Title

Some Generalizations of Bounded-Confidence Models of Opinion Dynamics

Author

Li, Grace

Publication Date

2024

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA
Los Angeles

Some Generalizations of Bounded-Confidence Models of Opinion Dynamics

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Mathematics

by

Grace Jingying Li

2024

ABSTRACT OF THE DISSERTATION

Some Generalizations of Bounded-Confidence
Models of Opinion Dynamics

by

Grace Jingying Li
Doctor of Philosophy in Mathematics
University of California, Los Angeles, 2024
Professor Mason Alexander Porter, Chair

Bounded-confidence models (BCMs) are a type of model of opinion dynamics with continuous-valued opinions. The two most popular BCMs are the Deffuant–Weisbuch (DW) model [DNA00] and the Hegselmann–Krause (HK) model [HK02]. In a BCM, interacting pairs of agents only influence each other if their opinions differ by less than some "confidence bound" between them. One major challenge in the study of opinion dynamics is making models more realistic and applicable to real-world data and situations. In this dissertation, we develop and study three generalizations of BCMs. They each incorporate some mechanism to make them more realistic, while maintaining tractability.

We first develop and study a generalization of the DW model that uses node weights to model heterogeneous agent-activity levels. The node weights in this BCM allow us to consider individuals in a social network that share their ideas and opinions more frequently than others. Using numerical simulations, we systematically investigate the effects of node weights, which we assign uniformly at random to the nodes. We demonstrate that introducing heterogeneous node weights results in longer convergence times and more opinion

fragmentation than in an associated baseline DW model.

We then investigate BCMs in which each pair of agents has a distinct confidence bound that changes when the pair interacts. The confidence bounds in these BCMs encode the mutual willingness of agents to consider each other's opinions. We demonstrate numerically that our adaptive BCMs tend to promote consensus and yield longer convergence times than the associated baseline BCMs. We also show that these adaptive BCMs can have neighboring agents that converge to the same opinion but are not receptive to each other. This qualitative behavior does not occur in the associated baseline BCMs.

Finally, we study BCMs that have multi-dimensional opinions that consist of multiple interdependent topics. When a pair of agents interact on a topic, whether or not they compromise their opinions depends on the differences in their opinions on all topics. Using numerical simulations, we demonstrate for these BCMs that the choice of initial opinion distribution has a large effect on the amount of opinion fragmentation.

The dissertation of Grace Jingying Li is approved.

Deanna Needell

Jacob G. Foster

Andrea Bertozzi

Mason Alexander Porter, Committee Chair

University of California, Los Angeles

2024

TABLE OF CONTENTS

1 Introduction **1**

2 Background on Networks **6**

2.1 Some network definitions 6

2.2 Communities in networks 8

2.3 Some examples of networks 9

2.3.1 Complete graphs 9

2.3.2 Erdős–Rényi (ER) graphs 9

2.3.3 Stochastic-block-model (SBM) graphs 10

3 Background on Bounded-Confidence Models (BCMs) of Opinion Dynamics **12**

3.1 The Deffuant–Weisbuch (DW) model 12

3.2 The Hegselmann–Krause (HK) model 14

3.3 Characterizing opinions in BCMs 15

3.3.1 Convergence of opinions 15

3.3.2 Effective graphs and opinion clusters 16

3.3.3 Opinion consensus and fragmentation 18

3.3.4 Quantifying opinion fragmentation 19

3.3.5 Convergence time 22

3.4 Some generalizations of BCMs 24

4 A Bounded-Confidence Model with Heterogeneous Node-Activity Levels 28

4.1 Introduction and motivation 28
4.2 A BCM with heterogeneous node-selection probabilities 32
4.2.1 Relating node weights to edge weights 34
4.3 Details of numerical simulations 34
4.3.1 Network structures 34
4.3.2 Node-weight distributions 37
4.3.3 Simulation specifications 40
4.4 Numerical simulations and results 42
4.4.1 Simulations on a complete graph 43
4.4.2 ER graphs 53
4.4.3 SBM graphs 55
4.4.4 Caltech network 57
4.4.5 Finite-size effects 61
4.5 Conclusions and discussion 65
4.5.1 Summary of our main results 65
4.5.2 Relating node weights to network structure 66
4.5.3 Relating node weights to node opinions 67
4.5.4 Edge-based heterogeneous activities 68
4.5.5 Importance of node weights 69

5 Bounded-Confidence Models with Adaptive Confidence Bounds 72

5.1 Introduction and motivation 72
5.1.1 Related work 73
5.2 BCMs with adaptive confidence bounds 76

5.2.1 Our HK model with adaptive confidence bounds 76
5.2.2 Our DW model with adaptive confidence bounds 78
5.3 Theoretical results 79
5.3.1 Adaptive-confidence HK model 79
5.3.2 Adaptive-confidence DW model 82
5.4 Details of our numerical simulations 89
5.4.1 Network structures 89
5.4.2 Simulation specifications 90
5.4.3 Quantifying model behaviors 92
5.5 Results of our numerical simulations 95
5.5.1 Adaptive-confidence HK model 96
5.5.2 Adaptive-confidence DW model 109
5.5.3 Network of network-scientist coauthorships 115
5.6 Conclusions and discussion 117
5.6.1 Summary and discussion of our results 117
5.6.2 Future work 118

6 Bounded-Confidence Models with Topic-Weighted Opinion Discordance 120

6.1 Introduction and motivation 120
6.2 BCMs with topic-weighted opinion discordance 122
6.2.1 Our HK model with topic-weighted opinion discordance 123
6.2.2 Our DW model with topic-weighted opinion discordance 124
6.3 Preliminary theoretical results 125
6.3.1 Convergence of opinions 125

6.3.2 Region of absorption 127
6.4 Details of our numerical simulations 131
6.4.1 Initial opinion distributions 131
6.4.2 Simulation specifications 133
6.5 Preliminary results of our numerical simulations 137
6.5.1 Topic-weighted HK model 137
6.5.2 Topic-weighted DW model 144
6.6 Discussion and next steps 146

7 Conclusions 149

A A Linear Analysis of Shock–Grain Interaction for Inertial Confinement Fusion (ICF) Ablators 151

B Detection of Overlapping Communities in Undirected Graphs with Personalized PageRank 153

References 157

LIST OF FIGURES

3.1 Sample trajectories of agent opinions in a DW model. 20

4.1 Convergence times in simulations of our node-weighted BCM on a 500-node complete graph with various node-weight distributions. 46

4.2 Steady-state numbers of major opinion clusters in simulations of our node-weighted BCM on a 500-node complete graph with various node-weight distributions. . . 48

4.3 Shannon entropies of the steady-state opinion-cluster profiles in simulations of our node-weighted BCM on a 500-node complete graph with various node-weight distributions. 50

4.4 Steady-state numbers of minor opinion clusters in simulations of our node-weighted BCM on a 500-node complete graph with various node-weight distributions. . . 51

4.5 Sample trajectories of agent opinions versus time t in a single simulation of our node-weighted BCM on a 500-node complete graph. 52

4.6 Shannon entropies of the steady-state opinion-cluster profiles in simulations of our node-weighted BCM on $G(500, p)$ ER random graphs for various node-weight distributions and several values of p. 54

4.7 Steady-state numbers of minor opinion clusters in simulations of our node-weighted BCM on the Caltech Facebook network with various distributions of node weights. 58

4.8 Steady-state numbers of major opinion clusters in simulations of our node-weighted BCM on the Caltech Facebook network with various distributions of node weights. 59

4.9 Shannon entropies of the steady-state opinion-cluster profiles in simulations of our node-weighted BCM on the Caltech Facebook network with various node-weight distributions. 60

4.10 Histogram of the node degrees of the Caltech Facebook network. 61

4.11 Convergence times in simulations of our node-weighted BCM on complete graphs of various sizes. 62

4.12 Shannon entropies of the steady-state opinion-cluster profiles in simulations of our node-weighted BCM on complete graphs of various sizes. 63

5.1 Examples of final effective graphs with $W(T_f) < 1$. 94

5.2 The numbers of major clusters in simulations of our adaptive-confidence HK model on a 1000-node complete graph for various combinations of the BCM parameters γ, δ, and c_0. 99

5.3 The weighted-average edge fraction $W(T_f)$ (see equation Equation (5.4.1)) in simulations of our adaptive-confidence HK model on a 1000-node complete graph for various combinations of the BCM parameters γ, δ, and c_0. 100

5.4 The convergence times (in terms of the number of time steps) on a logarithmic scale in simulations of our adaptive-confidence HK model on a 1000-node complete graph for various combinations of the BCM parameters γ, δ, and c_0. 102

5.5 The numbers of major clusters in simulations of our adaptive-confidence HK model on $G(1000, p)$ ER random graphs with (A–E) $p = 0.1$ and (F–J) $p = 0.5$ for various combinations of the BCM parameters γ, δ, and c_0. 103

5.6 The numbers of major clusters in simulations of our adaptive-confidence HK model on 1000-node SBM random graphs with connection probabilities $P_{aa} = P_{bb} = 1$ and $P_{ab} = 0.01$ for various combinations of the BCM parameters γ, δ, and c_0. 106

5.7 The numbers of major clusters in simulations of our adaptive-confidence HK model on the UC Santa Barbara network for various combinations of the BCM parameters γ, δ, and c_0. 107

5.8 The Shannon entropies in simulations of our adaptive-confidence HK model on the UC Santa Barbara network for various combinations of the BCM parameters γ, δ, and c_0. 108

5.9 The numbers of major clusters in simulations of our adaptive-confidence HK model on the Reed College network for various combinations of the BCM parameters γ, δ, and c_0. 109

5.10 The numbers of major clusters in simulations of (A) the baseline DW model and (B–J) our adaptive-confidence DW model on a 100-node complete graph for various combinations of the BCM parameters γ, δ, c_0, and m. 112

5.11 The weighted-average edge fraction $W(T_f)$ (see equation Equation (5.4.1)) in simulations of (A) the baseline DW model and (B–D) our adaptive-confidence DW model on a 100-node complete graph for various combinations of the BCM parameters γ, δ, c_0, and m. In (E), in which we show our simulations with $(\gamma, \delta) = (0.1, 0.5)$, we run all of our simulations to convergence (i.e., we ignore the bailout time) and use the resulting final opinion clusters. 114

5.12 The numbers of major clusters in simulations of (A) the baseline DW model and (B, C) our adaptive-confidence DW model on the NetScience network for various combinations of the BCM parameters γ, δ, c_0, and m. 116

5.13 The numbers of minor clusters in simulations of (A) the baseline DW model and (B, C) our adaptive-confidence DW model on the NetScience network for various combinations of the BCM parameters γ, δ, c_0, and m. 116

6.1 Boundaries of the region of absorption in 2D for $c = 0.225$ and various values of ω 130

6.2 Opinion on topic 2 given the opinion on topic 1 in our wedge distribution. . . . 132

6.3 The probability density functions for the Gaussian($\sigma = 0.22, \rho = 0$) and Gaussian($\sigma = 0.22, \rho = 0.8$) distributions of initial opinions in our simulations of our topic-weighted BCMs. 134

6.4 Numbers of major limit opinion clusters in simulations of our topic-weighted HK model on a 2000-node complete graph with various initial opinion distributions and various values of the BCM parameters c and ω. 138

6.5 Numbers of minor limit opinion clusters in simulations of our topic-weighted HK model on a 2000-node complete graph with various initial opinion distributions and various values of the BCM parameters c and ω. 139

6.6 Limit order parameters $Q(T_f)$ in simulations of our topic-weighted HK model on a 2000-node complete graph with various initial opinion distributions and various values of the BCM parameters c and ω. 141

6.7 The limit opinion clusters in simulations of our topic-weighted HK model and the baseline HK model with $c = 0.1$ and Gaussian initial opinion distributions. . . . 143

6.8 The limit opinion clusters in simulations of our topic-weighted HK model and the baseline HK model with $c = 0.2$ for the independent uniform and wedge initial opinion distributions. 145

6.9 Limit order parameters $Q(T_f)$ in simulations of our topic-weighted DW model on a 500-node complete graph with various values of the BCM parameters c and ω. 147

LIST OF TABLES

4.1 The networks on which we simulate our node-weighted BCM. 35

4.2 The names and specifications of our node-weight distributions. 37

4.3 Summary of the trends in our simulations of our node-weighted BCM. 44

5.1 The real-world networks on which we simulate our adaptive-confidence BCMs. . 90

5.2 The BCM parameter values that we examine in simulations of our adaptive-confidence BCMs. 91

5.3 Summary of the observed trends in our adaptive-confidence HK model. 97

5.4 Summary of the observed trends in our adaptive-confidence DW model. 111

5.5 Summary of the numbers of simulations of our adaptive-confidence DW model that reach the bailout time of 10^6 time steps. 113

ACKNOWLEDGMENTS

People are connected in many ways, and we are influenced by our interactions with one another. This journey would not have been possible without my interactions with and the support of many people, including my family, friends, mentors, and collaborators.

I am extremely thankful to my incredible advisor and mentor, Mason Porter. Thank you for being so generous with your time and spending hours discussing research ideas, deliberating ways to visualize and convey results, giving detailed feedback on writing, and providing advice on career paths and application materials. Mason has been incredibly supportive and kind, and he genuinely cares about the well-being of his students. Thank you for helping me become a better writer and more thoughtful researcher.

I thank my mentors, Seth Davidovits and Trevor Steil, at Lawrence Livermore National Laboratory. Thank you to Seth for your guidance and for your patience explaining a physics application that was completely new to me. Thank you to Trevor for your guidance and patience in teaching me skills in distributed computing, your help resolving cryptic error messages, and your feedback on proposal writing. Appendices A and B concern projects that were mentored by Seth Davidovits and Trevor Steil, respectively. I also thank Roger Pearce, Min Priest, and Geoff Sanders for their many ideas and helpful suggestions.

