

A HIERARCHY OF NETWORK MODELS GIVING BISTABILITY UNDER TRIADIC CLOSURE*

STEFANO DI GIOVACCHINO[†], DESMOND J. HIGHAM[‡], AND KONSTANTINOS
C. ZYGALAKIS[‡]

Abstract. Triadic closure describes the tendency for new friendships to form between individuals who already have friends in common. It has been argued heuristically that the triadic closure effect can lead to bistability in the formation of large-scale social interaction networks. Here, depending on the initial state and the transient dynamics, the system may evolve towards either of two long-time states. In this work, we propose and study a hierarchy of network evolution models that incorporate triadic closure, building on the work of Grindrod, Higham, and Parsons [Internet Math., 8 (2012), pp. 402–423]. We use a chemical kinetics framework, paying careful attention to the reaction rate scaling with respect to the system size. In a macroscale regime, we show rigorously that a bimodal steady state distribution is admitted. This behavior corresponds to the existence of two distinct stable fixed points in a deterministic mean-field ODE. The macroscale model is also seen to capture an apparent metastability property of the microscale system. Computational simulations are used to support the analysis.

Key words. birth and death processes, graph, Langevin equation, mean field, metastability, stochastic simulation algorithm

MSC codes. 91D30, 91C20, 65C20

DOI. 10.1137/21M1461290

1. Motivation and background. Network science is built on the study of pairwise interactions between individual elements in a system. However, it is becoming more widely recognized that higher-order motifs involving groups of elements are also highly relevant [2, 3, 35]. In this work we focus on the formation of triangles—connected triples of nodes—over time.

Naturally occurring networks are often observed to have an overabundance of triangles [25, 26]. In the case of social networks, where nodes represent people and interactions describe friendships, *triadic closure* is widely regarded as a key triangle-forming mechanism. The concept of triadic closure dates back to the work of Simmel [32] and gained attention after the widely cited article [20]. Suppose two people, B and C, are not currently friends but have a friend, A, in common. The triadic closure principle states that the existence of the edges A-B and A-C increases the likelihood that the edge B-C will form at some future time. In other words, having a friend in common increases the chance of two people becoming friends. As discussed

*Received by the editors November 22, 2021; accepted for publication (in revised form) July 8, 2022; published electronically December 5, 2022.

<https://doi.org/10.1137/21M1461290>

Funding: The work of the first author was supported by the GNCS-INDAM project and by the PRIN2017-MIUR project 2017JYCLSF “Structure preserving approximation of evolutionary problems.” The work of the second author was supported by the Engineering and Physical Sciences Research Council under grants EP/P020720/1 and EP/V015605/1. The work of the third author was supported by the Leverhulme Trust (RF/2020-310) and the Engineering and Physical Sciences Research Council under grant EP/V006177/1.

[†]Department of Engineering and Computer Science and Mathematics, University of L’Aquila, L’Aquila, 67100 Italy (stefano.digiovacchino@graduate.univaq.it).

[‡]School of Mathematics, University of Edinburgh, Edinburgh, EH9 3FD United Kingdom (d.j.higham@ed.ac.uk, k.zygalakis@ed.ac.uk).

in [10, Chapter 2], there are three reasons why the chance of the B-C connection increases. First, B and C both socialize with A and hence have a greater *opportunity* to meet. Second, A can vouch for B and C and hence raise the level of *trust* between them. Third, A may be *incentivized* to encourage the B-C friendship if maintaining a single triadic friendship is viewed as more efficient than maintaining a separate pair of dyadic friendships with B and C separately. Building on these ideas it is also natural, as in [13, 22], to argue that the increased likelihood that B and C will become connected via triadic closure grows in proportion to the number of friends they have in common.

With the advent of large-scale time-stamped data sets that record online human interactions, it has become possible to test for the presence of triadic closure [2, 25, 28]. Moreover, the triadic closure principle can be used as the basis of link prediction algorithms, which attempt to anticipate the appearance of new edges over time [2, 12, 27, 36].

A time-dependent random graph model that incorporates triadic closure was proposed and analyzed in [22], and also calibrated to real cellphone data. The model, from the general class defined in [21], takes the form of a discrete-time Markov chain where edges may appear or disappear independently at random over each time step. Numerical simulations revealed bistable dynamics—different paths of the same stochastic process were observed to evolve towards either a sparse regime or a rich, well-triangulated regime. A heuristic, deterministic mean-field approximation was put forward to explain the behavior.

In this work, our aim is to develop and study a hierarchy of time-dependent network models that incorporate triadic closure and admit bistable behavior. We use a chemical kinetics framework to describe the system at a microscale level and then consider macroscale approximations. This allows us to clarify the assumptions that go into the micro-to-macro step and to be clear about the scaling of the model parameters with respect to system size. At the macroscale level, the resulting process can be shown rigorously to admit bistable behavior, in the sense of possessing a bimodal steady state distribution for a particular choice of model parameters. The macroscale steady state also gives insights into the apparent metastability behavior observed in the microscale simulations for large system size, where the switching time between sparse and well-triangulated regimes becomes extremely long. We study this effect by computing the mean time to transition between regimes. We also introduce the corresponding diffusion, or Langevin, model and compute its mean transition times. Here, reaction rate theory [31] gives a good approximation to the observed growth in mean passage time with respect to system size. Finally, we show that the underlying deterministic reaction rate ODE, arising in the thermodynamic limit, gives a continuous-time analogue of the mean-field equation in [22].

The rest of the manuscript is organized as follows. In section 2 we set up the notation and define the microscale model. We also report on computational experiments that motivate the subsequent material. In section 3 we derive a macroscopic approximation and analyze its steady states and mean exit times. Section 4 introduces the corresponding Langevin and ODE models. We finish with a brief discussion in section 5.

2. Microscopic model. We suppose now that a network has a fixed set of n nodes, which at each point in time may be connected by undirected edges. For example, the nodes may represent individuals in an online social media platform and the edges may represent mutual friendships. To be concrete, we assume that the

nodes are labeled from 1 to n , and we let $N = n(n-1)/2$ denote the total number of edges that may be present at any time. We are interested in the case where n is large, say $n \geq 30$, and we find it useful to think of a *one-parameter family* of models, parameterized by N .

Our modeling framework allows edges to be created or deleted in continuous time. We will use a chemical kinetics setting where the creation or deletion of an edge is represented as a reaction between “edge” and “no edge” molecules. To be precise, consider any pair of distinct nodes, i and j . Without loss of generality, we take $i < j$. Then, with the corresponding undirected edge between nodes i and j we associate two species, E_{ij} and O_{ij} . Existence of E_{ij} represents the presence of the edge connecting nodes i and j and the existence of O_{ij} represents the absence of this edge. So exactly one of E_{ij} and O_{ij} exists at any given time t .

