

Bounded-confidence models of opinion dynamics with neighborhood effectsSanjukta Krishnagopal^{*}*Department of Computer Science, University of California, Santa Barbara, California 93106, USA*Mason A. Porter[†]*Department of Mathematics, University of California, Los Angeles, California 90095, USA;**Department of Sociology, University of California, Los Angeles, California 90095, USA;**and Santa Fe Institute, Santa Fe, New Mexico 87501, USA*

(Received 7 February 2024; revised 22 May 2025; accepted 8 July 2025; published 19 November 2025)

People's opinions evolve through social interactions, and their social networks evolve as their opinions change. In this paper, we generalize bounded-confidence models of opinion dynamics by incorporating neighborhood effects. In our neighborhood bounded-confidence models (NBCMs), interacting agents are influenced both by each other's opinions and by the opinions of the agents in each other's neighborhoods (through so-called "transitive influence"). We extend our NBCMs to adaptive network models that incorporate "transitive homophily", with agents rewiring to connect preferentially to agents whose mean neighbor opinions are close to their own opinions. We numerically simulate an asynchronously updating adaptive NBCM on a variety of networks, and we explore how its opinion dynamics and structural properties change as we adjust the relative importances of dyadic and transitive influence. We observe that the transitive influence of neighborhoods can play an important role in shaping the qualitative features of opinion dynamics.

DOI: [10.1103/4zcl-34xy](https://doi.org/10.1103/4zcl-34xy)**I. INTRODUCTION**

Human opinions, human behavior, and social interactions influence each other in inextricable ways. Human opinions and their dynamics play an important role in many real-world settings, such as decision-making [1], misinformation and disinformation campaigns [2], referendums [3], and political polarization [4]. A major factor in all of these realms is homophily, which encapsulates the tendency of individuals to interact more with individuals that are similar to themselves than with individuals that differ from them [5].

There are myriad models of opinion dynamics [6–8]. Some opinion models (e.g., voter models [9]) have discrete-valued opinions, and other opinion models have continuous-valued opinions. Models with continuous-valued opinions include the DeGroot consensus model [10,11], the Friedkin–Johnson model [12], bounded-confidence models (BCMs) [13,14], and models with explicit radicalization dynamics [15]. Much of the literature on opinion dynamics has focused on examining conditions for consensus [8], in which the opinions of the agents in a network converge to a single value (or, more loosely, when there is a single major opinion cluster), but it is also important to study polarization (in which there are two major opinion clusters), fragmentation (in which there are three or more major opinion clusters), and other situations.

In a traditional BCM, interacting agents change their opinions if and only if their current opinions are sufficiently similar. Specifically, two agents influence each other if the

difference between their opinions lies within a confidence bound. The two foundational BCMs are the Hegselmann–Krause (HK) model (which has synchronous updates of agent opinions) [16] and the Deffuant–Weisbuch (DW) model (which has asynchronous opinion updates) [17]. In the HK model, interactions take place in groups. In an interaction, each agent compromises its opinion with all neighboring agents whose opinions lie within a confidence bound. By contrast, in the DW model, interactions are pairwise (i.e., dyadic), with a single pair of agents interacting at each discrete time. If the difference between the opinions of two interacting agents lies within a confidence bound, then they compromise their opinions by some amount. Extensions of BCMs include models that incorporate leaders [18], stubborn agents [19], heterogeneous confidence bounds [20], smooth interactions (e.g., through sigmoidal functions) [21], coevolving opinions and networks [22,23], and polyadic agent interactions (i.e., with three or more agents interacting simultaneously) [24,25].

By studying BCMs on networks, researchers examine how network structure influences opinion dynamics [26]. Most BCMs incorporate only direct influence, with interacting agents considering only each other's opinions when determining whether or not to update their opinions. However, humans are influenced not only directly but also through transitive influence (e.g., through friends of friends) [27]. They are thus influenced by the neighbors of adjacent agents. Such transitive influence is notably prevalent in interactions on social media [28], which has low barriers to interacting with friends of friends. Moreover, individuals can be influenced not only by the friends of their friends (i.e., individuals that are two steps away), but also by individuals that are even farther away in a network. In their studies of influence on social

^{*}Contact author: sanjukta@ucsb.edu[†]Contact author: mason@math.ucla.edu

networks, Christakis and Fowler [29] posited the idea of “three degrees of influence” to illustrate the distance from an agent in a network that can still yield robust influence on human behavior and opinions. Along similar lines, Miranda *et al.* [30] studied innovation adoption and observed that the adoption rate of an innovation is influenced significantly both by nearest neighbors and by extended neighborhoods. Researchers have also highlighted the importance of extended neighborhoods in empirical network data (e.g., in Facebook [31]) and in the design of algorithms (e.g., for local community detection [32]).

In the present paper, we study opinion models that incorporate transitive influence from neighbors of neighbors. We formulate *neighborhood BCMs* (NBCMs), in which an agent can change its opinion based both on dyadic interactions with adjacent nodes and on the mean opinion of the neighbors of those adjacent nodes. We refer to this mean opinion as the “mean neighbor opinion”. For the dyadic influence, we employ the usual confidence bound of BCMs, so two interacting agents compromise their opinions by some amount if they are sufficiently similar to each other (i.e., if their opinions differ by less than a confidence bound). The two agents also experience transitive influence. In such influence, an agent i is influenced by a neighboring agent with whom it interacts if the neighboring agent’s mean neighbor opinion is within agent i ’s confidence bound (i.e., if the mean opinion of the neighbor’s neighboring nodes is sufficiently similar to agent i ’s opinion). Although the mean neighbor opinion influences whether or not an agent updates its opinion, the magnitude of opinion changes are affected only by the opinions of the interacting agents themselves. We explore the qualitative behavior of an asynchronously updating NBCM for different relative proportions of direct influence (i.e., through dyadic interactions) and transitive influence (i.e., through neighborhood effects). Because of the transitive influence (and unlike in conventional BCMs), a node in our NBCMs can update its opinion due to an interaction with an adjacent node without requiring that the adjacent node also update its opinion, even when both nodes have the same confidence bound. This asymmetry can arise even when all nodes have the same confidence bound. Additionally, because of transitive influence, two interacting nodes can influence each other even when their opinions differ by more than the confidence bound.

We are not aware of any existing BCMs that consider the above neighborhood effects. However, there does exist a DeGroot model in which agents update their opinions to become closer to the opinions of agents that are neighbors of their neighbors [33], and a very recent paper incorporated neighborhood effects into a voter model [34]. The neighborhood effects in our NBCMs model “transitive homophily” (i.e., neighborhood-based homophily) and differ in a key way from the neighborhood effects in Ref. [33]. In our model, changes in agent opinions are based only on the opinions of their neighbors (i.e., by agents that are one step away in a network and hence are adjacent to them). However, neighbors of neighbors (i.e., agents that are two steps away) influence whether or not an agent changes its opinion in the first place. Our NBCMs also inherit the desirable property that individuals have heterogeneous tolerances toward others’ opinions.

Notably, this heterogeneity is present even for a constant, homogeneous confidence bound.

We also study an NBCM that coevolves with network structure. We consider *neighborhood-based adaptation*, in which an agent can form or sever connections with another agent based both on the similarity of the other agent’s opinions to their own and on the mean opinion of the neighbors of that other agent. Changes in opinions can lead to changes in relationships between agents, which in turn can lead to changes in the structure of a network. See Ref. [35] for a review of coevolving (i.e., adaptive) network models. A variety of adaptive network models have been developed to study the coevolution of opinions and network structure. Examples include adaptive voter models [36,37], adaptive BCMs and related opinion models [22,23,38–40], and others. In an adaptive opinion model, there typically is a mechanism for agents to sever ties to agents with whom they disagree (or with whom they disagree sufficiently) and then establish ties to other agents, perhaps (for convenience) in a way that preserves the density of ties (i.e., edges).

Our adaptive NBCM has two notions of homophily: (1) *dyadic homophily*, in which agents with more similar opinions are more likely to have social ties with each other (i.e., “birds of a feather flock together”); and (2) *transitive homophily*, in which agents are more likely to have social ties with agents whose mean neighbor opinions are closer to their own opinions (i.e., “you are who you know”). We employ an edge-rewiring strategy that is based on transitive homophily and is similar to the strategy in Ref. [23]. When an agent breaks an edge, it rewires this edge to form a tie with another agent with a probability that is proportional to the similarity between its opinion and that agent’s mean neighbor opinion. This rewiring strategy explicitly considers the neighborhoods of other nodes. Given network adaptation through transitive homophily, we study how network structure and opinions coevolve for different relative importances of dyadic influence (which arises from opinion similarities between adjacent agents) and transitive influence (which arises from opinion similarities between agents and the mean neighbor opinions of adjacent agents). We quantify this relative importance through a neighborhood-tuning parameter. Our NBCM simulations reveal unexpected and nonmonotonic dependence of various network measures (such as degree assortativity and the adjacency matrix’s spectral gap) on the neighborhood-tuning parameter.

Our paper proceeds as follows. In Sec. II, we discuss BCMs on networks. In Sec. III, we present our NBCMs both without and with network adaptation. In Sec. III A, we introduce our NBCMs. We formulate both a neighborhood DW (NDW) model and a neighborhood HK (NHK) model. In Sec. III B, we incorporate an edge-rewiring mechanism for network adaptation into our NDW and NHK models. In this rewiring mechanism, the mean neighbor opinions of nodes affect the severing and establishment of edges. In Sec. IV, we examine the behavior of our adaptive NDW model. We present the results of simulations of our nonadaptive NDW model on time-independent networks in Sec. IV A, present the results of simulations of our adaptive NDW model in Sec. IV B, discuss the effects of parameters on opinion dynamics in Sec. IV C, discuss simulations of our adaptive

NDW model with homophilic rewiring on different types of networks in Sec. IV D, and examine the effects of network size on the qualitative dynamics of our adaptive NDW model in Sec. IV E. In Sec. V, we summarize our results and discuss various extensions of our work. Our code is available at Ref. [41].

II. BOUNDED-CONFIDENCE MODELS (BCMs) ON NETWORKS

In a BCM, an agent is receptive to the opinions of another agent if and only if the opinions of these two agents are sufficiently similar (i.e., they lie within some confidence bound) [13,14]. Consider an unweighted and undirected network (i.e., graph) G with N agents. Each agent i is a node of the network. We denote the network's associated adjacency matrix by A , where the entry $A_{i,j} = 1$ if there is an edge between agent i and agent j and $A_{i,j} = 0$ if there is no edge between them. We assume that the network G is unweighted for simplicity, but we can consider weighted networks by letting $A_{i,j}$ be the weight of the edge between nodes i and j . We refer to a pair of nodes as a “dyad”. Suppose that each agent i has a continuous-valued opinion $x_i(t) \in [0, 1]$ at discrete time t . The vector of opinions of the N agents is the “opinion profile” $X(t) = (x_1(t), \dots, x_N(t))$ at time t . A subset of nodes in which the opinions of consecutive nodes (which we order by their opinion values) are all within a specified value ϵ of each other is an “opinion cluster”. An opinion cluster is in “consensus” when all of its nodes have the same opinion. However, it is also possible for an opinion cluster to be in a different state, such as a pseudoconsensus (see Sec. IV A).