I thank my collaborators, Jerry (Jiajie) Luo and Weiqi Chu, that I met at UCLA. Working with you has been very fun and rewarding. We have had many insightful discussions on models of opinion dynamics, and you have been wonderful explaining your perspectives on problems. I also wish to thank the other members and alumni of my research group and the UCLA Networks Journal Club for helpful discussion and comments. I give special thanks to Heather Zinn Brooks, Phil Chodrow, Theodore Faust, Gillian Grindstaff, Abigail Hickok, Micheal Johnson, Leah Keating, Kaiyan Peng, and Sarah Tymochko.

Chapter 4 is adapted with permission from [LP23]. The work for that project was under the mentorship of Mason Porter. We wrote the paper together. Chapter 5 is adapted from

[LLP23], which is joint work with Jiajie Luo and under the mentorship of Mason Porter. In Chapter 5, the models were formulated jointly with Jiajie Luo, and many of the proofs of our theoretical results were developed primarily by Jiajie Luo. All authors wrote the paper together. Chapter 6 concerns a project that is ongoing work in collaboration with Weiqi Chu and Jiajie Luo. In Chapter 6, the idea for the models was developed primarily by Jiajie Luo and the convergence results are joint work with Weiqi Chu and Jiajie Luo.

I thank my dissertation committee members, Andrea Bertozzi, Deanna Needell, and Jacob Foster, for their helpful suggestions, comments, and ideas. I also thank Andrea and Deanna for their career guidance and advice.

I am grateful for funding support from the UCLA Dissertation Year Fellowship, the National Science Foundation Grant No. 1829071, and the National Science Foundation (Grant No. 1922952) through the Algorithms for Threat Detection (ATD) program. I additionally want to thank Brenda Buenrostro, Helen Chung, Martha Contreras, Yim Neang, Kaley Schuster, and Corinne Smith, for their administrative support.

I wish to thank my undergraduate professors at Cooper Union who helped me decide to attend graduate school in mathematics. Thank you to Partha Debroy for believing in me and convincing me that I could succeed in graduate school and that switching from engineering to mathematics was possible. Thank you to Benjamin Davis for your engaging classes in numerical methods and optimization that gave me a glimpse at the field of applied mathematics. Thank you to Stanislav Mintchev for your rigorous linear-algebra class that helped prepare me for grad school in mathematics and for your advice on application materials.

I thank the friends that I met at UCLA for their support and good company. Special thank you to Jason Brown, Erin George, Cecelia Higgins, Jerry Luo, Andrew Sack, and Yotam Yaniv for your help studying for qualifying exams. Thank you to Bon-Soon Lin for being a stellar bootcamp instructor and for all your time helping with the basic qualifying exam.

I thank my parents and siblings for their steadfast love and support. Thank you to my dad for instilling in me a level-headed approach to challenges. Thank you to my mom for being the biggest believer in my ability to succeed.

Last, but certainly not least, I thank my partner, Jason Schuchardt for his endless love and support and for being with me every step of the journey through graduate school. During the COVID-19 lockdowns, you were the best company I could ever hope to ask for. I am truly lucky that you are my partner and best friend.

VITA

2013–2017 B.E. Chemical Engineering with Minor in Mathematics, The Cooper Union for the Advancement of Science and Art

2019–2021 M.A. Mathematics, UCLA

2021–2024 Summer intern and external collaborator, Lawrence Livermore National Laboratory

PUBLICATIONS AND PREPRINTS

Grace J. Li and Mason A. Porter. "Bounded-Confidence Model of Opinion Dynamics with Heterogeneous Node-Activity Levels." *Physical Review Research*, 5(2): 023179, 2023.

Grace J. Li*, Jiajie Luo*, and Mason A. Porter. (*joint first author). "Bounded-Confidence Models of Opinion Dynamics with Adaptive Confidence Bounds." *arXiv:2303.07563*, 2024.

Grace J. Li, Jiajie Luo, Kaiyan Peng, and Mason A. Porter. "Using Mathematics to Study How People Influence Each Other's Opinions." *arXiv:2307.01915*, 2024. (Expository article submitted to *Frontiers for Young Minds*.)

Grace J. Li and Seth Davidovits. "Microphysics of Shock-Grain Interaction for ICF Ablators in a Fluid Approach." (submitted), 2024.

CHAPTER 1

Introduction

Humans are connected in numerous ways, and our many types of interactions with each other influence what we believe and how we act. To model how opinions spread between people or other agents, researchers across many disciplines have developed a variety of models of opinion dynamics [CFL09, SLS17, LA18, NVT20, Noo20, GOD21, PKI22, BAP24]. However, in part because of the difficulty of gathering empirical data on opinions, much of the research on opinion dynamics has focused on theory and model development, with little empirical validation [CFL09, GOD21, PKI22, Vaz22]. Some researchers have examined how human opinions change in controlled experimental settings with questionnaires [CZ15, VMG16, TFM16], and others have examined empirical opinion dynamics using data from social-media platforms [MDB20, Koz22, Koz23]. One of the many difficulties in empirically validating models of opinion dynamics is the potential sensitivity of model outcomes to measurement errors of real-life opinion values [CQ22]. See [Mas19] for a discussion of some challenges of validating models in the social sciences. Even with the difficulty of validating models of opinion dynamics, it is valuable to formulate and study such models. Developing mechanistic models forces researchers to clearly define assumptions, variables, and the relationships between variables; such models provide frameworks to explore and generate testable hypotheses about complex social phenomena [HL15, Vaz22].

In an agent-based model (ABM) of opinion dynamics, each agent is endowed with an opinion from some opinion space. The opinion space can be discrete or continuous, bounded or unbounded, and one-dimensional (1D) or higher-dimensional [CFL09, SLS17, NVT20].

Agents are placed on an underlying network that governs which agents can interact with each other. We view the nodes in a network as representing agents. In a discrete-time ABM of opinion dynamics, at each time step, one selects which agents are activated for interaction and then an update rule determines if and how their opinions change. We say that an opinion model updates 'synchronously' if one activates all agents at each time step. We say that a model updates 'asynchronously' if one selects a subset of agents to activate at each time step.

Bounded-confidence models (BCMs) are a popular type of ABM of opinion dynamics with continuous-valued opinions [NVT20, BAP24]. In a BCM, interacting agents influence each other only when their opinions are sufficiently similar. This situation is reminiscent of the psychological idea of selective exposure, which asserts that people tend to seek information or conversations that support their existing views and avoid those that challenge their views [CM11]. Under this assumption, an agent's views are influenced directly only by agents with sufficiently similar views. For example, social-media platforms have polarizing posts, but individuals can choose whether or not to engage with such content. They are not persuaded by everything in their social-media feeds. The two most popular BCMs of opinion dynamics are the Deffuant–Weisbuch (DW) model [DNA00] and the Hegselmann–Krause (HK) model [HK02]. The DW model updates asynchronously and the HK model updates synchronously.

In Chapter 2, we provide a brief background on definitions and concepts from network science that we use in later chapters. We also give some examples of networks on which we study our models of opinion dynamics. In Chapter 3, we give some background on BCMs. We introduce the DW and HK models, discuss ways to characterize opinions in BCMs, and discuss some generalizations of the standard DW and HK models.

It is challenging to incorporate the complex factors that play a role in opinion dynamics into a single tractable model [GOD21]. Researchers often choose a small number of aspects to study and incorporate into a model of opinion dynamics [CFL09, NVT20, Noo20]. In

this dissertation, we generalize BCMs to incorporate ideas that make them more realistic, while maintaining tractability. In Chapters 4–6, we develop and study three generalizations of BCMs, each of which incorporates a specific modification inspired by social interactions in real life. The BCMs in this dissertation are simplistic, and we make no claims about their ability to represent real-life opinion changes. However, by contributing to the mathematical theory of models of opinion dynamics, we seek to improve understanding of how relatively simple models can inform us about real-world scenarios. We study each of our models using numerical simulations, and we analytically prove results for our models in Chapters 5 and 6.

In Chapter 4, we study an asynchronous BCM that uses node weights in a network to model heterogeneous node-activity levels. The standard DW model [DNA00] updates asynchronously with a pair of agents selected uniformly at random to activate for interaction. Social interactions in real life are not uniformly random. In social networks, some individuals share their ideas and opinions more frequently than others. These disparities can arise from heterogeneous sociabilities, heterogeneous activity levels, different prevalences to share opinions when engaging in a social-media platform, or something else. To examine the impact of such heterogeneities on opinion dynamics, we generalize the DW model by incorporating node weights. The node weights allow us to model agents with different probabilities of interacting. Using numerical simulations, we systematically investigate (using a variety of network structures and node-weight distributions) the effects of node weights, which we assign uniformly at random to the nodes. We demonstrate that introducing heterogeneous node weights results in longer convergence times and more opinion fragmentation than in a baseline DW model. The node weights in our BCM allow one to consider a variety of sociological scenarios in which agents have heterogeneous probabilities of interacting with other agents. This chapter consists of work that was done in collaboration with Mason Porter [LP23].

In Chapter 5, we investigate BCMs in which each pair of agents has a distinct confidence bound that changes when the pair interacts. The quality of an interaction between indi-

viduals can affect how much they trust each other [GAP13, LMM17, CS20]. Rather than considering trust, our BCMs use a notion of "receptiveness", which encodes the willingness of an individual to consider the future opinions of other individuals. We generalize both the DW and HK models to incorporate heterogeneous, time-dependent confidence bounds. The confidence bounds in the resulting BCMs encode mutual receptiveness of agents, and they change when agents interact with each other. In our adaptive-confidence BCMs, when two nodes successfully compromise their opinions in an interaction (i.e., they have a "positive interaction"), they become more receptive to each other. Likewise, when two nodes interact but do not change their opinions (i.e., they have a "negative interaction"), they become less receptive to each other. We analytically and numerically explore the limiting behaviors of our adaptive-confidence BCMs. For a variety of networks and a wide range of values of the parameters that control the increase and decrease of confidence bounds, we demonstrate numerically that our adaptive BCMs result in fewer major opinion clusters (i.e., sets of nodes that converge to the same opinion) and longer convergence times than the baseline (i.e., nonadaptive) BCMs. We also show that these adaptive BCMs can have neighboring nodes that converge to the same opinion but are not receptive to each other. This qualitative behavior does not occur in the associated baseline BCMs. This chapter consists of work that was done in collaboration with Jiajie Luo and Mason Porter [LLP23].

In Chapter 6, we examine BCMs that have multi-dimensional opinions that consist of multiple interdependent topics. In real life, people can have interdependent opinions (e.g., political views) on different topics [Con06]. We generalize both the DW and HK models to K-dimensional opinions and use non-Euclidean topic-weighted opinion "discordance" functions to determine whether or not agents are receptive to each other on a particular topic. When a pair of agents interact on a topic, whether or not they compromise their opinions on that topic depends on the differences in their opinions on all topics. This chapter consists of ongoing work that is in collaboration with Weiqi Chu and Jiajie Luo. In our preliminary numerical simulations, we demonstrate for our topic-weighted BCMs that the choice of initial

opinion distribution has a large effect on (1) whether our models reaches a consensus state or a fragmented state and (2) the amount of opinion fragmentation.

In Chapter 7, we summarize and discuss our contributions. In Appendices A and B, we briefly discuss two projects that are in collaboration with researchers at Lawrence Livermore National Laboratory. The work in Appendix A is a collaboration with Seth Davidovits. The density variations from high-density carbon grains in ablators for inertial-confinement fusion (ICF) affect shock propagation in a target capsule. By combining asymptotic analysis of the linearized governing partial differential equations (building from work by Velikovich et al. [VWL07]) with Fourier analysis, we model the effects of such grains on shock propagation. Appendix B is about ongoing work in collaboration with Trevor Steil. We study a method [ACL06, WGD13] based on personalized PageRank to detect overlapping communities in undirected graphs.

CHAPTER 2

Background on Networks

In this chapter, we briefly describe some definitions and concepts from network science that we use in later chapters. In Section 2.1, we give some basic network definitions. In Section 2.2, we briefly describe communities in networks. In Section 2.3, we give some examples of networks on which we study our models of opinion dynamics.

2.1 Some network definitions

Networks consist of nodes (which represent entities) and the edges between them (which represent ties between nodes). Networks can encode relationships between entities in a wide variety of disciplines [New18]. In this dissertation, we focus on examining networks with nodes that represent individuals and edges that represent social relationships. In our models of opinion dynamics, an edge between nodes indicates that they can have a social interaction with each other and potentially influence each other's opinions.

We now present some standard definitions. For a review of standard network definitions, see [New18]. The simplest type of network is an undirected graph $G = (V, E)$, where V is a set of nodes and E is a set of edges between those nodes. We let $N = |V|$ denote the size (i.e., the number of nodes) of a network. When two nodes i and j have an edge (i, j) between them, we say that nodes i and j are *adjacent*. When a graph has edges from a node to itself, we refer to those edges as self-edges; if two nodes have multiple edges between them, we say that the graph has multi-edges. We refer to the nodes that are adjacent to a node i as

the neighbors of node i. We refer to the set of neighbors of a node i as the neighborhood of node i and denote it by $\mathcal{N}(i)$. The *degree* of a node i is the number of neighbors of node i, and we denote it by $\deg(i)$. The degree distribution of a graph G is a probability distribution that indicates the fraction of nodes with each degree. Many real-world networks have heavy-tailed degree distributions [BB03, BC19]. A "walk" on a network is a sequence of nodes and edges (with repeats allowed) that starts at a node, goes along an edge to an adjacent node, and so on. A *connected component* of an undirected network G is a maximal subgraph with a walk between each pair of nodes. There is a walk between any two nodes in the same connected component. There does not exist a walk between any two nodes in different connected components.

