lambda_+ + \lambda_- = 2$ (Figure 3 center).

In Figure 3 we plot the optimal drifts as s ranges from -0.2 to 0.2. This range captures all of the qualitative behaviour of the optimal drift, including four bifurcations. The key observation is that there are three persistent regions of maximum drift (in fact, these regions may also be identified in Figure 2): the lower left, upper right, and upper left. See Figure 3b where these three region are distinct. As s varies through the critical bifurcation points, these regions merge and divide. At the extremes, $s \approx \pm 0.2$, the regions' boundaries align with the dominant direction of diffusion (see Figure 3 center). For example, for $s \lessapprox -0.2$, the upper regions merge to form a strip approximately parallel to the boundary of the lower left region which is in turn aligned with the $(-1, 1)$ diagonal (Figure 3a). This alignment makes the y-drift strategy v_2 almost identical to the x-drift strategy v_1 (i.e. $v_2(x, y) = v_1(y, x) \approx v_1(x, y)$). Analogously for $s \gtrapprox 0.2$ but in the transverse direction: the y-drift is approximately the complement of the x-drift (one is maximised approximately where the other is minimised, see Figure 3f). These tightly coupled strategies are highly non-trivial, arising from the intricate interplay between the pay-off, the resetting, and the correlated dynamics.

Additional control constraints allow us to capture more realistically constrained scenarios. For example, if the first process was constructed to operate based only on its own state without knowledge of the full system, then we must relax the assumption of complete observability (algebraically, $v_1(x)$ has no y-dependence). How does that influence the optimal strategy of the second process $v_2(x, y)$? Alternatively, one process might depend on the full system state but nonetheless tailor its drift to favour certain regions of system states (for example, perhaps it greedily maximises profit in some subregion of state space which it favours despite this strategy being sub-optimal with respect to the global objective). What influence does this have on the optimal control of the second process?

Surprisingly, for all the cases we explored in two dimensions, the optimal strategy of the second process was nearly unchanged regardless of how we constrained the first process. This suggests that the objective landscape is effectively separable, rendering the optimal strategy of one process robust to sub-optimal strategies of the other process. This observation could be exploited for dimension reduction.

We have hitherto considered independent resets. However, when one system component fails it may induce the failure of other components. Such cascading failures may be accounted for in the general setting by resetting multiple processes upon the failure of some components, as detailed in section 3.1.

In Figure 4, we depict the optimal drifts for various reset costs C, in analogy with

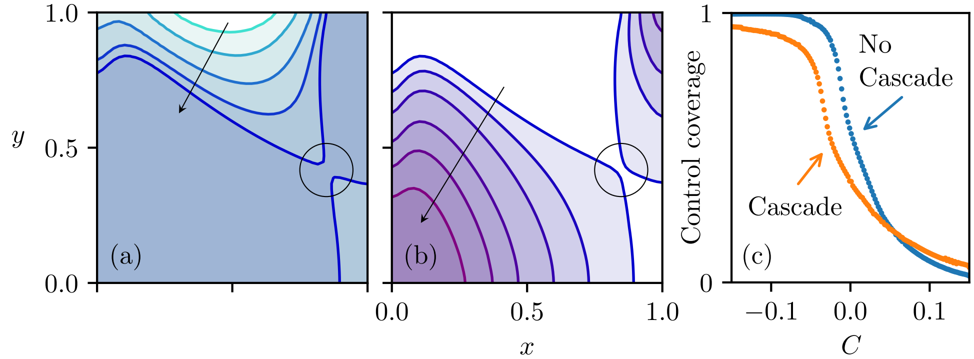

Fig. 4: (a,b) Optimal drift $v_1(x, y)$ for cascading failures and various resetting costs: (a) $C \in \{-0.1, -0.08, -0.06, -0.045, -0.0356\}$, (b) $C \in \{-0.0354, -0.02, 0, 0.03, 0.06, 0.1\}$. Maximal drift ($v_1 = U$) is shaded, minimum drift ($v_1 = u$) is unshaded. Black arrows show directions of increasing C. Both plots share the same (x, y) axes. Circles highlight the topology-changing bifurcations. Changes in colour are for visual aid but of no quantitative relevance. (c) Control coverage, the fraction of the state space on which there is maximum drift, defined in (5.4), versus reset cost C for the independent- and cascading-failure models.

Figure 2 and using the same parameters but with cascading failures. The optimal control is qualitatively similar: two regions of maximum drift (lower left and upper right) dominate the strategies, whose extents increases for decreasing resetting costs. A third region in the upper left, as in Figure 2, is recoverable in a different region of parameter space.

It is interesting to study the fraction of the state space on which there is maximum drift, which we call the control coverage. Since the optimal solutions are bang–bang and symmetric, this is captured by the quantity

$$\frac{\int_\Omega v_1(\boldsymbol{x}) - u \, \mathrm{d}\boldsymbol{x}}{(U - u)|\Omega|}. \tag{5.4}$$

When comparing the control coverage for the independent-failure model (Figure 2) and the cascading-failure model (Figure 4), one notable difference is that, for $C \lessapprox 0.06$, the control coverage is less extensive with cascading failures (Figure 4c). This aligns with the intuition that failure is more costly when it unleashes a cascade of component resets, and thus an optimal strategy might be less willing to incur reset. Curiously, above this critical reset cost $C \gtrapprox 0.06$ the phenomenon appears to invert: the control coverage associated with cascading failures is (slightly) more extensive. Despite the discrepancy being somewhat small, we conjecture that it is not a numerical artefact, but instead reflects the fact that, with cascading failures, reset sends all trajectories to the low-pay-off origin making a more extensive maximum drift in this vicinity more optimal.