In the DW model, opinion updates are asynchronous. At each discrete time, one chooses an edge uniformly at random; the nodes that are attached (i.e., incident) to that edge potentially compromise their opinions. Suppose that nodes i and j interact at time t . They update their opinions according to the rule

$$\begin{aligned} x_i(t+1) &= x_i(t) + \rho(x_j(t) - x_i(t))\mathbb{1}_{d_{BC}(i,j) < \epsilon_{i,j}}, \\ x_j(t+1) &= x_j(t) + \rho(x_i(t) - x_j(t))\mathbb{1}_{d_{BC}(i,j) < \epsilon_{i,j}}, \end{aligned} \quad (1)$$

where the distance between the opinions of agents i and j is $d_{BC}(i, j) = |x_i(t) - x_j(t)|$, the indicator function $\mathbb{1}_b = 1$ if condition b holds and $\mathbb{1}_b = 0$ if condition b does not hold, $\epsilon_{i,j}$ is the confidence bound, and $\rho \in (0, 0.5]$ is a constant (which is sometimes called a “convergence parameter”). The value $\rho = 0.5$ corresponds to an exact opinion compromise; in this case, two interacting agents that update their opinions both adopt the mean of their opinions. The confidence bound $\epsilon_{i,j}$ is symmetric in most studies, but one can make it asymmetric (i.e., $\epsilon_{i,j} \neq \epsilon_{j,i}$) if one desires. For a symmetric confidence bound, the opinion updates of nodes i and j are symmetric because $d_{BC}(i, j) = d_{BC}(j, i)$. For simplicity, we suppose that the confidence bound is homogeneous, so $\epsilon_{i,j} = \epsilon$ for all $i, j \in \{1, 2, \dots, N\}$. Therefore, when we discuss a confidence bound in the rest of our paper, we refer to it as “the” confidence bound. When adjacent nodes have opinions that differ by less than the confidence bound, we say that these nodes are “directly receptive” to each other.

In the HK model, opinion updates are synchronous and hence deterministic. At each discrete time, an agent interacts with all of its neighboring agents. Each agent updates its

opinion to the mean of the opinions of all neighbors with opinions within the confidence bound. That is,

$$x_i(t+1) = x_i(t) + \frac{\rho}{|\Gamma_i|} \sum_j A_{i,j} x_j(t) \mathbb{1}_{d_{BC}(i,j) < \epsilon}, \quad (2)$$

where $\Gamma_i = \{j | A_{i,j} = 1 \text{ and } |x_i - x_j| < \epsilon\}$ is the set of adjacent agents (i.e., neighbors) of agent i whose opinions are within the confidence bound. Given a network and an initial opinion profile, the HK model is deterministic. This is convenient for mathematical analysis.

Consider the HK model with a homogeneous confidence bound ϵ on a connected network (i.e., a network in which there exists a path from any node to any other node). In matrix form, Eq. (2) is $X(t+1) = W(t)X(t)$, where the opinion-update matrix W is related to the adjacency matrix and $X(t)$ is an opinion profile. If the matrix W is row-stochastic and type-symmetric (i.e., $\text{sgn}(W_{i,j}) = \text{sgn}(W_{j,i})$ and $W_{i,j} = 0$ if and only if $W_{j,i} = 0$) with nonzero diagonal entries and non-negative entries (which is true of the adjacency matrix A), then the HK model has the following properties [16,42]:

(1) Opinion updates do not change the order of the node opinions. That is, $x_i(t) \leq x_j(t)$ at any time t implies that $x_i(\tau) \leq x_j(\tau)$ for all times $\tau > t$. See Sec. 3D of Ref. [16].

(2) If the opinion difference between two nodes exceeds the confidence bound ϵ at time t , those two nodes cannot subsequently have opinions that differ by less than the confidence bound. See Sec. 3D of Ref. [16].

(3) A necessary condition to achieve a consensus opinion, which entails that $|x_i(T) - x_j(T)| \rightarrow 0$ for all nodes i and j at a finite time T , is that removing all edges between adjacent nodes that are not directly receptive to each other (i.e., neighboring nodes whose opinions differ by at least the confidence bound) results in a connected network. If this “pruned” network is disconnected at time t , then it remains disconnected at all subsequent times.

(4) By Theorem 2 of Ref. [42], the system converges to an opinion profile with consensus in each opinion cluster.

These properties are nice mathematically, but they differ markedly from typical real-life observations of opinion dynamics. Additionally, the opinion updates in the HK and DW models do not incorporate neighborhood effects. In Sec. III A, we generalize the HK and DW models to incorporate such effects. The resulting neighborhood-based HK model has rather different properties than the ones that we just discussed. Introducing neighborhood effects also yields interesting phenomena in DW models. In Sec. IV, we explore our neighborhood-based DW model in detail.

III. NEIGHBORHOOD-BASED OPINION MODELS AND NETWORK ADAPTATION

A. Neighborhood bounded-confidence models

We introduce neighborhood bounded-confidence models (NBCMs) that generalize the DW and HK models. Incorporating neighborhood effects allows us to investigate how an agent's proclivity to change its opinion depends both on the opinions of its adjacent agents and on the opinions of the neighbors of those adjacent agents. This idea encompasses

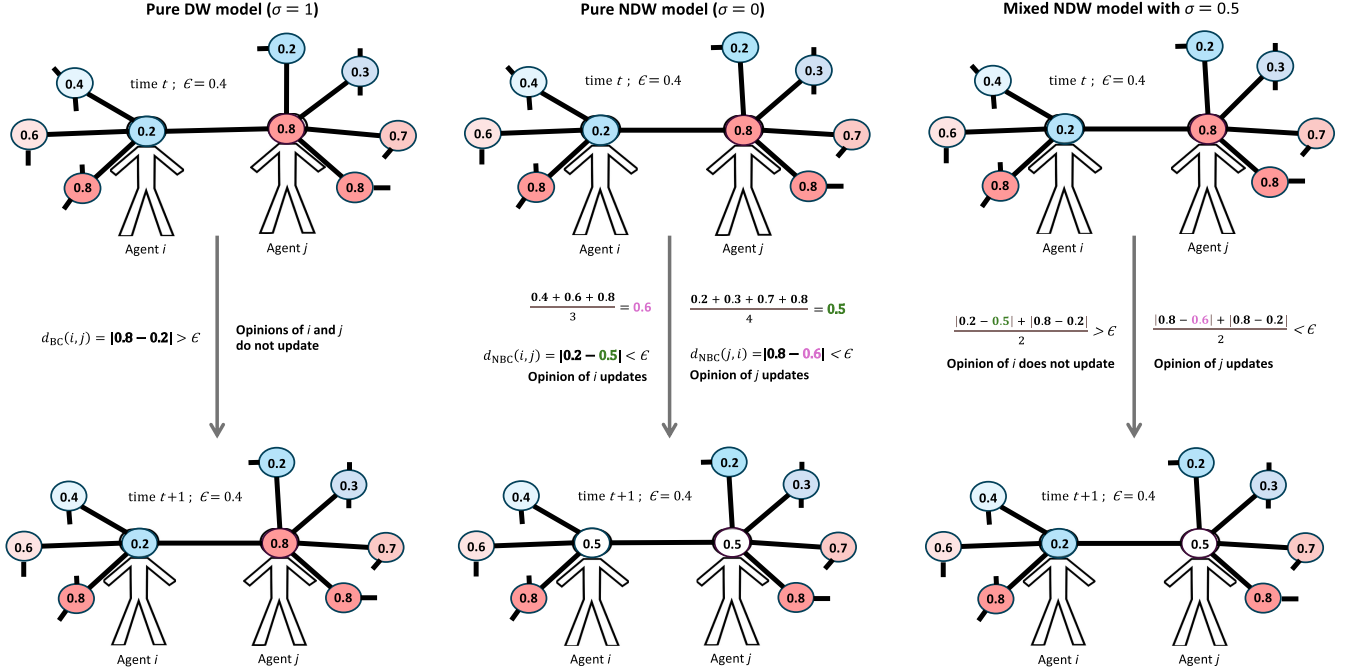


FIG. 1. A schematic illustration of opinion updates of two adjacent nodes, i and j , in the pure (i.e., baseline) DW model, pure NDW model, and the NDW model with neighborhood-tuning parameter $\sigma = 0.5$ (i.e., a “mixed” NDW model). All edges have weight 1 and convergence parameter $\rho = 1$.

the notion that “you are who you know”. For example, consider a mildly liberal agent i that is adjacent to two agents, j and k , that have the same centrist political opinion as each other but have different neighborhoods. Suppose that most nodes in j ’s neighborhood are liberals but that most nodes in k ’s neighborhood are conservatives. It seems plausible that agent i is influenced differently by agents j and k , perhaps with greater influence from agent j because its neighborhood has a political opinion that is closer to i ’s opinion. We extend both the HK and DW models to incorporate such neighborhood effects.

In our neighborhood DW (NDW) model, opinion updates are asynchronous. At each discrete time t , we choose $f = nN$ edges uniformly at random without replacement. If nN is not an integer, then we take $f = \lceil nN \rceil$ (where $\lceil \cdot \rceil$ is the ceiling function) to round up to the nearest integer. (For all reported numerical computations, nN is an integer.) The two incident nodes of each edge interact dyadically. We perform opinion updates sequentially; a node that updates its opinion uses its new opinion for subsequent interactions at time t .

Suppose that i and j are the two nodes that are incident to a chosen edge. They update their opinions according to the rule

$$\begin{aligned} x_i(t+1) &= x_i(t) + \rho(x_j(t) - x_i(t)) \mathbb{1}_{\sigma d_{BC}(i,j) + (1-\sigma)d_{NBC}(i,j) < \epsilon}, \\ x_j(t+1) &= x_j(t) + \rho(x_i(t) - x_j(t)) \mathbb{1}_{\sigma d_{BC}(i,j) + (1-\sigma)d_{NBC}(j,i) < \epsilon}, \end{aligned} \quad (3)$$

where the distance $d_{NBC}(i, j) = |x_i(t) - \frac{\sum_k A_{jk} w_k x_k(t)}{|\Gamma_j|}|$ and w_k is the weight of node k . Node weights can take a variety of possible values, and heterogeneous node weights can capture ideas like heterogeneous importances or activity levels [43].

