Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Endogenous Resets in Interacting Birth-Death Processes: A Graphon Mean-Field Framework

Sindhu Chaudhary

sindhuchaudhary55@gmail.com

Ghent University

Arnaud Devos

Ghent University

Dieter Fiems

Ghent University

Koen De Turck

Ghent University

Research Article

Keywords:

Posted Date: November 20th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11788934/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Endogenous Resets in Interacting Birth–Death Processes: A Graphon Mean-Field Framework

Sindhu Chaudhary[1*], Arnaud Devos[1], Dieter Fiems[1], Koen De Turck[1]

[1*]Department of Telecommunication and Information Processing, Ghent University, Ghent, 9000, Belgium.

*Corresponding author(s). E-mail(s): sindhuchaudhary55@gmail.com;
Contributing authors: Arnaud.Devos@UGent.be;
Dieter.Fiems@ugent.be; Koen.DeTurck@ugent.be;

Abstract

We study interacting birth–death processes in which an agent at zero sends reset events to other agents through a directed network. These endogenous catastrophes can stabilize coordinates that are supercritical (i.e. unstable) in isolation. For finite systems, we prove positive recurrence when every source strongly connected component of the reset graph contains a subcritical agent. For the case of two agents, we derive the stationary functional equation and a kernel construction of its boundary generating functions. For dense heterogeneous systems, we prove convergence on fixed finite time intervals to a type-dependent nonlinear process driven by a graphon reset kernel. We also characterize stationary solutions of the limiting equations by a fixed-point problem and prove existence and moment bounds under quantitative influence from subcritical types. We illustrate the mathematical results by several experiments.

Keywords: Interacting birth–death processes, Endogenous catastrophes, Graphon mean-field limit, Network-induced stabilization

1 Introduction

In many population systems, a component that collapses forces another back to the baseline state as well. A natural model for this is a collection of such birth–death processes, or agents, coupled by a reset mechanism. When an agent is at zero, it

sends reset events that instantaneously move selected receiving agents to zero. The event rate therefore depends on the current boundary occupancy of the senders. This state-dependent feedback is the central feature of the model: recurrent senders can repeatedly clear receivers that would drift to infinity without interaction.

One motivation comes from social trust and belief dynamics, a recurring theme in the statistical physics of social systems [1]. The state can represent a discrete stock of trust, with upward and downward jumps recording favorable and unfavorable experiences, while a reset represents a crisis that destroys accumulated trust. An agent at zero then models an individual or group whose loss of trust increases the risk of abrupt trust loss elsewhere. In previous work, Chaudhary et al. studied a homogeneous mean-field version of this idea using infinite-server queues with breakdowns. The breakdown rate depended on the population fraction at zero, leading to nonlinear stationary regimes and metastable switching [2]. The present model changes the local dynamics to $M/M/1$-type birth–death processes and resolves who affects whom through a directed reset graph or graphon kernel. It thereby permits heterogeneous rates, pairwise sender–receiver asymmetry, and finite-network stability questions that are absent from the homogeneous aggregate model.

The same mathematics also admits queueing and reliability interpretations. A reset may clear accumulated work in a service agent, or a failed agent may trigger a protective shutdown or restart elsewhere. In a social or reliability interpretation a reset is adverse, whereas in a workload interpretation it is stabilizing. Probabilistically, both are jumps to the boundary generated by the boundary occupancy of another agent. Other possible settings include cascading reputational shocks [3], defaults, or failures that propagate between interdependent systems [4]. Our results concern the common stochastic mechanism rather than any specific application.

1.1 Relation to the literature

The idea of randomly resetting a process to zero has appeared in several fields. It appears as catastrophes in birth–death and population processes [5] and as catastrophes, disasters and clearing in queues [6–11]. In statistical physics it is known as stochastic resetting [12–14]. In these works the reset acts on a single process and is driven either by an external clock or by the state of the agent being cleared. The closest setting with interacting components is that of G-networks [15, 16], where negative customers generated at one node remove work at another, and the stationary distribution often has a product form. In our model a receiving agent is cleared completely, and the reset rate is determined by how often its senders are at zero. The resets are therefore driven by the state of other agents, and product form is generally lost.

The finite-population model is also related to interacting queueing systems. Classical supermarket models demonstrate how weak interaction changes queue length distributions and gives rise to propagation of chaos [17, 18], also when servers interact only along the edges of a graph [19]. General mean-field methods for interacting Markov processes and communication systems are developed in [20, 21], but the interaction is of a different nature (through routing or selection of the shortest queue, etc.) than the boundary-generated jumps to zero in our model.

Our dense-network formulation belongs to the literature on dynamics over graph limits [22]. Graphons have been used for deterministic interacting equations and for stochastic particle systems on dense and inhomogeneous networks [23–27]. The directed kernel in this paper is not assumed symmetric: its first argument is the sender type and its second the receiver type. The proof must therefore retain the type coordinate and handle a reset drift that is quadratic in the empirical measure.

Finally, the two-agent calculation uses the kernel and boundary-value methods for random walks in the quarter plane developed in [28–30]. The closest classical model is the coupled-processor queue of [28], in which the service rate of one queue changes when the other queue is empty, in our model an empty agent instead resets the other one. Fourier elimination along a kernel contour is close in spirit to the processor-sharing calculation of [31], but endogenous resets introduce unknown boundary masses which can be solved by an additional system of equations.

1.2 Contributions and scope

The contributions are as follows.

(i) For finite N, we prove non-explosion and irreducibility and give a graph-theoretic sufficient condition for positive recurrence: every source strongly connected component contains a coordinate with $\lambda_i < \mu_i$. For two agents without critical coordinates, we show that this condition is also necessary.

(ii) For two agents, we derive the stationary functional equation and kernel integral identities and linear consistency systems for its boundary generating functions in the noncritical construction regime of Section 3.4. These give direct reconstruction when the constants system is nonsingular.

(iii) For dense graphon-kernel systems, we prove a fixed-finite-time mean-field limit and propagation of chaos under stated initialization hypotheses.

(iv) We characterize stationary solutions of the limiting nonlinear equations by an admissible fixed-point problem. Quantitative direct influence from subcritical types yields existence and positive reset pressure; a uniform version gives finite aggregate first moments.

(v) We show that the limits $N \to \infty$ and $t \to \infty$ cannot in general be interchanged: in the uniform supercritical model, the limiting equations have a positive stationary fixed point although every finite-N system is transient.

General uniqueness and attraction of stationary fixed points, and convergence of finite-N stationary laws, remain open.

1.3 Organization

Section 2 defines the finite model and states its stability result. Section 3 treats the two-agent chain. Section 4 introduces the dense graphon system, and Section 5 proves its finite-time mean-field limit. Section 6 studies stationary fixed points and moment bounds. Section 7 presents numerical illustrations, and Section 8 discusses limitations and open problems.

2 Finite-Agent Model

2.1 N-Agent Dynamics

We consider a system of N interacting agents indexed by $i \in \{1, \dots, N\}$. The state of agent i at time t is denoted by $X_i(t) \in \mathbb{Z}_{\geq 0}$ and evolves as a continuous-time Markov process. In the absence of interactions, each agent follows $M/M/1$-type birth–death dynamics with arrival rate $\lambda_i > 0$ and service rate $\mu_i > 0$. Specifically, when $X_i(t) = k \geq 1$, the state increases to $k+1$ at rate λ_i and decreases to $k-1$ at rate μ_i. When $X_i(t) = 0$, arrivals occur at rate λ_i, and there are no departures.

Interactions between agents are mediated through reset events. When an agent reaches the boundary state 0, it generates reset pressure that may instantaneously pull other agents to the boundary. To encode the reset interaction structure, we fix a nonnegative $N \times N$ matrix $\Gamma^{(N)} = (\Gamma^{(N)}_{ij})_{1 \leq i,j \leq N}$, where $\Gamma^{(N)}_{ij}$ denotes the reset intensity exerted by agent i on agent j when agent i is at the boundary state.

Throughout the paper, the first index of a reset kernel denotes the sender and the second the receiver, so the total reset rate experienced by an agent involves the transposed entries of $\Gamma^{(N)}$.

The resulting process $\boldsymbol{X}(t) = (X_1(t), \dots, X_N(t))$ is a continuous-time Markov chain on $\mathbb{Z}^N_{\geq 0}$. Its generator $\mathcal{L}^{(N)}$ acts on bounded test functions f as follows:

$$(\mathcal{L}^{(N)} f)(\boldsymbol{x}) = \sum_{i=1}^{N} \lambda_i \big[f(\boldsymbol{x} + \boldsymbol{e}_i) - f(\boldsymbol{x}) \big] + \sum_{i=1}^{N} \mu_i \mathbf{1}\{x_i > 0\} \big[f(\boldsymbol{x} - \boldsymbol{e}_i) - f(\boldsymbol{x}) \big] + \sum_{i=1}^{N} r_i^{(N)}(\boldsymbol{x}) \big[f(\boldsymbol{x} - x_i \boldsymbol{e}_i) - f(\boldsymbol{x}) \big].$$

Here, $\boldsymbol{e}_i$ denotes the ith N-dimensional unit vector and as such state $\boldsymbol{x} - x_i \boldsymbol{e}_i$ is the configuration obtained by resetting the ith coordinate to zero. Moreover, $r_j^{(N)}(\boldsymbol{x})$ denotes the total reset rate experienced by agent j in system state $\boldsymbol{X}(t) = \boldsymbol{x}$,

$$r_j^{(N)}(\boldsymbol{x}) = \sum_{i=1}^{N} \Gamma^{(N)}_{ij} \mathbf{1}\{x_i = 0\}. \tag{1}$$

Here $\mathbf{1}\{\cdot\}$ is the indicator function, which evaluates to 1 if its argument is true and to 0 otherwise. Each such reset event instantaneously resets $X_j(t)$ to 0, regardless of its current value.

2.2 Stability Mechanism

Reset rates depend on the current global configuration of the system, since $r_j^{(N)}$ depends on the states of all senders. The agents are therefore dependent even though their birth–death dynamics are independent. A key role in the stability of the system is played by the *criticality* of agents. An agent j is called subcritical iff $\lambda_j < \mu_j$, supercritical iff $\lambda_j > \mu_j$, and critical when they are equal.

In isolation, a subcritical agent is a positive recurrent birth–death process and returns to the boundary state 0 infinitely often. In the interacting system, additional reset events can only decrease the coordinate, and hence cannot make such a coordinate less stable. Subcritical agents therefore provide recurrent opportunities to generate reset events. The stability condition below requires that every part of the reset graph which receives no reset influence from outside contains at least one such source of boundary visits.

To state a sufficient condition for positive recurrence, we encode the reset interaction structure as a directed graph, writing $j \to k$ if $\Gamma_{jk}^{(N)} > 0$, i.e., if agent j can reset agent k when $X_j(t) = 0$. We call a set $A \subset \{1, \dots, N\}$ *externally reset-isolated* if there are no $j \notin A$ and $k \in A$ with $j \to k$, that is, if no agent outside A can reset any agent in A. Such a set is immune to reset pressure from the rest of the system, although reset interactions may still occur inside it.

Proposition 1 (Finite-N sufficient stability condition) *Assume that $\lambda_i > 0$ and $\mu_i > 0$ for all $i = 1, \dots, N$. Let G be the directed reset graph on $\{1, \dots, N\}$, with an edge $j \to k$ whenever $\Gamma_{jk}^{(N)} > 0$. Suppose that every nonempty externally reset-isolated set $A \subseteq \{1, \dots, N\}$ contains at least one subcritical agent:*

$$\exists j \in A \quad \text{such that} \quad \lambda_j < \mu_j.$$

Equivalently, every source strongly connected component of the reset graph contains at least one subcritical agent. Then the Markov chain $X(t)$ is positive recurrent.

We prove this proposition in Appendix A. In words, every source component of the reset graph must contain an agent that returns to the boundary often enough to start reset cascades, which then reach every other component. We do not claim that the condition is necessary for arbitrary finite reset graphs. The two-agent model of the next section shows the mechanism in its simplest form.

3 The Two-Agent Endogenous Reset Model

This section studies the smallest nontrivial instance of the interacting birth–death model with endogenous resets. Despite its simplicity, the two-agent system exhibits the key stabilization mechanism that later reappears in the mean-field and graphon limits.

3.1 Model definition

Let $X_i(t) \in \mathbb{Z}_{\geq 0}$ denote the level of agent i at time t, $i = 1, 2$. $X_i(t)$ increases by one at rate λ_i, and decreases by one at rate μ_i whenever $X_i(t) > 0$. Whenever $X_i(t) = 0$, agent i generates reset events. For the two-agent model we write $\gamma_1 := \Gamma_{12}$ and $\gamma_2 := \Gamma_{21}$, so that γ_i is the rate at which agent i, when at zero, resets the other agent. The process

$(X_1(t), X_2(t))$ is a continuous-time Markov chain on $\mathbb{Z}^2_{\geq 0}$ with transition rates

$$(X_1, X_2) \to \begin{cases} (X_1 + 1, X_2) & \text{rate } \lambda_1, \\ (X_1, X_2 + 1) & \text{rate } \lambda_2, \\ (X_1 - 1, X_2) & \text{rate } \mu_1 \mathbf{1}\{X_1 > 0\}, \\ (X_1, X_2 - 1) & \text{rate } \mu_2 \mathbf{1}\{X_2 > 0\}, \\ (0, 0) & \text{rate } \gamma_1 \mathbf{1}\{X_1 = 0\} + \gamma_2 \mathbf{1}\{X_2 = 0\}. \end{cases}$$

This model is a direct specialization of the finite-N dynamics introduced in Section 2.

3.2 Irreducibility and stability

Proposition 2 (Irreducibility) *If $\lambda_i > 0$ and $\mu_i > 0$ for $i = 1, 2$, the Markov chain $(X_1(t), X_2(t))$ is irreducible on $\mathbb{Z}^2_{\geq 0}$.*

Proposition 3 (Two-agent stability) *The two-agent process is positive recurrent if each agent is subcritical or is reset by a subcritical agent. For agent 1 this means that $\lambda_1 < \mu_1$, or that $\gamma_2 > 0$ and $\lambda_2 < \mu_2$. The condition for agent 2 is symmetric.*

If, in addition, $\lambda_i \neq \mu_i$ for $i = 1, 2$, then this condition is also necessary. Case by case, the condition reads:

$$\begin{aligned} \gamma_1 = \gamma_2 = 0 : &\quad \lambda_1 < \mu_1 \text{ and } \lambda_2 < \mu_2, \\ \gamma_1 > 0,\ \gamma_2 = 0 : &\quad \lambda_1 < \mu_1, \\ \gamma_1 = 0,\ \gamma_2 > 0 : &\quad \lambda_2 < \mu_2, \\ \gamma_1 > 0,\ \gamma_2 > 0 : &\quad \lambda_1 < \mu_1 \text{ or } \lambda_2 < \mu_2. \end{aligned}$$

Remark 1 The condition separates the role of the reset rates from the birth–death drifts. If $\gamma_1 > 0$ and $\gamma_2 = 0$, then agent 1 is never reset, so stability requires agent 1 to be subcritical, while agent 2 may be subcritical, critical, or supercritical. The case $\gamma_1 = 0$, $\gamma_2 > 0$ is symmetric. If $\gamma_1 > 0$ and $\gamma_2 > 0$, one subcritical agent is sufficient to stabilize the pair. The case of a critical agent that is not reset by a subcritical agent is left open.

Proof For two agents, the condition says exactly that every source component of the reset graph, with an edge $1 \to 2$ if $\gamma_1 > 0$ and $2 \to 1$ if $\gamma_2 > 0$, contains a subcritical agent. Sufficiency is therefore the case $N = 2$ of Proposition 1.