These simulations provide a taste of the mosaic complexity of drift strategies that optimally steer noisy complex systems. This richness is born out of the delicate balance needed in the trade-off between high performance and system crash while accounting for the interdependence of system components, in terms of their dynamics, knock-on failure, and coupling in the objective.

6. Conclusions.

We introduce a complex system of resetting stochastic processes, seeking to optimally control the system with respect to an objective. We study the interdependence of the processes due to the objective function, as well as interaction via correlated noise and cascading failures. Formulating in terms of the PDE governing the process law, we obtain a PDE-constrained optimisation problem via the application of renewal theory and asymptotic analysis. An adjoint analysis shows that the optimal controls are of type bang–bang. Numerical simulations, also leveraging the adjoint perspective, reveal bifurcation in the solution structure as system parameters are varied.

The present formulation may be readily extended. More general reset laws (e.g. nonrectangular state spaces) could capture more sophisticated failure models. Asymmetries abound in reset costs, reset boundary, control constraints, and diffusion coefficients. Various other objective function forms may be useful in different applications.

We explore only the infinite time horizon. The large-time perspective confers an enormous advantage as it allows us to consider an elliptic problem rather than a parabolic problem, and thereby achieves dimension reduction by eliminating time, while remaining applicable over large time horizons. Moreover, the renewal approach makes analysing the reset count tractable, which would otherwise be a highly nontrivial complication. The analysis for an objective on a finite-time horizon cannot rely on the large-time limit, and therefore the formulation remains in the parabolic setting and the renewal theory is not directly applicable. Nevertheless, studying the temporal convergence towards the steady state could justify the use of the infinite-time case as a relevant model even over finite time horizons. This problem requires care in the numerical scheme to guarantee conservation and stability.

The finite-difference schemes implemented are subject to the curse of dimensionality, however, machine-learning-based approaches make high dimensional systems tractable [31], where emergent structure and long-range properties typical of complex systems can be expected.

Appendix A. Discretisation scheme.

In this appendix we detail the discretisation employed for the PDEs describing the independent-failure model (3.16) and the cascading-failure model (3.18). The dependent variable is $\widetilde{p}_0$, but throughout this section when we refer to discretised quantities, we replace this with p to simplify the notation. This is emphasised so as not to be confused with the original use of p in the section 3.

Since the domain Ω is rectangular, the problem lends itself to a consistent finite-difference discretisation. On a grid of equal spacing h_i in each dimension $i = 1, \dots,$ d, we consider the discrete domain indexed by $\boldsymbol{j} \in \Gamma \subset \mathbb{Z}^d$ where

$$\Gamma = \{(j_1, \dots, j_d)^\top \mid 0 \le j_i \le N_i \text{ for all } i = 1, \dots, d\}, \tag{A.1}$$

corresponding to the points $\boldsymbol{x_j} = (j_1 h_1, \dots, j_d h_d) \in \Omega$ where $N_i h_i \approx L_i$. We denote spatial evaluation by subscripting: $p_{\boldsymbol{j}} \approx p(\boldsymbol{x_j}) = p(j_1 h_1, \dots, j_d h_d)$. All numerical results in this work use $h_i = 0.02$ and $N_i = 51$.

It is useful to introduce the boundary sets

$$\begin{aligned} \partial\Gamma_0 &= \{\boldsymbol{j} \in \Gamma \mid j_i = 0 \text{ for some } i\}, \\ \partial\Gamma_N &= \{\boldsymbol{j} \in \Gamma \mid j_i = N_i \text{ for some } i\}, \\ \mathring{\Gamma} &= \Gamma \setminus (\partial\Gamma_N \cup \partial\Gamma_0). \end{aligned} \tag{A.2}$$

The sets $\partial\Gamma_0$ and $\partial\Gamma_N$ are the discrete counterparts of the continuum $\partial\Omega_0$ and $\partial\Omega_L$, respectively. Whereas in the continuum setting the intersection $\partial\Omega_0 \cap \partial\Omega_L$ is a null set, this is no longer the case in the discrete setting: different boundary conditions are imposed at these two boundaries (see (3.16b) and (3.16c)), thus, in the discrete setting, we need to decide which of the two boundary conditions are applied to points in $\partial\Gamma_0 \cap \partial\Gamma_N$. We choose to assign points in the intersection to the N-boundary, $\partial\Gamma_N$, by defining $\widehat{\partial\Gamma}_0 := \partial\Gamma_0 \setminus \partial\Gamma_N$, and calling $\widehat{\partial\Gamma}_0$ the 0-boundary.

Another issue concerns the appropriate boundary conditions for points on two (or more) 0-boundaries, that is, $\boldsymbol{j} \in \widehat{\partial\Gamma}_0$ such that $j_i = j_k = 0$ for some $i \neq k$. To deal with this, we distinguish between $\boldsymbol{j}$ for which a single index is zero, and $\boldsymbol{j}$ for which more than one index is zero, which we denote $\widehat{\partial\Gamma}_0^1$ and $\widehat{\partial\Gamma}_0^{>1}$, respectively: $\widehat{\partial\Gamma}_0 = \widehat{\partial\Gamma}_0^1 \cup \widehat{\partial\Gamma}_0^{>1}$. We may set the solution value in $\widehat{\partial\Gamma}_0^{>1}$ arbitrarily, as we will design a consistent scheme whose support does not enter $\widehat{\partial\Gamma}_0^{>1}$.