We suppose for simplicity that $w_k = 1$ for all k . The *neighborhood-tuning parameter* $\sigma \in [0, 1]$ controls the relative weights for an agent i to potentially compromise its opinion with the opinion of its neighbor j due to (1) the similarity between node i ’s opinion and node j ’s opinion and (2) the similarity between node i ’s opinion and node j ’s mean neighbor opinion (i.e., a neighborhood-based transitive influence) [44]. With $\sigma = 1$, we recover the baseline DW model (1). With $\sigma = 0$, we obtain a “pure” NDW model. Intermediate values of σ combine the influence of opinion similarity between individual agents and neighborhood-based opinion similarity. We sometimes refer to this situation as a “mixed” NDW model. The convergence parameter ρ controls the amount that interacting agents compromise in one opinion-update step. The value $\rho = 0.5$ corresponds to an exact opinion compromise. The distance d_{BC} depends symmetrically on agents i and j , but the distance d_{NBC} depends asymmetrically on the two agents. Agent i looks at agent j ’s neighbors to determine whether or not it is influenced [45] by j , whereas agent j looks at agent i ’s neighbors to determine whether or not it is influenced by i . In Fig. 1, we show a schematic illustration of opinion updates (with exact opinion compromises) between two interacting agents for the pure (i.e., baseline) DW model, the pure NDW model, and a mixed NDW model with $\sigma = 0.5$.

In our neighborhood HK (NHK) model, opinion updates are synchronous. The opinions of the nodes update according to the equation

$$x_i(t+1) = x_i(t) + \frac{\rho}{|\Gamma_i|} \sum_j A_{i,j} x_j(t) \mathbb{1}_{\sigma d_{BC}(i,j) + (1-\sigma)d_{NBC}(i,j) < \epsilon}. \quad (4)$$

The opinion updates between two nodes, i and j , can be asymmetric (or even unidirectional) because the distances $d_{\text{NBC}}(i, j)$ and $d_{\text{NBC}}(j, i)$ need not be equal even if nodes i and j have equal confidence bounds. In fact, it is typically true that $d_{\text{NBC}}(i, j) \neq d_{\text{NBC}}(j, i)$. Setting $\sigma = 0$ yields a “pure” NHK model.

Given a network and an initial opinion profile, the opinion updates in our NHK model are deterministic. Properties of our NHK model include the following:

(1) The order of the node opinions can change with time. That is, $x_i(t) \leq x_j(t)$ for some time t does not imply that $x_i(t + \tau) \leq x_j(t + \tau)$ for all times $\tau > 0$. Switches in node-opinion order can arise from two nodes compromising their opinions when their opinions differ by more than the confidence bound (i.e., when $|x_i - x_j| > \epsilon$ for a homogeneous confidence bound ϵ) if a neighborhood’s mean opinion lies within the confidence bound. In the NHK model, it is also possible for a node’s opinion to “jump” from one opinion cluster to another. Such “opinion jumping” also arises in BCMs with polyadic interactions [24], although it occurs for a different reason.

(2) If the difference between the opinions of two nodes does not lie within a confidence bound at a certain time, these nodes can still influence each other through their neighborhoods. Their opinions can then evolve so that the difference between them lies within the confidence bound.

(3) The connectedness of a pruned network (i.e., a network that we obtain by removing edges between nodes whose opinions differ by more than the confidence bound) is no

longer a necessary condition for consensus (as it was in the HK model). It is possible for a pruned network to be disconnected at some time but then become connected later.

(4) Define the “neighborhood-pruned network” of a network to be the subnetwork that we obtain by removing all edges that are incident to at least one node whose mean neighbor opinion is not within the confidence bound of the other incident node. In the pure NHK model, a necessary condition for consensus is that the neighborhood-pruned network is connected at all times.

(5) The opinion-update matrix W is not type-symmetric because neighborhood-based influence is not symmetric (as node i can influence node j even when node j does not influence node i). That is, $d_{\text{NBC}}(i, j) = 0$ does not imply that $d_{\text{NBC}}(j, i) = 0$. Therefore, the NHK model does not satisfy the conditions in Theorem 2 of Ref. [42] that guarantee convergence to an opinion profile with consensus in each opinion cluster.

B. Neighborhood-based network adaptation

We now incorporate network adaptation into our NBCMs. In our adaptive NBCMs, agents consider both the opinions of adjacent agents and the mean neighbor opinions of adjacent agents when determining whether to establish or sever a tie. At each discrete time, we choose $f = nN$ edges uniformly at random without replacement. (We take $f = \lceil nN \rceil$ if nN is not an integer.) We then update the opinions of the nodes that are incident to each chosen edge using Eq. (3). We rewire the edge if

$$\min\{\sigma d_{\text{BC}}(i, j) + (1 - \sigma)d_{\text{NBC}}(i, j), \sigma d_{\text{BC}}(j, i) + (1 - \sigma)d_{\text{NBC}}(j, i)\} > \zeta, \quad (5)$$

where i and j are the two nodes that are incident to the edge. (If the inequality (5) holds after updating node opinions, we rewire the edge even if its incident nodes just compromised their opinions.) An edge that satisfies the inequality (5) is a “discordant” edge, so $\zeta \in [0, 1]$ is a discordance threshold. A larger value of ζ entails less tolerance of different opinions and a more stringent requirement to maintain an edge. If $\zeta \geq \epsilon$ (which is the case for all of our simulations), then it is only possible to sever an edge if the distance between the opinions of its incident nodes is at least as large as the confidence bound ϵ . When $\zeta = \epsilon$, we rewire any chosen edge in which at least one of the incident nodes does not compromise its opinion. The asymmetry in Eq. (5) is important. It is possible to sever an edge between nodes i and j if node i ’s opinion is sufficiently different from node j ’s opinion, even if node j ’s opinion is not far enough from i ’s opinion to exceed the discordance threshold. It is also

possible to sever an edge between nodes i and j if node i ’s opinion is sufficiently different from node j ’s mean neighbor opinion, even if node j ’s mean neighbor opinion is not far enough from i ’s opinion to exceed the discordance threshold.

When we remove an edge, we select one of its incident nodes, with equal probability of each, to rewire to a new node. Suppose that we select node i . Node i considers each node k in the set K of nodes that are not currently in its neighborhood, including the node from which it just detached. We employ one of the following two rewiring strategies:

(1) *Random rewiring*: Node i attaches to a random node (irrespective of the opinion of that node). In this random rewiring strategy, node i chooses a node $k \in K$ uniformly at random. Henceforth, whenever we use the term “random rewiring”, we mean this uniformly random rewiring strategy.

(2) *Homophilic rewiring*: Node i attaches to a node $k \in K$ with probability

$$P(i \rightarrow k) = \frac{1 - (\sigma d_{\text{BC}}(i, k) + (1 - \sigma)d_{\text{NBC}}(i, k))}{\sum_k (1 - (\sigma d_{\text{BC}}(i, k) + (1 - \sigma)d_{\text{NBC}}(i, k)))}. \quad (6)$$

We set $P(i \rightarrow i) = 0$ to prevent self-edges. For a progressively larger opinion distance $\tilde{d} := \sigma d_{\text{BC}}(i, k) + (1 - \sigma)d_{\text{NBC}}(i, k)$

between nodes i and k , there is a progressively smaller probability that node i rewires to attach to k . The neighborhood-

ALGORITHM 1: Pseudocode for our algorithm to simulate the NDW model with network adaptation.

```

Consider a graph  $G = (V, E)$  with a set  $V$  of nodes and a set  $E$  of edges.
# Initialization
Initialize the opinion of each node to a uniformly random value in  $[0, 1]$ .
Choose  $\text{rewire} \in \{\text{random}, \text{homophilic}\}$ .
for  $t < t_{\max}$  or convergence = False do
  # Opinion spreading
  Choose  $f = nN$  edges.      ★ (If  $nN$  is not an integer, then  $f = \lceil nN \rceil$ .)
  for each chosen edge  $e \in E$ :
    Consider nodes  $v_1$  and  $v_2$  that are incident to  $e$ .
    Update the opinions of nodes  $v_1$  and  $v_2$  using Eq. (3).
  # Opinion-driven network adaptation
  if  $e$  is discordant [see Eq. (5)]:
    Remove the edge  $e$ .
    Select a node that is incident to the edge  $e$ , with equal probability of the two incident nodes. Denote the chosen node by  $z$ .
    Let  $V_z = \{i \in V \text{ such that } A_{i,z} = 0\}$ .
    if  $\text{rewire} = \text{random}$ :
      Choose a node  $k \in V_z$  uniformly at random.
      Add an edge so that  $A_{k,z} = 1$ .
    if  $\text{rewire} = \text{homophilic}$ :
      Choose a node  $k \in V_z$  with probability given by Eq. (6).
      Add an edge so that  $A_{k,z} = 1$ .
  end for

```

tuning parameter σ allows us to interpolate between two different types of homophily:

(a) Dyadic homophilic rewiring: Node i is more likely to attach to nodes whose opinions are closer to its opinion (as in the rewiring strategy in Ref. [23]).

(b) Transitive homophilic rewiring: Node i is more likely to attach to nodes with a mean neighbor opinion that is closer to its opinion.

We consider purely dyadic homophilic rewiring by setting $\sigma = 0$. Such rewiring captures the idea that people are more likely to befriend somebody that they perceive as similar to themselves. We consider purely transitive homophilic rewiring by setting $\sigma = 1$. Such rewiring captures the idea that people are more likely to befriend somebody with friends that they perceive as similar to themselves.

In Algorithm 1, we outline our algorithm for our NDW model with neighborhood-based network adaptation.

IV. NUMERICAL SIMULATIONS AND BEHAVIOR OF THE ADAPTIVE NDW MODEL

We conduct numerical simulations to examine the qualitative behavior of our adaptive NDW model.

A. Opinion states in our adaptive NDW model with homophilic rewiring

In our simulations, we consider the following types of opinion states:

- (1) Consensus: A system is in consensus at a given time if all agents have the same opinion at that time.
- (2) Polarization: A system is polarized at a given time if there are two opinion clusters at that time. This notion of

polarization does not consider the sizes of these clusters (no matter how small they are).

(3) Fragmentation: A system is fragmented at a given time if there are three or more opinion clusters at that time. We again do not consider the sizes of these clusters.

(4) Pseudoconsensus: Consider a set of nodes with at least two distinct opinion values. If the difference in the opinions between two consecutive nodes (which we order by their opinion values) is always less than the confidence bound, then the set of nodes is in pseudoconsensus. (For example, a set of nodes with opinion values $\{0.1, 0.2, 0.3\}$ is in pseudoconsensus for $\epsilon = 0.2$.) Polarized and fragmented states can have opinion clusters in pseudoconsensus. In our simulations, we choose numerical tolerances to quantify what it means for node opinions to differ from each other. We categorize nodes as having the “same” opinion when their opinion values differ by less than 10^{-6} . In practice, in our computations, all nodes with “different” opinions differ in their opinion values by at least 10^{-2} .