For necessity, let $\lambda_i \neq \mu_i$ and suppose the condition fails, say for agent 1. Then $\lambda_1 > \mu_1$, and either $\gamma_2 = 0$ or $\lambda_2 > \mu_2$. A birth–death process with rates $\lambda > \mu$ that starts at 1 never hits 0 with probability $1 - \mu/\lambda > 0$. If $\gamma_2 = 0$, agent 1 is never reset, so from $(1, 0)$ it never returns to 0 with probability $1 - \mu_1/\lambda_1$. If $\lambda_2 > \mu_2$, no reset can occur from $(1, 1)$ until a coordinate hits zero, so with probability $(1 - \mu_1/\lambda_1)(1 - \mu_2/\lambda_2)$ neither coordinate ever does. With positive probability, the chain jumps from $(0, 0)$ to $(1, 0)$ and then to $(1, 1)$ without visiting $(0, 0)$ in between, so it returns to $(0, 0)$ with probability less than one. By irreducibility (Proposition 2), it is transient and in particular not positive recurrent. □

3.3 Stationary distribution and functional equation

Assume that the chain is positive recurrent. Let

$$\pi_{i,j} = \lim_{t\to\infty} \mathbb{P}(X_1(t) = i, X_2(t) = j)$$

denote the stationary distribution, and define the joint probability generating function

$$P(x,y) = \sum_{i=0}^{\infty}\sum_{j=0}^{\infty} \pi_{i,j} x^i y^j.$$

Theorem 4 (Functional equation) *For $0 < |x| < 1$ and $0 < |y| < 1$, the stationary generating function $P(x,y)$ satisfies the functional equation*

$$\begin{aligned}\Big(\lambda_1 + \lambda_2 + \mu_1 + \mu_2 - \lambda_1 x - \lambda_2 y - \frac{\mu_1}{x} - \frac{\mu_2}{y}\Big) P(x,y) \\ = \Big(\mu_1 - \gamma_1 - \frac{\mu_1}{x}\Big) P(0,y) + \Big(\mu_2 - \gamma_2 - \frac{\mu_2}{y}\Big) P(x,0) + \gamma_1 P(0,1) + \gamma_2 P(1,0). \quad (2)\end{aligned}$$

After multiplication by xy, the identity extends continuously to zero coordinates and also holds throughout any larger domain of absolute convergence used below.

3.4 Derivation of the stationary solution

Under the sufficient stability conditions of Proposition 3, the chain has a unique stationary law. Its pgf is the unique proper bivariate pgf satisfying the full functional equation: coefficient extraction recovers the stationary balance equations.

For the integral construction below, assume $\lambda_i, \mu_i > 0$ and $\gamma_i \geq 0$, the sufficient stability condition, and, after interchanging the agent labels if necessary,

$$\lambda_1 < \mu_1, \qquad \lambda_2 \neq \mu_2. \quad (3)$$

If $\lambda_2 > \mu_2$, stability requires $\gamma_1 > 0$. Individual zero reset rates are allowed. If $\gamma_1 = \gamma_2 = 0$, both coordinates must be subcritical, and their stationary pgf is simply

$$P(x,y) = \frac{1-\rho_1}{1-\rho_1 x}\frac{1-\rho_2}{1-\rho_2 y}, \qquad \rho_i := \lambda_i/\mu_i < 1. \quad (4)$$

We therefore use the integral construction only when $\gamma_1 + \gamma_2 > 0$.

The critical case $\lambda_2 = \mu_2$ is excluded from the construction only, not from Proposition 3: the second kernel contour then lies on the unit circle, and closing the constants system would require a separate boundary argument, which we do not supply.

Solving the functional equation (2) requires finding the boundary generating functions $P(x,0)$ and $P(0,y)$. As demonstrated in the seminal works of [29, 30], these boundary functions can be found as the solution of a boundary-value problem for analytic functions. For the functional equation considered here, however, a slightly more direct method can be applied. As demonstrated in Appendix B, this approach

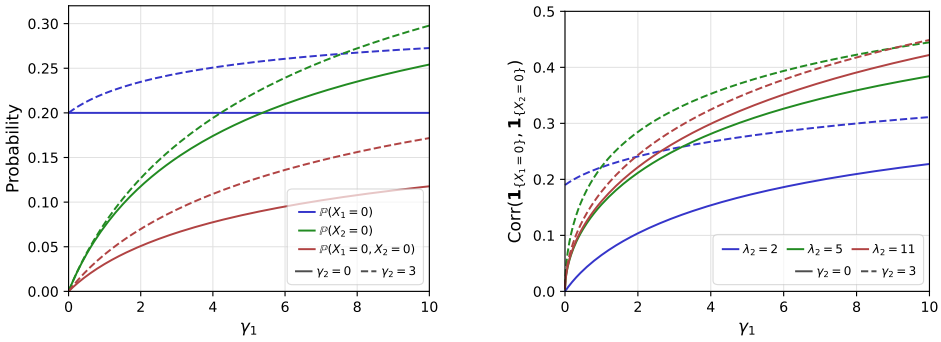

(a) Stationary probabilities for state 0, $\lambda_2 = 5$ (b) Correlation of $\mathbf{1}_{\{X_1=0\}}$ and $\mathbf{1}_{\{X_2=0\}}$

Fig. 1: Two-agent model versus γ_1, with $\gamma_2 = 0$ (solid) and $\gamma_2 = 3$ (dashed). Parameters: $\lambda_1 = 8, \mu_1 = 10, \mu_2 = 3$

exploits a Fourier-type orthogonality relation, comparable to [31], and yields $P(x, 0)$ and $P(0, y)$ as integrals. Nevertheless, due to the explicit appearance of $P(0, 1)$ and $P(1, 0)$ in (2)—which are not present in [31]—it is necessary to solve an additional linear system to determine these unknown constants.

The identities in Appendix B are stated for products of the boundary pgfs with linear prefactors. Zeros of those prefactors impose removable-singularity compatibility conditions rather than poles. The displayed constants systems give direct reconstruction when nonsingular; in singular cases additional compatibility relations are provided, but a universally nonsingular closure is not claimed. A formal solution must still reconstruct a normalized pgf with nonnegative coefficients satisfying the full functional equation.

3.5 Numerical example

To illustrate the effect of the reset rates in the two-agent model, we consider a setting in which agent 1 is stabilizing, with $\lambda_1 = 8$ and $\mu_1 = 10$. For agent 2, we take $\mu_2 = 3$ and $\lambda_2 = 5 > \mu_2$, so that it would be unstable in the absence of agent 1. Figure 1a shows the stationary probability of an empty state as a function of the reset rate γ_1 generated by agent 1. Solid and dashed lines correspond to $\gamma_2 = 0$ and $\gamma_2 = 3$, respectively, where γ_2 is the reset rate associated with agent 2. The figure clearly demonstrates the stabilizing effect of γ_1 on agent 2: increasing γ_1 leads to a higher probability that agent 2 is in state 0. In the limit $\gamma_1 \to 0$, the process governing agent 2 becomes unstable.

In the case $\gamma_2 = 0$ (solid lines), the dynamics of agent 1 reduce to those of a standard $M/M/1$ birth–death process, with stationary probability of an empty system given by $1 - \lambda_1/\mu_1$. When $\gamma_2 > 0$ (dashed lines), we observe that the probability that agent 1 is in state 0 increases with γ_1. This can be explained as follows: a higher γ_1 increases the probability that agent 2 is reset to state 0, which, in turn, raises the frequency of resets of agent 1 when $\gamma_2 > 0$. Conversely, the more frequent resets of

agent 1 by agent 2 also lead to more resets of agent 2, which explains why $\mathbb{P}(X_2 = 0)$ is larger for $\gamma_2 = 3$.

Since interaction between the agents is exclusively triggered when an agent is in the zero state, we quantify statistical dependence through the joint probability of simultaneous zero states. In particular, we consider the correlation of the indicator variables $\mathbf{1}_{\{X_1=0\}}$ and $\mathbf{1}_{\{X_2=0\}}$, that is,

$$\mathrm{Corr}\left(\mathbf{1}_{\{X_1=0\}}, \mathbf{1}_{\{X_2=0\}}\right) = \frac{\mathbb{P}(X_1 = 0, X_2 = 0) - \mathbb{P}(X_1 = 0)\mathbb{P}(X_2 = 0)}{\sqrt{\mathbb{P}(X_1 = 0)(1 - \mathbb{P}(X_1 = 0))\mathbb{P}(X_2 = 0)(1 - \mathbb{P}(X_2 = 0))}}.$$

Figure 1b shows the correlation as a function of the reset rate γ_1 generated by agent 1, for several values of the arrival rate λ_2 of agent 2. The figure clearly shows that the correlation increases as γ_1 increases. For large γ_1, the correlation is higher when agent 2 is more strongly supercritical. For $\lambda_2 = 2 < \mu_2$, agent 2 is stable on its own, and with $\gamma_2 = 3$ it still resets agent 1 when $\gamma_1 = 0$. This explains why the corresponding dashed curve starts at a positive value.

4 Graphon Interaction Framework

In this section we extend the finite-agent reset model to a network of connected and interacting agents. We focus on dense graphon-kernel interactions, in which every ordered pair of agents interacts through a fixed weight determined by their latent types. This formulation yields a tractable mean-field limit while allowing for heterogeneous interaction and rate structures.

4.1 Dense graphon-kernel model

Let $\Gamma : [0,1]^2 \to [0,\infty)$ be a bounded reset kernel. Here $\Gamma(u,v)$ denotes the reset intensity from a sender of type u to a receiver of type v. For each $N \geq 1$, we sample latent types $U_1, \ldots, U_N$ uniformly in $[0,1]$ and in an i.i.d. fashion.

Given (U_i), define interaction weights

$$\Gamma_{ij}^{(N)} = \begin{cases} \Gamma(U_i, U_j), & i \neq j, \\ 0, & i = j. \end{cases}$$

The diagonal convention suppresses self-resets and, as we will show later, does not affect the large-N limit. Each agent i carries a level $X_i^N(t) \in \mathbb{Z}_{\geq 0}$ and evolves according to type-dependent birth–death dynamics. Specifically, agent i has arrival rate $\lambda(U_i)$ and service rate $\mu(U_i)$, where $\lambda, \mu : [0,1] \to [0,\infty)$ are measurable functions.

4.2 Reset dynamics

Reset events are generated by agents at state 0 and propagate through the interaction network: conditional on the system state $X^N(t)$, agent j experiences reset events at

rate

$$r_j^N(t) = \frac{1}{N}\sum_{i=1}^{N} \Gamma_{ij}^{(N)}\, \mathbf{1}\{X_i^N(t) = 0\} = \frac{1}{N}\sum_{\substack{i=1 \\ i\neq j}}^{N} \Gamma(U_i, U_j)\mathbf{1}\{X_i^N(t) = 0\}. \tag{5}$$

Each such event instantaneously resets $X_j^N(t)$ to 0, regardless of its current value. The $1/N$ normalization keeps reset hazards of order $O(1)$ as the population size grows. We define the type-tagged empirical measure

$$\nu_t^N = \frac{1}{N}\sum_{i=1}^{N} \delta_{(U_i, X_i^N(t))} \quad \text{on } [0,1] \times \mathbb{Z}_{\geq 0},$$

where δ_a is the Dirac measure at a.

5 Graphon Mean-Field Limit

This section derives the mean-field limit of the dense graphon-kernel system of Section 4.

5.1 Mean-field equations

In the large-population limit, let $p_k(t,u)$ denote the probability that an agent of type $u \in [0,1]$ has level $k \in \mathbb{Z}_{\geq 0}$ at time t. A type-u agent experiences resets at the rate given by the weighted boundary mass of the population,

$$r(t,u) = \int_0^1 \Gamma(v,u)\, p_0(t,v)\, dv. \tag{6}$$

The forward equations for $(p_k(t,u))_{k\geq 0}$ are:

$$\partial_t p_k(t,u) = \lambda(u)p_{k-1}(t,u) + \mu(u)p_{k+1}(t,u) - \big(\lambda(u) + \mu(u) + r(t,u)\big)p_k(t,u), \quad k \geq 1, \tag{7}$$

$$\partial_t p_0(t,u) = \mu(u)p_1(t,u) - \lambda(u)p_0(t,u) + r(t,u)\big(1 - p_0(t,u)\big). \tag{8}$$

Equations (6)–(8) form a coupled nonlinear system indexed by the latent type u. We prove the following theorem and the propagation-of-chaos corollary in Appendix C.

Theorem 5 (Graphon mean-field limit) *Let $\Gamma : [0,1]^2 \to [0,\infty)$ be a bounded reset kernel that is Lebesgue-a.e. continuous, and let $\lambda, \mu : [0,1] \to [0,\infty)$ be bounded and Lebesgue-a.e. continuous. For each $N \geq 1$, let $(U_i)_{i=1}^N$ be i.i.d.* Unif$[0,1]$*, and let $(X_i^N(t))_{i=1}^N$ evolve according to the dense graphon-kernel reset model of Section 4, with reset rates (5) and type-tagged empirical measures ν_t^N as in Section 4. Assume that the initial empirical measures satisfy*

$$\nu_0^N \Rightarrow \nu_0 \quad \textit{in probability}, \qquad \nu_0(du, dk) = du\, p_k(0,u),$$

and that the first empirical moments are uniformly bounded:

$$\sup_{N\geq 1} \mathbb{E}\Big[\frac{1}{N}\sum_{i=1}^{N} X_i^N(0)\Big] < \infty.$$

Then, for every $T < \infty$, the processes $(\nu_t^N)_{t\in[0,T]}$ converge in probability in $D([0,T], \mathcal{P}([0,1]\times\mathbb{Z}_{\geq 0}))$ to the deterministic path $(\nu_t)_{t\in[0,T]}$, where

$$\nu_t(du, dk) = du\, p_k(t,u).$$

For every $t \in [0,T]$ and a.e. u, $(p_k(t,u))_{k\geq 0}$ is a probability distribution on $\mathbb{Z}_{\geq 0}$. Moreover, p is the unique solution of the graphon mean-field equations (6)–(8) with initial condition $p_k(0,u)$ that is continuous in integrated ℓ^1 distance:

$$\lim_{s\to t} \int_0^1 \sum_{k\geq 0} |p_k(s,u) - p_k(t,u)|\, du = 0.$$

Corollary 6 (Propagation of chaos) *In addition to the assumptions of Theorem 5, suppose that, conditional on the full type vector $(U_1,\dots,U_N)$, the initial levels are independent and satisfy*

$$\mathbb{P}\big(X_i^N(0) = k \mid U_1,\dots,U_N\big) = p_k(0,U_i), \qquad i = 1,\dots,N, \quad k \geq 0.$$

For each fixed $m \geq 1$ and $T < \infty$, the augmented processes

$$\big((U_1, X_1^N),\dots,(U_m, X_m^N)\big)$$

converge jointly in distribution on $([0,1]\times D([0,T],\mathbb{Z}_{\geq 0}))^m$ to independent copies of a typed nonlinear Markov process. Conditional on a limiting type u, this process has birth rate $\lambda(u)$, death rate $\mu(u)$ away from zero, and reset rate $r(t,u) = \int_0^1 \Gamma(v,u)p_0(t,v)\, dv$.

Remark 2 (Finite-time scope) Theorem 5 is a finite-time mean-field limit. It does not assert convergence of stationary distributions of the finite-N systems. See Section 6.1.