We implemented the schemes detailed in Refs. [27, 30, 55], and obtained the most stable convergence with the asymmetric scheme from Ref. [30], which we adopted in the interior of the domain, $\boldsymbol{j} \in \mathring{\Gamma}$, via

$$\sum_{k=1}^{d} \frac{(v_k p)_{\boldsymbol{j}} - (v_k p)_{\boldsymbol{j}-\boldsymbol{e}_k}}{h_k} - \frac{(\phi_k)_{\boldsymbol{j}+\boldsymbol{e}_k/2} - (\phi_k)_{\boldsymbol{j}-\boldsymbol{e}_k/2}}{h_k} = 0, \tag{A.3a}$$

where

$$\begin{aligned}(\phi_k)_{\boldsymbol{j}+\boldsymbol{e}_k/2} = {} & \frac{(D_{kk}p)_{\boldsymbol{j}+\boldsymbol{e}_k} - (D_{kk}p)_{\boldsymbol{j}}}{h_k} \\ & + \sum_{k\neq i=1}^{d} \frac{(D_{ki}p)_{\boldsymbol{j}+\boldsymbol{e}_i} + (D_{ki}p)_{\boldsymbol{j}+\boldsymbol{e}_k+\boldsymbol{e}_i} - (D_{ki}p)_{\boldsymbol{j}-\boldsymbol{e}_i} - (D_{ki}p)_{\boldsymbol{j}+\boldsymbol{e}_k-\boldsymbol{e}_i}}{4h_i},\end{aligned} \tag{A.3b}$$

with $\boldsymbol{e}_i$ denoting the ith basis vector (of zero entries except the ith entry of one) and denoting the diffusivity tensor $\boldsymbol{D} = (D_{ik})$ in the Cartesian frame, for all $\boldsymbol{j}$ such that the scheme is supported on $\mathring{\Omega} \setminus \widehat{\partial\Gamma}_0^{>1}$. For $\boldsymbol{j} \in \mathring{\Omega}$ where scheme (A.3) relies of evaluation in $\widehat{\partial\Gamma}_0^{>1}$, we use the scheme

$$\begin{aligned}&\sum_{k=1}^{d} \frac{(v_k p)_{\boldsymbol{j}} - (v_k p)_{\boldsymbol{j}-\boldsymbol{e}_k}}{h_k} \\ &\quad - \sum_{i=1}^{d} \frac{1}{h_k}\left[\frac{(D_{ik}p)_{\boldsymbol{j}+\boldsymbol{e}_k} - (D_{ik}p)_{\boldsymbol{j}-\boldsymbol{e}_i+\boldsymbol{e}_k}}{h_i} - \frac{(D_{ik}p)_{\boldsymbol{j}} - (D_{ik}p)_{\boldsymbol{j}-\boldsymbol{e}_i}}{h_i}\right] = 0.\end{aligned} \tag{A.4}$$

Since we only *subtract* from at most one index in the stencil (A.4), and since $j_k \geq 1$ for all $\boldsymbol{j} \in \mathring{\Gamma}$, no point in the interior includes any indices from $\widehat{\partial\Gamma}_0^{>1}$ in its stencil.

Both schemes (A.3) and (A.4) are consistent. Moreover, since we will consider non-negative drifts in the optimisation problem, the first-order terms in both schemes (A.3) and (A.4) are of upwinded form.

We now detail the discrete boundary conditions. For the N-boundaries, we set

$$p_{\boldsymbol{j}} = 0, \qquad \boldsymbol{j} \in \partial\Gamma_N. \tag{A.5}$$

For $\boldsymbol{j} \in \widehat{\partial\Gamma}^1_0$, we set

$$
\begin{aligned}
(v_k p)_{\boldsymbol{j}} &- \frac{(D_{kk}p)_{\boldsymbol{j}+\boldsymbol{e}_k} - (D_{kk}p)_{\boldsymbol{j}}}{h_k} - \sum_{k\neq i=1}^{d} \frac{(D_{ik}p)_{\boldsymbol{j}+\boldsymbol{e}_i} - (D_{ik}p)_{\boldsymbol{j}-\boldsymbol{e}_i}}{2h_i} \\
&= (v_k)_{\boldsymbol{j}^k} p_{\boldsymbol{j}^k} - \frac{(D_{kk}p)_{\boldsymbol{j}^k} - (D_{kk}p)_{\boldsymbol{j}^k-\boldsymbol{e}_k}}{h_k} - \sum_{k\neq i=1}^{d} \frac{(D_{ik}p)_{\boldsymbol{j}^k+\boldsymbol{e}_i} - (D_{ik}p)_{\boldsymbol{j}^k-\boldsymbol{e}_i}}{2h_i} \\
&= \frac{(D_{kk}p)_{\boldsymbol{j}^k-\boldsymbol{e}_k}}{h_k}, \qquad \text{(A.6)}
\end{aligned}
$$

for the single k such that $j_k = 0$, with $\boldsymbol{j}^k = (j_1, \ldots, j_{k-1}, N_k, j_{k+1}, \ldots, j_d)$ denoting the point opposite $\boldsymbol{j}$, and the last equality following from (A.5) since $\boldsymbol{j}^k \in \partial\Gamma_N$ and $\boldsymbol{j}^k \pm \boldsymbol{e}_i \in \partial\Gamma_N$ for all $i \neq k$.