Our nonadaptive and adaptive NDW models can have a variety of qualitative dynamics. Examples of such dynamics are the following:

- (1) Convergence: There exists a time T such that the opinion profile $X(T) = X(T + t)$ for all $t \geq 0$.
- (2) Pseudoconvergence: There exists a time T such that the opinion profile $X(T) = X(T + t)$ for all $t \in [0, \tau)$ but $X(T) \neq X(T + \tau)$.
- (3) Temporary consensus: Nodes that have the same opinions at some time but then fall out of consensus at a later time are in a temporary consensus when they all have the same opinion.
- (4) Opinion jumping: This phenomenon occurs when a node’s opinion changes by a value that is more than the

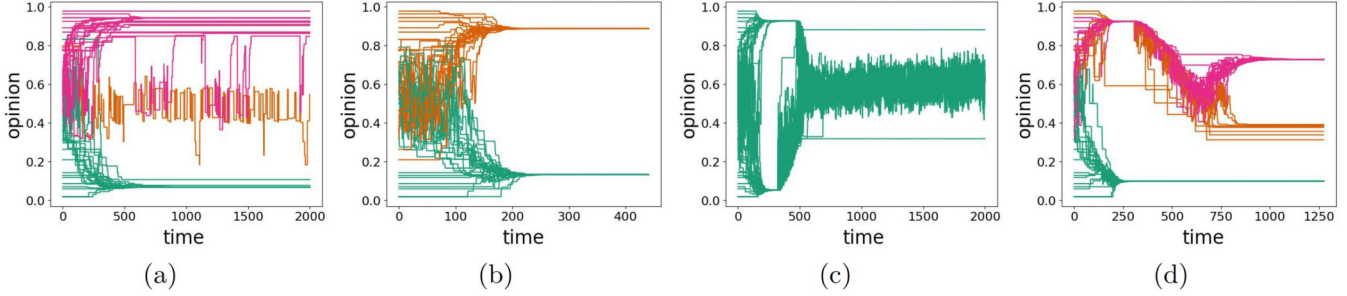


FIG. 2. Several examples of opinion dynamics in our neighborhood DW (NDW) model with homophilic rewiring. In each panel, we show one simulation on a network. The neighborhood-tuning parameter σ , the confidence bound ϵ , and the discordance threshold ζ have the values (a) $\sigma = 0$, $\epsilon = 0.1$, and $\zeta = 0.4$; (b) $\sigma = 0.1$, $\epsilon = 0.2$, and $\zeta = 0.2$; (c) $\sigma = 0$, $\epsilon = 0.2$, and $\zeta = 0.4$; and (d) $\sigma = 0.1$, $\epsilon = 0.1$, and $\zeta = 0.3$. At each discrete time, we consider $f = 0.2N$ dyads, where $N = 100$ is the size (i.e., the number of nodes) of the network. The convergence parameter is $\rho = 0.3$. We choose edges uniformly at random; if an edge is discordant, we rewire it using the homophilic rewiring strategy (6). It is possible for the same node to rewire multiple times. In each simulation, the initial network is a $G(N, p)$ Erdős-Rényi (ER) graph with an independent, homogeneous probability $p = 0.3$ of an edge between each pair of nodes. We initialize each node opinion to a uniformly random value in the interval $[0, 1]$. All depicted simulations use the same initial network and the same set of initial opinions. We terminate a simulation either when it reaches our stopping criterion or when $t_{\max} = 2000$ time steps have elapsed (whichever occurs first). We color the opinion trajectory of each node according to its final opinion cluster.

confidence bound ϵ in a single time step. When a node's opinion jumps, it can move from one opinion cluster to another. Opinion jumping can also occur in DW models with polyadic interactions [24].

(5) Opinion crossing: Consider two nodes i and j with opinions $x_i(t)$ and $x_j(t)$, respectively, at time t . If $x_i(t) > x_j(t)$ but $x_i(t + \tau) < x_j(t + \tau)$ for some time $\tau > 0$, then the opinions cross each other. Opinion crossing can occur in both our NDW model and our NHK model. It can also occur in the standard DW model, but it cannot occur in the standard HK model (under the mild assumptions that we stated in Sec. II).

In Fig. 2, we show examples of various behaviors in the NDW model (3) with homophilic rewiring. In each simulation, the initial network is a $G(N, p)$ Erdős-Rényi (ER) graph with $N = 100$ nodes and an independent, homogeneous probability $p = 0.3$ of an edge between each pair of nodes. We initialize the node opinions to uniformly random values in the interval $[0, 1]$. At each discrete time, we repeat a three-step process for $f = nN$ distinct edges (as before, we take $f = \lceil nN \rceil$ if nN is not an integer), which we choose uniformly at random without replacement. For each selected edge, we update the opinions of its incident nodes according to the opinion-update rule (3). If the edge is discordant, we rewire it using one of the strategies in Sec. III B. (In Fig. 2, we use the homophilic rewiring strategy (6).) We end a simulation either when it satisfies a stopping criterion or when $t_{\max} = 2000$ time steps have elapsed (whichever occurs first). Our stopping criterion is that no node changes its opinion by more than 10^{-3} in any time step for 200 consecutive time steps.

In Fig. 2(a), we show a fragmented state with pseudoconsensus in the pink opinion cluster (top) and green opinion cluster (bottom). The orange opinion cluster (middle) has not converged within 2000 time steps. For example, we observe opinion jumping at approximately times 600 and 1900; some opinion values (in the pink cluster) change by more than the confidence bound ϵ in a single time step. In Fig. 2(b), we show

an example of polarization. In this simulation, there are also several opinion crossings. In Fig. 2(c), two opinion clusters have converged (to opinions with values of approximately 0.9 and 0.3) and one opinion cluster has not converged within 2000 time steps. The opinions in this last cluster oscillate wildly between values of about 0.45 and 0.75. In Fig. 2(d), we observe a variety of interesting behaviors, including fragmentation into three opinion clusters, pseudoconsensus in the orange opinion cluster (middle), consensus in the pink (top) and green (bottom) opinion clusters, opinion crossing (which is most visible at early times), and opinion jumping (which is most visible in orange). Additionally, the pink opinion cluster has a temporary consensus between the approximate times 250 and 350.

B. Opinion dynamics in our adaptive NDW model with homophilic rewiring

In our adaptive NDW model, the structure of a network can change at each discrete time. These structural changes, in turn, affect the opinion dynamics on the network. In Fig. 3, we show the time evolution of the network properties in an adaptive network with homophilic rewiring (6). We compare our adaptive NDW model to a baseline adaptive DW model for different values of the neighborhood-tuning parameter $\sigma \in [0, 1]$, where $\sigma = 1$ corresponds to pure DW opinion updates and rewiring that depends only on dyadic homophily and $\sigma = 0$ corresponds to pure NDW opinion updates and rewiring that depends only on transitive homophily. As we will illustrate, the network properties depend in an interesting way on σ . In our simulations, the initial network is a $G(N, p)$ ER graph with $N = 50$ nodes and an independent, homogeneous probability $p = 0.3$ of an edge between each pair of nodes. We initialize the node opinions to uniformly random values in $[0, 1]$. We indicate the other parameter values in the caption of Fig. 3. Each curve is a mean of 20 simulations with the same parameter values but different randomizations,

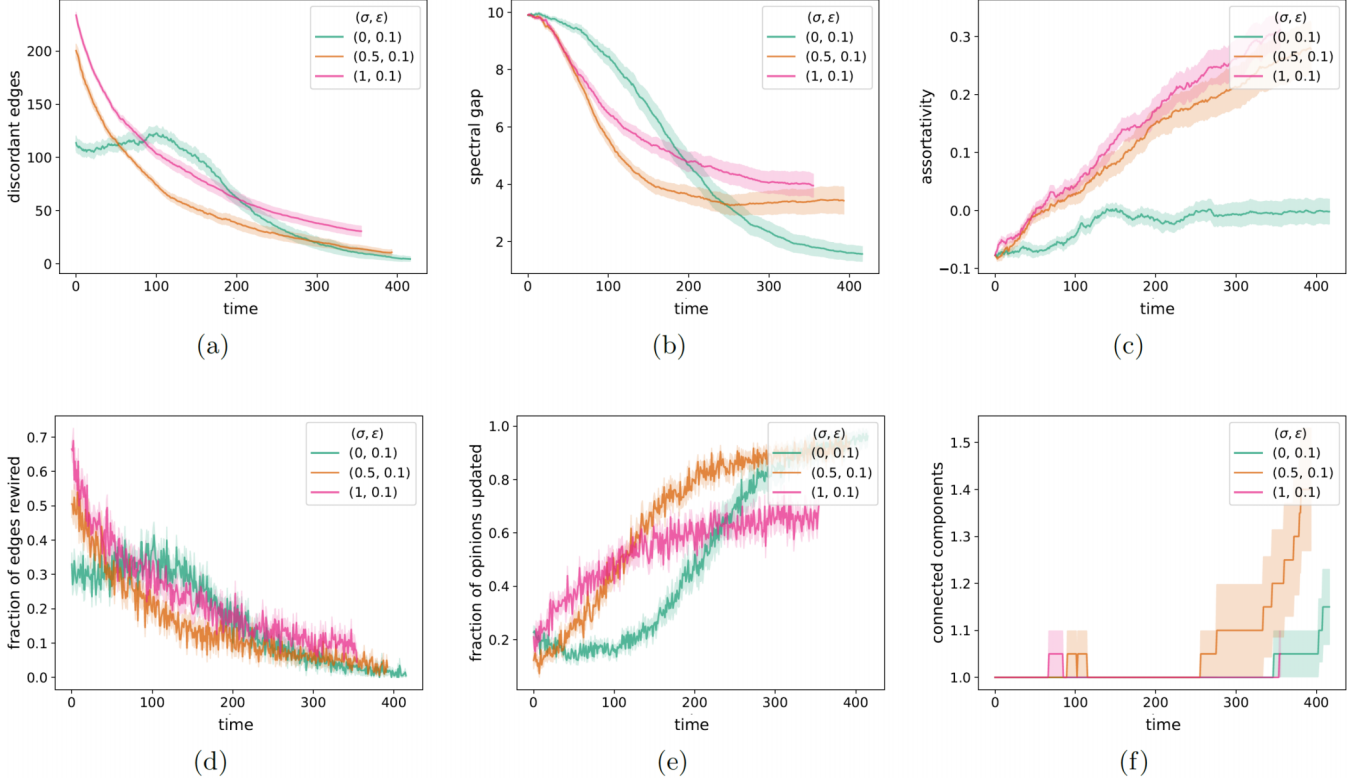


FIG. 3. We examine network adaptation as a function of time in our NDW model by plotting (a) the number of discordant edges, (b) the spectral gap of the network's adjacency matrix, (c) a degree-assortativity coefficient, (d) the fraction of edges in the chosen dyads that rewired (i.e., that are discordant), (e) the fraction of nodes in the chosen dyads that update their opinions, and (f) the mean number of connected components. We compare the network properties for a pure DW opinion-update rule (i.e., $\sigma = 1$), a mixed DW and NDW opinion-update rule with $\sigma = 0.5$, and a pure NDW opinion-update rule (i.e., $\sigma = 0$). In each simulation, the initial network is a $G(N, p)$ ER network with $N = 50$ nodes and an independent, homogeneous probability $p = 0.3$ of an edge between each pair of nodes. For each network, we initialize each node opinion to a uniformly random value in $[0, 1]$. We plot means of 20 simulations, with the same 20 initial networks and sets of initial opinions for each panel. The shaded regions indicate the standard errors. The confidence bound is $\epsilon = 0.1$, the discordance threshold is $\zeta = 0.2$, the number of edges that we choose at each discrete time for interaction is $f = 0.2N$, and the convergence parameter is $\rho = 0.3$. We terminate a simulation either when it reaches our stopping criterion or when $t_{\max} = 2000$ time steps have elapsed (whichever occurs first).

including both different ER graphs and different sets of initial opinions, with standard errors indicated by the shaded regions. All simulations in this subsection meet our stopping criterion.