The mechanisms governing the edge evolution will be modeled as the following set of reactions, with rate constants c_1 , c_2 , and \hat{c}_3 :



Here, reactions (2.1) and (2.2) represent spontaneous edge birth and spontaneous edge death, respectively. The third reaction, (2.3), captures the triadic closure effect: if i and j are not currently connected, then for every instance where there is a node k connected to both i and j , there is chance that i and j will become connected via triadic closure. Hence, we use the same principle as [22]: the overall chance of an edge arising from triadic closure is linearly proportional to the number of new triangles that this event would create.

We refer to (2.1)–(2.3) as the *microscopic model*. For the purpose of our work, this microscopic model is regarded as an exact description of the network evolution. Our aim is to study this model, and in later sections we will introduce approximations that allow us to gain insights.

In principle, we could apply the stochastic simulation algorithm (SSA), also known as Gillespie’s algorithm, the Doob–Gillespie algorithm, and the Bortz–Kalos–Lebowitz algorithm, to this system [6, 15, 16, 24]. The state vector, which records the number of E_{ij} and O_{ij} “molecules” at each time point will have $2N$ components. There are N reactions of type (2.1), N reactions of type (2.2), and $N(N-1)$ reactions of type (2.3). Hence, the stoichiometric vectors would have dimension $N(N+1)$. Given an initial state vector, at each step of the SSA we draw two random numbers: one to determine time of the next reaction and one to determine which reaction takes place.

Because exactly one of E_{ij} and O_{ij} can exist, the propensity functions for reactions (2.1)–(2.3) have simple forms:

- for reaction (2.1), the propensity function is c_1 if O_{ij} exists and zero otherwise;
- for reaction (2.2), the propensity function is c_2 if E_{ij} exists and zero otherwise;
- for reaction (2.3), the propensity function is \hat{c}_3 if O_{ij} , E_{jk} and E_{ik} exist and zero otherwise.

We note that reactions (2.1) and (2.2) involve individual edges, and it is reasonable to assume that the rate constants c_1 and c_2 do not depend on the system size. Reaction (2.3), however, involves interactions between edges, and the question of how \hat{c}_3 depends on the system size could be viewed as a modeling issue. We will argue in section 3 that \hat{c}_3 should scale like $1/n$, since this produces systems where the triadic

closure mechanism remains present but does not overwhelm the spontaneous birth and death of edges. More precisely, we will let

$$(2.4) \quad \hat{c}_3 = \frac{c_3}{n-2},$$

for a fixed constant c_3 (independent of n).

From the SSA perspective, we can reduce the size of the stoichiometric vectors, and thereby save considerably on storage and bookkeeping, by exploiting the special structure of this system. To do this, we draw three random numbers at each step: one to determine time of the next reaction, one to determine which *class* of reaction takes place, and one to pick a reaction from within this class. In this way, the computation can be performed directly with the symmetric adjacency matrix, $A(t) \in \mathbb{R}^{n \times n}$, where $(A(t))_{ij} = 1$ if the edge connecting i and j is present at time t and $(A(t))_{ij} = 0$ otherwise. (It would be sufficient to work with just the upper triangle of $A(t)$, but we find it more natural to use the whole matrix when describing the algorithm.)

To summarize this approach, we introduce the class-level propensities

$$(2.5) \quad a_{\text{birth}} = c_1 \sum_{i < j} \sum_j (1 - A(t)_{ij}),$$

$$(2.6) \quad a_{\text{death}} = c_2 \sum_{i < j} \sum_j A(t)_{ij},$$

$$(2.7) \quad a_{\text{triadic}} = \hat{c}_3 \sum_{i < j} \sum_j A_{\text{triadic}}(t)_{ij},$$

where $A_{\text{triadic}}(t) \in \mathbb{R}^{n \times n}$ is defined as

$$(2.8) \quad A_{\text{triadic}}(t)_{ij} = (A^2(t))_{ij}(1 - A(t)_{ij}).$$

In (2.5) we form the product of the rate constant c_1 and the total number of missing edges. Similarly, in (2.6) we form the product of the rate constant c_2 and the total number of current edges. To understand (2.7), note that the term $(A^2(t))_{ij}$ counts the number of nodes that are connected to both i and j . Hence, in (2.7) we form the product of the rate constant \hat{c}_3 and the total number of opportunities for triadic closure.

A step of SSA may then be performed as follows:

1. Compute a_{birth} , a_{death} , a_{triadic} , and form $a_{\text{sum}} = a_{\text{death}} + a_{\text{birth}} + a_{\text{triadic}}$.
2. Choose the time until the next reaction from an exponential distribution with parameter a_{sum} .
3. Choose the class of reaction: death, birth, or triadic, with probability proportional to a_{death} , a_{birth} , and a_{triadic} , respectively.
4. If the class is death, choose one of the nonzeros from the upper triangle of $A(t)$ uniformly at random, and set it to zero.
5. If the class is birth, choose one of the zeros from the upper triangle of $A(t)$ uniformly at random, and set it to one.
6. If the class is triadic, choose one of the nonzeros from the upper triangle of the matrix $A_{\text{triadic}}(t)_{ij}$ with probability proportional to its value, and set it to one.

It is worth pointing out here that the elements of the adjacency matrix represent chemical species, and they can only take the value 0 or 1. However, the Markov chain

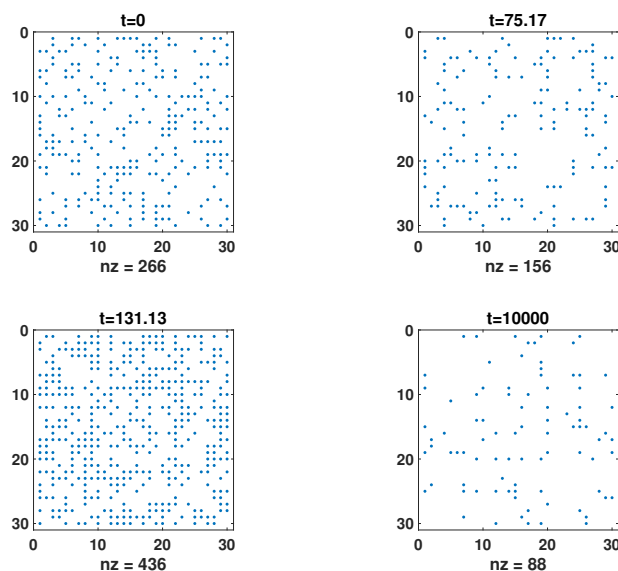


FIG. 1. Time evolution of the adjacency matrix $A(t)$. Nonzeros are represented by dots. We show the connectivity at the initial time (upper left), at a time of low density (upper right), a time of high density (lower left), and a later time where the density is again low (lower right). Here, we have $n = 30$ nodes and nz denotes the number of nonzeros in the adjacency matrix.

simulated by the algorithm described above has a state space of dimension 2^N (all the possible undirected graphs with n -nodes). This in turn implies that as the number of nodes n increases the computational cost of the algorithm will become a severe limiting factor in understanding the dynamics of the system. One possible way of accelerating the simulation is to use the τ -leaping algorithm [19]. In the case $\tau = 1$, this would essentially reduce to the model in [22].