In social networks, individuals tend to form social connections with people that are similar to themselves, such as people that have the same interests, have similar beliefs, or belong to the same demographic group [MSC01]. This phenomenon is known as *homophily* and can be summarized with the saying "birds of a feather flock together" [MSC01]. The opposite scenario, in which there tend to be more connections or interactions between individuals that are dissimilar to each other, is known as heterophily [MSC01, LVC14]. Researchers can measure the tendency of edges in a network to be between nodes that are similar to each other in some characteristic by calculating an "assortativity coefficient" [New03, New18, KO23]. A positive assortativity coefficient indicates "assortative mixing" and is used to measure homophily in social networks. Many real-world social networks have assortative mixing by node degree [New03].

In this dissertation, we study models of opinion dynamics on undirected graphs. Researchers have also studied other types of networks that can incorporate more complex information. We briefly describe some of them here. A directed graph is a graph in which each edge has a direction [New18]. If the edge (i, j) from node i to node j exists, then it is not necessary for the edge (j, i) in the other direction to exist. One can use directed edges to encode directed relationships such as which way a flight travels between airports, which

account follows another account on social media, or which way currency flows in online transactions. An *edge-weighted* graph (typically called a "weighted" graph) uses weights on edges to encode information such as the strength of a relationship [New04]. A *node-weighted* graph has node weights that incorporate node-level information; such graphs are less commonly studied than edge-weighted graphs [HDZ12]. In Chapter 4, we study a model of opinion dynamics on node-weighted graphs. A signed network [Hei46, CH56, LHK10] assigns each edge a positive or negative value to represent the nature of the relationship. For example, in a signed social network where the nodes represent people, one might use a positive edge to indicate friendship and a negative edge to indicate dislike for each other. A hypergraph consists of a set of nodes and a set of hyperedges, each of which connects an arbitrary number of nodes [Bre13]. One can use hypergraphs to study polyadic group relationships. A temporal network (which is also commonly known as a dynamic graph or an evolving graph) is a graph that changes with time [HS12].

2.2 Communities in networks

A "community" in a network is a set of nodes that are more densely connected to each other than to other communities in the network. Two nodes in the same community are in some sense more closely related or similar than two nodes in different communities. There is no single definition of what makes a "good" community in a network, and researchers have considered a variety of ways to quantify whether a set of nodes has good community structure and ways to detect communities in networks [FH16, CDM17]. See [FH16] for a review of community-detection methods for networks. See [CDM17] for an overview of calculations to quantitatively evaluate community structures in networks.

One notion of a good community is that there should be relatively more edges between nodes in the same community than between nodes in different communities [CDM17]. A way to quantify this is by calculating the *conductance* of a set of nodes [SM00, CDM17].

The conductance of a set of nodes C in a graph G is

$$\Phi(C) = \frac{|\partial C|}{\min\{\mathrm{vol}(C), 2|E| - \mathrm{vol}(C)\}}\,, \tag{2.2.1}$$

where $|\partial C|$ is the number of edges that are between a node in C and a node not in C, $\mathrm{vol}(C)$ is the number of edges that are between nodes in C, and $|E|$ is the number of edges in the graph. A set of nodes with smaller conductance can be considered a higher-quality community than a set of nodes with larger conductance. In Appendix B, we investigate a method for community detection that finds communities with small conductance values.

2.3 Some examples of networks

In this section, we describe some examples of undirected networks. In this dissertation, we simulate our models of opinion dynamics on these networks.

2.3.1 Complete graphs

A complete graph is a graph that has an edge between each pair of nodes. An undirected complete graph with N nodes has $\binom{N}{2} = \frac{1}{2}N(N-1)$ edges. When researchers study dynamical systems on a complete graph, they sometimes refer to the system as a "fully-mixed" or "well-mixed" system [DNA00]. When we study a model of opinion dynamics on a complete graph, every pair of nodes can interact in the model. In this dissertation, we examine each of our models of opinion dynamics on complete graphs before considering more complicated network structures.

2.3.2 Erdős–Rényi (ER) graphs

The Erdős–Rényi (ER) random-graph models [ER59, Gil59, ER60] are some of the simplest models for generating random graphs. There are two types of ER random-graph models:

the $G(N, m)$ and $G(N, p)$ ER models.

The $G(N, m)$ ER model [ER59, ER60] is an ensemble of random graphs with N nodes and m edges. Consider an N-node undirected graph. There are $\binom{N}{2} = \frac{1}{2}N(N-1)$ possible edges between pairs of nodes. One uniformly randomly chooses m of these edges to generate a $G(N, m)$ graph in the random-graph ensemble.

The $G(N, p)$ ER model [SR51, Gil59] is an ensemble of random graphs with N nodes and homogeneous, independent probability p of an edge between each pair of nodes. When $p = 1$, the $G(N, p)$ model yields a complete graph. For the $G(N, p)$ ER model, the expected number of edges is $\frac{1}{2}N(N-1)p$ and the expected mean degree is $(N-1)p$ [New18]. The $G(N, p)$ ER model has a Bernoulli degree distribution. In the limit as $N \to \infty$, the degree distribution approaches a Poisson degree distribution [New18]. The $G(N, p)$ ER model is also sometimes called the Bernoulli random-graph model, or "the" random-graph model [New18]. In this dissertation, we use $G(N, p)$ ER networks to simulate each of our models of opinion dynamics on networks with different edge densities.

2.3.3 Stochastic-block-model (SBM) graphs

A stochastic-block-model (SBM) [HLL83, New18] is a random-graph model that can incorporate specific mixing (such as assortative mixing) of groups of nodes. The simplest type of undirected SBM network consist of blocks that are $G(N, p)$ ER graphs (see Section 2.3.2). SBM networks are specified by the number N of nodes, a partition of the set of nodes into groups $S_1, S_2, \ldots S_r$, and an edge-probability matrix P that specifies the probability of on edge between any two nodes depending on their group memberships. For an undirected SBM with r groups, P is a symmetric $r \times r$ matrix with entries in the interval $[0, 1]$. The entry in the Ath row and Bth column of P is the independent probability of an edge between a node in group A and a node in group B.

We use SBM networks to simulate our models of opinion dynamics on networks with

underlying block structures. We consider undirected SBM networks with 2×2 blocks. We partition a network into two sets of nodes. Set A has 75% of the nodes of the network, and set B has 25% of the nodes. Our choice of partition sizes is inspired by the choices of Kureh and Porter [KP20]. We define a symmetric edge-probability matrix

$$P = \begin{bmatrix} P_{AA} & P_{AB} \\ P_{AB} & P_{BB} \end{bmatrix}, \tag{2.3.1}$$

where P_{AA} and P_{BB} are the probabilities that there is an edge between two nodes in set A and set B, respectively, and P_{AB} is the probability that there is an edge between a node in set A and a node in set B.

Inspired by the choices of Kureh and Porter [KP20], we consider two types of SBM networks. The first SBM has a two-community structure, in which there is a larger probability of edges within a community (see Section 2.2) than between communities. In a two-community SBM, the probabilities P_{AA} and P_{BB} are larger than P_{AB}, so edges between nodes in the same community exist with a larger probability than edges between nodes in different communities. The second SBM has a core–periphery structure, with a set of core nodes with a large probability of edges within the set, a set of peripheral nodes with a small probability of edges within the set, and an intermediate probability of edges between core nodes and peripheral nodes. Our core–periphery SBMs with core set A and periphery set B satisfy $P_{AA} > P_{AB} > P_{BB}$. In Chapter 4, we simulate our node-weighted BCM on both two-community and core–periphery SBMs. In Chapter 5, we simulate our BCMs with adaptive confidence bounds on a two-community SBM.

CHAPTER 3

Background on Bounded-Confidence Models (BCMs) of Opinion Dynamics

The two most popular bounded-confidence models (BCMs) of opinion dynamics are the Deffuant–Weisbuch (DW) model [DNA00] and the Hegselmann–Krause (HK) model [HK02]. In the standard DW and HK models, agent opinions take continuous values in a closed interval on the real line. This interval represents a continuous spectrum of views about something, such as an ideology (with the strongest views at the end points of the interval) or strength of support for a political candidate. We describe the DW model in Section 3.1 and the HK model in Section 3.2. In Section 3.3, we discuss ways to characterize opinions in BCMs. In Section 3.4, we discuss some generalizations of the DW and HK models.

3.1 The Deffuant–Weisbuch (DW) model

The standard DW model [DNA00] is a discrete-time asynchronous BCM. Consider a time-independent, unweighted, undirected network $G = (V, E)$ with no self-edges or multi-edges, where V is a set of nodes and E is a set of edges between them. Let $N = |V|$ denote the size (i.e., the number of nodes) of a network. The nodes of the network represent agents with opinions that lie in the closed interval $[0, 1]$. Let $x_i(t)$ denote the opinion of node i at time t. Let $\mathbf{x}(t) \in \mathbb{R}^N$ denote the opinion vector of the system; its ith entry is $x_i(t)$. One endows each node with an initial opinion, which one selects uniformly at random from the interval $[0, 1]$.

The standard DW model has two parameters. The *confidence bound* $c \in [0, 1]$ is a thresholding parameter; when two nodes interact, they compromise their opinions by some amount if and only if their opinions differ by less than c. The *compromise parameter* $m \in (0, 0.5]$ (which is also sometimes called a "convergence parameter" [DNA00] or a "cautiousness parameter" [MVP18]) parametrizes the amount that a node changes its opinion to compromise with the opinion of a node with which it interacts.

In the standard DW model, at each discrete time, one uniformly randomly selects an edge $(i, j) \in E$ for activation. At time t, suppose that one picks edge (i, j). Then, nodes i and j update their opinions through the following update rule:

$$\begin{aligned} x_i(t+1) &= \begin{cases} x_i(t) + m(x_j(t) - x_i(t)), & \text{if } |x_i(t) - x_j(t)| < c \\ x_i(t), & \text{otherwise}, \end{cases} \\ x_j(t+1) &= \begin{cases} x_j(t) + m(x_i(t) - x_j(t)), & \text{if } |x_i(t) - x_j(t)| < c \\ x_j(t), & \text{otherwise}. \end{cases} \end{aligned} \tag{3.1.1}$$

When $|x_i(t) - x_j(t)| < c$, we say that nodes i and j are *receptive* to each other at time t. When $|x_i(t) - x_j(t)| \geq c$, we say that nodes i and j are *unreceptive* to each other.

The confidence bound c controls the amount of "open-mindedness" of the nodes to different opinions. The extreme case $c = 0$ is degenerate (because no nodes update their opinions), and the case $c = 1$ allows all adjacent nodes to interact with each other (making the model a linear system). The compromise parameter controls how much a node changes its opinion after interacting with a node it is receptive to. When $m = 0.5$, two interacting nodes that are receptive to each other precisely average their opinions; when $m \in (0, 0.5)$, interacting nodes that are receptive to each other move towards each other's opinions, but they do not adopt the mean opinion. One can alternatively consider $m \in (0, 1)$ as in Meng et al. [MVP18], although this is an uncommon choice. When $m > 0.5$, nodes "overcompromise" when they change their opinions; they overshoot the mean opinion and change which side of the mean

opinion they are on.

3.2 The Hegselmann–Krause (HK) model

The standard HK model [Kra00, HK02] is a discrete-time synchronous BCM. The HK model was studied initially only on complete graphs [HK02]. Since then, the HK model has also been studied on networks (see, e.g., [For05, PFT18]). We consider the standard HK model to be its extension to networks. For a review of the HK model and its extensions, see [BAP24].

Consider a time-independent, unweighted, undirected network $G = (V, E)$ with no self-edges or multi-edges, where V is a set of nodes and E is a set of edges between them. As in the DW model, the nodes of a network represent agents with opinions that lie in the closed interval $[0, 1]$. Again, let $x_i(t)$ denote the opinion of node i at time t. Let $\mathbf{x}(t) \in \mathbb{R}^N$ denote the opinion vector of the system; its ith entry is $x_i(t)$. Each node is assigned an initial opinion, which one selects uniformly at random from the interval $[0, 1]$. The confidence bound c is the only parameter of the standard HK model. As in the DW model, c controls the amount of "open-mindedness" of the nodes to different opinions.

At each time t, we update the opinion of each node i by calculating

$$x_i(t+1) = |I(i, x(t))|^{-1} \sum_{j \in I(i,x(t))} x_j(t)\,, \tag{3.2.1}$$

where[1] $I(i, x(t)) = \{i\} \cup \{j \mid |x_i(t) - x_j(t)| < c \text{ and } (i, j) \in E\} \subseteq \{1, 2, \ldots, N\}$. We say that nodes i and j are receptive to each other at time t if their opinion difference is less than the confidence bound c (i.e., $|x_i(t) - x_j(t)| < c$). Thus, $I(i, x(t))$ consists of all adjacent nodes that node i is receptive to as well as node i itself.

[1] In [Kra00, HK02], $I(i, x(t)) = \{i\} \cup \{j \mid |x_i(t) - x_j(t)| \leq c \text{ and } (i, j) \in E\}$. We use a strict inequality to be consistent with the strict inequality in the DW model.