At this stage, we have $\prod_{i=1}^{d}(N_i + 1)$ unknowns and equations, which we write in matrix form via

$$\boldsymbol{M}\boldsymbol{p} = \boldsymbol{0}, \qquad \text{(A.7)}$$

where $\boldsymbol{p} = (p_{\boldsymbol{j}})$ is an $\prod_{i=1}^{d}(N_i + 1)$-dimensional vector and $\boldsymbol{M}$ is the finite-difference matrix. Since the system (A.7) is linear and homogeneous, any non-trivial solution may be multiplied by any nonzero scalar to yield another solution. This indeterminacy is resolved when the constraint (3.16d) is imposed, which may be discretised via

$$\sum_{\boldsymbol{j}\in\widehat{\partial\Gamma}^1_0} \left[(v_k p)_{\boldsymbol{j}} - \frac{(D_{kk}p)_{\boldsymbol{j}+\boldsymbol{e}_k} - (D_{kk}p)_{\boldsymbol{j}}}{h_k} - \sum_{k\neq i=1}^{d} \frac{(D_{ik}p)_{\boldsymbol{j}+\boldsymbol{e}_i} - (D_{ik}p)_{\boldsymbol{j}-\boldsymbol{e}_i}}{2h_i} \right] \frac{H}{h_k} = 1, \qquad \text{(A.8)}$$

where $k = k(\boldsymbol{j})$ denotes the index for which $j_k = 0$ and

$$H = \prod_{i=1}^{d} h_i. \qquad \text{(A.9)}$$

The right-hand side of (A.8) is positive since we have taken the *inward*-facing normal flux on the left-hand side. Denoting the discrete condition (A.8) by $\boldsymbol{L}^\top \boldsymbol{p} = 1$, we impose the condition by solving the forward system

$$\begin{pmatrix} \boldsymbol{M} & \boldsymbol{L} \\ \boldsymbol{L}^\top & 0 \end{pmatrix} \begin{pmatrix} \boldsymbol{p} \\ q \end{pmatrix} = \begin{pmatrix} \boldsymbol{0} \\ 1 \end{pmatrix}, \qquad \text{(A.10)}$$

where $\boldsymbol{0}$ denotes vectors of zeros and ones, respectively, and q plays the role of a Lagrange multiplier.

The cascading failure problem (3.18) is solved using the same scheme, without the inhomogeneous condition (A.8), but instead setting the entry, corresponding to the discrete location of the delta function, equal to $\prod_{i=1}^{d} 1/h_i$ on the right-hand side of (A.7). Results using both of these schemes were compared with the original scheme in dimension one (where the independent and cascading failure problems coincide, but the schemes differ), and both agreed with previous results from [41].