We visualize in Figure 1 the adjacency matrix $A(t)$ at different time instances along a path computed by the SSA algorithm described above. Here, a dot in the picture corresponds to the existence of an edge. We used reaction rate coefficients

$$(2.9) \quad c_1 = 0.025, \quad c_2 = 0.25, \quad c_3 = 0.91$$

with \hat{c}_3 defined in (2.4). Here, the initial configuration was a sample from the Erdős–Rényi class with parameter $p = 0.3$, i.e., every edge has independent probability p to exist at $t = 0$. As t increases the system is seen to transition between states of low and high connectivity. We can summarize the overall connectivity by the edge density

$$(2.10) \quad q(t) := \frac{1}{N} \sum_{i < j} \sum_j A(t)_{ij}.$$

In Figure 1, the upper and lower right states have edge density of around 0.18 and 0.1, respectively, and the lower left state has an edge density of around 0.5. The transitioning becomes clearer in Figure 2, where we plot the evolution of $q(t)$ for a range of different system sizes, again using an Erdős–Rényi initial condition with $p = 0.3$. As n is increased from 30 to 50, the edge density spends longer periods of time around each level. For $n = 80$ and $n = 100$, switching does not take place over the interval $0 \leq t \leq 10^4$. We will return to these observations in the next section, where we construct and study an explanatory macroscopic approximation.

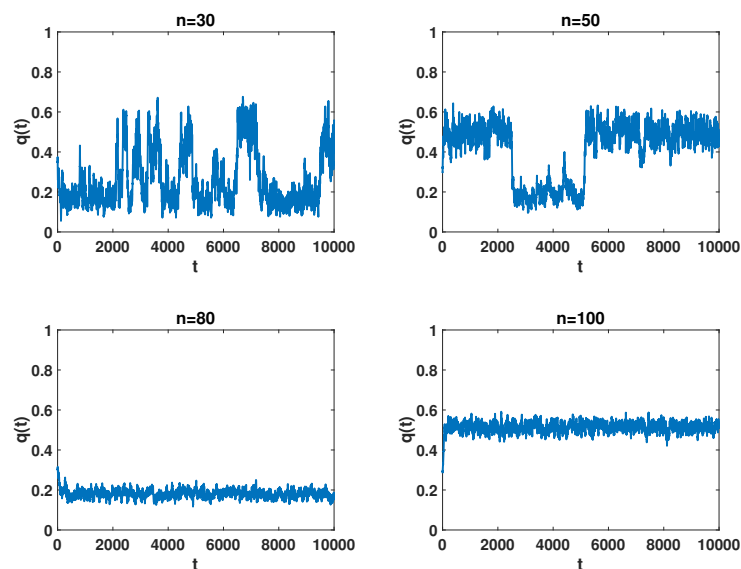


FIG. 2. Evolution of edge density for $n = 30, 50, 80, 100$, up to time $t = 10^4$.

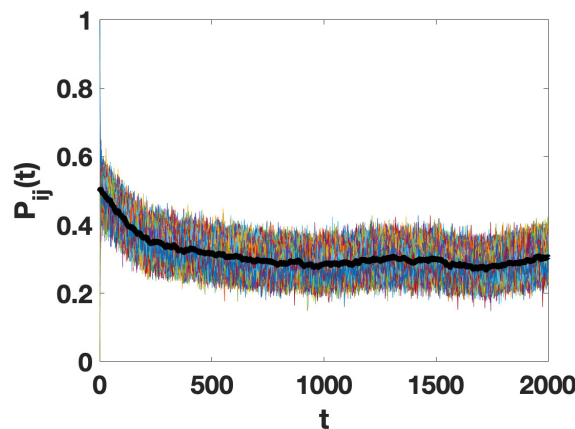


FIG. 3. Time evolution of all $P_{ij}(t)$, the probability that the edge $i \leftrightarrow j$ exists at time t . Thick black line is the mean of $P_{ij}(t)$ over all i and j .

Given that the overall edge density is a macroscopic quantity, it is natural to ask if it captures the effective behavior of the system. To investigate this issue, let $P_{ij}(t)$ denote the probability that the edge connecting i and j exists at time t . In Figure 3 we superimpose all the individual values $P_{ij}(t)$ for the case $n = 30$. Here, we used a Monte Carlo approach to estimate each $P_{ij}(t)$ by applying SSA to compute 250 independent paths. For each path, an initial state was chosen where half of the edges exist. We see that after a fast transient all probabilities fluctuate around a value that lies between the low and high edge density regimes that were seen in Figures 1 and 2. (The mean of $P_{ij}(t)$ over all i and j is superimposed as a thick black line.) This implies that after an initial transient phase the topology of the network is homogeneous in the sense that no nodes have distinguishable behavior.

3. Macroscopic model. We saw in section 2 that after an initial transient the specific topology of the network appears to be homogeneous. Motivated by this observation we now define a simplified stochastic model. In this *macroscopic* model we do not retain any information about the topological structure of the network; we simply keep track of the number of edges. This reduces the size of the state space dramatically; in the microscopic model there are 2^N distinct networks. In this macroscopic model the scalar edge count ranges between 0 and N . The simpler model allows long-time numerical simulation to be performed more efficiently, and, as we show in section 3.3, it is also amenable to rigorous analysis.

3.1. Macroscopic regime. Our simplified, macroscopic version of the microscale system (2.1)–(2.3) involves two species: E represents an edge and O represents a missing edge. The species undergo three types of reaction, representing birth, death, and triadic closure, respectively,



The state vector may be written $X(t) \in \mathbb{R}^2$, where $X_1(t)$ and $X_2(t)$ record the number of edges (species E) and missing edges (species O), respectively. Notice that $X_1(t) + X_2(t) = N$ for all $t \geq 0$. The stoichiometric vectors for the three reactions are

$$\nu_1 = \begin{bmatrix} 1 \\ -1 \end{bmatrix}, \quad \nu_2 = \begin{bmatrix} -1 \\ 1 \end{bmatrix}, \quad \nu_3 = \begin{bmatrix} 1 \\ -1 \end{bmatrix}.$$

Reactions (3.1) and (3.2) are first order reactions that model spontaneous birth and death of edges. These effects are independent of the network structure, and hence we reproduce the corresponding birth and death behavior from the microscale model exactly by taking propensity functions $a_1(X(t)) = c_1 X_2(t)$, and $a_2(X(t)) = c_2 X_1(t)$, where c_1 and c_2 are the rate constants from (2.1) and (2.2). For the reaction (3.3) representing triadic closure, we must introduce simplifying assumptions, because the macroscale regime does not keep track of individual node and edge labels. Based on our observation in section 2, we suppose that at the current time t edges are uniformly and independently distributed among the nodes—every possible edge $i \leftrightarrow j$ (with $i < j$) has a chance $X_1(t)/N$ of existing. (We note that this is effectively the mean-field assumption used in [22].) Our aim is now to approximate the number of *open wedges*; that is, node triples i, j, k such that edges $i \leftrightarrow j$ and $i \leftrightarrow k$ are present and edge $j \leftrightarrow k$ is absent. This quantity records the number of opportunities for the microscale reaction (2.3) to take place, and hence will be used in forming the propensity function at the macroscale level in (3.3).