3.3 Characterizing opinions in BCMs

In this section, we describe how we characterize opinions in BCMs. In BCMs with the update rules of the standard DW model (see Equation (3.1.1)) or the standard HK model (see Equation (3.2.1)), the opinions of nodes converge to steady-state values (see Section 3.3.1). When studying BCMs, researchers often study how long it takes to reach steady-state behavior and the steady-state opinions of nodes [MVP18, NVT20, HFQ20, PKI22]. To characterize these steady-state opinions, researchers commonly examine steady-state opinion clusters, which are sets of nodes that converge to the same opinion (see Section 3.3.2). The number and sizes of the steady-state opinion clusters determine whether a system reaches opinion "consensus", "polarization", or "fragmentation". Researchers sometimes make different choices for how they define consensus, polarization, and fragmentation; we discuss some possible choices and our definitions in Section 3.3.3. In Section 3.3.4, we discuss some ways we quantify opinion fragmentation. In Section 3.3.5, we discuss notions of the time required for nodes to reach their steady-state opinions and how we approximate "convergence time" in our numerical simulations.

3.3.1 Convergence of opinions

One can view each opinion update in a BCM as multiplying an opinion vector by a row-stochastic matrix. Consider the following theorem, which was stated and proved by Lorenz [Lor05].

Theorem 3.3.1 ([Lor05])**.** *Let $\{A(t)\}_{t=0}^{\infty} \in \mathbb{R}_{\geq 0}^{N \times N}$ be a sequence of row-stochastic matrices. Suppose that each matrix satisfies the following properties:*

(1) The diagonal entries of $A(t)$ are positive.

(2) For each $i, j \in \{1, \ldots, N\}$, we have that $[A(t)]_{ij} > 0$ if and only if $[A(t)]_{ji} > 0$.

(3) There is a constant $\alpha > 0$ such that the smallest positive entry of $A(t)$ for each t is

larger than α.

Given times t_0 and t_1 with $t_0 < t_1$, let

$$A(t_0, t_1) = A(t_1 - 1) \times A(t_1 - 2) \times \cdots \times A(t_0)\,. \tag{3.3.1}$$

If conditions (1)–(3) are satisfied, then there exists a time t' and pairwise-disjoint classes $\mathcal{I}_1 \cup \cdots \cup \mathcal{I}_p = \{1, \ldots, N\}$ such that if we reindex the rows and columns of the matrices to have the order $\mathcal{I}_1, \ldots, \mathcal{I}_p$, then

$$\lim_{t\to\infty} A(0,t) = \begin{bmatrix} K_1 & & 0 \\ & \ddots & \\ 0 & & K_p \end{bmatrix} A\,(0,t')\ , \tag{3.3.2}$$

where each K_q, with $q \in \{1, 2, \ldots, p\}$, is a row-stochastic matrix of size $|\mathcal{I}_q| \times |\mathcal{I}_q|$ whose rows are all the same.

As stated in [Lor05], Theorem 3.3.1 guarantees that the opinion of each node converges to a limit opinion value in the standard HK and DW models. For a node i, we define its *limit opinion* as $\lim_{t\to\infty} x_i(t)$, which we denote by x_i^*. In the field of dynamical systems, it is common to refer to x_i^* as the *steady-state opinion* of node i. We use both the terms "limit opinion" (in Chapters 5 and 6) and "steady-state opinion" (in Chapter 4) to refer to x_i^*.

3.3.2 Effective graphs and opinion clusters

Let $G = (V, E)$ be a time-independent, unweighted, and undirected graph without self-edges or multi-edges. We examine a BCM on G, and we associate to G a time-dependent "effective graph" $G_{\text{eff}}(t)$, which is a subgraph of G with edges only between nodes that are receptive

to each other at time t. That is,

$$G_{\text{eff}}(t) = (V, E_{\text{eff}}(t)), \tag{3.3.3}$$
$$E_{\text{eff}}(t) = \{(i,j) \in E \text{ such that } |x_i(t) - x_j(t)| < c\}.$$

Other researchers have referred to effective graphs as "confidence graphs" [BAP24], "communication graphs" [BBC13], and "corresponding graphs" [YDH14].

We refer to the connected components of the effective graph $G_{\text{eff}}(t)$ as the *opinion clusters* at time t. Equivalently, an opinion cluster $S_r(t)$ is a maximal connected set of nodes in which the pairwise differences in opinions are all strictly less than the confidence bound c. Adding any other node to $S_r(t)$ yields at least one pair of adjacent agents with an opinion difference of at least c. At each time t, nodes in different opinion clusters are unreceptive to each other and adjacent nodes in the same opinion cluster are receptive to each other. We say that the size of an opinion cluster is the number of nodes in the opinion cluster. We define the "diameter" $D(S_r(t))$ of an opinion clusters $S_r(t)$ as the maximum difference in opinions in that opinion cluster. That is,

$$D(S_r(t)) = \max\{|x_i(t) - x_j(t)| : \text{ nodes } i,j \in S_r(t)\}. \tag{3.3.4}$$

If the opinions of all nodes converge to some limit opinion, then it is insightful to examine the effective graph when all nodes have reached a limit opinion. We refer to the connected components of this effective graph as the *limit opinion clusters* or the *steady-state opinion clusters*. We use these two terms interchangeably. As we discussed in Section 3.3.1, for the standard DW and HK models, the opinion of each node converges to some limit opinion, so we can consider the limit opinion clusters. When the limit $\lim_{t\to\infty} G_{\text{eff}}(t)$ exists, we call it the *limit effective graph* and denote it by G^*_{eff}. As we show in Section 5.3, the limit effective graph exists for the standard HK model (see Theorem 5.3.4) and almost surely exists for the standard DW model (see Theorem 5.3.7). When the limit effective graph exists for the

standard HK and DW models, it equals the effective graph in which all nodes have attained their limit opinion values.

In the standard DW and HK models, one can determine the limit opinion clusters in finite time. If every pair of opinion clusters, S_a and S_b, is separated by a distance of at least c (i.e., $|x_i(t) - x_j(t)| \geq c$ for all $i \in S_a$ and all $j \in S_b$) at some time $\tilde{T}$, then (because c is fixed) no agent in S_a can influence the opinion of an agent in S_b (and vice versa) for all $t \geq \tilde{T}$. Meanwhile, agents in each opinion cluster continue to compromise their opinions with each other. In particular, if at some time $t \geq \tilde{T}$, the diameter of an opinion cluster (see Equation (3.3.4)) is less than the confidence bound c, then all pairs of adjacent nodes in that opinion cluster are receptive to each other and all nodes in that opinion cluster eventually converge to the same limit opinion. Therefore, for the standard DW and HK models, we observe the formation of distinct limit opinion clusters in finite time.[2] We use this feature of these standard BCMs to specify a stopping criterion for our numerical simulations (see Section 3.3.5).

3.3.3 Opinion consensus and fragmentation

In this dissertation, we investigate which situations in our BCMs yield consensus (specifically, they result in one "major" opinion cluster, which will characterize shortly) at steady state and which situations yield opinion fragmentation (when there are at least two distinct major clusters) at steady state. Some researchers use the term "polarization" to refer to the presence of exactly two opinion clusters (or of exactly two major opinion clusters) and "fragmentation" to refer to the presence of three or more opinion clusters (or of three or more major opinion clusters) [HK02, BGS16]. However, because we are interested in distinguishing between consensus states and any state that is not a consensus, we use the term "fragmentation" for any state with at least two major opinion clusters. We then quantify

[2]For the same reasons, we also observe the formation of distinct limit opinion clusters in finite time for our models in Chapters 4 and 6.

the extent of opinion fragmentation.

In some situations, an opinion cluster has very few agents. Consider a 500-node network in which 499 agents eventually have the same opinion, but the remaining agent (say, Agent 86, despite repeated attempts by Agent 99 and other agents to convince him) retains a distinct opinion at steady state. In applications, it may not be appropriate to think of this situation as opinion fragmentation. To handle such situations, we use a notion of "major clusters" and "minor clusters" [LAZ04, Lor07]. We characterize major and minor clusters in an ad hoc way. We define a "minor" opinion cluster in a network as an opinion cluster with at most $f_m N$ of the agents, where N is the number of nodes of a network and f_m is the fraction of nodes we use to determine if a cluster is minor. We consider different choices of f_m for our different BCMs. For example, we take $f_m = 0.02$ for our node-weighted BCM in Chapter 4, and we take $f_m = 0.01$ for our adaptive-confidence BCMs in Chapter 5. Any opinion cluster that is not a minor cluster is a "major" cluster. We account only for the number of major clusters when determining if a simulation reaches a consensus state (i.e., exactly one major cluster) or a fragmented state (i.e., more than one major cluster). However, we track the numbers and sizes of all major and minor clusters, and we use all clusters (i.e., both major and minor clusters) to quantify opinion fragmentation.

In this dissertation, we consider agent-based BCMs on finite networks, so finite-size effects may affect our results. (See our discussion in Section 4.4.5.) Alternatively, one can examine mean-field approximations as $N \to \infty$ [BKR03, GGL12, CPP21, FBQ21, DFB23]. With $N \to \infty$, one can suppose that opinion clusters that scale linearly with N are significant for determining whether there is opinion consensus or fragmentation.

3.3.4 Quantifying opinion fragmentation

Quantifying opinion fragmentation is much less straightforward than determining whether or not there is fragmentation. Researchers have proposed a variety of notions of fragmentation and polarization [BGS16], and they have also proposed several ways to quantify such

notions [BGS16, CW22, AWA22]. For example, for BCMs, Schawe et al. [SFH21] quantified how close a system is to consensus by determining the fraction of nodes in the largest opinion cluster.

In principle, a larger number of opinion clusters is one indication of more opinion fragmentation. However, as we show in Figure 3.1, there can be considerable variation in the sizes (i.e., the numbers of nodes) of the opinion clusters. For example, suppose that there are two opinion clusters. If the two opinion clusters have the same size, then one can view the opinions in the system as more polarized than if one opinion cluster has a large majority of the nodes and the other opinion cluster has a small minority. Additionally, although we use only major clusters to determine if a system reaches a consensus or fragmented state, we seek to distinguish quantitatively between the scenarios of opinion clusters (major or minor) with similar sizes from ones with opinion clusters with a large range of sizes.

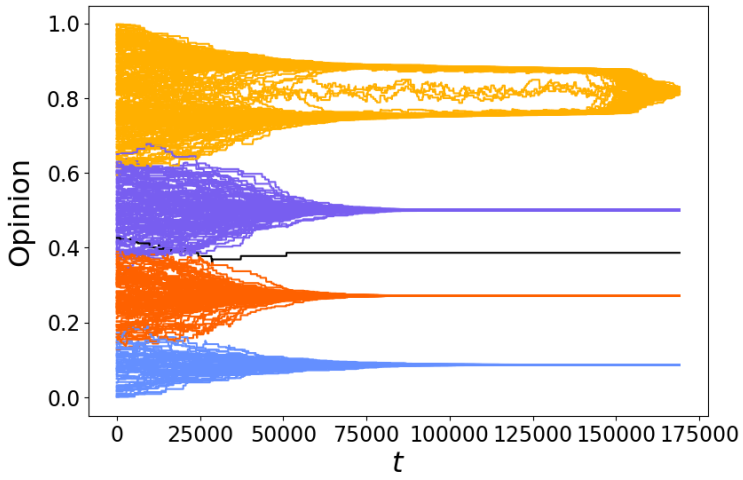

Figure 3.1: Sample trajectories of agent opinions versus time t in a single simulation of a DW model with confidence bound $c = 0.1$ and compromise parameter $m = 0.1$. We color the trajectory of each node by its steady-state opinion cluster. Observe that the steady-state opinion clusters have different sizes. There is one minor cluster (in black); it consists of a single node whose steady-state opinion is about 0.4. The opinion cluster that converges to the largest opinion value has about twice as many nodes as the other major clusters. (This figure originally appeared in [LP23].)

Wang et al. [WLE17] defined an order parameter Q for BCMs. We calculate[3] Q as

$$Q(t) = \frac{1}{N^2} \sum_{i=1}^{N} \sum_{j=1}^{N} \mathbf{1}_{|x_i(t) - x_j(t)| < c} \,, \tag{3.3.5}$$

where $\mathbf{1}_{|x_i - x_j| < c}$ is an indicator function that equals 1 when $|x_i - x_j| < c$ and equals 0 otherwise. The order parameter Q indicates the fraction of pairs of nodes that are receptive to each other. For a complete graph, the order parameter gives the number of edges in the effective graph. In a graph that is not complete, some nodes are not adjacent to each other, so the order parameter is more difficult to interpret. One can equivalently calculate $Q(t)$ as a normalized sum of the squares of the sizes of the opinion clusters. Suppose that there are R opinion clusters (see Section 3.3.2), which we denote by S_r for $r \in \{1, \ldots, R\}$. We refer to the set $\{S_r\}_{r=1}^{R}$ as an "opinion-cluster profile", which is a partition of the set of nodes of a network. Using this notation, one can calculate $Q(t)$ as

$$Q(t) = \frac{1}{N^2} \sum_{r=1}^{R} |S_r|^2 \,. \tag{3.3.6}$$

We use the order parameter Q to investigate the behaviors of our topic-weighted BCMs in Chapter 6.