In Fig. 3(a), we plot the time evolution of the number of discordant edges. We observe a marked difference between the pure DW model, in which the number of discordant edges decreases almost monotonically as a function of time, and the pure NDW model, in which the number of discordant edges tends to increase early in simulations (after a small initial dip) and subsequently decreases almost monotonically. In Fig. 3(b), we show the time evolution of the spectral gap of the adjacency matrix A . The spectral gap is the absolute value of the difference between the two largest eigenvalues of A . For a time-independent network, the spectral gap is inversely proportional to the relaxation time of a standard random walk on the network (and it is also related to network community structure) [48,49]. Intuitively, if nodes consider the neighborhoods of other nodes when adapting their opinions, there is an averaging effect that disincentivizes the clustering of opinions into distinct communities. Notably, the behavior of the NDW model with $\sigma = 0.5$, which assigns equal importance to transitive influence and dyadic influence, does not simply

interpolate between the pure DW model and the pure NDW model.

In Fig. 3(c), we plot the time evolution of the degree-assortativity coefficient, which is equal to the Pearson correlation coefficient between the node degrees of adjacent nodes [50]. A positive value of degree assortativity indicates a positive correlation between nodes of similar degrees. For all examined values of the neighborhood-tuning parameter σ and the confidence bound ϵ , we observe that degree assortativity tends to increase with time, as expected for homophilic rewiring. Additionally, the degree assortativity is much smaller for the pure NDW model (for which $\sigma = 0$) than for the other examined situations. In our NDW model, one tends to select nodes with larger degrees for possible opinion updates. Consequently, for simulations that start with uniformly distributed opinions, when we uniformly random select edges for a potential compromise between their incident nodes, we observe that larger-degree nodes' opinions are more likely than smaller-degree nodes' opinions to approach their neighbors' opinions. For the pure DW model (for which $\sigma = 1$), one can thus expect large-degree nodes to often attach to other large-degree nodes, leading to a large degree asso-

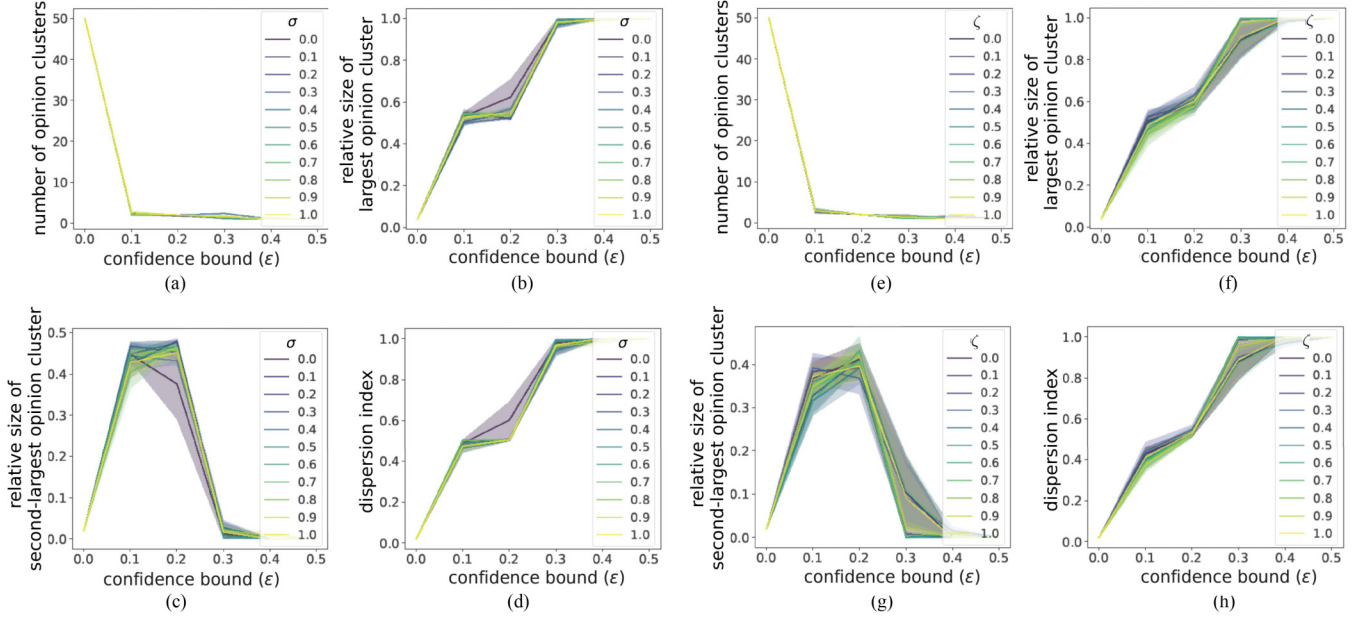


FIG. 4. We illustrate the dependence of several final opinion-profile properties in our NDW model with homophilic rewiring on the confidence bound ϵ for different values of (a, b, c, d) the neighborhood-tuning parameter σ and (e, f, g, h) the discordance threshold ζ . We show (a, e) the number of opinion clusters, (b, f) the fraction of nodes in the largest opinion cluster, (c, g) the fraction of nodes in the second-largest opinion cluster, and (d, h) the dispersion index Δ . In each simulation, the initial network is a $G(N, p)$ ER graph with $N = 50$ nodes and an independent, homogeneous probability $p = 0.3$ of an edge between each pair of nodes. In panels (a)–(d), the discordance threshold is $\zeta = 0.2$. In panels (e)–(h), the neighborhood-tuning parameter is $\sigma = 0.5$. For each network, we initialize each node opinion to a uniformly random value in $[0, 1]$. The number of edges that we select at each discrete time for agents to interact is $f = 0.2N$, and the convergence parameter is $\rho = 0.3$. We plot means of 5 simulations, with the same 5 initial networks and the same sets of initial opinions for each panel. The shaded regions indicate the standard errors.

ciativity. This feature is tempered in the NDW model, where neighborhoods also play a role in determining the adjacencies between nodes, so one can expect the pure NDW model with purely transitive homophilic rewiring to have smaller degree assortativity than the pure DW model with purely dyadic homophilic rewiring. Once again, the mixed NDW model does not simply interpolate between the pure DW model and the pure NDW model.

In Fig. 3(d), we plot the time evolution of the fraction of the f chosen edges in a time step that we rewire due to discordance. Recall that we consider f edges (i.e., f dyads) at each discrete time, but we only rewire the discordant edges. We observe a pattern that resembles the one in Fig. 3(a), although Fig. 3(d) is much noisier. In Fig. 3(e), we examine the time evolution of the fraction of the nodes in the chosen dyads that update their opinions. It is unsurprising but still worth remarking that the fraction of edges that are discordant is not inversely correlated with the fraction of opinion updates. Edge discordance and opinion updates are governed by different parameters. Interestingly, although the pure NDW model initially has the fewest opinion updates, it eventually has the most opinion updates. We observe noticeable differences in the trajectories of the fraction of opinion updates for the three values of the neighborhood-tuning parameter σ .

In Fig. 3(f), we plot the time evolution of the mean number of connected components of the networks. The pure NDW model and the pure DW model both typically yield the same

mean number of connected components (usually 1), although the pure NDW model ultimately yields more components than the pure DW model. Interestingly, we eventually obtain the largest mean number of connected components for $\sigma = 0.5$, which uses a mixture of the NDW and standard DW opinion-update rules.

C. Effect of parameters on opinion clusters in the NDW model with homophilic rewiring

We calculate a few different quantities to study the properties of the opinion clusters in the NDW model with homophilic rewiring. We examine how the number of opinion clusters, the relative sizes of the two largest opinion clusters, and the dispersion index [51] change as we vary the confidence bound ϵ , the neighborhood-tuning parameter σ , and the discordance threshold ζ . All simulations in this subsection meet our stopping criterion.

The dispersion index is [51]

$$\Delta = \frac{\sum_i s_i^2}{(\sum_i s_i)^2}, \quad (7)$$

where s_i is the size of the i^{th} opinion cluster. An opinion profile with r equal-sized clusters yields $\Delta = 1/r$.

In Fig. 4, we show the dependence of several final opinion-profile properties on the confidence bound ϵ for different values of the neighborhood-tuning parameter σ . In Fig. 4(a),

we see for all values of σ that the number of opinion clusters decreases dramatically as we increase the confidence bound from $\epsilon = 0$ to $\epsilon = 0.1$. When $\epsilon = 0$, nodes do not change their opinions. In Figs. 4(b) and 4(c), we show the relative sizes (with respect to the total number of nodes) of the largest and second-largest opinion clusters. The relative size of the largest opinion cluster tends to grow with ϵ and appears to reach 1 for $\epsilon \geq 0.4$ for all values of σ . By contrast, the relative size of the second-largest cluster depends nonlinearly on ϵ , with a peak size between $\epsilon = 0.1$ and $\epsilon = 0.2$. The dependence on σ is less obvious, although the curve corresponding to $\sigma = 0$ (i.e., the pure NDW model) appears to differ noticeably from those for the other values of σ . The sum of the relative sizes of the two largest opinion clusters [see Figs. 4(b) and 4(c)] is smaller than 1 for some parameter values, so there are more than two opinion clusters in these situations. However, these additional opinion clusters are very small for most parameter values. In Fig. 4(d), we observe that the dispersion index Δ tends to increase with ϵ . We also see that the dispersion index follows a very similar trend as the size of the largest opinion cluster. In Figs. 4(e)–4(h), we observe similar trends with the confidence bound ϵ as we did in Figs. 4(a)–4(d). Additionally, we do not observe a clear dependence on the discordance threshold ζ , which does not directly influence the opinions themselves. Instead, ζ influences which edges are rewired.

As we can see in Fig. 4, the dynamics of our adaptive NDW model depend strongly on the confidence bound ϵ and depend weakly on the neighborhood-tuning parameter σ and the discordance threshold ζ . The dependence on ϵ is familiar from studies of ordinary BCMs [14,26]. For small confidence bounds ϵ , the opinions of neighboring nodes readily differ by more than the confidence bound, and then nodes are unlikely to influence each other. This situation tends to yield many opinion clusters. As we increase ϵ , more nodes influence each other and we obtain fewer opinion clusters. The observed nonmonotonic dependence on ϵ of the sizes of the largest and second-largest opinion clusters [see Figs. 4(b), 4(d), 4(f), and 4(h)] is interesting and worthy of further investigation.