To derive this approximation, we first note there are $X_1(t)$ ways to choose an edge $i \leftrightarrow j$. Then, fixing this i and j , for every other node k there is a chance of $\approx (X_1(t) - 1)/N$ that the edge $i \leftrightarrow k$ also exists. So, taking into account all $n - 2$ nodes, there are $\approx (n - 2)X_1(t)(X_1(t) - 1)/N$ chains of the form $j \leftrightarrow i \leftrightarrow k$. For any such chain, the chance of the third edge, $j \leftrightarrow k$, being absent is $\approx X_2(t)/N$. So the overall number of open wedges is $\approx (n - 2)X_1(t)(X_1(t) - 1)X_2(t)/N^2$. Hence, we find that a suitable propensity function for reaction (3.3) is

$$(3.4) \quad a_3(X(t)) = \hat{c}_3(n - 2)X_2(t)X_1(t)(X_1(t) - 1)/N^2,$$

where \hat{c}_3 is the rate constant in (2.3).

Since $X_1(t) + X_2(t) = N$, the macroscopic model (3.1)–(3.3) may be written as a scalar, continuous-time birth and death process for which $X_1(t)$ takes integer values between 0 and N . Using λ_i and μ_i to denote the overall birth and death rates, we have the general form

$$(3.5) \quad \mathbb{P}(X_1(t + \delta t) = i + 1 | X(t) = i) = \lambda_i \delta t + o(\delta t) \text{ for } i = 0, \dots, N - 1,$$

$$(3.6) \quad \mathbb{P}(X_1(t + \delta t) = i - 1 | X(t) = i) = \mu_i \delta t + o(\delta t) \text{ for } i = 1, \dots, N.$$

Based on the arguments above, these birth and death rates take the form

$$(3.7) \quad \lambda_i = c_1(N - i) + \frac{\hat{c}_3(n - 2)}{N^2}(N - i)i(i - 1),$$

$$(3.8) \quad \mu_i = c_2 i.$$

Using the scaling (2.4), we obtain the birth rate

$$(3.9) \quad \lambda_i = N \left[c_1 \left(1 - \frac{i}{N} \right) + c_3 \left(1 - \frac{i}{N} \right) \left(\frac{i}{N} \right) \left(\frac{i - 1}{N} \right) \right].$$

We then rewrite the death rate (3.8) as

$$(3.10) \quad \mu_i = N c_2 \frac{i}{N}.$$

We note in particular that the $\hat{c}_3 \propto 1/\sqrt{N}$ scaling from (2.4) has balanced the size of the separate terms in the birth and death rates (3.9) and (3.10).

We emphasize that c_1 , c_2 , and c_3 are regarded as positive constants and we have a family of birth and death processes parameterized by the system size, N , representing the number of edges.

3.2. Comparison between the macroscopic and microscopic models.

Having derived an approximation to the microscopic model in the previous subsection it is natural to compare individual paths of the two stochastic models. We will choose two sets of reaction constants c_1, c_2, c_3 . From now on, and for reasons that will become apparent, we will refer to the parameters in (2.9) as the bistable regime, and we will also use a second set of parameters with a larger birth rate, c_1 , given by

$$(3.11) \quad c_1 = 0.25, \quad c_2 = 0.25, \quad c_3 = 0.91.$$

We will refer to (3.11) as the monostable regime.

The behavior of paths in the monostable regime can be seen in the upper plots of Figure 4, for $n = 30$ on the left and $n = 100$ on the right. The initial state was an Erdős–Rényi sample with parameter $p = 0.2$. For both values of n , we see that the microscopic and macroscopic models are behaving in the same way, namely, having the edge density increasing towards a high value and oscillating around it. Additionally, we superimpose in these plots the solution of the deterministic mean-field model in [22] (see also the discussion in section 4). As we can see, both models agree closely with the mean-field approximation, with the size of fluctuations becoming smaller for larger n .

The behavior of the paths in the bistable parameter regime (2.9) can be seen in the lower plots of Figure 4. Again we have $n = 30$ on the left and $n = 100$ on the right and we used an Erdős–Rényi initial configuration with $p = 0.2$. We see again there is excellent qualitative agreement between the microscopic and macroscopic paths.

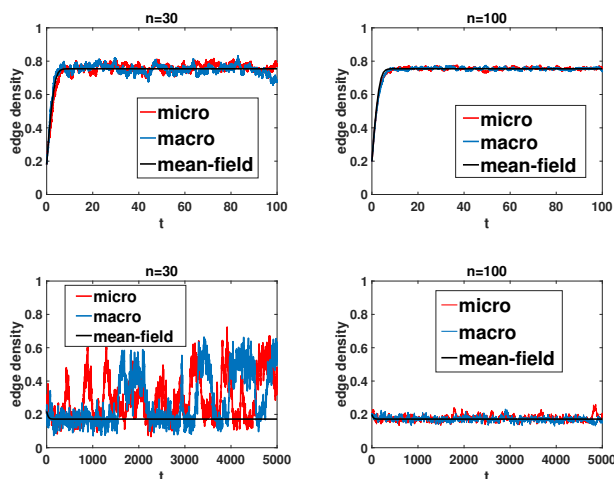


FIG. 4. Comparison between the microscopic, the macroscopic, and the mean-field model. An Erdős–Rényi initial configuration with $p = 0.2$ using $n = 30$ (left) and $n = 100$ (right). Upper plots: monostable regime (3.11). Lower plots: bistable regime (2.9).

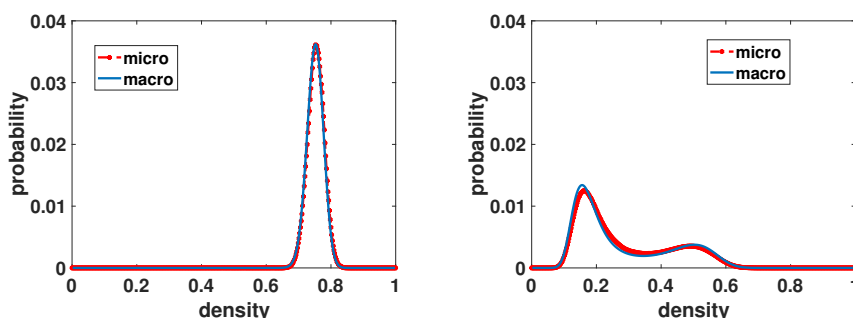


FIG. 5. Comparison between the steady states of the microscopic and macroscopic models in the monostable (left, parameter set (3.11)) and bistable (right, parameter set (2.9)) regimes for $n = 30$.