In our investigations in Chapter 4 and Chapter 5, we quantify opinion fragmentation by following Han et al. [HFQ20] and calculating Shannon entropy. Consider an opinion-cluster profile $\{S_r\}_{r=1}^{R}$. The fraction of agents in opinion cluster S_r is $|S_r|/N$. The Shannon entropy H of the opinion-cluster profile is

$$H = -\sum_{r=1}^{R} \frac{|S_r|}{N} \ln\left(\frac{|S_r|}{N}\right) \,. \tag{3.3.7}$$

Computing Shannon entropy allows us to use a scalar value to quantify the distribution

[3]Wang et al. [WLE17] used the inequality $|x_i - x_j| \leq c$. However, we use a strict inequality to be consistent with the strict inequality in the update rules in Equations (3.1.1) and (3.2.1).

of opinion-cluster sizes, with larger entropies indicating more opinion fragmentation. The Shannon entropy is larger when there are more opinion clusters. Additionally, for a fixed number R of opinion clusters, the Shannon entropy is larger when the opinion clusters are the same size than when the sizes are heterogeneous. We calculate the steady-state entropy $H(T_f)$ using all steady-state opinion clusters (i.e., both major and minor clusters). When there are sufficiently few minor clusters, we expect the number of major clusters to follow the same trend as the Shannon entropy.

3.3.5 Convergence time

Consider a BCM where nodes quickly reach a consensus for one set of parameters and initial conditions. Suppose that it takes twice as long to reach consensus for another set of parameters and initial conditions. Researchers are interested in differentiating between such situations and investigating notions of "convergence time", the time required for nodes to reach their steady-state opinions [Lor07, NVT20, PFT22, BAP24]. In this section, we describe some theoretical work on convergence time and then discuss numerical approximations for convergence time.

Dittmer [Dit01] proved that the opinions in the standard HK model on a complete graph converge in finite time. Researchers sometimes use the term "termination time" to refer to the maximum time that is required across all sets of initial opinions for all nodes to reach their steady-state opinions [EBN13, PFT22]. For the standard HK model on a complete graph, the termination time is finite. Bhattacharyya et al. [BBC13] studied an HK model on complete graphs with nodes that have multidimensional opinions (with the Euclidan norm for opinion distance) and proved lower and upper bounds on the termination time that are, respectively, polynomial in the number of nodes and polynomial in both the number of nodes and the opinion dimension. Etesami et al. [EBN13] proved for the multidimensional HK model on complete graphs that when a condition on the number of single-node opinion clusters is satisfied, the termination time has an upper bound that is polynomial in the

number of nodes. Hegselmann and Krause [HK19] also considered the multidimensional HK model on complete graphs and proved conditions that give geometrically fast convergence to consensus. Parasnis et al. [PFT22] showed that the nodes in the standard HK model can take arbitrarily long to reach their steady-state opinions (i.e., the termination time is infinite) for graphs that are not complete. They examined and proved bounds on the ε-convergence time, which is the time required for all nodes to be within ε of their steady-state opinions.

The stochastic updates of the DW model and its variants make it challenging to study the convergence time in those models. (It is easier to study convergence in the HK model.) Zhang and Chen [ZC15] derived upper bounds on the convergence rate for a synchronous and asymmetric variant of the DW model. In their model, every node interacts with (and is potentially influenced by) one other node at each time step. Their model is asymmetric because when a node i selects node j for interaction, j need not select i for interaction. Chen et al. [CSM20] examined a variant of the model by Zhang and Chen [ZC15] that has heterogeneous confidence bounds. They proved an upper bound on the convergence rate for this model when the compromise parameter $m \in [1/2, 1)$.

In this dissertation (see Chapters 4–6), we examine the convergence times of our BCMs using numerical simulations. For our numerical simulations, we need a stopping criterion, as it can potentially take arbitrarily long for nodes to reach their steady-state opinions in a BCM simulation. Our stopping criterion also gives us a proxy to compare how long it takes nodes in different simulations to converge to their steady-state opinions. In numerical investigations of BCMs, researchers commonly consider how many time steps it takes for a simulation to stop.

Our stopping criterion checks that the diameter (see Equation (3.3.4)) for each opinion cluster (see Section 3.3.2) is less than some tolerance. That is,

$$\max\{|x_i(t) - x_j(t)| \text{ such that } i, j \in S_r(t) \text{ for some } r\} < \text{tolerance}\,. \tag{3.3.8}$$

We refer to the time T_f at which we reach our stopping criterion as the "convergence time" of our simulations. This choice of stopping criterion is inspired by Meng et al. [MVP18]. Because of computational limitations, we set different tolerance values for our investigations of our different BCMs. As we discussed in Section 3.3.2, for the standard DW and HK models, by choosing a tolerance value that is less than the confidence bound c, we can determine the steady-state opinion clusters at our numerical convergence time. We choose tolerance values (specifically, values of 0.01 or 0.02 for our DW models and 1×10^{-6} for our HK models) that are considerably smaller than the confidence bounds c in our simulations. By doing so, we aim to obtain a numerical convergence time that gives a good sense of how long it takes for nodes to converge to their limit opinions.

Other researchers have made different choices to numerically examine convergence of BCMs. For example, researchers sometimes consider the magnitude of opinion changes over a certain number of time steps. Schawe [SFH21] numerically simulated the standard HK model and calculated the total opinion change for all nodes in a pair of consecutive time steps. They calculated

$$\Delta = \sum_{i=1}^{N} |x_i(t-1) - x_i(t)|$$

and stopped a simulation if Δ is less than their tolerance value. Sîrbu et al. [SPG19] and Kan et al. [KFP23] simulated variants of the DW model and stopped their simulations if the sum of the changes in node opinions over a certain number of consecutive time steps is smaller than a tolerance value.

3.4 Some generalizations of BCMs

There has been much research on the standard DW and HK models and their generalizations. See [Lor07, NVT20, Noo20, BAP24] for reviews.

The DW model was first studied on complete graphs and a square-lattice graph (i.e.,

graphs where the nodes are arranged in a grid pattern on a square) [DNA00]. The HK model was studied initially only on a complete graph [HK02]. Since then, to explore the effects of network structure on BCM dynamics, many researchers have examined DW models [MVP18] and HK models [For05, SFH21] on various time-independent graphs. Researchers have also examined BCMs on hypergraphs [HKB22, CP23] and on networks that coevolve with agent opinions [SLL14, KFP23].

Though relatively uncommon, some studies have considered initial agent opinions that arise from distributions other than uniform distributions, yielding initial conditions that differ from those in the standard DW and HK models. Several researchers [Jac06, Sha13, CTS13, Sob15, HKB22] used numerical simulations to examine the effect of nonuniform initial-opinion distributions on DW models. Kou et al. [KZP12] numerically investigated different initial-opinion distributions for their HK model with heterogeneous confidence bounds. Shang [Sha13] considered the DW model with an opinion space $\mathbb{R}$ on an infinite 1D lattice graph and derived a critical confidence bound, above which nodes almost surely reach a consensus at the mean of the initial opinions for any set of initial opinions. Gómez-Serrano et al. [GGL12] proved sufficient conditions on the initial-opinion distribution and confidence bound for a mean-field limit of the DW model to reach a consensus with all nodes converging to the same opinion. Hickok et al. [HKB22] developed a DW model on hypergraphs and proved sufficient conditions on the hypergraph and initial opinion distribution for nodes to almost surely reach a consensus. They also proved that a phase transition in the convergence time occurs when the variance of the initial opinion distribution equals the confidence bound of the model. Yang et al. [YDH14] proved a sufficient condition on the initial-opinion distribution for the HK model to reach consensus and proposed a variant of the HK model that guarantees convergence for more general initial opinions. Blondel et al. [BHT10] studied a continuous-time version of the HK model and proved that, for almost all initial conditions, the mean opinion is conserved and the opinion variance is nonincreasing. They also proposed a variant of the HK model with a continuum of agents and proved sufficient conditions on

the initial opinions to give existence and uniqueness of solutions to their model.

Other investigations have modified the confidence bound or compromise parameters. Such generalizations affect the opinion updates of interacting agents. Unlike the standard DW model, the standard HK model does not have a compromise parameter. Chazelle and Chu [CW17] studied an HK model with an inertia parameter that functions like a compromise parameter and determines how much a node moves from its own opinion when compromising. Several researchers [DAW02, Zha14, HDH18] considered DW models with heterogeneous compromise parameters.

Many researchers have generalized the DW and HK models by incorporating heterogeneity into the confidence bounds. Lorenz [Lor10] extended the DW and HK models so that each node has its own confidence bound, which can result in asymmetric influence and opinion updates. Using numerical simulations, Lorenz demonstrated that these BCMs are more likely than the standard BCMs to reach a consensus state when there are both open-minded and closed-minded nodes (which have large and small confidence bounds, respectively). By analyzing the heterogeneous-confidence DW model of [Lor10] on a complete graph, Chen et al. [CSM20] proved almost-sure convergence of opinions for certain parameter values and derived sufficient conditions for the nodes of a network to eventually reach a consensus. In a related work, Chen et al. [CSD20] examined a noisy HK model and showed that heterogeneous confidence bounds in this setting can yield larger differences in node opinions in the infinite-time limit. Su et al. [SGW17] examined the heterogeneous-confidence HK model of [Lor10] and proved that at least some nodes of a network converge to a steady-state opinion in finite time. Researchers have also incorporated edge-based heterogeneities in the confidence bounds of BCMs. Shang [Sha14] studied a DW model with edge-based confidence bounds that arise from independent and identically distributed Poisson processes. They derived sufficient conditions for consensus to occur almost surely for a 1D lattice graph. Etesami [Ete19] examined an HK model on networks with time-independent edge-heterogeneous confidence bounds and proved that their model is Lyapunov stable.

Researchers have also generalized BCMs and other opinion models by incorporating adaptivity and time-dependence into the parameters. Weisbuch et al. [WDA02] studied a generalized DW model in which each node has a heterogeneous, time-dependent confidence bound that is proportional to the standard deviation of the opinions that that node observed in all prior interactions. They also consider a variant of their model that places more weight on opinions in recent interactions. Deffuant et al. [DAW02] examined a "relative agreement" DW model in which each agent has a unique uncertainty parameter that determines whether interacting nodes influence each other and how much they compromise their opinions. A node changes both its opinion and its uncertainty values when it is influenced by another node. Bagnoli et al. [BCF07] considered a variant of the DW model in which each pair of nodes has an affinity parameter that changes based on the magnitude of their opinion difference each time they interact and determines whether they influence each other. Bernardo, Vasca, and Iervolino [VBI21, BVI22] developed variants of the HK model in which nodes have individual, time-dependent confidence bounds that depend on the opinions of neighboring nodes. In their models, nodes that differ too much from their neighbors widen their confidence bounds so that they seek different opinions.

CHAPTER 4

A Bounded-Confidence Model with Heterogeneous Node-Activity Levels

In this chapter, we study a generalization of the DW model that has node weights to model heterogeneous agent-activity levels. Using numerical simulations, we systematically investigate the effects of node weights, which we assign uniformly at random to the nodes. This chapter is adapted[1] from an original paper [LP23] that I co-authored with my advisor Mason A. Porter.[2] We include our code and all figures (including those that are not shown in this chapter) in our repository at `https://gitlab.com/graceli1/NodeWeightDW`.

4.1 Introduction and motivation

In this chapter, we generalize the DW model to incorporate heterogeneous node-activity levels. Although the DW model has been generalized in many ways [NVT20], few studies have modified the procedure to select which agents interact in a time step. The ones that have modified this procedure (see, e.g., [AC15, ZXL18, SPG19, PRM22]) have focused on specific scenarios, rather than on investigating the effects of introducing heterogeneities into agent-selection probabilities.

In the standard DW model (see Section 3.1), one uniformly randomly selects pairs of

[1] Adapted with permission from [LP23]. Copyright 2023, American Physical Society.

[2] I led this project with mentorship from Mason Porter. Mason Porter and I wrote the paper [LP23] together.

agents to interact, but social interactions in real life are not uniformly random. Some studies of DW models have modified the selection procedure that determines which agents interact with [AC15, ZXL18, SPG19, PRM22]. When selecting agents in a way that is not uniformly at random, one can think of the agents as having different activity levels that encode their interaction frequencies. (In a given time interval, we expect these agents to have different numbers of interactions.) The idea of heterogeneous node-activity levels plays an important role in activity-driven models of temporal networks [PGP12]. Activity-driven frameworks have also been used to model which agents can interact with each other in studies of opinion dynamics. Li et al. [LHM17] developed an activity-driven model of opinion dynamics on networks with nodes with assigned activity rates (i.e., assigned activation probabilities). At each time step of their model, one removes all existing edges and then the active agents randomly form connections to other agents. All agents then evaluate the mean opinions of their neighbors to determine if and how to update their own opinions [LHM17]. Researchers have also incorporated heterogeneous agent selection in voter models of opinion dynamics. Masuda et al. [MGR10] studied a voter model with heterogeneous "flip" rates, which one can interpret as heterogeneous node weights that encode activity levels. Baronchelli et al. [BCP11] studied a voter model with heterogeneous edge weights, which one can interpret as encoding heterogeneous edge activities.

Some researchers have generalized the DW model to incorporate heterogeneous agent selection. Alizadeh and Cioffi-Revilla [AC15] studied a modified DW model that incorporates a repulsion mechanism (which was proposed initially by Huet et al. [HDJ08]) in which interacting agents with opinions that differ by more than a cognitive-dissonance threshold move farther away from each other in the space of opinions. They used two-dimensional (2D) vector-valued opinions and placed their nodes on complete graphs. To model agents with different activity levels, Alizadeh and Cioffi-Revilla [AC15] implemented a Poisson node-selection probability, which one can interpret as independent internal "clocks" that determine agent activation. In comparison to selecting agent pairs uniformly at random

(as in the standard DW model), the Poisson node-selection probability can either lessen or promote the spread of extremist opinions, depending on which opinions are more prevalent in more-active agents.