Remark 3 (Regularity of the reset kernel) The a.e.-continuity assumption avoids technical pathologies of arbitrary bounded measurable kernels and is what the proof uses in Lemma 14. It still covers the standard finite-block case, which is discontinuous only on the block-boundary lines and is therefore Lebesgue-a.e. continuous.

5.2 Uniform mean-field as a special case

The uniform-mixing mean-field model is obtained from (6)–(8) by choosing a constant reset kernel and constant rate functions. Concretely, take $\Gamma(u,v) \equiv c$ and $\lambda(u) \equiv \lambda$, $\mu(u) \equiv \mu$. If, in addition, the solution is type-homogeneous, i.e., $p_k(t,u) \equiv p_k(t)$ for a.e. u, then the reset intensity (6) reduces to $r(t) = c\, p_0(t)$, and the mean-field equations become, for $k \geq 1$,

$$\dot{p}_k(t) = \lambda p_{k-1}(t) + \mu p_{k+1}(t) - \big(\lambda + \mu + r(t)\big)p_k(t), \tag{9}$$

together with the boundary equation

$$\dot{p}_0(t) = \mu p_1(t) - \lambda p_0(t) + r(t)\big(1 - p_0(t)\big). \tag{10}$$

The above equations constitute the homogeneous scalar mean-field model with endogenous resets, in which interaction acts only through a single dimension, namely, the fraction of agents at the boundary state.

6 Stationary Solutions of the Limiting Equations

This section studies stationary solutions of the limiting nonlinear equations (6)–(8), which should not be confused with limits of invariant distributions of the finite-N chains as $N \to \infty$.

Write $\pi_k(u)$ for a stationary type-u probability and set

$$r(u) = \int_0^1 \Gamma(v,u)\pi_0(v)\,dv. \tag{11}$$

For fixed u, the stationary equations are those of a reflected birth–death process with a constant reset-to-zero rate $r(u)$ from every positive state. We first characterize the situation in which there is only one type, writing λ, μ, and r without their type arguments.

Proposition 7 (One-type stationary law) *Assume $\lambda > 0$ and $\mu > 0$.*

(i) *If $r > 0$, the process has a unique stationary distribution.*

(ii) *If $r = 0$ and $\lambda < \mu$, it has the usual $M/M/1$ stationary distribution.*

(iii) *If $r = 0$ and $\lambda \geq \mu$, it has no stationary probability distribution.*

In cases (i) *and* (ii)*, let α be the smaller root of*

$$\mu\alpha^2 - (\lambda + \mu + r)\alpha + \lambda = 0. \tag{12}$$

Then $0 < \alpha < 1$, and

$$\pi_0 = h_{\lambda,\mu}(r) = 1 - \alpha, \tag{13}$$

$$\pi_k = (1-\alpha)\alpha^k, \qquad k \geq 0, \tag{14}$$

$$\mathbb{E}[X] = \frac{\alpha}{1-\alpha}. \tag{15}$$

For $r = 0$ and $\lambda < \mu$, these expressions reduce to $\alpha = \lambda/\mu$ and $\pi_k = (1 - \lambda/\mu)(\lambda/\mu)^k$.

Proof For $k \geq 1$, the stationary bulk equations give

$$\mu\pi_{k+1} - (\lambda + \mu + r)\pi_k + \lambda\pi_{k-1} = 0.$$

If $r > 0$, the characteristic polynomial has one root $\alpha \in (0,1)$ and one root larger than one. Summability eliminates the latter, so $\pi_k = C\alpha^{k-1}$ for $k \geq 1$. Normalization gives $C = (1-\pi_0)(1-\alpha)$, and the boundary equation $0 = \mu\pi_1 - \lambda\pi_0 + r(1-\pi_0)$ yields (13). This produces a probability distribution satisfying all balance equations. Irreducibility then gives uniqueness.

When $r = 0$, the roots are 1 and λ/μ. If $\lambda < \mu$, summability removes the unit root and gives the standard geometric law. If $\lambda \geq \mu$, no nonzero summable solution of the bulk and boundary equations can be normalized. Equivalently, the reflected birth–death process is null recurrent when $\lambda = \mu$ and transient when $\lambda > \mu$. □

For a single type, the model is the $M/M/1$ queue with catastrophes. For numerical evaluation, the smaller root used above is more stably written as

$$\alpha(r) = \frac{2\lambda}{\lambda + \mu + r + \sqrt{(\lambda + \mu + r)^2 - 4\lambda\mu}}.$$

Thus $\alpha(r)$ and $h_{\lambda,\mu}(r)$ are continuous for $r > 0$, and their limits as $r \downarrow 0$ give the following extension. Define

$$h_{\lambda,\mu}(0) := \begin{cases} 1 - \lambda/\mu, & \lambda < \mu, \\ 0, & \lambda \geq \mu. \end{cases} \tag{16}$$

For fixed positive λ and μ, the resulting map $r \mapsto h_{\lambda,\mu}(r)$ is continuous on $[0, \infty)$. At $r = 0$ with $\lambda \geq \mu$, however, the value in (16) is only a continuous extension: by Proposition 7, it does not represent a stationary law.

Define the positive integral operator

$$(Kb)(u) := \int_0^1 \Gamma(v, u) b(v)\, dv \tag{17}$$

and the response map

$$(\mathcal{T}b)(u) := h_{\lambda(u),\mu(u)}\big((Kb)(u)\big). \tag{18}$$

A stationary solution necessarily has boundary mass $b = \pi_0$ satisfying

$$b = \mathcal{T}b. \tag{19}$$

Conversely, a fixed point b determines a stationary solution through (14) provided it is *admissible*, meaning

$$(Kb)(u) > 0 \quad \text{for a.e. } u \text{ such that } \lambda(u) \geq \mu(u). \tag{20}$$

No positivity condition is required on subcritical types, for which the $r = 0$ stationary law is normalizable. Sections 6.1 and 6.2 treat two models without subcritical types, where the existence of a stationary solution depends on the strength of the coupling. Section 6.3 gives general existence under influence from subcritical types.

6.1 Finite-population versus limiting stationarity

The finite-time mean-field limit and the stationary fixed-point problem involve different limiting operations. This distinction is particularly important in the uniform model when $\lambda > \mu$. For every fixed N, take the complete dense interaction with reset weight c/N between distinct agents. Starting from a state with all coordinates equal to one, no reset is generated until a coordinate first reaches zero. Up to that time, the

coordinates evolve as independent birth–death processes. A supercritical birth–death process started from one avoids zero with probability $1 - \mu/\lambda$. Consequently,

$$\mathbb{P}_{\mathbf{1}}\big(X_i^N(t) > 0 \text{ for every } i \text{ and } t \geq 0\big) = \Big(1 - \frac{\mu}{\lambda}\Big)^N > 0. \tag{21}$$

On this event no reset ever occurs and the origin is never reached. The finite chain is irreducible, so it is not recurrent and is therefore transient. In particular, it has no invariant probability distribution. The finite-N sufficient condition of Proposition 1 does not apply: the sole source component contains no subcritical agent.

The limiting nonlinear equations can nevertheless have a positive stationary fixed point. To see this, let b denote the stationary boundary mass in the uniform model. Its reset rate is $r = cb$, so the fixed-point equation is

$$b = h_{\lambda,\mu}(cb). \tag{22}$$

When $\lambda > \mu$, substituting $r = cb$ into (12) and (13), and eliminating α, gives

$$b\,\big[b(c+\mu) - c + \lambda - \mu\big] = 0.$$

The root $b = 0$ is the non-admissible continuous-extension fixed point discussed below (16). A positive, and therefore admissible, stationary fixed point exists exactly when $c > \lambda - \mu$; it is then unique and equals

$$b_* = \frac{c - \lambda + \mu}{c + \mu}. \tag{23}$$

For $(\lambda, \mu, c) = (8, 5, 10)$, this gives $b_* = 7/15$ and stationary reset rate $cb_* = 14/3$, even though every fixed finite system is non-recurrent.

There is no contradiction. Theorem 5 takes $N \to \infty$ on a fixed time interval, whereas stationarity concerns $t \to \infty$. The theorem establishes neither uniform-in-time convergence nor that the limits $N \to \infty$ and $t \to \infty$ may be interchanged.

A large finite system may display equilibrium-like behavior near a positive mean-field fixed point over an observed finite horizon before a rare loss of all boundary senders switches off the reset mechanism. Establishing such metastability, including escape-time asymptotics, requires a separate analysis and is not claimed here.

6.2 Exact threshold for two cross-coupled blocks

The stationary fixed-point equations admit a sharper characterization for two equal-sized blocks with no within-block interaction. We use a per-block intensity convention: if $b_i = \pi_0^{(i)}$ is the boundary mass in block i, then

$$r_1 = \Gamma_{21} b_2, \qquad r_2 = \Gamma_{12} b_1. \tag{24}$$

Thus Γ_{ij} is the reset intensity sent from block i to block j after normalization by the size of the sending block. For equal blocks, the corresponding graphon kernel

has cross-block values $2\Gamma_{ij}$ because the graphon integral is normalized by the whole population.

Proposition 8 (Two-block stationary threshold) *Suppose that $\lambda_i \geq \mu_i > 0$ for $i = 1, 2$ and that $\Gamma_{12}, \Gamma_{21} > 0$. The two-block limiting equations with reset rates (24) possess an admissible stationary solution if and only if*

$$\Gamma_{12}\Gamma_{21} > (\lambda_1 - \mu_1)(\lambda_2 - \mu_2). \tag{25}$$

When it exists, the admissible stationary solution is unique.

Proof Set $d_i := \lambda_i - \mu_i \geq 0$. For an admissible solution, $b_i \in (0, 1)$. Substitution of $\alpha_i = 1 - b_i$ in (12) gives the equivalent one-type relation

$$r_i(1 - b_i) = b_i(d_i + \mu_i b_i). \tag{26}$$

Combining this relation with (24) yields

$$\Gamma_{21}b_2 = b_1 f_1(b_1), \qquad \Gamma_{12}b_1 = b_2 f_2(b_2), \qquad f_i(b) := \frac{d_i + \mu_i b}{1 - b}. \tag{27}$$

Multiplying the first two equations and cancelling $b_1 b_2 > 0$ shows that every admissible solution must satisfy

$$\Gamma_{12}\Gamma_{21} = f_1(b_1)f_2(b_2) > d_1 d_2.$$

This proves necessity, including nonexistence at equality.

For sufficiency and uniqueness, define

$$g(b) := \frac{b f_1(b)}{\Gamma_{21}}, \qquad 0 \leq b < 1.$$

The functions f_i are continuous and strictly increasing, while g is continuous and strictly increasing from zero to infinity. Hence there is a unique $L \in (0, 1)$ such that $g(L) = 1$. On $0 < b < L$, the first equation in (27) is equivalent to $b_2 = g(b)$, and the second is equivalent to

$$f_1(b)f_2(g(b)) = \Gamma_{12}\Gamma_{21}. \tag{28}$$

The left-hand side is continuous and strictly increasing, tends to $d_1 d_2$ as $b \downarrow 0$, and tends to infinity as $b \uparrow L$. Equation (28) therefore has exactly one solution in $(0, L)$ precisely under (25). It determines unique $b_1, b_2 \in (0, 1)$ and positive reset rates. Proposition 7 then supplies the unique geometric stationary law in each block, so the fixed point is admissible and unique. □

6.3 Stability considerations

As in the finite-agent model of Section 2, stability is driven by types that both return to the boundary frequently and generate reset pressure. The obstruction is a measurable set of types that receives no reset influence from outside and contains no types capable of generating persistent boundary-triggered reset pressure.

6.3.1 Structural guidance

We first introduce a graphon-level analogue of the notion of an externally reset-isolated set from the finite-agent model.

Definition 1 (Externally reset-isolated measurable sets) A measurable set $A \subset [0,1]$ with positive Lebesgue measure is called *externally reset-isolated* if

$$\int_{A^c} \Gamma(v,u)\,dv = 0 \quad \text{for a.e. } u \in A.$$

Equivalently, types in A receive no reset influence from outside A.

Let

$$S_- := \{u \in [0,1] : \lambda(u) < \mu(u)\} \qquad \text{and} \qquad B := \{u \in [0,1] : \lambda(u) \geq \mu(u)\}$$

denote the intrinsically subcritical and non-subcritical regions, respectively; critical types with $\lambda(u) = \mu(u)$ belong to B.

A natural continuum analogue of the finite-N source-component condition would require every externally reset-isolated measurable set A of positive measure to satisfy $\mathrm{Leb}(A \cap S_-) > 0$. We use this qualitative condition only as structural intuition. In full graphon generality, such reachability assumptions do not provide quantitative lower bounds on the reset rates transmitted along indirect paths. We therefore turn to verifiable direct-influence conditions that ensure the existence and quantitative stability properties of stationary graphon fixed points.

6.3.2 Basic one-type estimates

The following lemmas collect the one-type estimates used below. Their proofs rely on the stationary generator identity applied to the truncated Lyapunov functions $V_K(k) = k \wedge K$.

Lemma 9 (Boundary mass of subcritical types) *Consider a stationary birth–death process with birth rate λ, death rate μ, and reset rate $r \geq 0$ that sends the process to zero. If $\lambda < \mu$, then every stationary distribution satisfies*

$$\pi_0 \geq 1 - \frac{\lambda}{\mu}.$$

Proof Using $V_K(k) = k \wedge K$ and the stationary generator equation, the reset contribution is nonpositive, while the birth and death terms yield $0 \leq \lambda\,\mathbb{P}(X < K) - \mu\,\mathbb{P}(0 < X \leq K)$. Letting $K \to \infty$ gives $0 \leq \lambda - \mu(1 - \pi_0)$ and hence the claim. □

Lemma 10 (Moment bound under positive reset) *For a stationary birth–death process with birth rate λ, death rate $\mu \geq 0$, and reset rate $r > 0$,*

$$\mathbb{E}[X] \leq \frac{\lambda}{r}.$$

Proof Applying the stationary generator identity to $V_K(k) = k \wedge K$ and noting that resets decrease V_K by $k \wedge K$, we obtain $0 \le \lambda - r\,\mathbb{E}[V_K(X)]$. Letting $K \to \infty$ proves the bound. □

Finally, when $\lambda < \mu$ and $r \ge 0$, the process is stochastically dominated by the classic $M/M/1$ queue with no resets, so that

$$\mathbb{E}[X] \le \frac{\lambda}{\mu - \lambda}. \tag{29}$$

6.3.3 Direct influence from subcritical types

The lemmas above imply quantitative stability whenever subcritical types send sufficient reset pressure to non-subcritical types.

Theorem 11 (Existence under direct subcritical influence) *Assume that* $\lambda(u) > 0$ *and* $\mu(u) > 0$ *for a.e.* u*, that* Γ *is bounded, and that*

$$\int_{S_-} \Gamma(v,u)\,dv > 0 \qquad \text{for a.e. } u \in B. \tag{30}$$

Then the limiting graphon equations possess at least one stationary solution. Moreover, every stationary solution satisfies $r(u) > 0 \qquad$ *for a.e.* $u \in B$.

Proof Work in $L^2([0,1])$ and let

$$\mathcal{C} := \{b \in L^2([0,1]) : 0 \le b \le 1 \text{ a.e.}\}.$$

This set is nonempty, closed, bounded, and convex. Since $\Gamma \in L^2([0,1]^2)$, the operator K in (17) is Hilbert–Schmidt and hence compact on L^2. The map $\mathcal{T}$ in (18) sends $\mathcal{C}$ into itself.