D. Simulations of our adaptive NDW model with homophilic rewiring on different types of initial networks

We now simulate our adaptive NDW model on several additional types of initial networks. These simulations complement our previous simulations on ER graphs. All simulations in this subsection meet our stopping criterion. We examine the following types of networks:

(1) Holme–Kim power-law cluster graphs: A Holme–Kim graph [52] is a generalization of the standard Barabási–Albert (BA) preferential-attachment graph [50,53] that also adds triangles. One starts with a set of isolated nodes. In each preferential-attachment step, one adds a new node and connects it to m existing nodes with probabilities that are proportional to their degrees. With probability \tilde{p} , for each new edge from the preferential-attachment step, one adds another edge and forms a triangle by connecting the new node to a neighbor (which one selects uniformly at random) of the previously linked node. With probability $1 - \tilde{p}$ or if no such new edges are possible, one instead performs another preferential-

attachment step. By adding triangles, one increases the triadic clustering in the network.

(2) Newman–Watts–Strogatz (NWS) small-world graphs [54]: To create an NWS graph, one starts with a ring network. One then connects each node of the ring to its r nearest neighbors, where r is an even positive integer. One then adds “shortcut” edges as follows. For each edge of the graph, with probability \tilde{p} , one adds a new edge from one of its incident nodes (which one selects uniformly at random) to a uniformly randomly chosen node.

(3) The Zachary Karate Club (ZKC) graph [55]: The ZKC graph is an astoundingly popular 34-node social network of friendships in a karate club. This real-world network has well-studied community structure [56].

In Fig. 5, we show the results of our simulations of our adaptive NDW model with homophilic rewiring on these different networks. As one can see in Figs. 5(a), 5(d), and 5(g) and in Fig. 3(a) (where we considered ER networks), the time evolution of the number of discordant edges is qualitatively similar for all of the examined networks. The number of discordant edges decreases with time for all examined networks, and the pure adaptive NDW model (i.e., with $\sigma = 0$) appears to have a plateau-like structure at early times. Additionally, of the examined scenarios, the spectral gap decreases to the smallest value for our pure adaptive NDW model. See Figs. 5(b), 5(e), and 5(h) [and also see Fig. 3(b)]. The spectral gap for the Holme–Kim graph decreases with time, whereas the spectral gaps for the NWS and ZKC graphs increase initially before they too tend to decay. The spectral gap for the ZKC graph has additional structure, perhaps due to its strong community structure. Finally, as one can see in Figs. 5(c), 5(f), and 5(i), the degree-assortativity plateaus at the smallest value for the pure adaptive NDW model, the second-smallest value for the mixed NDW model, and the largest value for the baseline DW model. We observe the same order of degree assortativities for ER graphs [see Fig. 3(c)]. The approach to the final degree-assortativity value appears to be qualitatively different for the NWS graph than for the other graphs.

E. Effect of network size for ER graphs

We now study the effect of network size on the qualitative dynamics of our adaptive NDW model. We simulate the adaptive NDW model with homophilic rewiring on $G(N, p)$ ER networks with connection probability $p = 0.3$ and $N = 100$, $N = 200$, and $N = 400$ nodes. In Fig. 6, we show the results of these simulations. Unsurprisingly, we observe that larger networks take longer to reach a steady state. In our simulations with $N = 50$ nodes (see Fig. 3), we observe convergence (up to a 10^{-3} error tolerance), with simulations achieving our stopping criterion before our simulations run for $t_{\max} = 2000$ time steps. However, for $N = 400$ nodes, our simulations of the pure adaptive NDW model with purely transitive homophilic rewiring (i.e., with $\sigma = 0$) run for $t_{\max} = 2000$ time steps without achieving our stopping criterion.

We briefly summarize a few qualitative features of Fig. 6. In the left column of Fig. 6, we observe that the number of discordant edges in our pure adaptive NDW model with

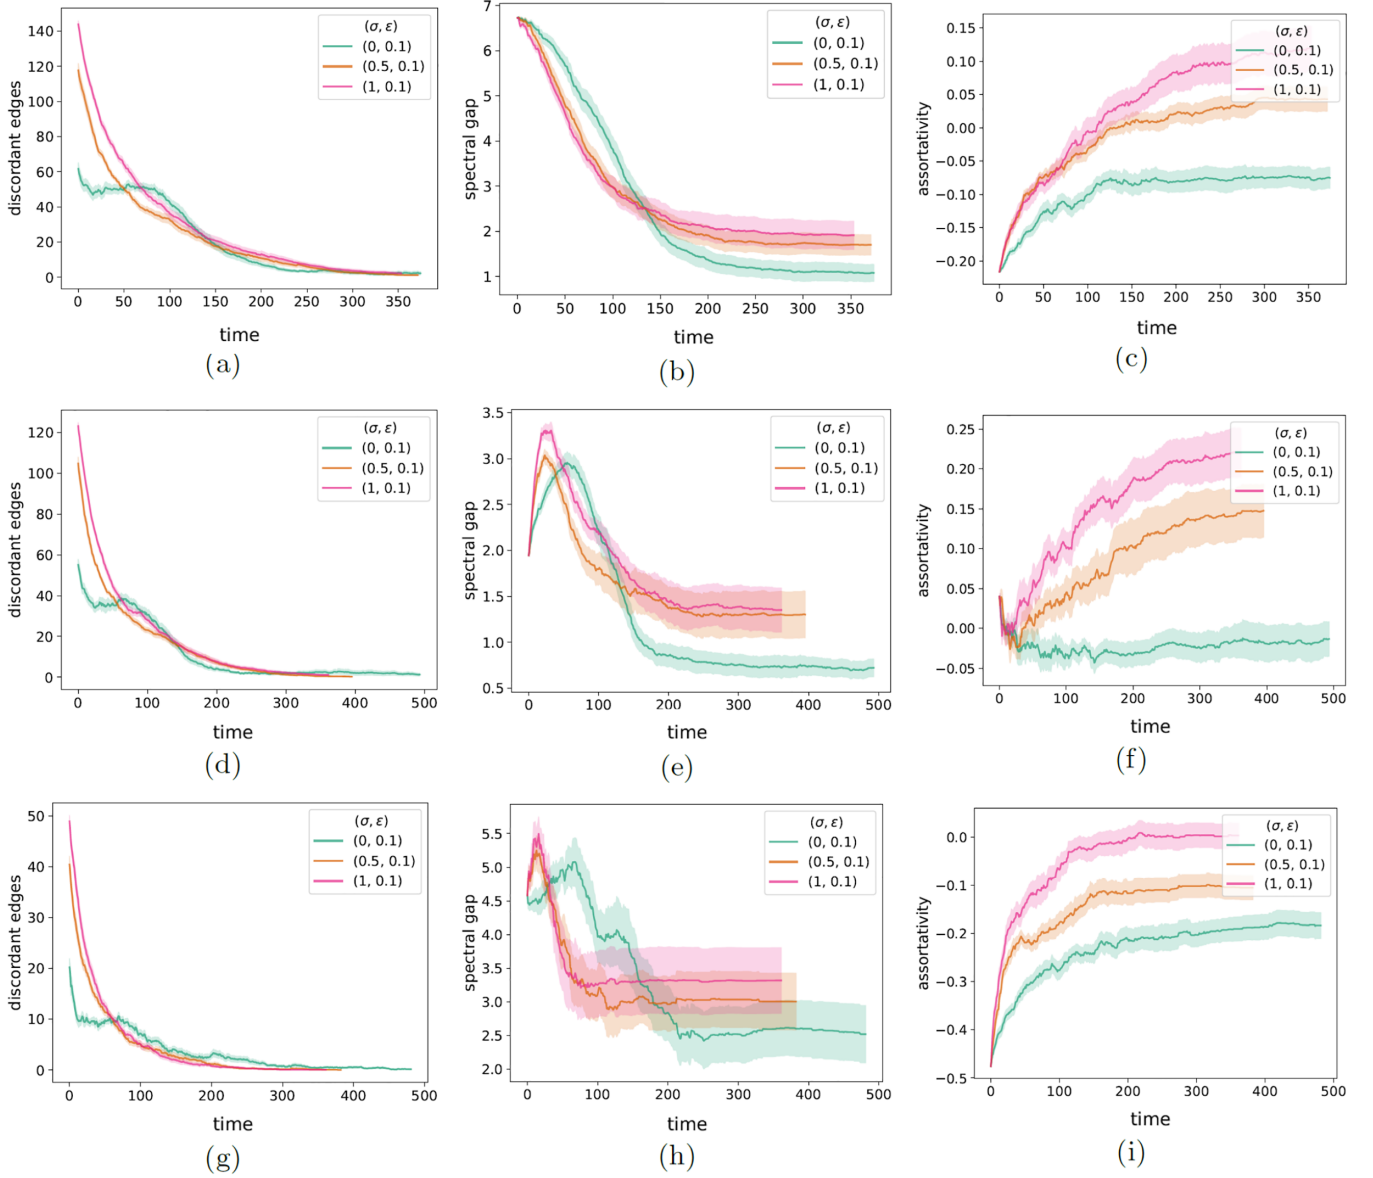


FIG. 5. We simulate our adaptive NDW model with homophilic rewiring on three types of networks. We consider (top) a Holme–Kim power-law graph with clustering, (middle) a Newman–Watts–Strogatz (NWS) small-world graph, and (bottom) the Zachary Karate Club graph. For each type of graph, we plot (a, d, g) the number of discordant edges in the network, (b, e, h) the spectral gap of the adjacency matrix, and (c, f, i) a degree-assortativity coefficient. We compare the network properties of the baseline adaptive DW model (i.e., the neighborhood-tuning parameter is $\sigma = 1$), a mixed adaptive NDW model (with $\sigma = 0.5$), and a pure adaptive NDW model (i.e., $\sigma = 0$). For the Holme–Kim graph, we start with 5 isolated nodes, add a new node to attach to 5 existing nodes using linear preferential attachment, and incorporate additional edges for triadic closure with probability $\tilde{p} = 0.3$. For the NWS small-world graph, each node is adjacent to 6 nearest neighbors and the probability of adding a new edge for each edge is $\tilde{p} = 0.3$. The Holme–Kim and NWS graphs have $N = 50$ nodes, and the Zachary Karate Club graph has $N = 34$ nodes. We initialize each node opinion to a uniformly random value in the interval $[0, 1]$. Each curve in each panel is a mean of 5 simulations with different initial opinions. We use the same Holme–Kim graph and the same NWS graph for all simulations. The shaded regions indicate the standard errors. The confidence bound is $\epsilon = 0.1$, the discordance threshold is $\zeta = 0.2$, the number of edges that we choose at each discrete time for interaction is $f = 0.2N$, and the convergence parameter is $\rho = 0.3$. We terminate a simulation either when it reaches our stopping criterion or when $t_{\max} = 2000$ time steps have elapsed (whichever occurs first).