Zhang et al. [ZXL18] examined a modified DW model with asymmetric updates on activity-driven networks. In their model, each node has a fixed activity potential, which one assigns uniformly at random from a distribution of activity potentials. The activity potential of an agent gives its probability of activating. At each time step, each active agent i randomly either (1) creates a message (e.g., a social-media post) or (2) boosts a message that was created by a neighboring agent j. If agent i boosts a message from agent j, then i updates its opinion using the standard DW update mechanism. Zhang et al. [ZXL18] simulated their model on a social network from Tencent Weibo and found that the distribution of activity potentials influences the location of the transition between opinion consensus and fragmentation. The node weights in our BCM are similar in spirit to the activity potentials of Zhang et al. [ZXL18]; they can encode the social activity levels of individuals, such as their frequencies of posting or commenting on social media. However, the way that we incorporate node weights in our BCM differs fundamentally from what Zhang et al. did in [ZXL18]. We consider a time-independent network G, and we select a single pair of neighboring agents to interact at each time step. We first randomly select one agent with a probability that is proportional to its node weight, and then we randomly select a second neighboring agent with a probability that depends on its node weight. The two selected agents then update their opinions using the DW update mechanism.

Heterogeneities in which interactions occur in a social network arise not only because some individuals are more likely to have interactions, but also because some pairs of individuals are more likely to interact than other pairs [BCP11]. The curation of content in social-media feeds is affected by homophily, which is the idea that individuals have a tendency to connect with others that are similar to themselves (e.g., perhaps they have similar ideas or beliefs) [MSC01]. Social-media feeds tend to show content to users that closely matches

their profiles and past activities [Spo17]. To examine the effect of such algorithmic bias on opinion dynamics, Sîrbu et al. [SPG19] studied a modified DW model that includes a homophily-promoting activation mechanism. At each time step, one agent is selected uniformly at random, and then one of its neighbors is selected with a probability that depends on the magnitude of the opinion difference between that neighbor and the first agent. The simulations by Sîrbu et al. of their model on complete graphs suggest that more algorithmic bias yields slower convergence times and more opinion fragmentation [SPG19]. Pansanella et al. [PRM22] applied the same algorithmic-bias model to a variety of network topologies (specifically, Erdős–Rényi, Barabási–Albert, and Lancichinetti–Fortunato–Radicchi (LFR) graphs), and they found similar trends as Sîrbu et al. did on complete graphs.

From the investigations in [AC15, ZXL18, SPG19, PRM22], we know that incorporating heterogeneous agent-selection probabilities into a DW model can influence opinion dynamics. Each of these papers examined a specific implementation of heterogeneous agent selection. We are not aware of any systematic investigations of the effects of heterogeneous agent selection on opinion dynamics in asynchronous BCMs. In the present chapter, we study a BCM with heterogeneous agent-selection probabilities, which we implement using node weights. In general terms, we are studying a dynamical process on node-weighted networks. We use node weights to model agents with different probabilities of interacting. These probabilities can encode heterogeneities in individual behavior, such as in sociability or activity levels. We conduct a methodical investigation of the effects of incorporating heterogeneous node weights, which we draw from various distributions, into our generalization of the DW model. We examine these effects on a variety of networks. In our study, we consider fixed node weights that we assign in a way that disregards network structure and agent opinions. However, one can readily adapt the node weights in our BCM to consider a variety of sociological scenarios in which agents have heterogeneous selection probabilities. We find that introducing heterogeneous node weights into our node-weighted BCM results in longer convergence times and more opinion fragmentation than selecting nodes uniformly at random. Moreover,

when studying models with heterogeneous node selection, our results illustrate that it is important to consider the baseline influence of assigning node-selection probabilities uniformly at random before drawing conclusions about more specific mechanisms such as algorithmic bias [SPG19]. More generally, our model illustrates the relevance of incorporating node weights into network analysis and dynamics.

4.2 A BCM with heterogeneous node-selection probabilities

We now introduce our BCM with heterogeneous node-selection probabilities. Consider an undirected network $G = (V, E)$. The nodes in the network represent agents that have opinions that lie in the closed interval $[0, 1]$. Let each node i have a time-dependent opinion $x_i(t) \in [0, 1]$. As in the standard DW model (see Section 3.1), there is a confidence bound $c \in [0, 1]$ and compromise parameter $m \in (0, 0.5]$. In our BCM, each node also has a fixed node weight $w_i > 0$ that encodes sociability, how frequently it engages in conversations, or simply the desire to share its opinions. One can think of a node's weight as a quantification of how frequently it communicates with its friends or posts on social media. By incorporating network structure, the standard DW model can include nodes with different numbers of friends (or other social connections). However, selecting interacting node pairs uniformly at random is unable to capture the heterogeneous interaction frequencies of individuals. By introducing node weights, we encode such heterogeneity and then examine how it affects opinion dynamics in a BCM. Although we employ fixed node weights, one can adapt our model to include time-dependent node weights, such as through purposeful strategies (e.g., posting on social media more frequently as one's opinions become more extreme).

In our node-weighted BCM, at each discrete time, we first select a node i with a probability that is proportional to its weight. Node i then interacts with a neighbor j, which we select with a probability that is equal to its weight divided by the sum of the weights of i's

neighbors. That is, the probabilities of first selecting node i and then selecting node j are

$$P_1(i) = \frac{w_i}{\sum_{k=1}^{N} w_k}, \quad P_2(j|i) = \frac{w_j}{\sum_{k \in \mathcal{N}(i)} w_k}, \tag{4.2.1}$$

where $\mathcal{N}(i)$ denotes the neighborhood (i.e., the set of neighbors) of i and $j \in \mathcal{N}(i)$. Once we select the pair of interacting agents, we update their opinions following the DW opinion update rule in Equation (3.1.1). For convenience, we repeat the opinion update rule:

$$\begin{aligned} x_i(t+1) &= \begin{cases} x_i(t) + m(x_j(t) - x_i(t)), & \text{if } |x_i(t) - x_j(t)| < c \\ x_i(t), & \text{otherwise}, \end{cases} \\ x_j(t+1) &= \begin{cases} x_j(t) + m(x_i(t) - x_j(t)), & \text{if } |x_i(t) - x_j(t)| < c \\ x_j(t), & \text{otherwise}. \end{cases} \end{aligned} \tag{4.2.2}$$

To capture the fact that some agents have more frequent interactions (such as from greater sociability or a stronger desire to share their opinions) than others, we use a node-based agent-selection procedure. For our node-based selection, we first randomly select one node and then randomly selects one of its neighbors. In the standard DW model (see Section 3.1), one uses an edge-based agent-selection procedure. The choice between edge-based and node-based agent selection can have substantial effects on the dynamics of voter models of opinion dynamics [KP20], and we expect that this is also true for other types of opinion-dynamics models. We are not aware of a comparison of edge-based and node-based agent selection in asynchronous BCMs (and, in particular, in DW models), and it seems both interesting and relevant to explore this issue. Most past research on the DW model has considered edge-based selection [NVT20]. However, [AC15, SPG19, PRM22] used a node-based selection procedure to model heterogeneous activities of agents.

4.2.1 Relating node weights to edge weights

Our BCM incorporates heterogeneous node-selection probabilities with node weights that model phenomena such as the heterogeneous sociability of individuals. One can also study heterogeneous selection probabilities of pairwise (i.e., dyadic) interactions, instead of focusing on the probabilities of selecting individuals. For instance, an individual may discuss their ideological views with a close friend more frequently than with a work colleague. One can use edge weights to determine the probabilities of selecting the dyadic interactions in a BCM. For example, one can assign each edge (i, j) a weight p_{ij} such that the sum of the weights is 1 (i.e., $\sum_{(i,j)\in E} p_{ij} = 1$). Then, at each discrete time, one can select edge (i, j) for interaction in a DW model with probability p_{ij}.

One can relate such edge selection to the node selection in our BCM. For example, one can select the edge between nodes i and j either by selecting node i and then node j or by selecting node j and then node i. In particular, our node-weighted BCM is equivalent to a DW model in which, at each discrete time, the edge (i, j) is selected for interaction with probability

$$p_{ij} = P_1(i)P_2(j|i) + P_1(j)P_2(i|j)\,, \tag{4.2.3}$$

where P_1 and P_2 are defined in Equation (4.2.1).

4.3 Details of numerical simulations

In this section, we discuss the setup of our numerical simulations of our node-weighted BCM.

4.3.1 Network structures

We now describe the details of the networks on which we simulate our node-weighted BCM. We summarize these networks in Table 4.1.

We simulate our BCM on complete graphs as a baseline scenario that lets us examine

how incorporating heterogeneous node-selection probabilities affects opinion dynamics. To examine finite-size effects from our networks, we consider complete graphs with sizes $N \in \{10, 20, 30, 45, 65, 100, 150, 200, 300, \ldots, 1000\}$. For all other synthetic networks, we consider networks with $N = 500$ nodes.

Table 4.1: The networks on which we simulate our node-weighted BCM.

Network	**Description**	**Parameters**
$C(N)$	Complete graph with N nodes	$N \in \{10, 20, 30, 45, 65, 100, 150, 200, 300, \ldots, 1000\}$
$G(N, p)$	Erdős–Rényi (ER) random-graph model with N nodes and homogeneous, independent edge probability p	$p \in \{0.1, 0.3, 0.5, 0.7\}$
Two-Community SBM[1]	Stochastic block model with 2×2 blocks. Edges between nodes in the same set (A or B) exist with a larger probability than edges between nodes in different sets; the block probabilities satisfy $P_{BB} > P_{AA} > P_{AB}$.	$P_{AA} = 49.9/374$ $P_{BB} = 49.9/124$ $P_{AB} = 1/500$
Core–Periphery SBM[1]	Stochastic block model with 2×2 blocks. Set A is a set of core nodes and set B is a set of peripheral nodes. The block probabilities satisfy $P_{AA} > P_{AB} > P_{BB}$.	$P_{AA} = 147.9/374$ $P_{BB} = 1/174$ $P_{AB} = 1/25$
Caltech Network	The largest connected component of the Facebook friendship network at Caltech on one day in fall 2005. This network, which is part of the FACEBOOK100 data set [RKM11, TMP12], has 762 nodes and 16,651 edges.	

[1]Our SBM networks have $N = 500$ nodes. We partition an SBM network into two sets of nodes; set A has 75% of the nodes, and set B has 25% of the nodes.

To consider networks with different edge densities, we generate synthetic networks using the $G(N, p)$ ER random-graph model (see Section 2.3.2). When $p = 1$, this yields a complete

graph. We examine $G(500, p)$ graphs with $p \in \{0.1, 0.3, 0.5, 0.7\}$.

To determine how a network with an underlying block structure affects the dynamics of our node-weighted BCM, we consider SBM networks with 2×2 blocks (see Section 2.3.3). Inspired by the choices of Kureh and Porter [KP20], we consider two types of SBM networks: the first has a two-community structure, and the second has a core–periphery structure. To construct our 2×2 SBMs, we partition a network into two sets of nodes; set A has 375 nodes (i.e., 75% of the network) and set B has 125 nodes (i.e., 25% of the network).

Consider the symmetric edge-probability matrix in Equation (2.3.1). In a two-community SBM, the probabilities P_{AA} and P_{BB} are larger than P_{AB}, so edges between nodes in the same community exist with a larger probability than edges between nodes in different communities. For our two-community SBM, we choose P_{AA} and P_{BB} so that the expected mean degree matches that of the $G(500, 0.1)$ ER model if we consider only edges within set A or only edges within set B. A network from the $G(N, p)$ model has an expected mean degree of $p(N-1)$ [New18], so we want the two communities of these SBM networks to have an expected mean degree of $49.9 = 0.1 \times 499$. We thus use the edge probabilities $P_{AA} = 49.9/374$ and $P_{BB} = 49.9/124$. To ensure that there are few edges between the sets A and B, we choose $P_{AB} = 1/500$.

We want our core–periphery SBM with core set A and periphery set B to satisfy $P_{AA} > P_{AB} > P_{BB}$. We chose P_{AA} so that the expected mean degree matches that of the $G(500, 0.3)$ model (i.e., it is 147.9) if we only consider edges within the set A. We thus choose the edge probability $P_{AA} = 147.9/374$. To satisfy $P_{AA} > P_{AB} > P_{BB}$, we choose $P_{AB} = 1/25$ and $P_{BB} = 1/174$.

Finally, we investigate our node-weighted BCM on a real social network from Facebook friendship data. We use the Caltech network from the FACEBOOK100 data set [RKM11, TMP12]. Its nodes encode individuals at Caltech, and its edges encode Facebook "friendships" between them on one day in fall 2005. We only consider the network's largest connected component, which has 762 nodes and 16,651 edges.

4.3.2 Node-weight distributions

In Table 4.2, we give the parameters and probability density functions of the node-weight distributions that we examine in our BCM. In this subsection, we discuss our choices of distributions.

Table 4.2: The names and specifications of our node-weight distributions. We show both the general mathematical expressions for the means and the specific values of the means for our parameter values. For the Pareto distributions, we truncate the distribution means to four digits after the decimal point. For all other distributions, the means are exact. Additionally, for our Pareto-80-43 distribution, 80% of the node weights are distributed among approximately 42.7723% of nodes (see Footnote 3), rather than among 43% of nodes.