However, unlike the monostable regime, for $n = 30$ the paths now do not agree with the mean-field model, which, by its deterministic nature, does not switch between the low and high density regimes.

Our second comparison of the two models concerns their long-time behavior. In the case of the macroscopic model, we will see in the next section that it is possible to obtain an analytic expression for the steady state. For the microscopic model it is not possible to derive an analytic expression for the steady state, so we used a long-time simulation ($T = 5 \times 10^6$ in the bistable case, $T = 10^4$ in the monostable case). We plot this comparison in Figure 5. Here, the horizontal axis represents the scaled state values $0, 1/N, 2/N, \dots, 1$. We observe excellent agreement between the two steady states both in the monostable and bistable regime.

3.3. Steady state of the macroscopic model. We now show that the bistable behavior observed in Figures 1, 4, and 5 can be explained by analyzing the long-term behavior of the macroscopic model. Standard theory [34] shows that in our ergodic setting a process of the form (3.5)–(3.6) has a stationary distribution

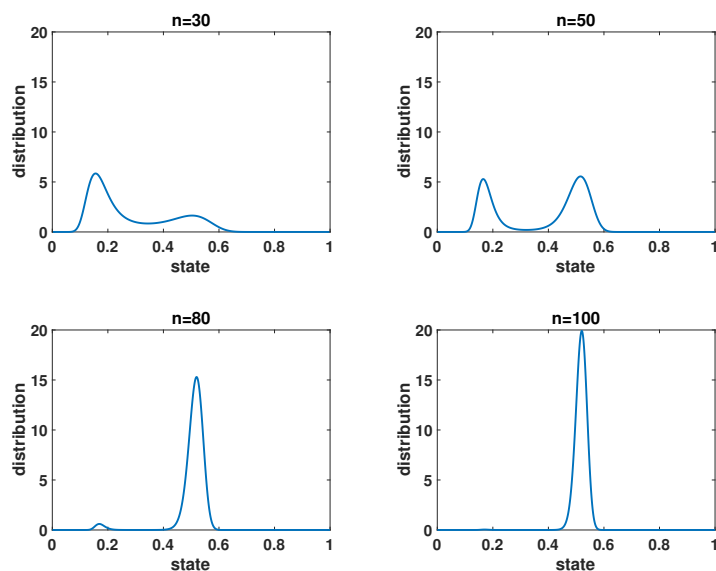


FIG. 6. Steady state distribution of the macroscopic model (3.5)–(3.6) for $n = 30, 50, 80, 100$ in the bistable regime (2.9).

$$\pi_j = \mathbb{P} \left(\lim_{t \rightarrow \infty} X_1(t) = j \right), \quad j = 0, 1, \dots, N,$$

that satisfies

$$(3.12) \quad \pi_j = \frac{\prod_{i=0}^{j-1} \lambda_i}{\prod_{i=1}^j \mu_i} \pi_0.$$

In Figure 6 we show the steady state distribution in the bistable regime (2.9) when the number of nodes is $n = 30, 50, 80$, and 100 . This figure may be compared with Figure 2. For internal consistency we have regarded these four curves as continuous-valued probability density functions over $[0, 1]$, thereby normalizing them to have unit area. We see in Figure 6 that the bimodal steady state distribution has peaks at around 0.2 and 0.5 , in agreement with the low and high connectivity regimes seen in Figure 2. We also see that although the location of the two peaks seems to be fixed, their relative heights vary. For $n = 30$, the lower mode in Figure 6 captures more probability, which is consistent with the trajectory of $q(t)$ in Figure 2 spending more time in the lower connectivity regime. As n increases the higher connectivity regime begins to dominate, and indeed the bimodality is barely visible for $n = 100$ in Figure 6. This is consistent with the less frequent switching observed in Figure 2.

The following result shows that the bimodality of the macroscopic steady state is connected to the roots of the cubic equation (3.13), which depend only on c_1 , c_2 , and c_3 . We note that these roots are $p_1^* \approx 0.17$, $p_2^* \approx 0.31$, and $p_3^* \approx 0.52$ for the bistable regime used in Figures 2 and 6. (We also note that the cubic equation (3.13) has a single real root of $p^* \approx 0.75$ in the monostable regime (3.11) used in the upper plots of Figure 4.)

THEOREM 3.1. *Consider the scalar cubic equation*

$$(3.13) \quad (1 - p)(c_1 + c_3 p^2) - c_2 p = 0$$

for which any real root must necessarily lie in $(0, 1)$. Then

- (a) the cubic equation has one real root $0 < p^* < 1$ if and only if for sufficiently large N the family of birth and death processes (3.5)–(3.6) with (3.9) and (3.10) has a unimodal steady state with maximum probability at state j , where $j/N = p^* + O(1/N)$,
- (b) the cubic equation has three real roots $0 < p_1^* < p_2^* < p_3^* < 1$ if and only if for sufficiently large N the family of birth and death processes (3.5)–(3.6) with (3.9) and (3.10) has a bimodal steady state with locally maximum probability at state j_1 , where $j_1/N = p_1^* + O(1/N)$ and at state j_3 , where $j_3/N = p_3^* + O(1/N)$, and with locally minimum probability at state j_2 , where $j_2/N = p_2^* + O(1/N)$.

Proof. Let

$$(3.14) \quad f(p) = \frac{(1-p)(c_1 + c_3 p^2)}{c_2},$$

so a root of (3.13) corresponds to $f(p) = p$. Because $f(p) > 0$ for $p < 0$ and $f(p) < 0$ for $p > 1$, we see that any real root must lie in $(0, 1)$. We know from (3.12) that $\pi_{j+1}/\pi_j = \lambda_j/\mu_{j+1}$. Hence, from (3.9), (3.10), and (3.14), we see that for any $j = 0, 1, \dots, N-1$,

$$(3.15) \quad \left| \frac{\pi_{j+1}}{\pi_j} - \frac{f(j/N)}{j/N} \right| \leq \frac{C}{N},$$

where C is a constant that is uniform in j .

For case (a), suppose that the cubic equation has one root p^* in $(0, 1)$. Then, by construction,

$$\frac{f(p)}{p} > 1 \text{ for } p \in (0, p^*) \quad \text{and} \quad \frac{f(p)}{p} < 1 \text{ for } p \in (p^*, 1).$$

Let j^* be the largest j such that $j/N \leq p^*$. Then, for sufficiently large N , we see from (3.15) that

$$\frac{\pi_{j+1}}{\pi_j} > 1 \text{ for } j < j^* \quad \text{and} \quad \frac{\pi_{j+1}}{\pi_j} < 1 \text{ for } j > j^*.$$

Hence, $\pi_0 < \pi_1 < \dots < \pi_{j^*}$ and $\pi_{j^*+1} > \pi_{j^*+2} > \dots > \pi_N$. So one or both of π_{j^*} and π_{j^*+1} give a maximum value.