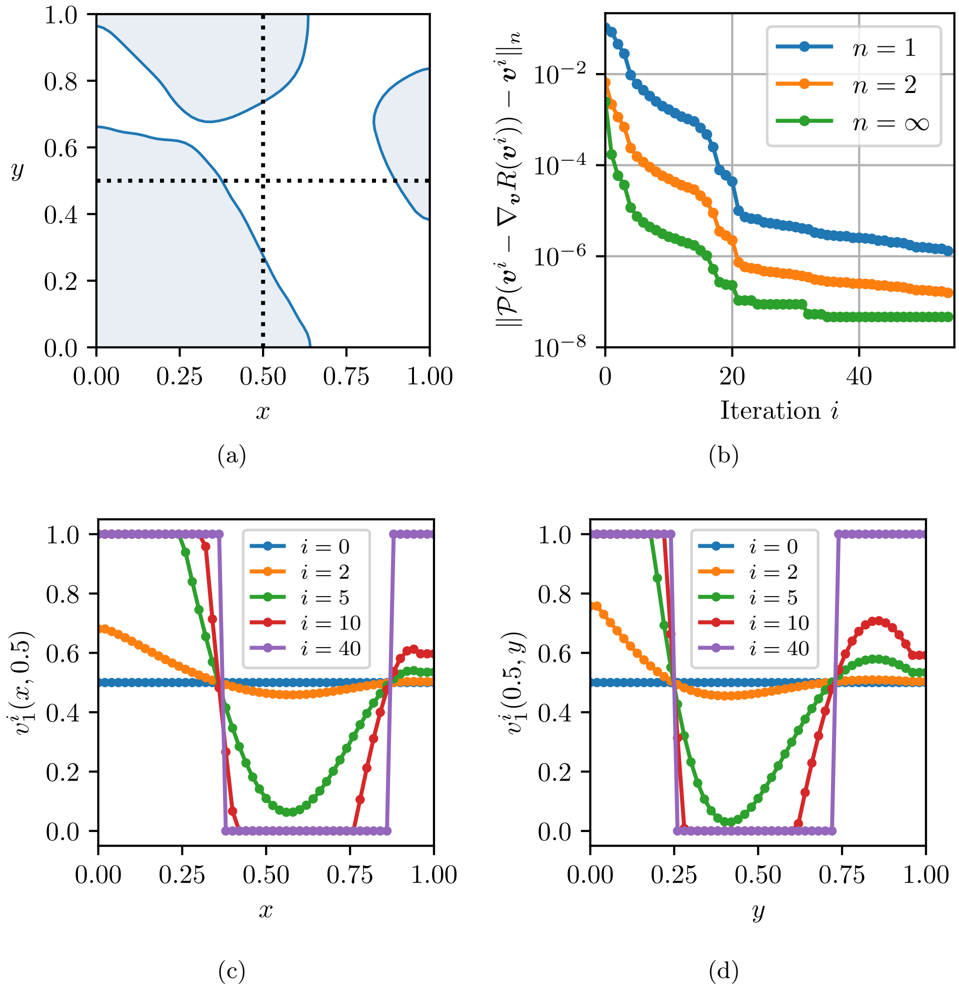

Fig. 5: Results of the numerical optimisation scheme for the independent-failure model and correlated noise (5.2) with $s = -0.11$. At iteration i, the controllable drift is denoted by $v^i(x, y)$ (and its discrete counterpart by $\boldsymbol{v}_i$). (a) Optimal drift $v_1^{55}(x, y)$ for the final iterate $i = 55$, where the maximum drift ($v_1^{55} = U$) is shaded and minimum drift ($v_1^{55} = u$) is unshaded. The dotted black lines show slices of the (x, y)-domain on which we plot the drifts v_1^i in panels (c) and (d). (b) Convergence of the ℓ^n norm of the projected gradient. (c,d) The drift $v_1^i(x, y)$ on the slices depicted in panel (a) at different iterations i of the optimisation algorithm.

This completes the discretisation of the governing PDEs (3.16) and (3.18).

Appendix B. Numerical convergence.

In this appendix, we demonstrate the numerical convergence achieved by the L-BFGS-B algorithm based on the adjoint approach described in section 4.1 and the discrete scheme detailed in Appendix A. We choose the case of independent failures (maximising R in (3.15) subject to (3.16)) and correlated noise, corresponding to the results depicted in Figure 3.

We start with the drift $\boldsymbol{v}^i$ after $i = 55$ iterations in Figure 5a. Our aim is to demonstrate that the objective evaluated with this drift, $R(\boldsymbol{v}^i)$, is (approximately) a local maximum, and thus this drift is (approximately) optimal. Unconstrained first-order optimality conditions require that the gradient vanishes at a local optimum. In our case, the control drift is bounded $\boldsymbol{v}(\boldsymbol{x}) \in [u, U]^d$, in which case the projected gradient plays the role of the gradient as we now describe.

The projection $\mathcal{P}$ maps each component of the drift vector $\boldsymbol{v}^i$ onto the bounds $[u, U]$, that is, for each component z in $\boldsymbol{v}^i$

$$z \overset{\mathcal{P}}{\mapsto} \begin{cases} 0, & z < u, \\ z, & z \in (u, U), \\ 1, & z > U. \end{cases} \tag{B.1}$$

The projected gradient of $R(\boldsymbol{v}^i)$ is given by the difference between the projection of the control perturbed by the negative gradient, $\mathcal{P}(\boldsymbol{v}^i - \nabla_{\boldsymbol{v}} R(\boldsymbol{v}^i))$, and the control, $\boldsymbol{v}^i$. To see why the first-order optimality conditions are simply the vanishing projected gradient, consider a local optimum at $\boldsymbol{v}^*$ and separate those components strictly within the bounds from those at the bounds. The gradient $\nabla_{\boldsymbol{v}} R(\boldsymbol{v}^*)$ must vanish in the components where $\boldsymbol{v}^*$ is within the bounds (otherwise the objective could be increased by a small admissible perturbation in the direction of the negative gradient, thus the point is not an optimum):

$$v_k^* \in (u, U) \implies -[\nabla_{\boldsymbol{v}} R(\boldsymbol{v}^*)]_k = 0. \tag{B.2}$$

Therefore, these components must also vanish in the projected gradient. For components of $\boldsymbol{v}^*$ at a bound, $\boldsymbol{v}_k^* \in \{u, U\}$, the negative gradient does not point inward to the interval $[u, U]$ (otherwise the objective could be increased by a small admissible perturbation):

$$v_k^* = u \implies -[\nabla_{\boldsymbol{v}} R(\boldsymbol{v}^*)]_k \le 0, \qquad v_k^* = U \implies -[\nabla_{\boldsymbol{v}} R(\boldsymbol{v}^*)]_k \ge 0. \tag{B.3}$$

Therefore, the projection of the perturbed control remains at the same boundary for these components, thus vanishing in the projected gradient. Ultimately, we find that

$$\mathcal{P}(\boldsymbol{v}^* - \nabla_{\boldsymbol{v}} R(\boldsymbol{v}^*)) - \boldsymbol{v}^* = \boldsymbol{0}, \tag{B.4}$$

as we set out to demonstrate.

In Figure 5b we show the ℓ^n norm of the projected gradient at each iteration. There projected gradient rapidly decays until iteration $i \approx 20$ after which time the convergence is slower. In Figures 5c and 5d we show the control $v_1^i(x, y)$, that is, the control of the first process at iteration i, along one-dimensional slices of the domain corresponding to the black dotted lines in Figure 5a. The rapid convergence is evident: by iteration $i = 10$ the drift closely matches its final configuration. The notable exceptions are regions near the L-boundary, where convergence is slower due to the relative insensitivity of the objective function to the drift. Nonetheless, from the very earliest iterations we see that the $(u + U)/2$ level set is clearly established and preserved, motivating its choice as the contour delineating regions of maximum and minimum optimal drift, as discussed in Appendix C.