Distribution	**Probability density function**	**Parameter values**	**Domain**	**Mean**	
Constant	$\delta(x-1)$	N/A	$\{1\}$	1	1
Pareto-80-43	$\frac{\alpha}{x^{\alpha+1}}$	$\alpha = \log_{4.5}(10)$	$[1,\infty)$	$\frac{\alpha}{\alpha-1}$	2.8836
Pareto-80-20		$\alpha = \log_4(5)$			7.2126
Pareto-90-10		$\alpha = \log_9(10)$			21.8543
Exp-80-43	$\frac{1}{\beta}\exp\left(\frac{-(x-1)}{\beta}\right)$	$\beta = 1.8836$	$[1,\infty)$	$\beta+1$	2.8836
Exp-80-20		$\beta = 6.2125$			7.2125
Exp-90-10		$\beta = 20.8543$			21.8543
Unif-80-43	$\frac{1}{b-1}$	$b = 4.7672$	$[1,b]$	$\frac{1}{2}(1+b)$	2.8836
Unif-80-20		$b = 13.4250$			7.2125
Unif-90-10		$b = 42.7086$			21.8543

To study the effects of incorporating node weights in our BCM, we compare our model to a baseline DW model. To ensure a fair comparison, we implement a baseline DW model that selects interacting agents uniformly at random using a node-based selection process. As we discussed at the end of Section 4.2, it is much more common to employ an edge-based selection process. We refer to the case in which all node weights are equal to 1 (that is, $w_i = 1$

for all nodes i) as the "constant weight distribution". The constant weight distribution (and any other situation in which all node weights equal the same positive number) results in a uniformly random selection of nodes for interaction. This is what we call the "baseline DW model"; we compare our DW model with heterogeneous node weights to this baseline model. We reserve the term "standard DW model" for the DW model with uniformly random edge-based selection of agents. When all of the nodes of a network have the same degree, our baseline DW model is equivalent to the standard DW model.

The node weights in our BCM encode heterogeneities in interaction frequencies, such as when posting content online. The majority of online content arises from a minority of user accounts [GTC09]. A "90-9-1 rule" has been proposed for such participation inequality. In this rule of thumb, about 1% of the individuals in online discussions (e.g., on social-media platforms) account for most contributions, about 9% of the individuals contribute on occasion, and the remaining 90% of the individuals are present online (e.g., they consume content) but do not contribute to it [Nie06]. Participation inequality has been documented in a variety of situations, including in the numbers of posts on digital-health social networks [Mie14], posts on internet support groups [CCG14], and contributions to open-source software-development platforms [GCB20]. Inequality in user activity has also been examined on Twitter (which recently was rebranded as $\mathbb{X}$)[AMS19]. For example, Xiong and Liu [XL14] used a power-law distribution to model the number of tweets about different topics. A few years ago, a survey by the Pew Research Center found that about 10% of the accounts of adult Twitter users in the United States generate about 80% of the tweets of such accounts [WH19].

One can interpret the node weights in our BCM as encoding the participation frequencies of individuals who contribute content to a social-media platform. We model online participation inequality by using a Pareto distribution for the node weights. This choice of distribution is convenient because of its simple power-law form. It has also been used to model inequality in a variety of other contexts, including distributions of wealth, word frequencies, website visits, and numbers of paper citations [New05]. When representing

social-media interactions, we care only about accounts that make posts or comments; we ignore inactive accounts. Therefore, we impose a minimum node weight in our model. We use the Pareto type-I distribution, which is defined on $[1, \infty)$, so each node has a minimum weight of 1. This positive minimum weight yields reasonable computation times in our simulations of our BCM. Nodes with weights near 0 would have very small probabilities of interacting, and allowing such weights would prolong simulations.

Let Pareto-X-Y denote the continuous Pareto distribution in which (in theory) $X\%$ of the total node weight is distributed among $Y\%$ of the nodes. In practice, once we determine the N node weights for a simulation with a Pareto node-weight distribution, it is not true that precisely $X\%$ of the total weight is distributed among $Y\%$ of the N nodes. The Pareto principle (which is also known as the "80-20 rule") is a popular rule of thumb that suggests that 20% of a population of individuals have 80% of the available wealth [New05]. Accordingly, we consider a Pareto-80-20 distribution. As an example of a node-weight distribution with a more extreme inequality, we also consider a Pareto-90-10 distribution. Additionally, as an example of a node-weight distribution with less inequality, we also consider a Pareto-80-43 distribution.[3]

We also examine uniform and exponential distributions of node weights. To match the domain of our Pareto distributions, we shift the uniform and exponential distributions so that their minimum node weight is also 1. We also choose their parameters to approximate the means of our Pareto distributions. We use Exp-X-Y and Unif-X-Y as shorthand notation to denote exponential and uniform distributions, respectively, with means that match that

[3] Inspired by the results of the aforementioned Pew Research Center survey of Twitter (now rebranded as $\mathbb{X}$) users [WH19], we intended to consider a Pareto-80-10 distribution, in which one expects 80% of the total weight to be distributed among 10% of nodes. The Pareto distribution family that we consider has the probability density function $\alpha/x^{\alpha+1}$ (see Table 4.2). A Pareto-80-10 distribution entails that $\alpha = \log_{4.5}(5)$. However, we mistakenly used $\alpha = \log_{4.5}(10)$ in our simulations. This choice corresponds to a Pareto-80-43 distribution. To calculate the percentage $Y\%$ of nodes that one expects to have $X\%$ of the total node weight, one calculates $Y/100 = 1 - (1 - X/100)^{(1-1/\alpha)}$. In our Pareto distribution with $\alpha = \log_{4.5}(10)$, we have that 80% of the node weights are distributed among approximately 42.7723% of the nodes. As a shorthand, we refer to this distribution as a "Pareto-80-43" distribution.

of the Pareto-X-Y distribution to four decimal places (see Table 4.2). When we examine the results of our numerical simulations, we want to compare distributions with similar means. We use the phrase "80-20 distributions" to refer to the Pareto-80-20, Exp-80-20, and Unif-80-20 distributions. We analogously use the phrases "90-10 distributions" and "80-43 distributions." In total, we examine three different families of distributions (Pareto, exponential, and uniform) with tails of different heaviness. In Table 4.2, we show the details of the probability density functions and the parameters of our node-weight distributions.

4.3.3 Simulation specifications

In our node-weighted BCM, agents have opinions in the 1D opinion space $[0, 1]$. Accordingly, we suppose that the confidence bound $c \in (0, 1)$. We suppose that the compromise parameter $m \in (0, 0.5]$, which is the typically studied range for DW models [MVP18, NVT20].

In our node-weighted BCM, the generation of the graphs in a random-graph ensemble, the sets of node weights, the sets of initial opinions, and the selection of pairs of agents to interact at each time step are all stochastic. We use Monte Carlo simulations to reduce these sources of noise in our simulation results. For each of our random-graph models (i.e., the ER and SBM graphs), we generate five graphs. For each graph and each node-weight distribution, we randomly generate ten sets of node weights. For each set of node weights, we generate ten sets of initial opinions that are distributed uniformly at random. In total, we consider 100 distinct sets of initial opinions and node weights for the Monte Carlo simulations of each individual graph. When we compare simulations from different distributions of node weights for the same individual graph, we reuse the same 100 sets of initial opinions.

The update rule (see Equation (4.2.2)) for our node-weighted BCM is the same as the update rule in the standard DW model. Therefore, using the same argument as in the standard DW model (see Section 3.3.1), each opinion in our node-weighted BCM converges to a steady-state value. Furthermore, as in the standard DW model, we can determine the steady-state opinion clusters in finite time (see the discussion at the end of Section 3.3.2) for

our node-weighted BCM. In our numerical simulations in this chapter, we use the stopping criterion in Equation (3.3.8). We specify that one of our simulations has converged if each opinion cluster has a diameter that is less than a tolerance of 0.02. That is, for each opinion cluster S_r, we have $\max_{i,j \in S_r} |x_i - x_j| < 0.02$. We denote the time step in which a simulation reaches the stopping criterion by T_f.

It is computationally expensive to numerically simulate a DW model. Additionally, as we will show in Section 4.4, our node-weighted DW model with heterogeneous node weights often converges to a steady state even more slowly than the baseline DW model. To reduce the computational burden of checking for convergence, we compute the convergence time to three significant digits. We thereby avoid checking for convergence at each time step. To guarantee that each simulation stops in a reasonable amount of time, we set a bailout time of 10^9 time steps. In our simulations, the convergence time T_f is always shorter than the bailout time. We thus report the results of our simulations as steady-state results.

To characterize the opinions in our node-weighted BCM, we calculate the quantities that we described in Sections 3.3.3 and 3.3.4 at steady state. Namely, we calculate the steady-state numbers of major and minor opinion clusters (see Section 3.3.3) and steady-state Shannon entropy $H(T_f)$ (see Equation (3.3.7)) of entire opinion-cluster profiles (see Section 3.3.4). In this chapter, we say that an opinion cluster is a "minor opinion cluster" if it has at most 2% of the nodes in a network. That is, an opinion cluster S_r is a minor opinion cluster if $|S_r| \leq 0.02N$. We say that an opinion cluster that is not a minor cluster is a "major opinion cluster".

Additionally, we calculate a quantity that we call the *mean local receptiveness*, which we define shortly. One way to quantify opinion fragmentation is to look at a local level and consider individual nodes of a network. As Musco et al. [CW22] pointed out, if an individual node has many neighbors with similar opinions to it, then it may be "unaware" of other opinions in the network. For example, most of the neighbors of a node may hold an opinion that is uncommon in the network. This phenomenon is sometimes called a "majority

illusion" [LYW16]. If a set of adjacent nodes tend to have neighbors with similar opinions as theirs, they may be in an "echo chamber" [FGR16], as it seems that they are largely exposed only to conforming opinions. To quantify the local observations of nodes, Musco et al. [CW22] examined a notion of local agreement by calculating the fraction of a node's neighbors with opinions on the same side of the network's mean opinion as that node. In our simulations, we often observe opinion fragmentation with three or more opinion clusters. Therefore, we need to look beyond the mean opinion of an entire network. To do this, we introduce the *local receptiveness* of an node. At time t, a node i with neighborhood $\mathcal{N}(i)$ has a local receptiveness of

$$L_i(t) = \frac{|\{j \in \mathcal{N}(i) : |x_i(t) - x_j(t)| < c\}|}{|\mathcal{N}(i)|}\,. \tag{4.3.1}$$

That is, $L_i(t)$ is the fraction of the neighbors of agent i at time t to which it is receptive (see Section 3.1). In the present chapter, we only consider connected networks, so each agent i has $|\mathcal{N}(i)| \geq 1$ neighbors. If one wants to consider isolated nodes, one can assign them a local receptiveness of 0 or 1. In our numerical simulations, we calculate the local receptiveness of each agent of a network at the convergence time T_f. We then calculate the mean $\langle L_i(T_f)\rangle$ of all agents in the network. This is the steady-state mean local receptiveness, as it is based on edges in the steady-state effective graph $G_{\text{eff}}(T_f)$ (see Equation (3.3.3)). When consensus is not reached, a smaller mean local receptiveness is an indication of greater opinion fragmentation. As we will discuss in Section 4.4, computing Shannon entropy and mean local receptiveness can give insight into the extent of opinion fragmentation when one considers them in concert with the number of opinion clusters.

4.4 Numerical simulations and results

In this section, we present results of our numerical simulations of our node-weighted BCM. In our numerical experiments, we consider compromise-parameter values $m \in \{0.1, 0.3, 0.5\}$.

For the confidence bound, we first consider $c \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$, and we then examine additional values of c near regions with interesting results. As we discussed in Section 4.3.3, for each individual graph, we use 100 distinct sets of initial opinions and node weights in Monte Carlo simulations of our BCM. For each of the random-graph models (i.e., ER and SBM graphs), we generate five graphs. For the 500-node complete graphs, we use the 10 weight distributions in Table 4.2. Because of computation time, we consider the 90-10 distributions only for the 500-node complete graph. For the other networks in Table 4.1, we consider seven distributions: the constant weight distribution, the three 80-43 distributions, and the three 80-20 distributions.

In Table 4.3, we summarize the trends that we observe in the examined networks. In the following subsections, we discuss details of our results for each type of network. For each network, we plot the numbers of major and minor clusters, Shannon entropies, and values of mean local receptiveness at steady state. We include our code and all figures (including those not shown in this chapter) in our repository at `https://gitlab.com/graceli1/NodeWeightDW`.

4.4.1 Simulations on a complete graph

The simplest underlying network structure on which we simulate our node-weighted BCM is a complete graph. A complete graph gives a baseline setting to examine how heterogeneous node-selection probabilities affect opinion dynamics. In our numerical simulations on a 500-node complete graph, we consider all three means (which we denote by 80-43, 80-20, and 90-10) for each of the uniform, exponential, and Pareto node-weight distribution families.

The standard DW model on a complete graph with agents with opinions in the interval $[0, 1]$ eventually reaches consensus if the confidence bound $c \geq 0.5$. As one decreases c from 0.5, there are progressively more steady-state opinion clusters (both major and minor clusters) [BKR03, Lor07]. Lorenz [Lor07] showed using numerical simulations that the number of major clusters is approximately $\lfloor \frac{1}{2c} \rfloor$ for the standard DW model. Therefore, there is a

Table 4.3: Summary of the trends in our simulations of our node-weighted BCM. Unless we note otherwise, we observe these trends for each of the networks that we examine (complete graphs, ER and SBM random graphs, and the Caltech Facebook network).

Quantity	**Trends**
Convergence Time	• For fixed values of the confidence bound c and compromise parameter m, the heterogeneous weight distributions have longer convergence times than the constant weight distribution.
Opinion Fragmentation[1]	• For fixed values of $c \in [0.1, 0.4]$ and m, the heterogeneous weight distributions usually have more opinion fragmentation than the constant weight distribution. • For fixed values of c and m and a fixed distribution mean, there usually is more opinion fragmentation when the tail of a distribution is heavier. • For fixed values of c and m and a fixed family of distributions, there usually is more opinion fragmentation when a distribution has a larger mean.
Number of Major Clusters	• A larger minimum value of c is required to always reach consensus for a heterogeneous weight distribution than for the constant weight distribution. • For fixed values of c and m and a fixed distribution mean, there usually are more major clusters when a distribution tail is heavier. • For fixed values of c and m and a fixed family of distributions, there usually are more major clusters when a distribution has a larger mean.
Number of Minor Clusters	• For the constant weight distribution and for fixed c, there usually are more minor clusters when the compromise parameter $m \in \{0.3, 0.5\}$ than when $m = 0.1$. The heterogeneous weight distributions do not follow this trend.[2]

[1]We quantify opinion fragmentation using Shannon entropy and mean local receptiveness. We observe clearer trends for the Shannon entropy than for the mean local receptiveness.