purely transitive homophilic rewiring (i.e., with $\sigma = 0$) increases with time at first and then eventually decreases. By contrast, for $\sigma = 1$ (i.e., for the baseline adaptive DW model with purely dyadic homophilic rewiring), the number of discordant edges consistently decreases with time. In the middle

column of Fig. 6, we observe that the spectral gap decreases much more sharply for the pure DW model (with $\sigma = 1$) and the mixed NDW model ($\sigma = 0.5$) than for the pure NDW model (with $\sigma = 0$). Finally, in the right column of Fig. 6, we observe that the pure NDW model with purely transitive

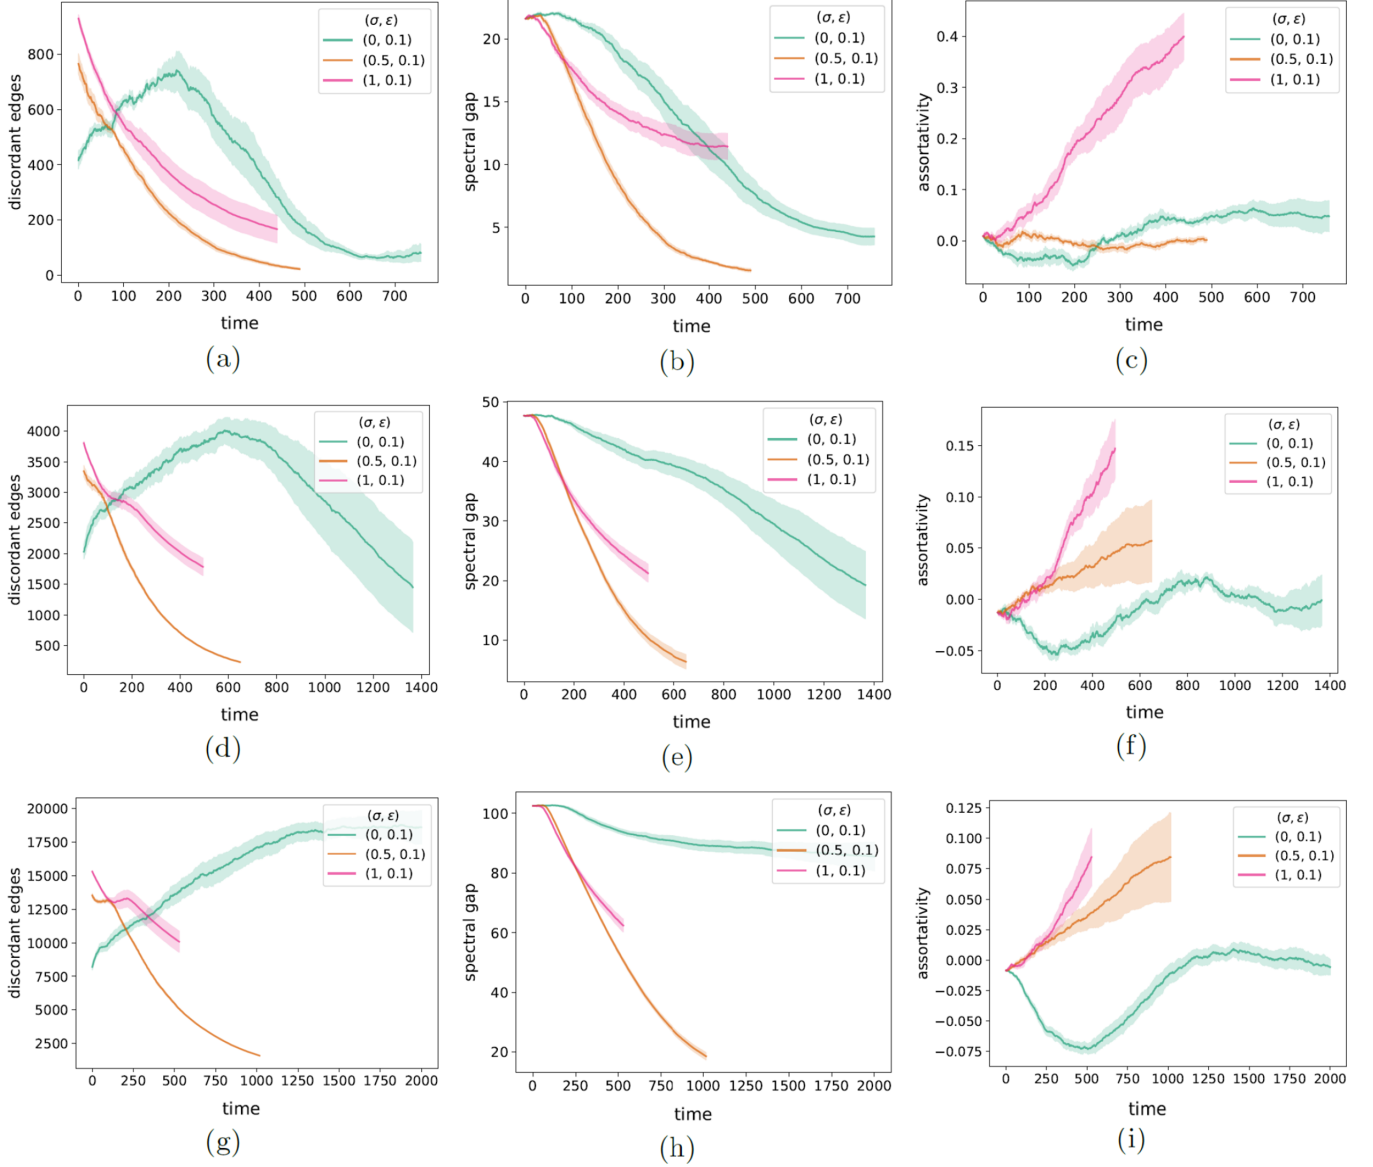


FIG. 6. We simulate our adaptive NDW model with homophilic rewiring on $G(N, p)$ ER graphs with (top) $N = 100$ nodes, (middle) $N = 200$ nodes, and (bottom) $N = 400$ nodes. For each of these situations, we plot (a, d, g) the number of discordant edges, (b, e, h) the spectral gap of the adjacency matrix, and (c, f, i) a degree-assortativity coefficient. We compare the network properties of a pure DW model with purely dyadic homophilic rewiring (i.e., $\sigma = 1$), a mixed NDW model (specifically, with $\sigma = 0.5$), and a pure NDW model with purely transitive homophilic rewiring (i.e., with $\sigma = 0$). In each simulation, the initial network is a $G(N, p)$ ER network (with $N = 100$, $N = 200$, and $N = 400$ nodes in the top, middle, and bottom rows, respectively) and an independent, homogeneous probability $p = 0.3$ of an edge between each pair of nodes. We initialize each node opinion to a uniformly random value in $[0, 1]$. We plot means of 5 simulations, with the same 5 initial networks and sets of initial opinions for each row. The shaded regions indicate the standard errors. The confidence bound is $\epsilon = 0.1$, the discordance threshold is $\zeta = 0.2$, the number of edges that we choose at each discrete time for interaction is $f = 0.2N$, and the convergence parameter is $\rho = 0.3$. We terminate a simulation either when it reaches our stopping criterion or when $t_{\max} = 2000$ time steps have elapsed (whichever occurs first).

homophilic rewiring has the smallest degree assortativity of the three examined situations for the ER networks with $N = 200$ and $N = 400$ nodes. For the ER networks with $N = 100$ nodes, the pure and mixed NDW models have similar degree assortativities.

In Fig. 7, we show the dependence of several final opinion-profile properties on the confidence bound ϵ for ER networks of different sizes. We observe that the relative sizes of the opinion clusters (as a fraction of the total network size) do not depend strongly on network size.

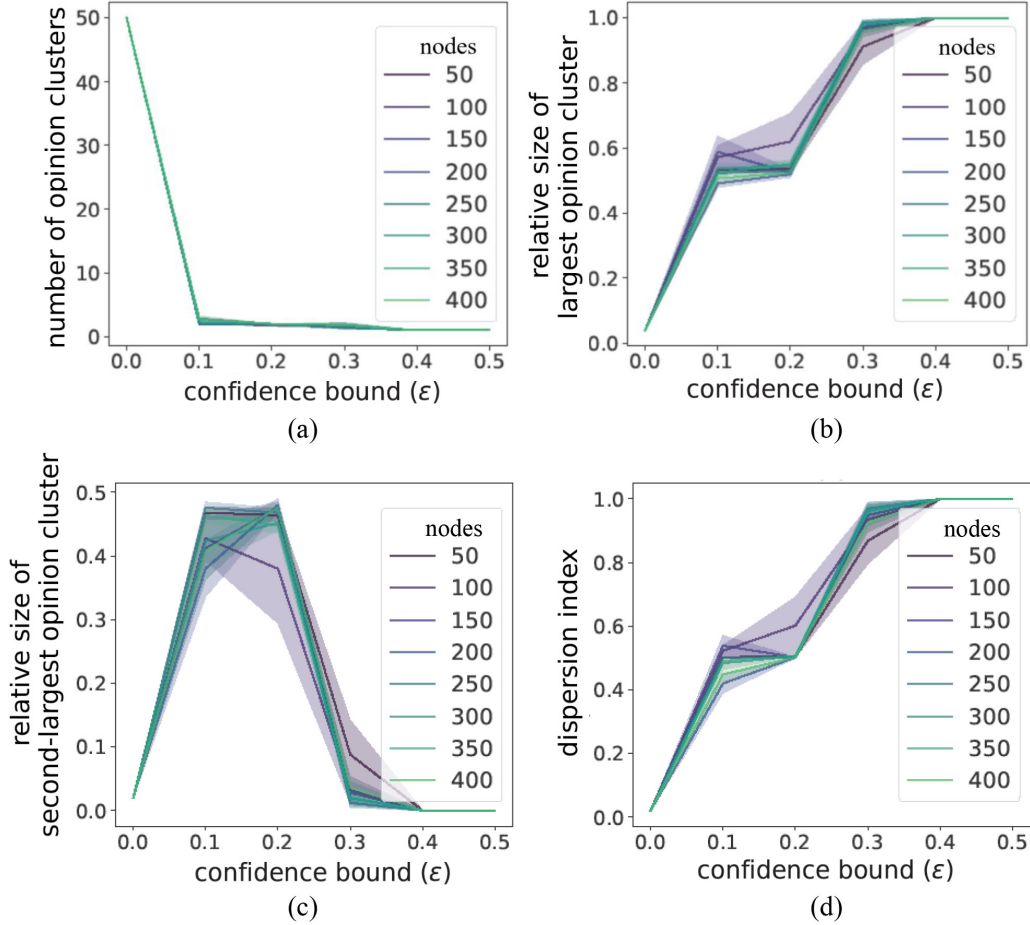


FIG. 7. We illustrate the dependence of several final opinion-profile properties in our adaptive NDW model with homophilic rewiring on the confidence bound ϵ for $G(N, p)$ ER networks of different sizes N . We show (a) the number of opinion clusters, (b) the fraction of nodes in the largest opinion cluster, (c) the fraction of nodes in the second-largest opinion cluster, and (d) the dispersion index Δ . In each simulation, the initial network is a $G(N, p)$ ER graph with N nodes and an independent, homogeneous probability $p = 0.3$ of an edge between each pair of nodes. For each network, we initialize each node opinion to a uniformly random value in $[0, 1]$. The neighborhood-tuning parameter is $\sigma = 0.5$, the discordance threshold is $\zeta = 0.2$, the number of edges that we select at each discrete time for agents to interact is $f = 0.2N$, and the convergence parameter is $\rho = 0.3$. We plot means of 5 simulations, with the same 5 initial networks and sets of initial opinions for each panel. The shaded regions indicate the standard errors.