We record why $\mathcal{T}$ is continuous and compact despite the type-dependent parameters. If $b_n \to b$ in L^2, then $Kb_n \to Kb$ in L^2. Every subsequence has a further subsequence converging pointwise a.e. Pointwise continuity of $h_{\lambda(u),\mu(u)}$ and the bound $0 \le h \le 1$ then give $\mathcal{T}b_n \to \mathcal{T}b$ in L^2 by dominated convergence. The same subsequence argument, starting from the relative compactness of $\{Kb : b \in \mathcal{C}\}$, shows that $\mathcal{T}(\mathcal{C})$ is relatively compact. Schauder's fixed-point theorem therefore yields some $b \in \mathcal{C}$ with $b = \mathcal{T}b$.

For $v \in S_-$, Proposition 7 makes the response a genuine one-type stationary boundary mass. Lemma 9 therefore gives

$$b(v) = h_{\lambda(v),\mu(v)}((Kb)(v)) \ge 1 - \frac{\lambda(v)}{\mu(v)} > 0.$$

Consequently, for a.e. $u \in B$,

$$(Kb)(u) \ge \int_{S_-} \Gamma(v,u) b(v)\,dv > 0$$

by (30). Thus the fixed point is admissible in the sense of (20), and Proposition 7 constructs a stationary solution. Finally, consider any stationary solution. By Lemma 9, its boundary mass satisfies $\pi_0(v) > 0$ for a.e. $v \in S_-$. Condition (30) then implies $r(u) > 0$ for a.e. $u \in B$. □

6.3.4 Uniform influence and finite first moments

Define, for $\delta > 0$, $S_\delta := \{v \in [0,1] : \mu(v) - \lambda(v) \geq \delta\}$, and impose two quantitative assumptions:

$$\int_{S_-} \frac{\lambda(u)}{\mu(u) - \lambda(u)}\, du < \infty, \tag{SM}$$

$$\exists\, \delta, \eta > 0 : \qquad \int_{S_\delta} \Gamma(v,u)\, dv \geq \eta \quad \text{for a.e. } u \in B. \tag{UDI}$$

The first condition guarantees integrable tails for the intrinsically subcritical types; the second provides a uniform lower bound on the reset pressure transmitted from a uniformly subcritical subset.

Proposition 12 (Uniform direct influence and finite moment) *Suppose that* (SM) *and* (UDI) *hold, that* λ, μ, *and* Γ *are bounded, and that* $\lambda(u), \mu(u) > 0$ *for a.e.* u. *Then at least one stationary solution exists. Every stationary solution satisfies*

$$r(u) \geq r_* := \eta \frac{\delta}{\|\mu\|_\infty} \qquad \text{for a.e. } u \in B.$$

Consequently, every stationary solution has finite aggregate first moment and obeys

$$\int_0^1 \mathbb{E}_{\pi(u)}[X]\, du \leq \int_{S_-} \frac{\lambda(u)}{\mu(u) - \lambda(u)}\, du + \frac{\|\lambda\|_\infty}{r_*} \mathrm{Leb}(B) < \infty.$$

Proof Existence follows from Theorem 11, because (UDI) implies (30). For $v \in S_\delta$, Lemma 9 gives

$$\pi_0(v) \geq \frac{\mu(v) - \lambda(v)}{\mu(v)} \geq \frac{\delta}{\|\mu\|_\infty}.$$

Therefore, for a.e. $u \in B$ satisfying (UDI),

$$r(u) \geq \int_{S_\delta} \Gamma(v,u)\pi_0(v)\, dv \geq \frac{\delta}{\|\mu\|_\infty} \int_{S_\delta} \Gamma(v,u)\, dv \geq r_*.$$

Lemma 10 then yields $\mathbb{E}_{\pi(u)}[X] \leq \|\lambda\|_\infty / r_*$ on B, while (29) and (SM) control the integral over S_-. Combining the bounds proves the claim. □

Remark 4 Critical types with $\lambda(u) = \mu(u)$ lie in B and require positive reset influence in every stationary solution. Sharp continuum criteria based only on externally reset-isolated sets, as well as uniqueness and long-time convergence for the nonlinear equations, remain open.

7 Numerical Experiments

We report numerical experiments in three settings. First, Section 7.1 illustrates the homogeneous model of Section 5.2 and the order-of-limits example of Section 6.1. Second, a finite-block population in Section 7.2 compares a directed two-group interaction with the two-agent model of Section 3. Finally, Section 7.3 treats position-dependent graphon kernels. This section illustrates the mean-field limit of Theorem 5 and the propagation of chaos of Corollary 6, and tests the stationary results of Section 6.

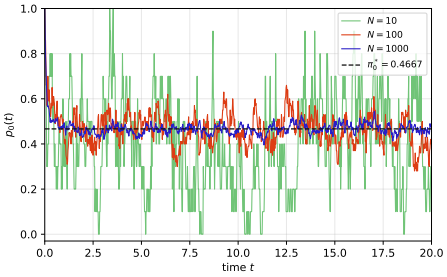

(a) Boundary mass $p_0(t)$ for $N = 10, 100, 1000$, all starting from $x_0 = 0$.

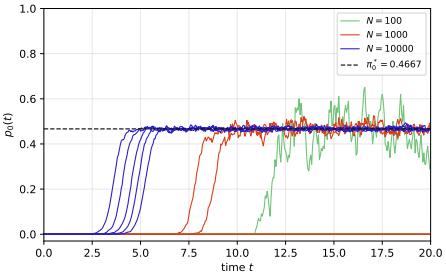

(b) Boundary mass $p_0(t)$ at fixed $x_0 = 15$ for $N = 100, 1000, 10^4$, five realizations each.

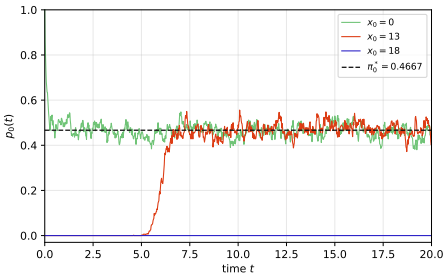

(c) Boundary mass $p_0(t)$ at $N = 500$ for three initial levels

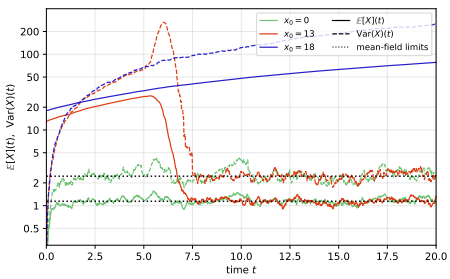

(d) Mean and variance

Fig. 2: Uniform Mean Field: Special case. Parameters $\lambda = 8, \mu = 5, c = 10$

7.1 Homogeneous population

We illustrate the uniform mean-field model of Section 5.2 in the supercritical regime $\lambda > \mu$. As shown in Section 6.1, every corresponding finite-N chain is transient and has no stationary distribution, even though the limiting nonlinear equations possess a positive stationary fixed point when $c > \lambda - \mu$.

We use $\lambda = 8$, $\mu = 5$ and $c = 10$, for which (23) gives the positive mean-field fixed point $p_0^* = 7/15$. The experiments therefore concern finite-time tracking of this mean-field regime, not convergence to a finite-N stationary distribution.

Figure 2 shows that the finite systems track this mean-field regime over the simulated horizon. Started from $x_0 = 0$ (panel (a)), the boundary mass fluctuates around p_0^*, with fluctuations that shrink as N grows. Started from a high level where no agent is initially at zero in panels (b) and (c), no resets occur and the levels drift upwards. Once agents start to reach zero, the reset mechanism switches on and the boundary mass moves quickly to p_0^*. Panel (d) shows that the mean and variance first grow and then settle near their mean-field values, rather than growing linearly.

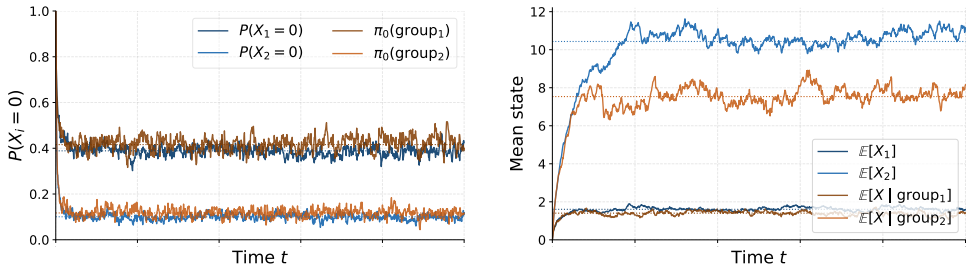

(a) Zero state occupancy proportion

(b) Mean state

Fig. 3: Two-subgroup model compared with the two-agent model, $N = 500$. Two-agent curves averaged over 500 runs while subgroups curves are a single realization. The dotted lines are the stationary predictions.

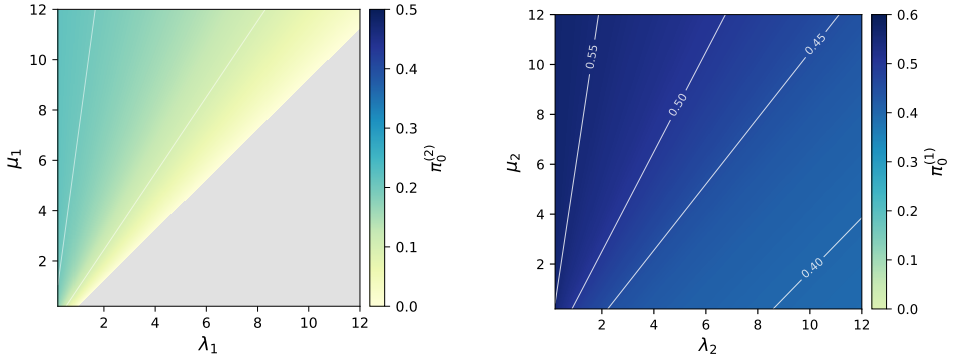

(a) $\pi_0^{(2)}$ as groups 1's rates (λ_1, μ_1) vary

(b) $\pi_0^{(1)}$ as groups 2's rates (λ_2, μ_2) vary

Fig. 4: Cross-influence between the groups, each group's own rates held fixed. The grey area signifies no stationary solution

7.2 Two Subgroups versus Two Agent Model

In this section we work with two groups of equal population size, group 1 is subcritical while group 2 is non-subcritical, with positive reset intensities Γ_{12} and Γ_{21}. The direct-subcritical influence condition in Theorem 11 reduces to $\Gamma_{12} > 0$.

Throughout, we use $\lambda_1 = 5$, $\mu_1 = 8$, $\lambda_2 = 10$, $\mu_2 = 6$, $\Gamma_{12} = 1.5$, $\Gamma_{21} = 2$ and $\Gamma_{11} = \Gamma_{22} = 0$, so there is no within-group interaction.

We use the per-block convention of Section 6.2. At $N = 2$ each block holds a single agent, and the model reduces to the two-agent chain of Section 3 with $\gamma_1 = \Gamma_{12}$ and $\gamma_2 = \Gamma_{21}$.

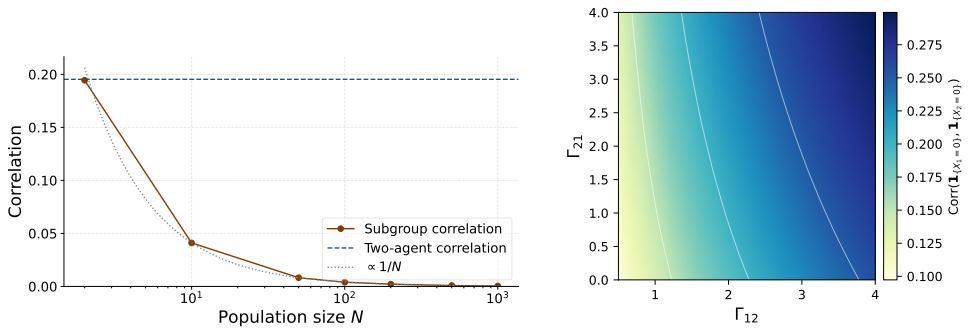

(a) correlation against population size N

(b) exact two-agent correlation over the reset intensities

Fig. 5: Correlation of the indicators $\mathbf{1}_{\{X_1=0\}}$ and $\mathbf{1}_{\{X_2=0\}}$

Figure 3 comes from the exact Gillespie simulations over the time interval $[0, 50]$. For the two-agent curve the simulation is averaged over 500 replications and the subgroup curve is a single realization at $N = 500$. Figure 5(a) uses the same simulator at each $N \in \{2, 10, 50, 100, 200, 500, 1000\}$, the correlation being time-averaged over $t \in [25, 50]$ to exclude the transient behavior.

Group 2 is non-subcritical and would drift to infinity in isolation, yet Figure 3(b) shows the mean level rising and then stabilizing. The mechanism of section 2.2 is realized for a population as boundary visits from group 1 generate reset events that capture the drift of group 2, which has no stabilizing mechanism of its own. The simulated curves settle at the fixed points values and satisfy lemma 9, which requires $\pi_0^{(1)} \geq 1 - \lambda_1/\mu_1$ and lemma 10 which requires $\mathrm{E}[\mathrm{X}|\mathrm{group}_2] \leq \lambda_2/r_2$.

The two panels of Figure 3 differ in an informative way. The boundary occupancies of the two systems almost coincide, while the mean levels separate clearly for the non-subcritical case. In the limit the stationary law is geometric by Proposition 7, and the mean is highly sensitive to the boundary mass, so a small gap in the boundary mass is magnified in the mean graph. For the two-agent case reset only occurs when agent 1 is at zero, so agent 2 passes through time stretches where it receives no reset pressure and grows freely, thus giving a higher mean value.

Figure 4 shows the stationary boundary mass of one group as the other group's parameters vary, computed from the fixed point without simulation. In the left panel, as group 1 becomes more subcritical its own boundary mass increases, which raises the reset pressure on group 2 and thereby $\pi_0^{(2)}$. The grey region marks parameters with no admissible stationary fixed point; when both groups are non-subcritical, Proposition 8 shows that this happens precisely when (25) fails. With $\Gamma_{12} = 1.5$, $\Gamma_{21} = 2$ and $\lambda_2 - \mu_2 = 4$, the boundary of the grey region is the line $\lambda_1 - \mu_1 = 0.75$, parallel to the diagonal. In the right panel, group 2 has only a small effect on the boundary mass of group 1, which stays above the floor of Lemma 9.

Figure 5(a) shows the correlation of indicators $\mathbf{1}_{\{X_1=0\}}$ and $\mathbf{1}_{\{X_2=0\}}$ for two agents in distinct groups, as a function of population size. The correlation decays at $1/N$

rate typical of mean-field systems, which is consistent with Corollary 6. At $N = 2$, the block model reduces exactly to the two-agent chain and the simulated value agrees with the exact solution.

Figure 5(b) shows the same correlation for the two-agent model across reset intensities. Correlation increases with both reset intensities, since stronger coupling makes the two agents more likely to be at zero at the same time. This is the regime in which the independence of the mean-field limit is expected to be least accurate for small populations.

7.3 Position Dependent graphon

In this section the latent type is a position $u \in [0, 1]$, and both the rate functions and the reset kernel vary with it. The birth rate ramps inside each of two regions, so that the load $\lambda(u)/\mu(u)$ changes across every type and not only between the two regions:

$$\lambda(u) = \begin{cases} 5 + 5u, & u < 0.5, \\ 7 + 6u, & u \geq 0.5, \end{cases} \qquad \mu(u) = \begin{cases} 8, & u < 0.5, \\ 6, & u \geq 0.5. \end{cases} \tag{31}$$

The load rises from 0.63 to 0.94 on $[0, 0.5)$ and from 1.67 to 2.17 on $[0.5, 1]$, so $S_- = [0, 1/2), \quad B = [1/2, 1]$, and both have positive Lebesgue measure. The drift gap $\mu - \lambda$ falls from 3 to 0.5 across S_-, so S_δ genuinely depends on δ, while staying away from zero, so that (SM) holds. The rates jump only at $u = 0.5$ and are therefore Lebesgue a.e. continuous, as Theorem 5 requires.