In this example, despite an intricate optimal drift topology close to a bifurcation and strong anisotropy in the diffusion operator, we observe good convergence to a

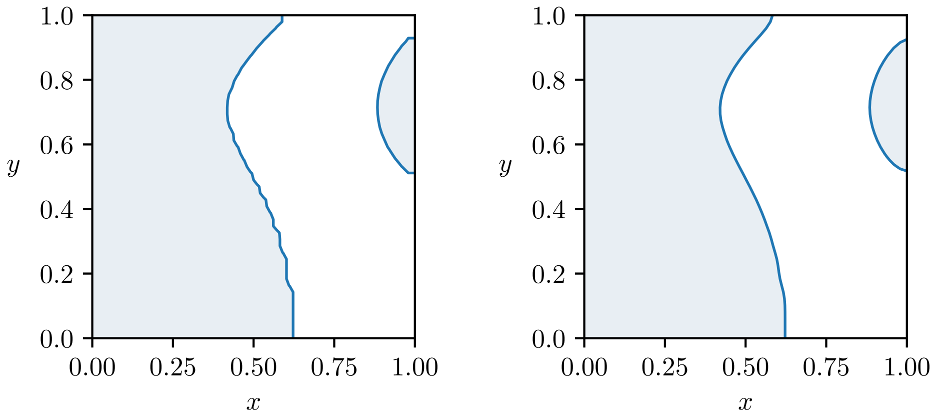

(a) Not smoothed. (b) Smoothed.

Fig. 6: Optimal drift strategy $v_1(x, y)$ (shaded regions show maximum drift, $v_1 = U$, unshaded regions show minimum drift, $v_1 = u$) for defaults parameters given in (5.1) with and without smoothing.

local optimum. The convergence criterion used to produce all the other figures was a threshold of 2×10^{-6} for the ℓ^∞-norm of the projected gradient. The initial control guess used in the optimisation was the uniform $\boldsymbol{v}^0(\boldsymbol{x}) \equiv (u + U)/2$. All calculations were performed on a standard laptop computer with simulation times on the order of seconds per optimal drift.

Appendix C. Contour plots.

As detailed in section 4.2, the optimal drifts are of bang–bang type, taking either their upper bound or lower bound but not intermediate values. Therefore, to plot an optimal drift field $v_i(\boldsymbol{x})$ it suffices to draw the region(s) of maximum and minimum drift. We choose to draw the contour(s) of the $(u + U)/2$ level set, obtained by interpolation of the drift field $v_i(\boldsymbol{x})$, which we consider to separate these two regions and shade only the region of maximum drift. Since the drift field is given only at grid points, the resulting contour tends to be unnaturally rectangular (Figure 6a). We thus choose to smooth the drift field with a Gaussian filter, resulting in a smoother contour (Figure 6b). In Figures 2 to 4 we use various colours to plot the contours and shade the regions of maximal drift. These are merely for visual aid to show the continuity of the optimal drift with respect to the changing parameter values away from a bifurcation, but have no quantitative relevance.

REFERENCES

[1] R. Albert, H. Jeong, and A.-L. Barabási, *Error and attack tolerance of complex networks*, Nature, 406 (2000), pp. 378–382, https://doi.org/10.1038/35019019.

[2] S. V. Anulova, H. Mai, and A. Y. Veretennikov, *On iteration improvement for averaged expected cost control for one-dimensional ergodic diffusions*, SIAM J. Control Optim., 58 (2020), pp. 2312–2331, https://doi.org/10.1137/19M1271944.

[3] A. Arapostathis, V. S. Borkar, and M. K. Ghosh, *Ergodic Control of Diffusion Processes*, Encyclopedia of Mathematics and its Applications, Cambridge University Press, 2011, https://doi.org/10.1017/CBO9781139003605.