V. CONCLUSIONS AND DISCUSSION

We studied adaptive bounded-confidence models (BCMs) of opinion dynamics that incorporate neighborhood effects into both opinion dynamics and network adaptation. In addition to the usual dyadic influence of BCMs, these neighborhood BCMs (NBCMs) incorporate transitive influence in determining whether or not agents compromise their opinions when they interact with other agents. For such transitive influence to occur, the mean opinions of the neighbors of the neighbors (e.g., the friends of friends) of agents need to be sufficiently similar to their own opinions. They thereby incorporate a notion of “you are who you know” into opinion updates. We formulated neighborhood-informed generalizations of both the Hegselmann–Krause (HK) model and the Deffuant–Weisbuch (DW) model, and we argued that the neighborhood HK (NHK) model includes qualitative behavior (such as changes in the order of node opinions with time) that cannot occur in the standard HK model.

We then developed adaptive NDW and NHK models with neighborhood-informed network adaptation, and we examined the qualitative behavior of the adaptive NDW model in various scenarios. Our neighborhood-informed homophilic rewiring strategy is based on transitive homophily and yields interesting network properties. For example, in our adaptive NDW model, we typically observed a smaller degree assortativity and often observed a smaller spectral gap than in our baseline adaptive DW model. In the adaptive NDW model, we also observed nonmonotonic behavior in the number of discordant edges as a function of time.

Neighborhood-based transitivity—which can arise through both transitive homophily and transitive influence—exerts notable effects on human behavior [27–29,31]. It is thus important to incorporate such ideas into models of opinion dynamics. However, as with all other models of opinion dynamics, our NBCMs have several notable limitations. In particular, we made many simplistic assumptions about

human behavior. For example, all of the agents in our models are identical and have homogeneous confidence bounds. It is relevant to explore heterogeneities in these and other features.

There are a variety of ways to extend our NBCMs. These extensions include both commonly noted possibilities (such as the incorporation of multidimensional opinions, heterogeneous confidence bounds, and polyadic interactions [8]) and generalizations that align directly with our model's neighborhood focus. In particular, we considered only the two-step neighbors of agents (e.g., friends of friends), and it will be interesting to examine generalizations of our NBCMs that include more expansive neighborhoods. Additionally, it is worthwhile to extend other types of opinion models to explore neighborhood effects, and our study of NBCMs provides an illustrative example to help guide such efforts. Another interesting direction is to combine our NBCMs with models of

disease spread to explore how neighborhood-based opinion dynamics and disease dynamics affect each other.

ACKNOWLEDGMENTS

M.A.P. acknowledges financial support from the National Science Foundation (Grant No. 1922952) through the Algorithms for Threat Detection (ATD) program. S.K. acknowledges support from the UC Presidential Postdoctoral Fellowship. We thank two anonymous referees for their many helpful comments.

DATA AVAILABILITY

The only data in the present paper is simulation data, which is not available. The code that we used to produce our simulations is available at Ref. [41].

-
- [1] M. A. Milburn, *Persuasion and Politics: The Social Psychology of Public Opinion* (Thomson Brooks/Cole Publishing Company, Pacific Grove, CA, USA, 1991).
 - [2] Z. Guo, J. Valinejad, and J.-H. Cho, *IEEE Trans. Netw. Sci. Eng.* **9**, 3775 (2022).
 - [3] T. T. C. Lin, *Asian J. Commun.* **32**, 152 (2022).
 - [4] J. T. Jost, D. S. Baldassarri, and J. N. Druckman, *Nat. Rev. Psychol.* **1**, 560 (2022).
 - [5] M. McPherson, L. Smith-Lovin, and J. M. Cook, *Annu. Rev. Sociol.* **27**, 415 (2001).
 - [6] H. Xia, H. Wang, and Z. Xuan, *Int. J. Knowl. Syst. Sci.* **2**, 72 (2011).
 - [7] H. Noorazar, K. R. Vixie, A. Talebanpour, and Y. Hu, *Int. J. Mod. Phys. C* **31**, 2050101 (2020).
 - [8] M. Starnini, F. Baumann, T. Galla, D. Garcia, G. Iñiguez, M. Karsai, J. Lorenz, and K. Sznajd-Weron, *arXiv:2507.11521*.
 - [9] S. Redner, *C. R. Phys.* **20**, 275 (2019).
 - [10] J. R. P. French Jr., *Psychol. Rev.* **63**, 181 (1956).
 - [11] M. H. DeGroot, *J. Am. Stat. Assoc.* **69**, 118 (1974).
 - [12] N. E. Friedkin and E. C. Johnsen, *J. Math. Sociol.* **15**, 193 (1990).
 - [13] J. Lorenz, *Int. J. Mod. Phys. C* **18**, 1819 (2007).
 - [14] C. Bernardo, C. Altafini, A. Proskurnikov, and F. Vasca, *Automatica* **159**, 111302 (2024).
 - [15] F. Baumann, P. Lorenz-Spreen, I. M. Sokolov, and M. Starnini, *Phys. Rev. Lett.* **124**, 048301 (2020).
 - [16] R. Hegselmann and U. Krause, *J. Artif. Soc. Soc. Simul.* **5**(3), 2 (2002).
 - [17] G. Weisbuch, G. Deffuant, F. Amblard, and J.-P. Nadal, *Complexity* **7**, 55 (2002).
 - [18] Y. Zhao, L. Zhang, M. Tang, and G. Kou, *Comput. Oper. Res.* **74**, 205 (2016).
 - [19] Y. Tian and L. Wang, *Automatica* **96**, 213 (2018).
 - [20] P. Sobkowicz, *Front. Phys.* **3**, 17 (2015).
 - [21] H. Z. Brooks, P. S. Chodrow, and M. A. Porter, *SIAM J. Appl. Dyn. Syst.* **23**, 1442 (2024).
 - [22] M. Brede, *Complexity* **2019**, 1486909 (2019).
 - [23] U. Kan, M. Feng, and M. A. Porter, *J. Complex Netw.* **11**, cnac055 (2023).
 - [24] A. Hickok, Y. Kureh, H. Z. Brooks, M. Feng, and M. A. Porter, *SIAM J. Appl. Dyn. Syst.* **21**, 1 (2022).
 - [25] H. Schawe and L. Hernández, *Commun. Phys.* **5**, 32 (2022).
 - [26] X. F. Meng, R. A. Van Gorder, and M. A. Porter, *Phys. Rev. E* **97**, 022312 (2018).
 - [27] G. Liu, Y. Wang, and M. Orgun, *Proceedings of the AAAI Conference on Artificial Intelligence* (Association for the Advancement of Artificial Intelligence, Washington, DC, USA, 2011), Vol. 25, pp. 1222–1229.
 - [28] N. A. Christakis and J. H. Fowler, *Connected: The Surprising Power of our Social Networks and How They Shape our Lives* (Little, Brown Spark, New York, NY, USA, 2009).
 - [29] N. A. Christakis and J. H. Fowler, *Stat. Med.* **32**, 556 (2013).
 - [30] M. Miranda, M. Pereda, A. Sánchez, and E. Estrada, *PNAS Nexus* **3**, pgae409 (2024).
 - [31] J. Ugander, B. Karrer, L. Backstrom, and C. Marlow, *arXiv:1111.4503*.
 - [32] L. G. S. Jeub, P. Balachandran, M. A. Porter, P. J. Mucha, and M. W. Mahoney, *Phys. Rev. E* **91**, 012821 (2015).
 - [33] Q. Zhou, Z. Wu, A. H. Altalhi, and F. Herrera, *Info. Sci.* **519**, 363 (2020).
 - [34] F. J. Muñoz and J. C. Nuño, *arXiv:2411.16671*.
 - [35] R. Berner, T. Gross, C. Kuehn, J. Kurths, and S. Yanchuk, *Phys. Rep.* **1031**, 1 (2023).
 - [36] P. Holme and M. E. J. Newman, *Phys. Rev. E* **74**, 056108 (2006).
 - [37] R. T. Durrett, J. P. Gleeson, A. L. Lloyd, P. J. Mucha, F. Shi, D. Sivakoff, J. E. S. Socolar, and C. Varghese, *Proc. Natl. Acad. Sci. USA* **109**, 3682 (2012).
 - [38] M. Del Vicario, A. Scala, G. Caldarelli, H. E. Stanley, and W. Quattrociocchi, *Sci. Rep.* **7**, 40391 (2017).
 - [39] B. Kozma and A. Barrat, *Phys. Rev. E* **77**, 016102 (2008).
 - [40] A. Parravano, A. Andina-Diaz, and M. A. Meléndez-Jiménez, *PloS ONE* **11**, e0164323 (2016).
 - [41] See <https://bitbucket.org/neighborhood-bounded-confidence-model-of-opinion-dynamics>.
 - [42] J. Lorenz, *Physica A* **355**, 217 (2005).
 - [43] G. J. Li and M. A. Porter, *Phys. Rev. Res.* **5**, 023179 (2023).
 - [44] Node i 's opinion contributes to this mean.
 - [45] In some prior studies, such as in investigations of DeGroot models [46,47], researchers have interpreted influence as a notion of trust. With such an interpretation, the amount of trust that agent i places in agent j depends on j 's neighborhood.

- [46] W. Xia, M. Cao, and K. H. Johansson, *IEEE Trans. Control Netw. Syst.* **3**, 46 (2015).
- [47] P. Liu, Y. Li, and P. Wang, *IEEE Trans. Fuzzy Syst.* **31**, 307 (2022).
- [48] N. Masuda, M. A. Porter, and R. Lambiotte, *Phys. Rep.* **716–717**, 1 (2017).
- [49] P. Van Mieghem, *Graph Spectra for Complex Networks*, 2nd ed. (Cambridge University Press, Cambridge, UK, 2023).
- [50] M. E. J. Newman, *Networks*, 2nd ed. (Oxford University Press, Oxford, UK, 2018).
- [51] B. Derrida and H. Flyvbjerg, *J. Phys. A: Math. Gen.* **19**, L1003 (1986).
- [52] P. Holme and B. J. Kim, *Phys. Rev. E* **65**, 066109 (2002).
- [53] A.-L. Barabási and R. Albert, *Science* **286**, 509 (1999).
- [54] M. E. J. Newman and D. J. Watts, *Phys. Lett. A* **263**, 341 (1999).
- [55] W. W. Zachary, *J. Anthropol. Res.* **33**, 452 (1977).
- [56] M. Girvan and M. E. J. Newman, *Proc. Natl. Acad. Sci. USA* **99**, 7821 (2002).