We use three reset kernels, each depending on the two types only through the distance $s = |u - v|$ between them:

$$\Gamma_{\sin}(v, u) = 5 + 3\sin(3\pi s), \qquad \Gamma_{\text{lin}}(v, u) = 8(1 - s), \qquad \Gamma_{\text{pow}}(v, u) = \frac{2}{(s + 1/4)^{4/5}}.$$

All three are bounded, continuous and nonnegative, and all three are symmetric, since they depend on $|u - v|$ alone. The linear and power kernels fall off with distance, so influence is strongest between nearby positions. The sine kernel does not fall off monotonically, which lets us separate the effect of proximity from the effect of how much boundary mass a sender carries. The three kernels also differ in overall size. Averaged over the unit square, the sine, linear, and power kernels have means approximately 5.64, 5.33, and 3.47, respectively. The power kernel therefore transmits less reset pressure overall, although square averages alone do not determine stationary pressure because sender occupancy and spatial placement also matter. The finite-N reset rates are given by (5), and the stationary limit by the fixed-point relation (19). with the sender index written first, and in the limit $r(u) = \int_0^1 \Gamma(v, u)\, \pi_0(v)\, \mathrm{d}v$. Together with the closed form for π_0 in terms of r, this gives the fixed point relation (19).

Every simulation is an exact Gillespie realization of the finite N chain with rates (31) and (5), started from $X_i^N(0) = 0$ for all i and run to $T = 25$. The latent positions are drawn once, as a single list of 1000 values, and then kept fixed for the whole study. The system of size N consists of the first N entries of that list, so smaller populations are nested inside larger ones and all three kernels act on exactly the same positions.

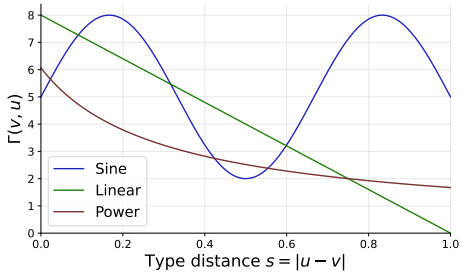

(a) Reset kernels against the type distance $s = |u - v|$

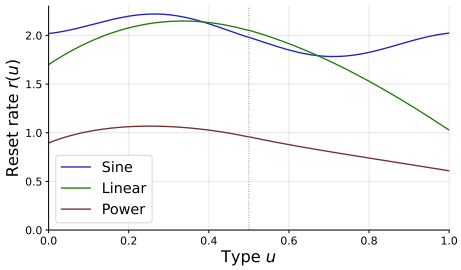

(b) Stationary reset rate $r(u)$

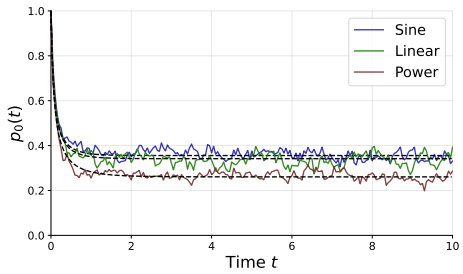

(c) Population boundary mass against time

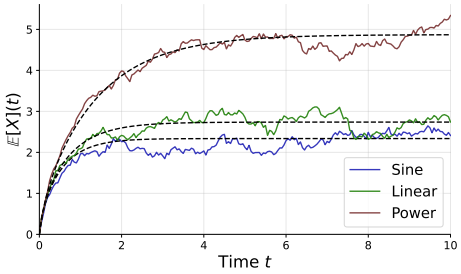

(d) Population mean level against time

Fig. 6: Position-dependent graphon: (a) the three reset kernels and (b) the resulting stationary reset rate; (c)–(d) a single realization at $N = 500$ compared with the mean-field solution.

This controls type-sampling variation across population sizes, although dynamical Monte Carlo variability remains. Stationary quantities are time averages over $[10, 25]$, and types are collected into ten equal bins.

The fixed point (19) is computed by iteration on a uniform grid of 400 types, started both from $r \equiv 0$ and from a large constant. For all three kernels both runs converged to the same profile, which we use below. We use this numerically observed profile below without interpreting it as a uniqueness proof.

Figure 6(a)–(b) shows the three kernels and the stationary reset rate they produce. The direct influence integral $\int_{S_-} \Gamma(v, u)\mathrm{d}v$ is bounded below on B for all three, so condition (30) holds and Theorem 11 applies. Every non-subcritical type receives strictly positive reset pressure and thus has a normalizable stationary distribution. The realized rates on B sit well above the bound $r_* = \eta\delta/\|\mu\|_\infty$ of Proposition 12, by roughly an order of magnitude in each case. That is what one expects, since the derivation of r_* throws away every sender outside S_δ and then replaces each boundary mass by the smallest value Lemma 9 permits.

The shape of $r(u)$ follows the shape of the kernel. For the linear and power kernels the pressure falls off across B with distance from S_-, because those are the types

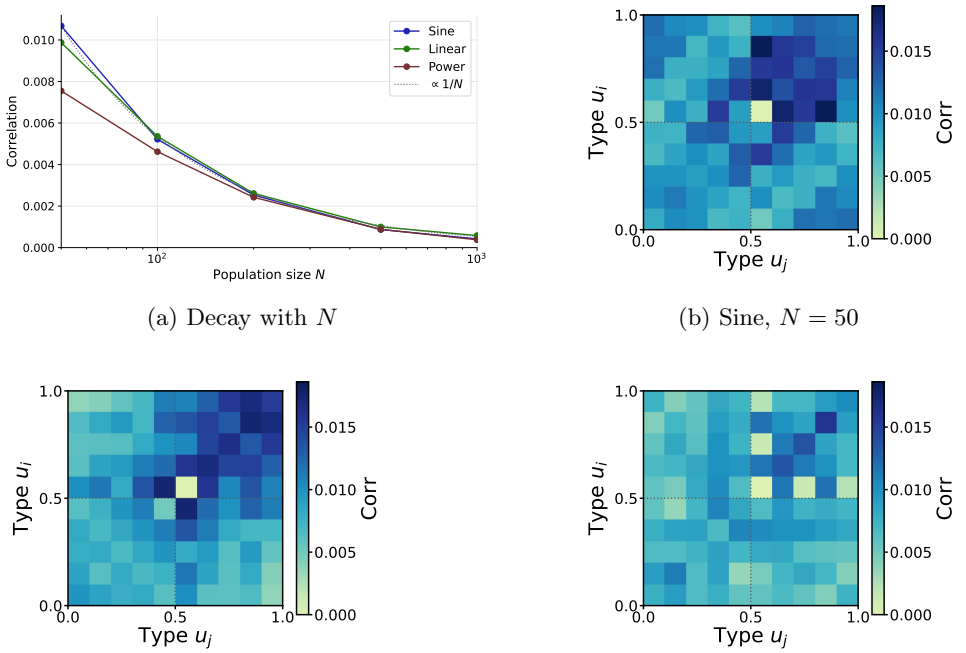

(a) Decay with N

(b) Sine, $N = 50$

(c) Linear, $N = 50$

(d) Power, $N = 50$

Fig. 7: Boundary-indicator correlation: (a) decay with population size N; (b)–(d) correlation by type pair for $N = 50$ under the sine, linear and power kernels.

furthest from the senders that spend time at the boundary. For the sine kernel it does not fall off, since the most influential senders are not the nearest ones. The size of $r(u)$ follows the size of the kernel. The sine and linear kernels deliver rates between about 1 and 2 on B, while the power kernel delivers between about 0.6 and 1. This ordering carries through the rest of the section.

Figure 6(c)–(d) compares a single realization at $N = 500$ with the solution of the forward equations (6) to (8), for the population boundary mass and for the mean level. The simulated curves follow the mean-field solution through the transient and settle near the predicted values, consistent with the finite-time convergence in Theorem 5. The figure also shows that the transient is over well before $t = 10$, which is why averaging over $[10, 25]$ gives a sensible estimate of the stationary quantities used later. The mean level separates the kernels in the way the reset rates predict. The power kernel, which transmits the least pressure, settles at the highest mean level, and the sine kernel at the lowest.

Figure 7(a) shows the correlation of the boundary indicators, averaged over pairs of agents of distinct type, against the population size. It decays like $1/N$ for all three kernels. Since the positions are fixed, this is the correlation conditional on the types, which is related to the convergence in Corollary 6. Agents interact only when one of

them is at the boundary, so this decay is the interaction itself washing out as the population grows.

Figures 7(b)– 7(d) resolve the same quantity by type pair at $N = 50$, the smallest population we consider and the one where the structure is easiest to see. For the sine and linear kernels the dependence is clearly strongest when both types lie in B, about half as strong again as for the other two blocks. A subcritical agent returns to the boundary under its own drift irrespective of the rest of the population, so relatively little of its behavior is shared with others. A non-subcritical agent has no autonomous positive-recurrent stationary law; reset pressure can nevertheless sustain its boundary occupancy, and that pressure is common to B. For the power kernel the three blocks are much closer together. That kernel transmits the least pressure, as the reset rate in Figure 6(a)–(b) show, so the common signal is weaker and the two groups behave more alike.

Figures 8(a)–8(c) show $\pi_k(u)$ over the type and level plane on a logarithmic color scale, at three population sizes together with the stationary solution of the limiting equations. Each finite-N panel comes from a single realization. At $N = 50$ a bin holds about five agents, so levels rarer than roughly 10^{-3} may never be observed. Such cells are drawn at the pale end of the color scale, so the pale upper-left region means too rare to resolve, not zero. It shrinks as N grows.

The geometric decay in k from Proposition 7 shows up as evenly spaced color bands, and the spacing $\log \alpha(u)$ changes continuously with position rather than only between the two regions, which is the effect of letting the load $\lambda(u)/\mu(u)$ change with position. The tails are heavier on B than on S_-, and heavier still for the power kernel than for the other two. This is the same ordering as in the reset rates and in the mean levels, and comes from the same place. On S_- the level is controlled by the agent's own drift. On B there is no such control and the level is held down only by resets, so a weaker kernel leaves a longer tail. The power panels are visibly darker at high k on the right hand side for exactly this reason. As N grows, the finite-N panels approach the limiting profile. This is consistent with Theorem 5, since the histograms are time averages over $t \in [10, 25]$, inside a finite horizon on which the limiting solution has already relaxed.

8 Discussion and Extensions

Endogenous boundary-generated resets provide a stabilization mechanism that is visible at three different levels. For finite directed graphs, recurrent source coordinates can initiate reset waves, yielding the source-component sufficient condition for positive recurrence. For two agents, the same mechanism appears in a stationary kernel equation with nonlocal boundary terms. For dense populations, sender boundary masses determine the reset pressure in a finite-time graphon mean-field limit.

The stationary results for the limiting equations are intentionally separate from that convergence theorem. Direct quantitative influence from subcritical types is sufficient for existence, and uniform influence controls aggregate first moments. Externally reset-isolated measurable sets suggest a sharper continuum criterion, but the present

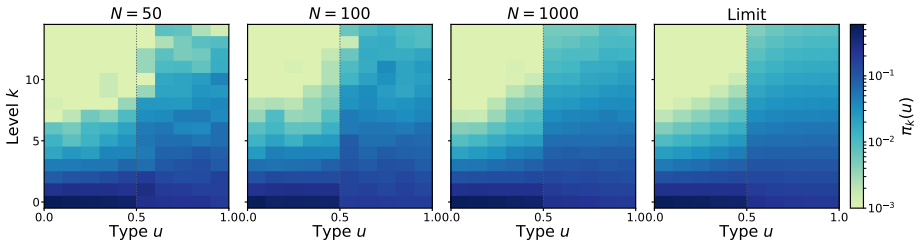

(a) Sine kernel

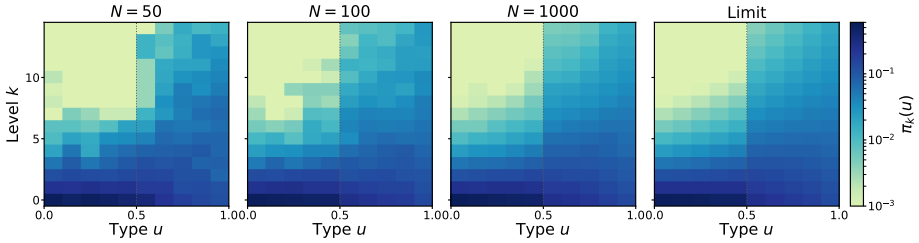

(b) Linear kernel

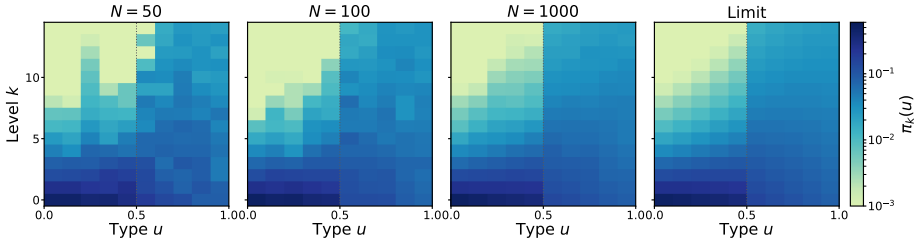

(c) Power kernel

Fig. 8: Stationary level distributions by type for the three reset kernels.

work does not prove that such a structural condition alone is sufficient. Nor does it prove general uniqueness or attraction of stationary fixed points.

The uniform supercritical example of Section 6.1 shows that a positive stationary fixed point of the limiting equations can coexist with transience of every finite system. Uniform-in-time approximation, escape-time asymptotics, and convergence of invariant measures require separate arguments.

Other open directions include the unclassified critical finite-source cases, a complete treatment of the critical two-agent kernel contour, sparse or time-dependent interaction graphs, and sharp conditions for uniqueness of the stationary graphon fixed point. Numerical priorities are independent validation of the two-agent formulas and reproducible error-controlled studies of finite-time mean-field convergence.

Declarations

Funding

Koen De Turck acknowledges funding from the Special Research Fund (BOF) of Ghent University, project BOF.STG.2022.0026.01

Ethics approval

This study is exempt from approval by an ethics committee, as it concerns purely mathematical research.

Appendix A Finite-N stability: proof of Proposition 1

We prove Proposition 1 by a regeneration argument. Throughout this appendix we suppress the superscript (N) and write

$$r_i(x) = \sum_{j=1}^{N} \Gamma_{ji} \mathbf{1}_{\{x_j=0\}}.$$

Recall that the reset graph has a directed edge $j \to i$ whenever $\Gamma_{ji} > 0$.

A.1 Basic properties of the Markov chain

The process is non-explosive. Indeed, its total rate of effective state changes is bounded uniformly in x by

$$q_* := \sum_{i=1}^{N} (\lambda_i + \mu_i) + \sum_{i=1}^{N} \sum_{j=1}^{N} \Gamma_{ij} < \infty.$$

Equivalently, one may construct the process from independent Poisson arrival, service, and reset-attempt clocks and dominate the number of effective jumps by a Poisson process of rate q_*. A reset attempt from j to i changes the state only when $X_j = 0$ and $X_i > 0$.

The chain is also irreducible on $\mathbb{Z}_{\geq 0}^N$. From every state, a finite sequence of service transitions reaches $\mathbf{0}$, and from $\mathbf{0}$ a finite sequence of arrival transitions reaches any prescribed state. All transitions in these paths have positive rates, and the bounded competing rates imply that each finite path has positive probability.