On the other hand, suppose we have a unimodal distribution, that is, $\pi_0 < \pi_1 < \dots < \pi_{j^*}$ and $\pi_{j^*+1} > \pi_{j^*+2} > \dots > \pi_N$. Then, by (3.15),

$$\frac{f(p)}{p} + O(1/N) > 1 \text{ for } p = \frac{j}{N} \text{ with } p < p^*$$

and

$$\frac{f(p)}{p} + O(1/N) < 1 \text{ for } p = \frac{j}{N} \text{ with } p > p^*.$$

For this to be true for all sufficiently large N , by continuity we must have a unique root in $(0, 1)$ at p^* .

Case (b) may be proved in a similar manner. \square

3.4. Exit times. Figures 2 and 6 suggest that for large n there is a metastability effect, where the process spends a long time in one regime before eventually switching to another. It is, of course, extremely challenging to verify such behavior by directly simulating a trajectory [5]. To gain further understanding of this type of behavior, we will study the expected time taken to move between modes in the bimodal case of the macroscopic model.

Let $G \in \mathbb{R}^{(N+1) \times (N+1)}$ be the tridiagonal matrix

$$(3.16) \quad \begin{bmatrix} -\lambda_0 & \lambda_0 & 0 & \dots & 0 \\ \mu_1 & -(\lambda_1 + \mu_1) & \lambda_1 & \ddots & \vdots \\ 0 & \mu_2 & -(\lambda_2 + \mu_2) & \lambda_2 & \vdots \\ \vdots & \ddots & \ddots & \ddots & \lambda_{N-1} \\ 0 & \dots & \dots & \mu_N & -\mu_N \end{bmatrix},$$

and, given $0 \leq i \leq N$, let $Q^{[i]} \in \mathbb{R}^{N \times N}$ be the matrix obtained from G by neglecting the i th row and column, where we index from 0 to N . Denote by 1_N the vector in \mathbb{R}^N of all ones, and consider the linear system

$$(3.17) \quad Q^{[i]} \tau^{[i]} = -1_N.$$

The component $\tau_j^{[i]}$ then gives the mean time spent by the process (3.5)–(3.6) to reach the state i starting from state j [33]; that is,

$$\tau_j^{(i)} = \mathbb{E}[\inf t : X(t) = i | X(0) = j], \quad j = 0, \dots, N, \quad j \neq i.$$

Let $\lfloor \cdot \rfloor$ denote the integer flooring operation. In the case where the cubic (3.13) has three real roots, $0 < p_1^* < p_2^* < p_3^*$, and for a fixed value n , and hence N , we let $N_1 = \lfloor p_1^* N \rfloor$, $N_2 = \lfloor p_2^* N \rfloor$, and $N_3 = \lfloor p_3^* N \rfloor$. Hence, by Theorem 3.1, N_1 and N_3 are the state values corresponding to the peaks of the bimodal density, and N_2 corresponds to the trough in-between. In Figure 7, on the left, by repeatedly solving appropriate systems of the form (3.17), we show the mean time taken for the process to move from N_1 to N_2 (dashed line) and from N_3 to N_2 (solid line), as a function of n . This characterizes for a specific system size n how long it takes on average to switch between the two regimes. Note that both axes in Figure 7 are logarithmically scaled. We see that as the system size increases, the time to transition between the two modes increases dramatically. Moreover, for small n , the $N_1 \rightarrow N_2$ transition takes longer than the $N_3 \rightarrow N_2$ transition, whereas this relation changes over as n increases. In Figure 7, on the right, we show the ratio of the $N_1 \rightarrow N_2$ and $N_3 \rightarrow N_2$ mean exit times. This behavior is consistent with the steady state plots in Figure 6, where for small (respectively, large) n there is more probability mass around N_1 (respectively, N_3).

4. Langevin and deterministic ODE models. The macroscale model (3.5)–(3.6) with (3.9) and (3.10) may be viewed as part of a hierarchy of models that includes a stochastic differential equation (SDE), or Langevin process, and a deterministic reaction rate ODE. We will first look at the Langevin version and test whether it retains the metastable behavior observed earlier.

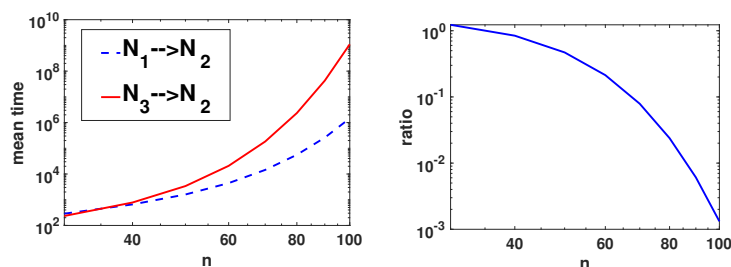


FIG. 7. Bistable regime: mean time for the macroscale model to reach the central trough of the steady state, N_2 , starting from each peak, N_1 and N_3 (left). Ratio between corresponding times (right).

Using standard arguments [17, 18] the Langevin equation may be written as an Itô SDE of the form

$$(4.1a) \quad dy(t) = \mu(y(t))dt + \beta(y(t))d\mathbf{W}(t),$$

$$(4.1b) \quad \mu(y) = c_1(1-y) - c_2y + c_3(1-y)y^2,$$

$$(4.1c) \quad \beta(y) = \frac{1}{\sqrt{N}}[\sqrt{c_1(1-y)}, -\sqrt{c_2y}, \sqrt{c_3(1-y)y^2}],$$

where $\mathbf{W}(t) = [W_1(t), W_2(t), W_3(t)]^T$ is formed of three independent Brownian motions. Here the scalar, real-valued, random variable $y(t)$ represents the edge density at time t . Because of the presence of the square roots in (4.1c), the process becomes ill-defined when $y(t)$ leaves the interval $[0, 1]$. This circumstance is common in Langevin formulations; intuitively, it arises from the fact that the approximations used to transfer between Poisson and diffusion processes break down when the molecule count for any species becomes small. Many approaches have been proposed to overcome this issue; see, for example, [1, 9, 29] and the references therein. We will avoid this difficulty by focusing on the mean first passage time for the SDE with a reflecting boundary at one end, thereby considering only paths that do not leave $[0, 1]$. Letting

$$\sigma^2(y) = \frac{1}{N} (c_1(1-y) + c_2y + c_3(1-y)y^2),$$

we may then define the following partial differential operator

$$(4.2) \quad Lu := \mu(y)\frac{du}{dy} + \frac{1}{2}\sigma^2(y)\frac{d^2u}{dy^2},$$

and the corresponding boundary value problem in the interval (a, b)

$$(4.3) \quad LT = -1,$$

subject to one of the two following pairs of boundary conditions:

$$(4.4a) \quad T'(0) = 0, \quad T(b) = 0,$$

$$(4.4b) \quad T(a) = 0, \quad T'(1) = 0.$$

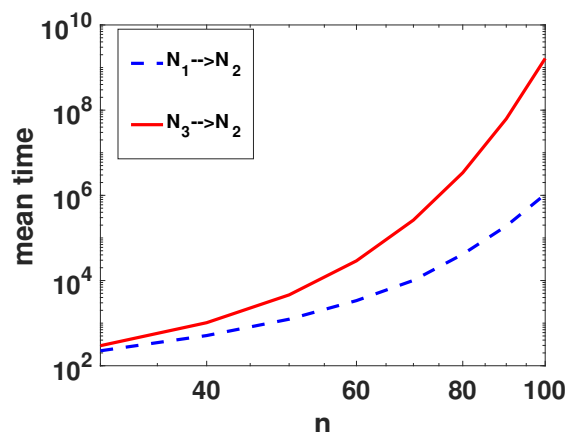


FIG. 8. Analogues of the mean exit times in Figure 7 for the Langevin equation (4.1).

The solution $T(x)$ with boundary conditions (4.4a) gives the mean time for paths of (4.1) starting from x to reach the value b when the SDE has a reflecting boundary condition at $y(t) = 0$ [14]. Similarly, with the boundary conditions (4.4b), $T(x)$ gives the mean time for the SDE paths to reach a starting from x when there is a reflecting boundary condition at $y(t) = 1$.

In the manner of Figure 7, in Figure 8 we show the mean times taken for the Langevin process to reach the central trough of the steady state distribution (scaled to represent edge count rather than edge density) from each peak. The curves were obtained by numerically solving the relevant boundary value problems. In separate computations, not reported here, we observed that it was very unlikely for an SDE path to leave the interval $[0, 1]$ over any reasonable timescale, and hence the reflecting boundary condition, which is introduced to avoid the technical issue of ill-defined square roots, has little effect on the results.

For further comparison, in Figure 9 we superimpose the mean exit times from Figure 7 (on the left) for the macroscale model and the mean first passage times from Figure 8 for the Langevin process. Here, we have plotted n^2 on the horizontal axis against the logarithm of the mean time on the vertical axis. We see that the SDE captures the extreme rate observed for the macroscale model, and for both models the mean time appears to scale exponentially in n^2 , in line with the reaction rate theory of Kramers [23, 31]. It would, of course, be of interest to pursue the rigorous analysis of (4.1) with respect to long-term behavior as the system size increases.

4.1. Deterministic model. The reaction rate ODE arising from (4.1) takes the form $dy/dt = \mu(y(t))$. We note that μ coincides with the cubic equation in (3.13), and hence the edge densities p^* in part (a) of Theorem 3.1 and p_1^*, p_2^*, p_3^* in part (b) of Theorem 3.1 coincide with fixed points of this ODE. We also note that the deterministic mean field from [22] is essentially an Euler discretization of this ODE with unit step size. In the upper plots of Figure 4 we see that, in the monostable regime, the deterministic mean-field approximation captures the long-term behavior of the macrolevel edge density. In the lower plots, corresponding to the bistable regime, the deterministic approximation is not able to switch, and hence settles on one of the two stable fixed points, p_1^* or p_3^* , depending on the initial condition.

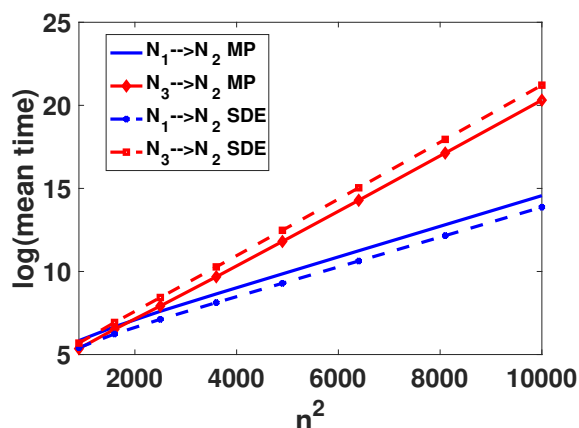


FIG. 9. Comparison between mean exit times for the SDE (4.1) (as in Figure 8) and the macroscale Markov process (3.5)–(3.6) with (3.9) and (3.10) (as in Figure 7).

5. Discussion. Our aim in this work was to perform multiscale stochastic modeling and analysis to add insights to the bistability effect observed in [22]. A major objective was to develop simplified stochastic models that (a) reproduce the switching and metastability behavior observed for the full microscale model, (b) make long-term simulation feasible, and (c) offer the potential for rigorous analysis. A key step in this work was to set up a suitable scaling for the triadic closure reaction rate as a function of system size, i.e., the total number of possible edges, N . Returning to the original microscale model (2.1)–(2.3), we note that \hat{c}_3 is the rate constant for each individual triadic closure reaction. So the chance of reaction (2.3) happening is \hat{c}_3 if each of O_{ij} , E_{jk} , and E_{ik} exist, and zero otherwise. Based on the way our macroscale approximations arose, it was then natural for \hat{c}_3 to have the scaling $\hat{c}_3 \propto 1/\sqrt{N}$ in (2.4). Here, *the chance of any individual triadic closure event taking place decreases as the system size grows*. This scaling is necessary if we wish to be in a regime where spontaneous birth, spontaneous death, and triadic closure all coexist for large system size. In particular, keeping \hat{c}_3 constant as N grows would lead to a system where triadic closure dominates, in general.

There are several directions in which this work could be pursued. For example, we note that although online social networking has the potential to increase the number of friendships an individual may form, there is evidence of an upper limit of around 150 in practice [4, 8]. This “Dunbar number” effect can be explained in part by cognitive constraints and in part by the time costs of maintaining relationships [7, 8]. In our model, in a bistable regime the low and high connectivity levels (p_1^*N and p_2^*N in Figure 6) lead to nodal degrees that become arbitrarily large in terms of N . Hence, in the case of very large networks it would be reasonable to investigate models that incorporate some sort of saturation or carrying capacity effect in order to limit the maximum number of friendships a person may maintain.

Furthermore, it is of course possible to move from triadic closure to more general mechanisms that encourage cliques [30] or other motifs to form in a network. Here new challenges include the incorporation of appropriate higher order matrix-level nonlinearities in expressions such as (2.7), and the justification of mean-field approximations across longer-range interactions. Moreover, the network evolution model could include mechanisms that produce heterogeneous degree distributions by,

for example, encouraging preferential attachment [11]. Finally, it would be of interest to develop calibration methods along the lines of [28] in order to fit parameters to real data and search for bistability effects.

Data Availability Statement. Code to perform the computational experiments reported here can be found at <https://www.maths.ed.ac.uk/~dhigham/algfiles.html>.