[4] R. R. Barton, *Computing forward difference derivatives in engineering optimization*, Eng. Optimiz., 20 (1992), pp. 205–224, https://doi.org/10.1080/03052159208941281.
[5] R. Bartoszynski, W. J. Bühler, W. Chan, and D. K. Pearl, *Population processes under the influence of disasters occurring independently of population size*, J. Math. Biol., 27 (1989), pp. 167–178, https://doi.org/10.1007/BF00276101.
[6] I. Ben-Ari and R. G. Pinsky, *Ergodic behavior of diffusions with random jumps from the boundary*, Stoch. Proc. Appl., 119 (2009), pp. 864–881, https://doi.org/10.1016/j.spa.2008.05.002.
[7] O. Bénichou, M. Moreau, P.-H. Suet, and R. Voituriez, *Intermittent search process and teleportation*, J. Chem. Phys., 126 (2007), p. 234109, https://doi.org/10.1063/1.2741516.
[8] S. Boccaletti, V. Latora, Y. Moreno, M. Chavez, and D.-U. Hwang, *Complex networks: Structure and dynamics*, Phys. Rep., 424 (2006), pp. 175–308, https://doi.org/10.1016/j.physrep.2005.10.009.
[9] P. J. Brockwell, J. Gani, and S. I. Resnick, *Birth, immigration and catastrophe processes*, Adv. Appl. Probab., 14 (1982), pp. 709–731, https://doi.org/10.2307/1427020.
[10] S. V. Buldyrev, R. Parshani, G. Paul, H. E. Stanley, and S. Havlin, *Catastrophic cascade of failures in interdependent networks*, Nature, 464 (2010), pp. 1025–1028, https://doi.org/10.1038/nature08932.
[11] X. Chao and Y. Zheng, *Transient analysis of immigration birth–death processes with total catastrophes*, Probab. Eng. Inform. Sc., 17 (2003), pp. 83—-106, https://doi.org/10.1017/S0269964803171057.
[12] B. De Bruyne, J. Randon-Furling, and S. Redner, *Optimization in first-passage resetting*, Phys. Rev. Lett., 125 (2020), p. 050602, https://doi.org/10.1103/PhysRevLett.125.050602.
[13] B. De Bruyne, J. Randon-Furling, and S. Redner, *Optimization and growth in first-passage resetting*, J. Stat. Mech.-Theory E., 2021 (2021), p. 013203, https://doi.org/10.1088/1742-5468/abcd33.
[14] I. Dobson, B. A. Carreras, V. E. Lynch, and D. E. Newman, *Complex systems analysis of series of blackouts: Cascading failure, critical points, and self-organization*, Chaos, 17 (2007), p. 026103, https://doi.org/10.1063/1.2737822.
[15] S. N. Dorogovtsev and J. F. F. Mendes, *Evolution of networks*, Adv. Phys., 51 (2002), pp. 1079–1187, https://doi.org/10.1080/00018730110112519.
[16] I. Eliazar, T. Koren, and J. Klafter, *Searching circular DNA strands*, J. Phys.-Condens. Mat., 19 (2007), p. 065140, https://doi.org/10.1088/0953-8984/19/6/065140.
[17] M. R. Evans and S. N. Majumdar, *Diffusion with optimal resetting*, J. Phys. A-Math. Theor., 44 (2011), p. 435001, https://doi.org/10.1088/1751-8113/44/43/435001.
[18] M. R. Evans and S. N. Majumdar, *Diffusion with stochastic resetting*, Phys. Rev. Lett., 106 (2011), p. 160601, https://doi.org/10.1103/PhysRevLett.106.160601.
[19] M. R. Evans and S. N. Majumdar, *Diffusion with resetting in arbitrary spatial dimension*, J. Phys. A-Math. Theor., 47 (2014), p. 285001, https://doi.org/10.1088/1751-8113/47/28/285001.
[20] M. R. Evans and S. N. Majumdar, *Run and tumble particle under resetting: a renewal approach*, J. Phys. A-Math. Theor., 51 (2018), p. 475003, https://doi.org/10.1088/1751-8121/aae74e.
[21] M. R. Evans, S. N. Majumdar, and G. Schehr, *Stochastic resetting and applications*, J. Phys. A-Math. Theor., 53 (2020), p. 193001, https://doi.org/10.1088/1751-8121/ab7cfe.
[22] E. Gelenbe, *Search in unknown random environments*, Phys. Rev. E, 82 (2010), p. 061112, https://doi.org/10.1103/PhysRevE.82.061112.
[23] M. B. Giles and N. A. Pierce, *An introduction to the adjoint approach to design*, Flow Turbul. Combust., 65 (2000), pp. 393–415, https://doi.org/10.1023/A:1011430410075.
[24] M. Girvan and M. E. J. Newman, *Community structure in social and biological networks*, P. Natl. Acad. Sci. USA, 99 (2002), pp. 7821–7826, https://doi.org/10.1073/pnas.122653799.
[25] C. P. Gomes and B. Selman, *Algorithm portfolios*, Artif. Intell., 126 (2001), pp. 43–62, https://doi.org/10.1016/S0004-3702(00)00081-3.
[26] C. P. Gomes, B. Selman, N. Crato, and H. Kautz, *Heavy-tailed phenomena in satisfiability and constraint satisfaction problems*, J. Autom. Reasoning, 24 (2000), pp. 67–100, https://doi.org/10.1023/A:1006314320276.
[27] D. Greenspan and P. C. Jain, *On non negative difference analogues of elliptic differential equations*, J. Frankl. Inst., 279 (1965), pp. 360–365, https://doi.org/10.1016/0016-0032(65)90423-0.
[28] I. Grigorescu and M. Kang, *Brownian motion on the figure eight*, J. Theor. Probab., 15 (2002), pp. 817–844, https://doi.org/10.1023/A:1016232201962.
[29] I. Grigorescu and M. Kang, *Ergodic properties of multidimensional brownian motion with*

rebirth, Electron. J. Probab., 12 (2007), pp. 1299–1322, https://doi.org/10.1214/EJP.v12-450.

[30] S. Günter, Q. Yu, J. Krüger, and K. Lackner, *Modelling of heat transport in magnetised plasmas using non-aligned coordinates*, J. Comput. Phys., 209 (2005), pp. 354–370, https://doi.org/10.1016/j.jcp.2005.03.021.

[31] J. Han, A. Jentzen, and E. Weinan, *Solving high-dimensional partial differential equations using deep learning*, P. Natl. Acad. Sci. USA, 115 (2018), pp. 8505–8510, https://doi.org/10.1073/pnas.1718942115.

[32] P. Holme, B. J. Kim, C. N. Yoon, and S. K. Han, *Attack vulnerability of complex networks*, Phys. Rev. E, 65 (2002), p. 056109, https://doi.org/10.1103/PhysRevE.65.056109.

[33] B. Krishna Kumar and D. Arivudainambi, *Transient solution of an M/M/1 queue with catastrophes*, Comput. Math. Appl., 40 (2000), pp. 1233–1240, https://doi.org/10.1016/S0898-1221(00)00234-0.