A.2 Graph structure

Let $\mathcal{C}$ be the set of strongly connected components of the reset graph. The condensation graph (the graph obtained by contracting each strongly connected component to a single vertex) on $\mathcal{C}$ is acyclic. A component $C \in \mathcal{C}$ is called a source component if no edge enters C from another component.

The externally reset-isolated-set condition of Proposition 1 is equivalent to requiring that each source component contain a subcritical index. Every source component is

itself externally reset-isolated. Conversely, suppose that a nonempty externally reset-isolated set A contains no subcritical index. The subgraph induced by A has a source strongly connected component C relative to A. No edge enters C from $A \setminus C$, and no edge enters A from A^c. Therefore no edge enters C from outside C, so C is a source component of the full graph with no subcritical index, a contradiction.

Choose, for each source component C, one subcritical index $s(C) \in C$. Let

$$S_0 := \{s(C) : C \text{ is a source component}\}.$$

By construction, $\lambda_s < \mu_s, \qquad s \in S_0$.

Moreover, every vertex $i \in \{1, \dots, N\}$ is reachable by a directed path from some $s \in S_0$. This follows because every component in a finite acyclic condensation graph is reachable from at least one source component, and within a strongly connected source component the chosen stabilizer can reach every other vertex of that component.

A.3 Return of the source stabilizers to zero

Define $B_0 := \{x \in \mathbb{Z}_{\geq 0}^N : x_s = 0 \text{ for all } s \in S_0\}$ and $\tau_{B_0} := \inf\{t \geq 0 : X(t) \in B_0\}$. We first show that $\mathbb{E}_x[\tau_{B_0}] < \infty$ for every initial state x.

Construct a family of independent autonomous birth–death processes $(Y_s(t))_{s \in S_0}$, where Y_s has birth rate λ_s, death rate μ_s, reflection at zero, and initial condition $Y_s(0) = x_s$. Couple Y_s and X_s using the same birth and death clocks, and let reset events act only on X_s. Then

$$X_s(t) \leq Y_s(t), \qquad t \geq 0.$$

Since $\lambda_s < \mu_s$, each Y_s is positive recurrent and hits zero in finite mean time. Because S_0 is finite, the product process $Y_{S_0}(t) = (Y_s(t))_{s \in S_0}$ is positive recurrent on $\mathbb{Z}_{\geq 0}^{S_0}$ and hits its all-zero state in finite mean time. Whenever $Y_s(t) = 0$ for all $s \in S_0$, the domination implies $X(t) \in B_0$. Therefore $\mathbb{E}_x[\tau_{B_0}] < \infty$ for every x.

A.4 A reset wave from B_0 to the all-zero state

We next show that, from every state in B_0, the process reaches the origin with a probability bounded below independently of the coordinates outside S_0.

Since every vertex is reachable from S_0, choose an ordering $i_1, \dots, i_m$ of the vertices outside S_0, together with parent vertices p_k, such that

$$p_k \in S_0 \cup \{i_1, \dots, i_{k-1}\}, \qquad p_k \to i_k.$$

Thus $\Gamma_{p_k i_k} > 0$. This is a directed spanning forest rooted at S_0. If $m = 0$, then $B_0 = \{\mathbf{0}\}$ and the conclusion is immediate; in this case set $p_* = 1$ and $T_* = 0$. Henceforth suppose $m \geq 1$.

Use the graphical construction from Section A.1, with an independent Poisson reset-attempt clock of rate Γ_{ji} for each edge $j \to i$. Starting from B_0, define $Z_{k-1} := S_0 \cup \{i_1, \dots, i_{k-1}\}$.

At the beginning of step k, all coordinates in Z_{k-1} are zero. If i_k is already zero, declare the step complete immediately. Otherwise, require the reset attempt $p_k \to i_k$ to occur before any arrival to a coordinate in Z_{k-1}. While the protected set remains at zero, the parent p_k is at zero, so its designated reset-attempt clock is active and has rate $\Gamma_{p_k i_k}$. Competing arrivals to the protected set have total rate $\Lambda_{k-1} := \sum_{\ell \in Z_{k-1}} \lambda_\ell$.

Other reset attempts and service completions can only help or are irrelevant, and events involving coordinates outside the protected set do not spoil the step. Hence, conditional on completion of the preceding steps, step k succeeds with probability at least

$$\frac{\Gamma_{p_k i_k}}{\Gamma_{p_k i_k} + \Lambda_{k-1}}.$$

After a successful step, i_k joins the protected zero set. Consequently, the probability of completing the entire reset wave and reaching $\mathbf{0}$ is at least

$$p_* := \prod_{k=1}^{m} \frac{\Gamma_{p_k i_k}}{\Gamma_{p_k i_k} + \Lambda_{k-1}} > 0. \tag{A1}$$

The bound is independent of the initial levels outside S_0.

A.5 Uniform return after a failed reset wave

A trial is one execution of the ordered construction in Section A.4. It succeeds if all m steps are completed. It fails at the first arrival to the current protected set Z_{k-1} that occurs before the designated reset attempt at a nontrivial step.

At a failure time, every source stabilizer is either still zero or is the coordinate that has just jumped from zero to one. Therefore $X_s \le 1, \qquad s \in S_0$. A failed arrival to a non-source coordinate leaves every source stabilizer at zero. No bound on the remaining coordinates is needed because membership in B_0 depends only on S_0.

Lemma 13 (Uniform return after failure) *There exists a constant $M < \infty$ such that*

$$\sup_{\{x:\, x_s \le 1 \ \forall s \in S_0\}} \mathbb{E}_x[\tau_{B_0}] \le M.$$

Proof Use the dominating product process Y_{S_0} constructed in Section A.3. If $x_s \le 1$ for every $s \in S_0$, then Y_{S_0} starts in the finite set $\{0,1\}^{S_0}$. Its expected hitting time of the all-zero state is finite from each state in this set because the product process is positive recurrent. Let M be the maximum of these finitely many expectations. Since $X_s(t) \le Y_s(t)$, the same uniform bound holds for the hitting time of B_0 by X. □

The duration of a trial also has a uniform conditional mean bound. At a nontrivial step k, the minimum of the designated reset-attempt clock and the arrival clocks from Z_{k-1} is exponential with rate $\Gamma_{p_k i_k} + \Lambda_{k-1}$. Other clock rings do not delay these clocks;

a step with $i_k = 0$ has duration zero. Hence

$$\sup_{x \in B_0} \mathbb{E}_x[\text{duration of one trial}] \le T_* := \sum_{k=1}^{m} \frac{1}{\Gamma_{p_k i_k} + \Lambda_{k-1}} < \infty. \tag{A2}$$

A.6 Regeneration and positive recurrence

Let $\tau_{\mathbf{0}} := \inf\{t \ge 0 : X(t) = \mathbf{0}\}$. By Section A.3, the process reaches B_0 in finite mean time from every initial state. Starting whenever the process enters B_0, perform the trial of Section A.4. If it fails, wait until the process returns to B_0 and begin a new trial. The strong Markov property, (A1), Lemma 13, and (A2) apply conditionally at every repetition.

Let K be the number of trials begun up to and including the first successful one. No independence of the trials is required: since every conditional success probability is at least p_*,

$$\mathbb{P}_x(K \ge n) \le (1 - p_*)^{n-1}, \qquad n \ge 1.$$

Thus $\mathbb{E}_x[K] \le 1/p_*$ and $\mathbb{E}_x[K - 1] \le (1 - p_*)/p_*$. Summing the conditional mean bounds for the trial durations and the returns following failed trials gives

$$\begin{aligned} \mathbb{E}_x[\tau_{\mathbf{0}}] &\le \mathbb{E}_x[\tau_{B_0}] + \frac{T_*}{p_*} + \frac{1 - p_*}{p_*} M \\ &= \mathbb{E}_x[\tau_{B_0}] + \left(\frac{1}{p_*} - 1\right)(T_* + M) + T_* < \infty. \end{aligned} \tag{A3}$$

It remains to verify finite mean return time. Starting from $\mathbf{0}$, the first effective transition is an arrival and occurs after an exponential time with rate $\sum_i \lambda_i > 0$. The resulting state belongs to the finite set $\{\boldsymbol{e}_1, \ldots, \boldsymbol{e}_N\}$, and (A3) gives a finite expected hitting time of $\mathbf{0}$ from every state in that set. Hence $\mathbf{0}$ has finite expected return time. By Section A.1, the chain is irreducible and non-explosive. An irreducible countable-state continuous-time Markov chain with a state of finite mean return time is positive recurrent. This completes the proof of Proposition 1.

Appendix B Exact Analysis of the Two-Agent Model

This appendix contains the detailed derivations underlying the exact stationary analysis of the two-agent endogenous reset model presented in Section 3. The analysis is based on the kernel method for random walks in the quarter plane and follows an approach that is comparable to [31], but modified to accommodate the endogenous reset terms that arise in the present model.

B.1 Balance equations and generating function

Let $P(x, y)$ denote the joint probability generating function of the stationary distribution $(\pi_{i,j})_{i,j \ge 0}$ as defined in Section 3 and consider the corresponding balance

equations:

$$
\begin{aligned}
(\lambda_1+\lambda_2+\mu_1+\mu_2)\pi_{i,j} &= \mu_1\pi_{i+1,j}+\mu_2\pi_{i,j+1}+\lambda_1\pi_{i-1,j}+\lambda_2\pi_{i,j-1}, && i\geq 1, j\geq 1;\\
(\lambda_1+\lambda_2+\mu_2+\gamma_1)\pi_{0,j} &= \mu_1\pi_{1,j}+\mu_2\pi_{0,j+1}+\lambda_2\pi_{0,j-1}, && j\geq 1;\\
(\lambda_1+\lambda_2+\mu_1+\gamma_2)\pi_{i,0} &= \mu_1\pi_{i+1,0}+\mu_2\pi_{i,1}+\lambda_1\pi_{i-1,0}, && i\geq 1;\\
(\lambda_1+\lambda_2)\pi_{0,0} &= \mu_1\pi_{1,0}+\mu_2\pi_{0,1}+\gamma_2\sum_{k=1}^{\infty}\pi_{k,0}+\gamma_1\sum_{k=1}^{\infty}\pi_{0,k}.
\end{aligned}
$$

Multiplying the balance equations by $x^i y^j$ and summing over all admissible states yields the functional equation

$$K(x,y)P(x,y)=A(x,y), \tag{B4}$$

where

$$K(x,y)=\lambda_1+\lambda_2+\mu_1+\mu_2-\lambda_1 x-\lambda_2 y-\frac{\mu_1}{x}-\frac{\mu_2}{y},$$

and

$$A(x,y)=\Big(\mu_1-\gamma_1-\frac{\mu_1}{x}\Big)P(0,y)+\Big(\mu_2-\gamma_2-\frac{\mu_2}{y}\Big)P(x,0)+\gamma_1 P(0,1)+\gamma_2 P(1,0).$$

This proves Theorem 4 and sets up the functions $K(x,y)$ and $A(x,y)$ which are used in the solution of this equation.

Throughout the construction we impose the scope of Section 3.4: $\lambda_i,\mu_i>0$, $\gamma_i\geq 0$, the sufficient stability condition, $\lambda_1<\mu_1$, $\lambda_2\neq\mu_2$, and $\gamma_1+\gamma_2>0$.

B.2 Kernel equation and reduced boundary relation

For fixed $x\in\mathbb{C}\setminus\{0\}$, the kernel equation $K(x,y)=0$ has two solutions $y_-(x)$ and $y_+(x)$ given by

$$y_\pm(x)=\frac{\lambda_1+\lambda_2+\mu_1+\mu_2-\lambda_1 x-\frac{\mu_1}{x}\pm\sqrt{\Delta(x)}}{2\lambda_2},$$

where

$$\Delta(x)=\Big(\lambda_1+\lambda_2+\mu_1+\mu_2-\lambda_1 x-\frac{\mu_1}{x}\Big)^2-4\lambda_2\mu_2.$$

Whenever $(x,y_\pm(x))$ lies in the domain of analyticity of P, vanishing of the kernel forces $A(x,y_\pm(x))=0$. Hence, for $y=y_\pm(x)$,

$$\Big(\mu_1-\gamma_1-\frac{\mu_1}{x}\Big)P(0,y_\pm(x))+\Big(\mu_2-\gamma_2-\frac{\mu_2}{y_\pm(x)}\Big)P(x,0)+\gamma_1 P(0,1)+\gamma_2 P(1,0)=0. \tag{B5}$$

B.3 Region of convergence

Write $p_i^{(1)} := \mathbb{P}(X_1 = i) = \sum_{j=0}^{\infty} \pi_{i,j}$. These marginal probabilities satisfy

$$p_i^{(1)} \leq (\lambda_1/\mu_1)^i p_0^{(1)}.$$

Indeed, applying the stationary generator identity to the bounded indicator of the upper set $\{X_1 \geq i\}$, for $i \geq 1$, gives the level-crossing relation

$$\mu_1 \mathbb{P}(X_1 = i) - \lambda_1 \mathbb{P}(X_1 = i-1) = -\gamma_2 \sum_{k=i}^{\infty} \pi_{k,0} \leq 0.$$

Iteration yields the claimed bound.

Consequently, for any (x, y) with $|x| < \frac{\mu_1}{\lambda_1}$ and $|y| \leq 1$

$$\sum_{i=0}^{\infty}\sum_{j=0}^{\infty} |\pi_{i,j} x^i y^j| \leq \sum_{i=0}^{\infty} \mathbb{P}(X_1 = i)|x|^i \leq \mathbb{P}(X_1 = 0) \sum_{i=0}^{\infty} \left(\frac{\lambda_1}{\mu_1}|x|\right)^i < \infty.$$

Thus P converges absolutely for $|x| < \mu_1/\lambda_1$ and $|y| \leq 1$. It is jointly analytic in the open region $|x| < \mu_1/\lambda_1$, $|y| < 1$, and continuous up to $|y| = 1$ locally uniformly in x in that region. This does not assert joint analyticity across the unit-circle boundary. If $\lambda_2 < \mu_2$, the symmetric level-crossing argument gives absolute convergence for $|y| < \mu_2/\lambda_2$ and $|x| \leq 1$, and analyticity in its interior. We use the reduced kernel equation only at points in these convergence domains, with boundary values obtained by continuity.

B.4 Integral representation for $P(x, 0)$

Assume $\lambda_1 < \mu_1$ and define $\tau_1 = \sqrt{\mu_1/\lambda_1}$. On $|x| = \tau_1$, writing $x = \tau_1 e^{i\theta}$, the kernel equation $K(x, y) = 0$ has real coefficients and a unique root in $(0, 1)$, which we denote by $y(\theta)$:

$$y(\theta) = \frac{\lambda_1 + \lambda_2 + \mu_1 + \mu_2 - 2\sqrt{\lambda_1\mu_1}\cos\theta - \sqrt{(\lambda_1 + \lambda_2 + \mu_1 + \mu_2 - 2\sqrt{\lambda_1\mu_1}\cos\theta)^2 - 4\lambda_2\mu_2}}{2\lambda_2}.$$

To check the contour domain, set $A_1(\theta) = \sum_{i=1}^{2}(\lambda_i + \mu_i) - 2\sqrt{\lambda_1\mu_1}\cos\theta$. Then

$$A_1(\theta) \geq \lambda_2 + \mu_2 + (\sqrt{\mu_1} - \sqrt{\lambda_1})^2 > \lambda_2 + \mu_2.$$

The quadratic in y is positive at zero and negative at one, so its smaller root satisfies $0 < y(\theta) < 1$. Since $1 < \tau_1 < \mu_1/\lambda_1$, the contour points lie in the analytic domain just established. Moreover, $\mu_2 - \gamma_2 - \mu_2/y(\theta) < 0$, even when $\gamma_2 = 0$.