REFERENCES

- [1] D. F. ANDERSON, D. J. HIGHAM, S. C. LEITE, AND R. J. WILLIAMS, *On constrained Langevin equations and (bio)chemical reaction networks*, Multiscale Model. Simul., 17 (2019), pp. 1–30.
- [2] A. R. BENSON, R. ABEBE, M. T. SCHAUB, A. JADBABAIE, AND J. KLEINBERG, *Simplicial closure and higher-order link prediction*, Proc. Natl. Acad. Sci. USA, 115 (2018), pp. E11221–E11230.
- [3] A. R. BENSON, D. F. GLEICH, AND J. LESKOVEC, *Higher-order organization of complex networks*, Science, 353 (2016), pp. 163–166.
- [4] J. BOASE, J. B. HERRIGAN, B. WELLMAN, AND L. RAINIE, *The Strength of Internet Ties*, Pew Internet Amer. Life Project, Pew Research Center, Washington, DC, 2006.
- [5] J. CARR, D. B. DUNCAN, AND C. WALSHAW, *Numerical approximation of a metastable system*, IMA J. Numer. Anal., 15 (1995), pp. 505–521.
- [6] J. L. DOOB, *Markoff chains – denumerable case*, Trans. Amer. Math. Soc., 58 (1945), pp. 422–438.
- [7] R. I. M. DUNBAR, *Neocortex size as a constraint on group size in primates*, J. Hum. Evol., 22 (1992), pp. 469–493.
- [8] R. I. M. DUNBAR, *Do online social media cut through the constraints that limit the size of offline social networks?*, R. Soc. Open Sci., 3 (2016), 3150292.
- [9] A. DUNCAN, R. ERBAN, AND K. ZYGALAKIS, *Hybrid framework for the simulation of stochastic chemical kinetics*, J. Comput. Phys., 326 (2016), pp. 398–419.
- [10] D. EASLEY AND J. KLEINBERG, *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*, Cambridge University Press, Cambridge, 2010.
- [11] N. EIKMEIER AND D. F. GLEICH, *Classes of preferential attachment and triangle preferential attachment models with power-law spectra*, J. Complex Netw., 8 (2019), cnz040.
- [12] E. ESTRADA AND F. ARRIGO, *Predicting triadic closure in networks using communicability distance functions*, SIAM J. Appl. Math., 75 (2015), pp. 1725–1744.
- [13] J. FOX, T. ROUGHGARDEN, C. SESHADRI, F. WEI, AND N. WEIN, *Finding cliques in social networks: A new distribution-free model*, SIAM J. Comput., 49 (2020), pp. 448–464.
- [14] C. W. GARDINER, *Handbook of Stochastic Methods, for Physics, Chemistry and the Natural Sciences*, 3rd ed., Springer, Berlin, 2004.
- [15] D. T. GILLESPIE, *A general method for numerically simulating the stochastic time evolution of coupled chemical reactions*, J. Comput. Phys., 22 (1976), pp. 403–434.
- [16] D. T. GILLESPIE, *Exact stochastic simulation of coupled chemical reactions*, J. Phys. Chem., 81 (1977), pp. 2340–2361.
- [17] D. T. GILLESPIE, *Markov Processes: An Introduction for Physical Scientists*, Academic Press, San Diego, CA, 1991.
- [18] D. T. GILLESPIE, *The chemical Langevin equation*, J. Chem. Phys., 113 (2000), pp. 297–306.
- [19] D. T. GILLESPIE, *Approximate accelerated stochastic simulation of chemically reacting systems*, J. Chem. Phys., 115 (2001), pp. 1716–1733.
- [20] M. S. GRANOVETTER, *The strength of weak ties*, Amer. J. Soc., 78 (1973), pp. 1360–1380.
- [21] P. GRINDROD AND D. J. HIGHAM, *Models for evolving networks: With applications in telecommunication and online activities*, IMA J. Manag. Math., 23 (2012), pp. 1–16.
- [22] P. GRINDROD, D. J. HIGHAM, AND M. C. PARSONS, *Bistability through triadic closure*, Internet Math., 8 (2012), pp. 402–423.
- [23] P. HÄNGGI, P. TALKNER, AND M. BORKOVEC, *Reaction-rate theory: Fifty years after Kramers*, Rev. Mod. Phys., 62 (1990), pp. 251–341.
- [24] D. J. HIGHAM, *Modeling and simulating chemical reactions*, SIAM Rev., 50 (2008), pp. 347–368.
- [25] J. KIM AND J. DIESNER, *Over-time measurement of triadic closure in coauthorship networks*, Soc. Netw. Anal. Min., 7 (2017), 9.

- [26] T. G. KOLDA, A. PINAR, T. PLANTENGA, C. SESHADHRI, AND C. TASK, *Counting triangles in massive graphs with MapReduce*, SIAM J. Sci. Comput., 36 (2014), pp. S44–S77.
- [27] T. LOU, J. TANG, J. HOPCROFT, Z. FANG, AND X. DING, *Learning to predict reciprocity and triadic closure in social networks*, ACM Trans. Knowl. Discov. Data, 7 (2013), 5.
- [28] A. V. MANTZARIS AND D. J. HIGHAM, *Infering and calibrating triadic closure in a dynamic network*, in Temporal Networks, P. HOLME, and J. SARMAKI, eds., Springer, Berlin, 2013, pp. 265–282.
- [29] M. MERRITT, A. ALEXANDERIAN, AND P. A. GREMAUD, *Multiscale global sensitivity analysis for stochastic chemical systems*, Multiscale Model. Simul., (2021), pp. 440–459.
- [30] K. PATTISELANNO, J. K. DIJKSTRA, C. STEGLICH, W. VOLLEBERGH, AND R. VEENSTRA, *Structure matters: The role of clique hierarchy in the relationship between adolescent social status and aggression and prosociality*, J. Youth Adolesc., 44 (2015), pp. 2257–2274.
- [31] G. A. PAVLIOTIS, *Stochastic Processes and Applications: Diffusion Processes, the Fokker-Planck and Langevin Equations*, Springer, Berlin, 2014.
- [32] G. SIMMEL, *Soziologie: Untersuchungen über die Formen der*, Suhrkamp, Frankfurt, Germany, 1992.
- [33] L. SZPRUCH AND D. J. HIGHAM, *Comparing hitting time behavior of Markov jump processes and their diffusion approximations*, Multiscale Model. Simul., 8 (2010), pp. 605–621.
- [34] H. M. TAYLOR AND S. KARLIN, *An Introduction To Stochastic Modeling*, 3rd ed., Academic Press, San Diego, CA, 1998.
- [35] L. TORRES, A. S. BLEVINS, D. BASSETT, AND T. ELIASSI-RAD, *The why, how, and when of representations for complex systems*, SIAM Rev., 63 (2021), pp. 435–485.
- [36] H. YIN, A. R. BENSON, AND J. UGANDER, *Measuring directed triadic closure with closure coefficients*, Netw. Sci., 8 (2020), pp. 551–573.