[34] B. Krishna Kumar and S. Pavai Madheswari, *Transient behaviour of the M/M/2 queue with catastrophes*, Statistica, 62 (2002), pp. 129–136, https://doi.org/10.6092/issn.1973-2201/1287.

[35] V. Kumar, O. Sadekar, and U. Basu, *Active brownian motion in two dimensions under stochastic resetting*, Phys. Rev. E, 102 (2020), p. 052129, https://doi.org/10.1103/PhysRevE.102.052129.

[36] L. Kusmierz, S. N. Majumdar, S. Sabhapandit, and G. Schehr, *First order transition for the optimal search time of Lévy flights with resetting*, Phys. Rev. Lett., 113 (2014), p. 220602, https://doi.org/10.1103/PhysRevLett.113.220602.

[37] E. G. Kyriakidis, *Stationary probabilities for a simple immigration-birth-death process under the influence of total catastrophes*, Stat. Probabil. Lett., 20 (1994), pp. 239–240, https://doi.org/10.1016/0167-7152(94)90048-5.

[38] E. G. Kyriakidis and A. Abakuks, *Optimal pest control through catastrophes*, J. Appl. Probab., 26 (1989), pp. 873—879, https://doi.org/10.2307/3214392.

[39] Y. J. Leung, W. V. Li, and Rakesh, *Spectral analysis of brownian motion with jump boundary*, P. Am. Math. Soc., 136 (2008), pp. 4427–4436.

[40] M. Luby, A. Sinclair, and D. Zuckerman, *Optimal speedup of Las Vegas algorithms*, Inform. Process. Lett., 47 (1993), pp. 173–180, https://doi.org/10.1016/0020-0190(93)90029-9.

[41] D. Lunz, *Optimal feedback control in first-passage resetting*, J. Phys. A-Math. Theor., 53 (2020), p. 44LT01, https://doi.org/10.1088/1751-8121/abbc7c.

[42] S. C. Manrubia and D. H. Zanette, *Stochastic multiplicative processes with reset events*, Phys. Rev. E, 59 (1999), pp. 4945–4948, https://doi.org/10.1103/PhysRevE.59.4945.

[43] M. Montero, A. Masó-Puigdellosas, and J. Villarroel, *Continuous-time random walks with reset events*, Eur. Phys. J. B, 90 (2017), p. 176, https://doi.org/10.1140/epjb/e2017-80348-4.

[44] J. L. Morales and J. Nocedal, *Remark on "Algorithm 778: L-BFGS-B: Fortran subroutines for large-scale bound constrained optimization"*, ACM T. Math. Software, 38 (2011), pp. 1–4, https://doi.org/10.1145/2049662.2049669.

[45] Y. Moreno, J. B. Gómez, and A. F. Pacheco, *Instability of scale-free networks under node-breaking avalanches*, Europhys. Lett., 58 (2002), p. 630, https://doi.org/10.1209/epl/i2002-00442-2.

[46] Y. Moreno, R. Pastor-Satorras, A. Vázquez, and A. Vespignani, *Critical load and congestion instabilities in scale-free networks*, Europhys. Lett., 62 (2003), p. 292, https://doi.org/10.1209/epl/i2003-00140-7.

[47] A. E. Motter, *Cascade control and defense in complex networks*, Phys. Rev. Lett., 93 (2004), p. 098701, https://doi.org/10.1103/PhysRevLett.93.098701.

[48] A. E. Motter and Y.-C. Lai, *Cascade-based attacks on complex networks*, Phys. Rev. E, 66 (2002), p. 065102, https://doi.org/10.1103/PhysRevE.66.065102.

[49] M. E. J. Newman, *The structure and function of complex networks*, SIAM Rev., 45 (2003), pp. 167–256, https://doi.org/10.1137/S003614450342480.

[50] E. Roldán, A. Lisica, D. Sánchez-Taltavull, and S. W. Grill, *Stochastic resetting in backtrack recovery by RNA polymerases*, Phys. Rev. E, 93 (2016), p. 062411, https://doi.org/10.1103/PhysRevE.93.062411.

[51] M. L. Sachtjen, B. A. Carreras, and V. E. Lynch, *Disturbances in a power transmission system*, Phys. Rev. E, 61 (2000), pp. 4877–4882, https://doi.org/10.1103/PhysRevE.61.4877.

[52] I. Santra, U. Basu, and S. Sabhapandit, *Run-and-tumble particles in two dimensions under stochastic resetting conditions*, J. Stat. Mech.-Theory E., 2020 (2020), p. 113206, https://doi.org/10.1088/1742-5468/abc7b7.

[53] S. H. Strogatz, *Exploring complex networks*, Nature, 410 (2001), pp. 268–276, https://doi.org/10.1038/35065725.

[54] S. Thurner, R. Hanel, and P. Klimek, *Introduction to the theory of complex systems*, Oxford University Press, 2018, https://doi.org/10.1093/oso/9780198821939.001.0001.

[55] B. van Es, B. Koren, and H. J. de Blank, *Finite-difference methods for extremely anisotropic diffusion*, in Advanced numerical methods for complex environmental models: needs and availability, I. Faragó, Á. Havasi, and Z. Zlatev, eds., Bentham Science Publishers, 2013, pp. 29–52, https://doi.org/10.2174/9781608057788113010006.