Now notice that $y(-\theta) = y(\theta)$. Combining the two relations corresponding to $x = \tau_1 e^{\pm i\theta}$ in (B5) and eliminating $P(0, y(\theta))$ from these two equations yields

$$(\mu_1-\gamma_1-\lambda_1 x)P(x,0)-(\mu_1-\gamma_1-\lambda_1\bar{x})P(\bar{x},0) = \mathrm{i}\,\frac{2(\gamma_1 P(0,1)+\gamma_2 P(1,0))\sqrt{\lambda_1\mu_1}\sin\theta}{\mu_2-\gamma_2-\mu_2/y(\theta)}\,.$$

Writing

$$(\mu_1-\gamma_1-\lambda_1 x)P(x,0) = \sum_{m=0}^{\infty} c_m x^m,$$

one obtains

$$\sum_{m=0}^{\infty} c_m \tau_1^m \sin(m\theta) = \frac{(\gamma_1 P(0,1)+\gamma_2 P(1,0))\sqrt{\lambda_1\mu_1}\sin\theta}{\mu_2-\gamma_2-\mu_2/y(\theta)}.$$

Using orthogonality of sine functions and summing the resulting series yields the integral representation

$$\begin{aligned}(\mu_1-\gamma_1-\lambda_1 z)P(z,0) &= (\mu_1-\gamma_1)P(0,0)\\ &+\frac{2}{\pi}(\gamma_1 P(0,1)+\gamma_2 P(1,0))\mu_1 z\int_0^{\pi}\frac{\sin^2\theta}{(\mu_2-\gamma_2-\mu_2/y(\theta))(\tau_1^2+z^2-2\tau_1 z\cos\theta)}\,d\theta,\end{aligned} \tag{B6}$$

valid for $|z| < \tau_1$.

B.5 Representation for $P(0, y)$

For the symmetric construction, set $\tau_2 = \sqrt{\mu_2/\lambda_2}$ and

$$x(\theta) = \frac{A_2(\theta)-\sqrt{A_2(\theta)^2-4\lambda_1\mu_1}}{2\lambda_1}, \quad A_2(\theta) := \sum_{i=1}^{2}(\lambda_i+\mu_i) - 2\sqrt{\lambda_2\mu_2}\cos\theta. \tag{B7}$$

Since $\lambda_2 \neq \mu_2$, $A_2(\theta) \geq \lambda_1+\mu_1+(\sqrt{\mu_2}-\sqrt{\lambda_2})^2 > \lambda_1+\mu_1$, and hence $0 < x(\theta) < 1$. If $\lambda_2 < \mu_2$, the contour $|y| = \tau_2 > 1$ lies in the symmetric extended convergence domain. If $\lambda_2 > \mu_2$, it lies inside the unit disc. Thus both contour constructions are justified in the stated noncritical scope, and $\mu_1-\gamma_1-\mu_1/x(\theta) < 0$ even when $\gamma_1 = 0$.

The same conjugate-point elimination and Fourier argument now yield, for $|z| < \tau_2$,

$$\begin{aligned}(\mu_2-\gamma_2-\lambda_2 z)P(0,z) &= (\mu_2-\gamma_2)P(0,0)\\ &+\frac{2}{\pi}(\gamma_1 P(0,1)+\gamma_2 P(1,0))\mu_2 z\int_0^{\pi}\frac{\sin^2\theta}{(\mu_1-\gamma_1-\mu_1/x(\theta))(\tau_2^2+z^2-2\tau_2 z\cos\theta)}\,d\theta,\end{aligned} \tag{B8}$$

where $x(\theta)$ is given by (B7). For either contour, $\tau_i^2+z^2-2\tau_i z\cos\theta = (z-\tau_i e^{i\theta})(z-\tau_i e^{-i\theta})$ has no zero when $|z| < \tau_i$, so the integral defines an analytic function on that disc.

B.6 Removable zeros of the boundary prefactors

To express the constants systems compactly, define

$$q_i(z) := \mu_i - \gamma_i - \lambda_i z, \qquad a_i := \mu_i - \gamma_i, \qquad C_\gamma := \gamma_1 P(0,1) + \gamma_2 P(1,0), \tag{B9}$$

$$D_2(\theta) := \mu_2 - \gamma_2 - \mu_2/y(\theta), \quad D_1(\theta) := \mu_1 - \gamma_1 - \mu_1/x(\theta). \tag{B10}$$

For $i = 1, 2$, set

$$H_i(z) := \frac{2\mu_i z}{\pi} \int_0^\pi \frac{\sin^2\theta}{D_{3-i}(\theta)(\tau_i^2 + z^2 - 2\tau_i z\cos\theta)}\, d\theta, \qquad |z| < \tau_i. \tag{B11}$$

The denominator bounds above justify differentiation under the integral:

$$H_i'(z) = \frac{2\mu_i}{\pi} \int_0^\pi \frac{(\tau_i^2 - z^2)\sin^2\theta}{D_{3-i}(\theta)(\tau_i^2 + z^2 - 2\tau_i z\cos\theta)^2}\, d\theta. \tag{B12}$$

Writing $P_1(z) := P(z,0)$ and $P_2(z) := P(0,z)$, the already derived identities (B6)–(B8) become

$$q_i(z)P_i(z) = a_i P(0,0) + C_\gamma H_i(z), \qquad |z| < \tau_i. \tag{B13}$$

These are identities for the products on their left-hand sides; they must not be divided by $q_i(z)$ at a zero of that factor. If $\zeta_i := a_i/\lambda_i$ lies in the corresponding disc, analyticity of the boundary pgf forces

$$a_i P(0,0) + C_\gamma H_i(\zeta_i) = 0. \tag{B14}$$

The apparent singularity is then removable, with value

$$P_i(\zeta_i) = -\frac{C_\gamma}{\lambda_i} H_i'(\zeta_i).$$

In particular, when $a_i = 0$, the value equation at zero is trivial, but its derivative gives the nontrivial constraint

$$-\lambda_i P(0,0) = C_\gamma H_i'(0). \tag{B15}$$

Thus $\gamma_i = \mu_i$ does not create a pole in the stationary pgf; it requires a derivative compatibility equation to recover its value at zero.

B.7 Determination of constants

Introduce the unknowns

$$A := P(0,1), \quad B := P(1,0), \quad P_{00} := P(0,0), \quad D_x := P_x(1,0), \quad D_y := P_y(0,1),$$

so that $C_\gamma = \gamma_1 A + \gamma_2 B$. Since $\tau_1 > 1$, the first boundary pgf and H_1 are analytic at 1. Substituting $y = 1$ into (2), taking the first-order limit as $x \to 1$, and using $P(1,1) = 1$ gives

$$\mu_1 A - \gamma_2 D_x = \mu_1 - \lambda_1. \tag{B16}$$

Differentiating (B13) for $i = 1$ at $z = 1$ yields

$$-\lambda_1 B + q_1(1) D_x = C_\gamma H_1'(1), \tag{B17}$$

while evaluation at $z = 1$ gives

$$q_1(1) B = a_1 P_{00} + C_\gamma H_1(1). \tag{B18}$$

All these relations are linear in the unknowns. They remain valid if $q_1(1) = 0$, since no division by that factor has been used.

Subcritical second coordinate: $\lambda_2 < \mu_2$**.**

Here $\tau_2 > 1$, so the symmetric limits are legitimate and give

$$\mu_2 B - \gamma_1 D_y = \mu_2 - \lambda_2, \tag{B19}$$

$$-\lambda_2 A + q_2(1) D_y = C_\gamma H_2'(1), \tag{B20}$$

$$q_2(1) A = a_2 P_{00} + C_\gamma H_2(1). \tag{B21}$$

Equations (B16), (B17), (B19), and (B20) constitute an explicit four-by-four system for A, B, D_x, D_y. Once these are determined, either value equation (B18) or (B21) with $a_i \neq 0$ gives P_{00}. If $a_1 = a_2 = 0$, use (B15) instead. The other value and removable-zero relations must also hold and provide compatibility checks.

Supercritical second coordinate: $\lambda_2 > \mu_2$**.**

Put

$$y_0 := \frac{\mu_2}{\lambda_2} = \tau_2^2 < \tau_2, \qquad Q := P(0, y_0).$$

Although 1 lies outside the disc of (B8), y_0 is inside it. At $x = 1$, the kernel vanishes at y_0. The finite pgf value there therefore forces the numerator in (B4) to vanish:

$$-\gamma_1 Q + (\mu_2 - \lambda_2) B + \gamma_1 A = 0. \tag{B22}$$

Evaluating (B13) for $i = 2$ at y_0 gives

$$-\gamma_2 Q = a_2 P_{00} + C_\gamma H_2(y_0). \tag{B23}$$

Together, (B16), (B17), (B18), (B22), and (B23) form a five-by-five linear system with unknown vector

$$(A, B, D_x, Q, P_{00}).$$

If $a_1 = a_2 = 0$, P_{00} disappears from both value equations and hence from this entire system. In that case augment it with (B15); retain the value equations as compatibility conditions rather than discarding them without a rank check. If exactly one a_i vanishes, the other value equation still contains P_{00}, and the removable-origin relation remains an additional check. For example, the stable parameters $(\lambda_1, \mu_1, \lambda_2, \mu_2, \gamma_1, \gamma_2) = (8, 10, 5, 3, 10, 3)$ have $a_1 = a_2 = 0$ and require such an additional relation.

Reconstruction and scope of the linear systems.

When the relevant coefficient matrix is nonsingular, it directly determines the constants. If singular, these equations remain necessary relations and may be augmented by the unused value and removable-zero constraints. We do not claim that a nonsingular closure has been proved for every parameter choice in (3).

The boundary identities determine the boundary pgfs on their stated discs, with removable values supplied by the derivative formulas. Any continuation needed in the full unit bidisc must remain analytic and agree with the constants $P(0,1)$ and $P(1,0)$. Once both boundary pgfs are fixed, (B4) recovers $P(x,y)$ away from the kernel, and kernel values are supplied by continuity. A formal algebraic candidate is a stationary pgf only if this reconstruction is normalized, has nonnegative coefficients, and satisfies the full functional equation and its boundary identities. Irreducibility and positive recurrence guarantee uniqueness of such a stationary pgf, not nonsingularity of every proposed constants system.

Appendix C Graphon mean-field limit: proof of Theorem 5

We prove the finite-time empirical-measure limit by a martingale-problem argument. The model-specific point is continuity of the reset drift: it is a quadratic functional of the empirical measure, and its continuity follows from the a.e.-continuity of the kernel and the Lebesgue type marginal of the limit. We then prove Corollary 6 by coupling.

C.1 Setup and notation

Let $E := [0,1] \times \mathbb{Z}_{\geq 0}$, where the level coordinate has the discrete topology, and equip $\mathcal{P}(E)$ with the topology of weak convergence. We work under the hypotheses of Theorem 5: Γ, λ and μ are bounded and Lebesgue-a.e. continuous.

For each N, the types $U_1, \ldots, U_N$ are i.i.d. Unif$[0,1]$ and $\nu_t^N = N^{-1} \sum_{i=1}^N \delta_{(U_i, X_i^N(t))}$. The initial conditions satisfy $\nu_0^N \Rightarrow \nu_0$ in probability, with $\nu_0(du, dk) = du\, p_k(0,u)$, together with the moment bound

$$\sup_{N \geq 1} \mathbb{E}\left[\frac{1}{N} \sum_{i=1}^N X_i^N(0)\right] < \infty. \tag{C24}$$

For $\nu \in \mathcal{P}(E)$, let ν_0 denote its boundary marginal, $\nu_0(B) := \nu(B \times \{0\})$, and set

$$r_\nu(u) := \int_{[0,1]} \Gamma(v,u)\, \nu_0(dv). \tag{C25}$$

For a bounded test function $\varphi : E \to \mathbb{R}$, define

$$\begin{aligned}(\mathcal{A}_\nu \varphi)(u,k) = {} & \lambda(u)[\varphi(u,k+1) - \varphi(u,k)] \\ & + \mu(u)\mathbf{1}_{\{k>0\}}[\varphi(u,k-1) - \varphi(u,k)] \\ & + r_\nu(u)\mathbf{1}_{\{k>0\}}[\varphi(u,0) - \varphi(u,k)].\end{aligned} \tag{C26}$$

C.2 Well-posedness of the limiting equations

Consider the Banach space

$$\mathcal{B} := L^1\big([0,1]; \ell^1(\mathbb{Z}_{\geq 0})\big), \qquad \|p\|_{\mathcal{B}} := \int_0^1 \sum_{k\geq 0} |p_k(u)|\, du.$$

For $p, q \in \mathcal{B}$, the associated reset rates satisfy

$$\|r_p - r_q\|_{L^1([0,1])} \leq \|\Gamma\|_\infty \|p - q\|_{\mathcal{B}}, \qquad \|r_p\|_\infty \leq \|\Gamma\|_\infty \tag{C27}$$

whenever p is probability-valued. Let $F(p)$ denote the right-hand side of (7)–(8). The shift operators corresponding to births and deaths have ℓ^1 operator norm at most two. For the reset term, adding and subtracting the reset operator with rate r_p applied to q and using (C27) gives, for probability-valued p and q,

$$\|F(p) - F(q)\|_{\mathcal{B}} \leq C\|p - q\|_{\mathcal{B}}, \qquad C := 2\|\lambda\|_\infty + 2\|\mu\|_\infty + 4\|\Gamma\|_\infty. \tag{C28}$$

More generally, F is locally Lipschitz on $\mathcal{B}$, so the Picard theorem gives a unique local solution. The positive probability simplex is invariant: summing over levels gives zero total mass derivative, while positivity follows from the forward-equation form with the nonnegative time-dependent rate $r_p(t,u)$. The solution therefore remains in the probability simplex. On this simplex, the bounded rates give a uniform bound on $\|F(p)\|_{\mathcal{B}}$, preventing finite-time blow-up. This gives a unique global probability-valued solution in $C([0,T], \mathcal{B})$ for every $T < \infty$.

C.3 Empirical martingale problem

Let $\mathcal{D}$ be a countable convergence-determining class of bounded continuous functions on E with finite support in the level coordinate. For $\varphi \in \mathcal{D}$, set

$$\begin{aligned}B_\varphi(\nu) := {} & \int_E \lambda(u)[\varphi(u,k+1) - \varphi(u,k)]\, \nu(du,dk) \\ & + \int_E \mu(u)\mathbf{1}_{\{k>0\}}[\varphi(u,k-1) - \varphi(u,k)]\, \nu(du,dk)\end{aligned}$$

$$+ \iint_{E\times E} \Gamma(v,u)\mathbf{1}_{\{\ell=0\}}\mathbf{1}_{\{k>0\}}[\varphi(u,0)-\varphi(u,k)] \times \nu(dv,d\ell)\nu(du,dk). \tag{C29}$$

Thus $B_\varphi(\nu) = \langle \nu, \mathcal{A}_\nu \varphi\rangle$. A direct computation with the finite-N generator shows that

$$M_t^{N,\varphi} := \langle \nu_t^N, \varphi\rangle - \langle \nu_0^N, \varphi\rangle - \int_0^t B_\varphi(\nu_s^N)\, ds \tag{C30}$$

is a martingale. In the reset part of (C29), the product measure expands into the required double sum over senders and receivers. Its diagonal terms vanish because $\mathbf{1}_{\{X_i^N=0\}}\mathbf{1}_{\{X_i^N>0\}} = 0$. Consequently, the convention $\Gamma_{ii}^{(N)} = 0$ introduces no correction in the empirical drift.

Each effective jump changes $\langle \nu_t^N, \varphi\rangle$ by at most $2\|\varphi\|_\infty/N$. The total system jump intensity is at most

$$N\big(\|\lambda\|_\infty + \|\mu\|_\infty + \|\Gamma\|_\infty\big).$$

It follows that

$$\mathbb{E}\big[\langle M^{N,\varphi}\rangle_T\big] \le \frac{C_\varphi T}{N}. \tag{C31}$$

By Doob's inequality,

$$\mathbb{E}\left[\sup_{t\le T} |M_t^{N,\varphi}|^2\right] \longrightarrow 0. \tag{C32}$$

C.4 Moment bounds and tightness

Let $A_i(T)$ be the number of arrivals to particle i by time T. Since services and resets only decrease levels,

$$\sup_{t\le T} X_i^N(t) \le X_i^N(0) + A_i(T).$$

Therefore (C24) implies

$$\sup_{N\ge 1}\mathbb{E}\left[\sup_{t\le T}\int_E k\,\nu_t^N(du,dk)\right] \le \sup_{N\ge 1}\mathbb{E}\left[\frac{1}{N}\sum_{i=1}^N X_i^N(0)\right] + T\|\lambda\|_\infty < \infty. \tag{C33}$$

Because the type space is compact, the sublevel sets of $\nu \mapsto \int_E k\, d\nu$ are compact in $\mathcal{P}(E)$. Markov's inequality and (C33) consequently give compact containment.

For every $\varphi \in \mathcal{D}$, the drift B_φ is uniformly bounded, while (C31) controls the martingale increments. The Aldous–Rebolledo criterion therefore gives tightness of the real-valued processes $\langle \nu^N, \varphi\rangle$. Compact containment and a standard measure-valued tightness criterion then give tightness of $\{\nu^N\}$ in $D([0,T],\mathcal{P}(E))$. Finally, every jump of ν^N has bounded-Lipschitz size at most $2/N$, so every limit point has continuous paths.

C.5 Continuity of the nonlinear drift

Lemma 14 (Continuity at measures with Lebesgue type marginal) *Let $\eta^n \Rightarrow \eta$ in $\mathcal{P}(E)$, and suppose that the type marginal of η is Lebesgue measure. Then, for every $\varphi \in \mathcal{D}$,*

$$B_\varphi(\eta^n) \longrightarrow B_\varphi(\eta).$$

Proof Consider first the birth integrand

$$(u,k) \longmapsto \lambda(u)[\varphi(u,k+1) - \varphi(u,k)].$$

It is bounded, and its discontinuity set is contained in $D_\lambda \times \mathbb{Z}_{\geq 0}$, where D_λ is the discontinuity set of λ. This set has η-measure zero because the type marginal of η is Lebesgue measure. The extended portmanteau theorem therefore gives convergence of the birth term. The same argument applies to the death term.

For the reset term, $\eta^n \otimes \eta^n \Rightarrow \eta \otimes \eta$. Its integrand is

$$((v,\ell),(u,k)) \longmapsto \Gamma(v,u)\mathbf{1}_{\{\ell=0\}}\mathbf{1}_{\{k>0\}}[\varphi(u,0) - \varphi(u,k)].$$

The level indicators are continuous because the level space is discrete. The remaining discontinuities are contained in the discontinuity set of Γ (up to the empty discontinuity set of φ). The type-pair marginal of $\eta \otimes \eta$ is two-dimensional Lebesgue measure, so this set has $\eta \otimes \eta$-measure zero. A second application of the extended portmanteau theorem proves convergence of the reset term. □

C.6 Identification of limit points

Let $\nu^{N_m} \Rightarrow \nu$ along a subsequence. The type marginal of ν_t^N is the time-independent empirical type measure $N^{-1}\sum_i \delta_{U_i}$, which converges almost surely to Lebesgue measure. Thus every ν_t has Lebesgue type marginal.

Using a Skorokhod representation for the convergent subsequence, we may pass to the limit almost surely. The limit path is continuous, so $\nu_s^{N_m} \Rightarrow \nu_s$ for every s. Lemma 14 and the uniform boundedness of B_φ then give, by dominated convergence,

$$\int_0^t B_\varphi(\nu_s^{N_m})\,ds \longrightarrow \int_0^t B_\varphi(\nu_s)\,ds.$$

Together with (C32), this shows that every limit point satisfies

$$\langle \nu_t, \varphi\rangle = \langle \nu_0, \varphi\rangle + \int_0^t B_\varphi(\nu_s)\,ds, \qquad \varphi \in \mathcal{D}. \tag{C34}$$

Since the type marginal is Lebesgue measure, disintegration gives $\nu_t(du,dk) = du\,p_k(t,u)$. Taking test functions of the form $\varphi(u,\ell) = g(u)\mathbf{1}_{\{\ell=k\}}$, first for g in a countable dense subset of $C([0,1])$ and then by approximation, shows that (C34) is equivalent to (6)–(8) in the $L^1([0,1];\ell^1)$ sense. Since $F(p(s))$ is uniformly bounded in $\mathcal{B}$, this is the Bochner integral form of the limiting equation and $p \in C([0,T],\mathcal{B})$.

C.7 Uniqueness and conclusion of the mean-field proof

Let p and $\widetilde{p}$ be two solutions with the same initial condition, and write $D(t) := \|p(t) - \widetilde{p}(t)\|_{\mathcal{B}}$

The integral form of the equations and (C28) imply $D(t) \le C \int_0^t D(s)\, ds$

Gronwall's inequality gives $D(t) = 0$. Thus every subsequential limit is the same deterministic solution constructed in the well-posedness subsection. Tightness therefore implies convergence of the full sequence in distribution in $D([0,T], \mathcal{P}(E))$. Because the limit is deterministic, this convergence is in probability. This completes the proof of Theorem 5.

C.8 Typed propagation of chaos

We now impose the additional initial-independence assumption of Corollary 6. For each i, construct a nonlinear process $\overline{X}_i$ with type U_i, initial state $\overline{X}_i(0) = X_i^N(0)$, and deterministic reset rate $r(t, U_i)$ from the mean-field solution. Use independent driving clocks across i, so the augmented processes $(U_i, \overline{X}_i)$ are i.i.d. Couple X_i^N and $\overline{X}_i$ by using the same birth and service clocks and a maximal coupling of their reset clocks. Let,

$$\tau_i^N := \inf\{t \ge 0 : X_i^N(t) \ne \overline{X}_i(t)\}, \qquad e_N(t) := \frac{1}{N} \sum_{i=1}^N \mathbb{P}(\tau_i^N \le t).$$

Before τ_i^N, only an unmatched reset can separate the two coordinates. Writing

$$Z_i^N(t) := \frac{1}{N} \sum_{j \ne i} \Gamma(U_j, U_i) \mathbf{1}_{\{\overline{X}_j(t) = 0\}},$$

we have

$$|r_i^N(t) - r(t, U_i)| \le \frac{\|\Gamma\|_\infty}{N} \sum_{j \ne i} \mathbf{1}_{\{X_j^N(t) \ne \overline{X}_j(t)\}} + |Z_i^N(t) - r(t, U_i)|. \tag{C35}$$

Conditional on U_i, the summands in $Z_i^N(t)$ are independent and bounded, and their common mean after averaging over the sender type is $r(t, U_i)$. Accounting for the omitted $j = i$ term and using the variance bound gives, uniformly in i and t,

$$\mathbb{E}|Z_i^N(t) - r(t, U_i)| \le \frac{\|\Gamma\|_\infty}{\sqrt{N}} + \frac{\|\Gamma\|_\infty}{N}. \tag{C36}$$

Integrating the unmatched-reset intensity in (C35), averaging over i, and using $\mathbf{1}_{\{X_j^N(t) \ne \overline{X}_j(t)\}} \le \mathbf{1}_{\{\tau_j^N \le t\}}$ yields

$$e_N(t) \le \|\Gamma\|_\infty \int_0^t e_N(s)\, ds + t\|\Gamma\|_\infty \left(N^{-1/2} + N^{-1}\right).$$

Hence Gronwall's inequality implies $e_N(T) \leq C_T N^{-1/2}$. The coupled system is exchangeable in the particle labels, so for fixed m,

$$\mathbb{P}\left(\exists i \leq m,\ \exists t \leq T : X_i^N(t) \neq \overline{X}_i(t)\right) \leq m e_N(T) \longrightarrow 0.$$

Since the augmented nonlinear processes are independent, this proves Corollary 6.

C.9 Relation to existing results

The argument follows the standard martingale-problem route for mean-field jump processes with bounded rates. General propagation-of-chaos results of this kind can be found in Sznitman [20] and Graham [18], and limits of interacting systems on dense graphs and large networks in Medvedev [23] and Delattre, Fournier and Hoffmann [24]. The model-specific step is Lemma 14: writing the reset drift as a double integral makes the a.e.-continuity of the graphon kernel sufficient for weak convergence. The coupling proof above supplies the additional initial-chaos hypothesis needed for tagged-particle convergence.

References

[1] Castellano, C., Fortunato, S., Loreto, V.: Statistical physics of social dynamics. Reviews of Modern Physics **81**(2), 591–646 (2009) https://doi.org/10.1103/RevModPhys.81.591

[2] Chaudhary, S., Fiems, D., De Turck, K.: Trust in crisis: A mean-field model of social dynamics. Physica A: Statistical Mechanics and its Applications **698**, 131803 (2026) https://doi.org/10.1016/j.physa.2026.131803

[3] Watts, D.J.: A simple model of global cascades on random networks. Proceedings of the National Academy of Sciences **99**(9), 5766–5771 (2002) https://doi.org/10.1073/pnas.082090499

[4] Buldyrev, S.V., Parshani, R., Paul, G., Stanley, H.E., Havlin, S.: Catastrophic cascade of failures in interdependent networks. Nature **464**(7291), 1025–1028 (2010) https://doi.org/10.1038/nature08932

[5] Brockwell, P.J., Gani, J., Resnick, S.I.: Birth, immigration and catastrophe processes. Advances in Applied Probability **14**(4), 709–731 (1982) https://doi.org/10.2307/1427020

[6] Di Crescenzo, A., Giorno, V., Nobile, A.G., Ricciardi, L.M.: On the $M/M/1$ queue with catastrophes and its continuous approximation. Queueing Systems **43**, 329–347 (2003) https://doi.org/10.1023/A:1023261830362

[7] Jain, G., Sigman, K.: A Pollaczek–Khintchine formula for $M/G/1$ queues with disasters. Journal of Applied Probability **33**, 1191–1200 (1996) https://doi.org/10.2307/3214996

[8] Boxma, O.J., Perry, D., Stadje, W.: Clearing models for $M/G/1$ queues. Queueing Systems **38**, 287–306 (2001) https://doi.org/10.1023/A:1010903605577

[9] Di Crescenzo, A., Giorno, V., Krishna Kumar, B., Nobile, A.G.: A double-ended queue with catastrophes and repairs, and a jump-diffusion approximation. Methodology and Computing in Applied Probability **14**(4), 937–954 (2012) https://doi.org/10.1007/s11009-011-9214-2

[10] Dimou, S., Economou, A.: The single server queue with catastrophes and geometric reneging. Methodology and Computing in Applied Probability **15** (2013) https://doi.org/10.1007/s11009-011-9271-6

[11] Barbhuiya, F.P., Kumar, N., Gupta, U.C.: Batch renewal arrival process subject to geometric catastrophes. Methodology and Computing in Applied Probability **21**, 69–83 (2019) https://doi.org/10.1007/s11009-018-9643-2

[12] Evans, M.R., Majumdar, S.N.: Diffusion with stochastic resetting. Physical Review Letters **106**, 160601 (2011) https://doi.org/10.1103/PhysRevLett.106.160601

[13] Evans, M.R., Majumdar, S.N., Schehr, G.: Stochastic resetting and applications. Journal of Physics A: Mathematical and Theoretical **53**, 193001 (2020) https://doi.org/10.1088/1751-8121/ab7cfe

[14] Abundo, M.: The first-passage area of Wiener process with stochastic resetting. Methodology and Computing in Applied Probability **25**(4) (2023) https://doi.org/10.1007/s11009-023-10069-4

[15] Gelenbe, E.: Product-form queueing networks with negative and positive customers. Journal of Applied Probability **28**(3), 656–663 (1991)

[16] Artalejo, J.R.: G-networks: A versatile approach for work removal in queueing networks. European Journal of Operational Research **126**(2), 233–249 (2000)

[17] Vvedenskaya, N.D., Dobrushin, R.L., Karpelevich, F.I.: Queueing system with selection of the shortest of two queues: An asymptotic approach. Problems of Information Transmission **32**(1), 15–27 (1996)

[18] Graham, C.: Chaoticity on path space for a queueing network with selection of the shortest queue among several. Journal of Applied Probability **37**(1), 198–211 (2000)

[19] Budhiraja, A., Mukherjee, D., Wu, R.: Supermarket model on graphs. The Annals of Applied Probability **29**(3), 1740–1777 (2019) https://doi.org/10.1214/18-AAP1437

[20] Sznitman, A.-S.: Topics in propagation of chaos. In: École d'Été de Probabilités

de Saint-Flour XIX – 1989. Lecture Notes in Mathematics, vol. 1464, pp. 165–251. Springer, Berlin (1991). https://doi.org/10.1007/BFb0085169

[21] Benaïm, M., Le Boudec, J.-Y.: A class of mean field interaction models for computer and communication systems. Performance Evaluation **65**(11–12), 823–838 (2008)

[22] Lovász, L., Szegedy, B.: Limits of dense graph sequences. Journal of Combinatorial Theory, Series B **96**(6), 933–957 (2006)

[23] Medvedev, G.S.: The nonlinear heat equation on dense graphs and graph limits. SIAM Journal on Mathematical Analysis **46**(4), 2743–2766 (2014)

[24] Delattre, S., Fournier, N., Hoffmann, M.: Hawkes processes on large networks. Annals of Applied Probability **26**(1), 216–261 (2016)

[25] Bhamidi, S., Budhiraja, A., Wu, R.: Weakly interacting particle systems on inhomogeneous random graphs. Stochastic Processes and their Applications **129**(6), 2174–2206 (2019) https://doi.org/10.1016/j.spa.2018.06.014

[26] Bayraktar, E., Chakraborty, S., Wu, R.: Graphon mean field systems. The Annals of Applied Probability **33**(5), 3587–3619 (2023) https://doi.org/10.1214/22-AAP1901

[27] Bayraktar, E., Wu, R.: Graphon particle system: Uniform-in-time concentration bounds. Stochastic Processes and their Applications **156**, 196–225 (2023) https://doi.org/10.1016/j.spa.2022.11.008

[28] Fayolle, G., Iasnogorodski, R.: Two coupled processors: the reduction to a Riemann–Hilbert problem. Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete **47**(3), 325–351 (1979) https://doi.org/10.1007/BF00535168

[29] Fayolle, G., Iasnogorodski, R., Malyshev, V.: Random Walks in the Quarter-Plane. Applications of Mathematics, vol. 40. Springer, Berlin (1999)

[30] Cohen, J.W., Boxma, O.J.: Boundary Value Problems in Queueing System Analysis. North-Holland Mathematics Studies, vol. 79. North-Holland, Amsterdam (1983)

[31] Konheim, A.G., Meilijson, I., Melkman, A.: Processor-sharing of two parallel lines. Journal of Applied Probability **18**(4), 952–956 (1981) https://doi.org/10.2307/3213071