[2]For the Caltech network, we usually observe more minor clusters when $m \in \{0.3, 0.5\}$ than when $m = 0.1$ for each of our heterogeneous weight distributions.

transition between consensus and opinion fragmentation for $c \in [0.25, 0.3]$. In our simulations, we observe that this transition occurs for $c \in [0.25, 0.4]$ in our node-weighted BCM. To examine this transition, we zoom in on these values of c. For the uniform and exponential distributions, we focus on $c \in [0.25, 0.3]$. For the Pareto distributions, the transition occurs for larger values of c than for the other distributions; we consider additional values of $c \in [0.3, 0.4]$. For the constant weight distribution, which gives our baseline DW model, we examine all values of c that we consider for any other distribution.

In Figure 4.1, we show the convergence times of our BCM simulations for various node-weight distributions. For fixed values of c and m, all of the heterogeneous weight distributions yield longer convergence times than the constant weight distribution. Additionally, for fixed c and m and a fixed family of distributions (uniform, exponential, or Pareto), the convergence time increases as we increase the mean of a distribution. Furthermore, for fixed c and for each heterogeneous weight distribution, the convergence time usually increases as we decrease the compromise parameter m. When calculating convergence time, we include time steps in which two nodes interact but do not change their opinions. To see if the heterogeneous weight distributions have inflated convergence times as a result of having more of these futile interactions, we also calculate the number of time steps to converge when we exclude such time steps. That is, we count the total number of opinion changes that it takes to converge. On a logarithmic scale, there is little difference between the total number of opinion changes and the total number of time steps to converge.

In Figure 4.2, we show the numbers of steady-state major opinion clusters in our BCM simulations for various node-weight distributions. For all weight distributions, consensus occurs in all of our simulations when the confidence bound $c \geq 0.5$. For fixed values of $c \in [0.1, 0.4]$ and m, the heterogeneous weight distributions usually yield more steady-state major clusters than the constant weight distribution. When we introduce heterogeneous node weights into our BCM, we need a larger confidence bound c than for the constant weight distribution to always reach consensus in our simulations. It appears that our BCM

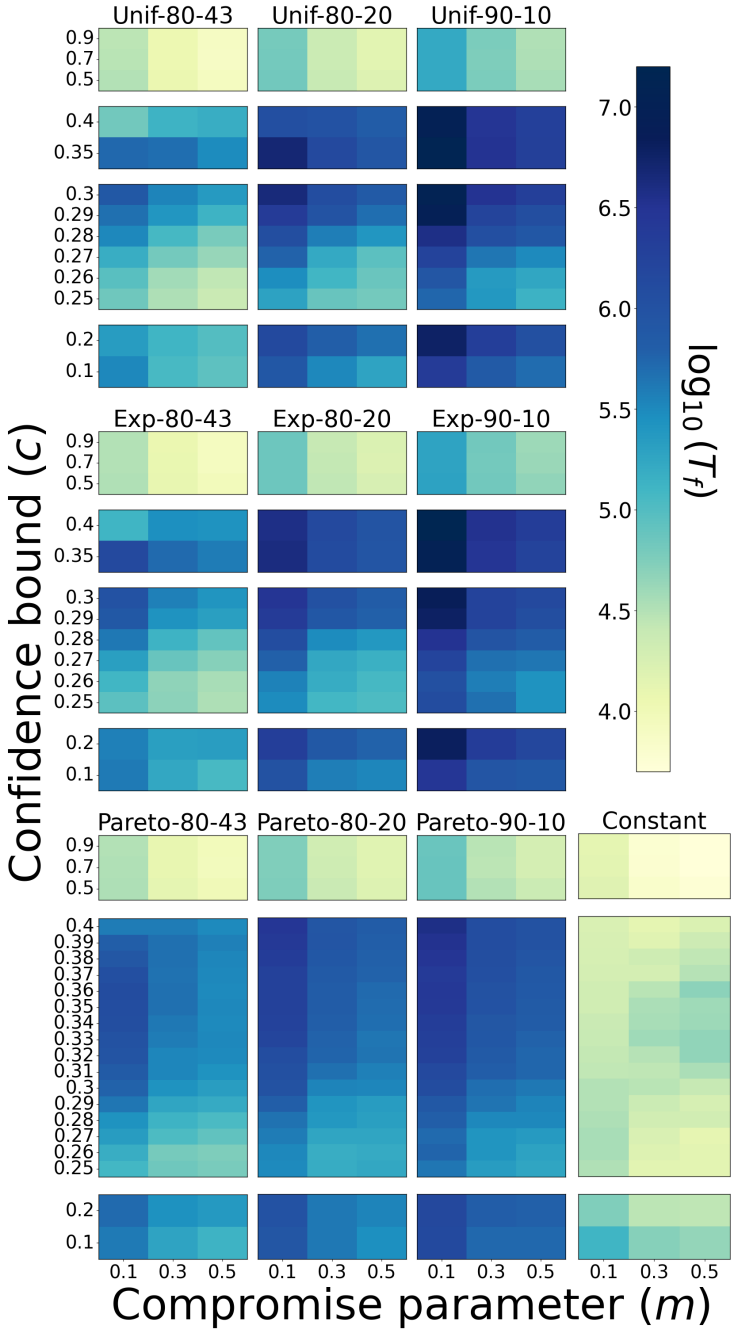

Figure 4.1: Convergence times (in terms of the number of time steps) in simulations of our node-weighted BCM on a 500-node complete graph with various node-weight distributions. If we consider only the time steps in which interacting nodes actually change their opinions, the convergence times are smaller; however, the trends are the same. For this heat map and all subsequent heat maps, the depicted values are means of our simulations of our BCM for each node-weight distribution and each value of the BCM parameter pair (c, m). (This figure originally appeared in [LP23].)

with heterogeneous node weights tends to have more opinion fragmentation than the baseline DW model. For fixed c and m, we observe for each distribution family that the mean number of steady-state major clusters increases as we increase the distribution mean. To see this, proceed from left to right in Figure 4.2 from the 80-43 distributions to the 80-20 distributions and then to the 90-10 distributions. Additionally, for fixed values of c and m and a fixed distribution mean, there are usually more steady-state major clusters as we proceed from a uniform distribution to an exponential distribution and then to a Pareto distribution.

To investigate how the node-weight distribution and the BCM parameters (i.e., c and m) affect the amount of opinion fragmentation, we calculate the Shannon entropy (see Equation (3.3.7)) and mean local receptiveness (see Equation (4.3.1)) at steady state. In Figure 4.3, we show the steady-state entropy values of our BCM simulations for various node-weight distributions. For all node-weight distributions, when there is opinion fragmentation instead of consensus, the steady-state entropy increases as we decrease the confidence bound c for fixed m. In line with our observations in Figure 4.2, when $c \in [0.1, 0.4]$, simulations of heterogeneous weight distributions usually yield larger entropies than the constant weight distribution. For fixed values of c and m and a fixed distribution mean, we also tend to observe a slightly larger entropy as we proceed from a uniform distribution to an exponential distribution and then to a Pareto distribution. For the Pareto distribution family and fixed values of c and m, the entropy increases as we increase the distribution mean. (Proceed from left to right in Figure 4.3.) The exponential and uniform distributions also have this trend, although it is less pronounced (i.e., the entropies tend to increase only slightly) than for the Pareto distributions. When we quantify fragmentation using Shannon entropy, we conclude that increasing the mean node weight has less effect on the amount of opinion fragmentation for the uniform and exponential distributions than it does for the Pareto distributions. Because Shannon entropy depends on the sizes of the opinion clusters, it provides more information about opinion fragmentation than tracking only the number of major opinion clusters. Our plot of the steady-state mean local receptiveness illustrates the same trends

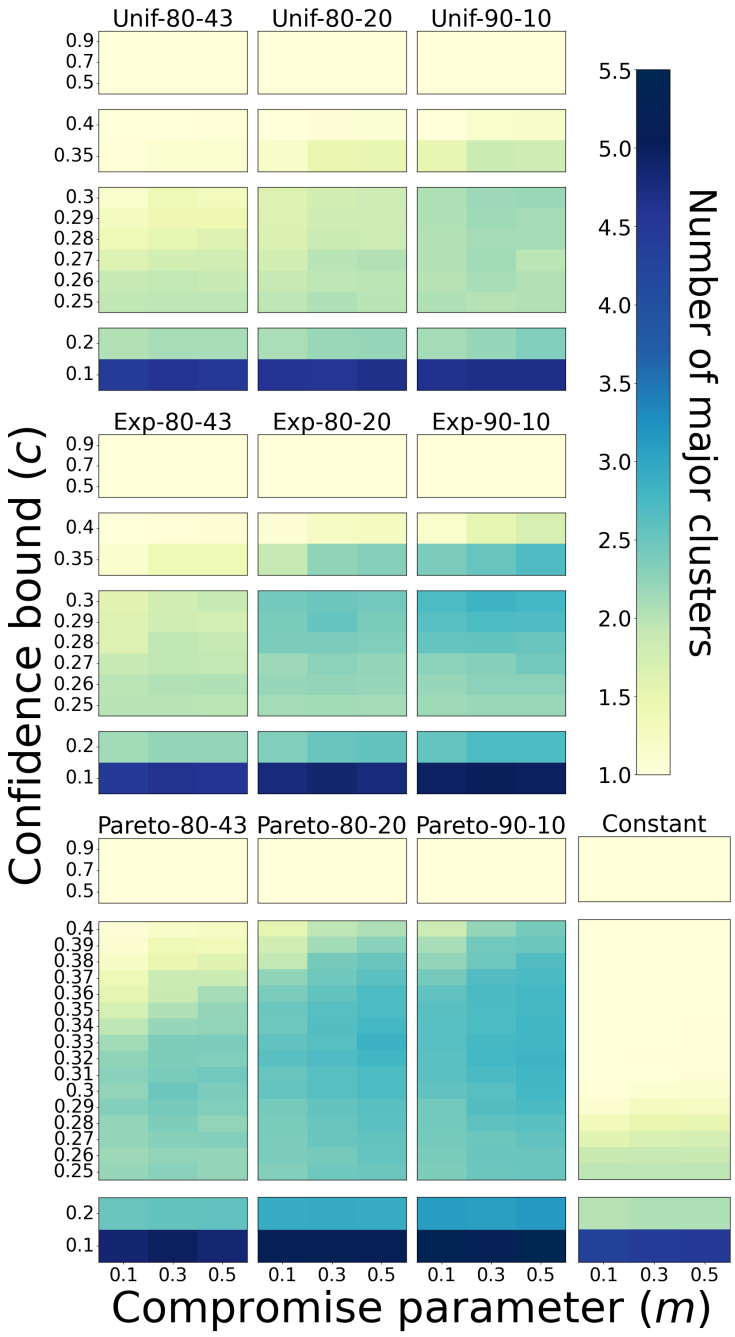

Figure 4.2: Steady-state numbers of major opinion clusters in simulations of our node-weighted BCM on a 500-node complete graph with various node-weight distributions. We consider a cluster to be major cluster if it has more than 2% of the nodes of a network. In this case, a major cluster must have at least 11 nodes. (This figure originally appeared in [LP23].)

as the entropy, but we do not show it here. (See our code repository for the relevant figure.) This suggests that both Shannon entropy and mean local receptiveness are useful for quantifying opinion fragmentation.

In Figure 4.4, we show the numbers of steady-state minor opinion clusters in our BCM simulations. For each node-weight distribution and each value of c and m, when we take the mean of our 100 simulations, we obtain at most three steady-state minor clusters. We observe the most minor clusters when $c \in \{0.1, 0.2\}$, which are the smallest confidence bounds that we examine. For the constant weight distribution, we typically observe more minor clusters when $m \in \{0.3, 0.5\}$ than when $m = 0.1$. However, we do not observe this trend for the heterogeneous weight distributions. For example, for the Pareto-80-43 distribution, when $c \in [0.34, 0.4]$, decreasing m results in more minor opinion clusters. For the three Pareto distributions, as we decrease m, we also observe that minor clusters tend to appear at smaller confidence bounds. Smaller values of m entail smaller opinion compromises for interacting agents; this may give more time for agents to interact before they settle into their steady-state opinion clusters. For the constant weight distribution, this may reduce the number of minor clusters by giving more opportunities for agents to assimilate into a major cluster. However, for our heterogeneous weight distributions, nodes with larger weights have larger probabilities of interacting with other nodes and we no longer observe fewer minor clusters as we decrease m.

We now propose a possible mechanism by which our node-weighted BCM may promote the trends in Table 4.3. In Figure 4.5, we show the trajectories of opinions versus time for a single simulation with node weights that we draw from a Pareto-80-43 distribution. To qualitatively describe our observations, we examine the large-weight and small-weight nodes (i.e., the nodes that are near and at the extremes of a set of node weights in a given simulation). Because our node-selection probabilities are proportional to node weights, we normalize the node weights in a simulation to sum to 1. In Figure 4.5, the large-weight nodes appear to quickly stabilize into their associated steady-state major opinion clusters,

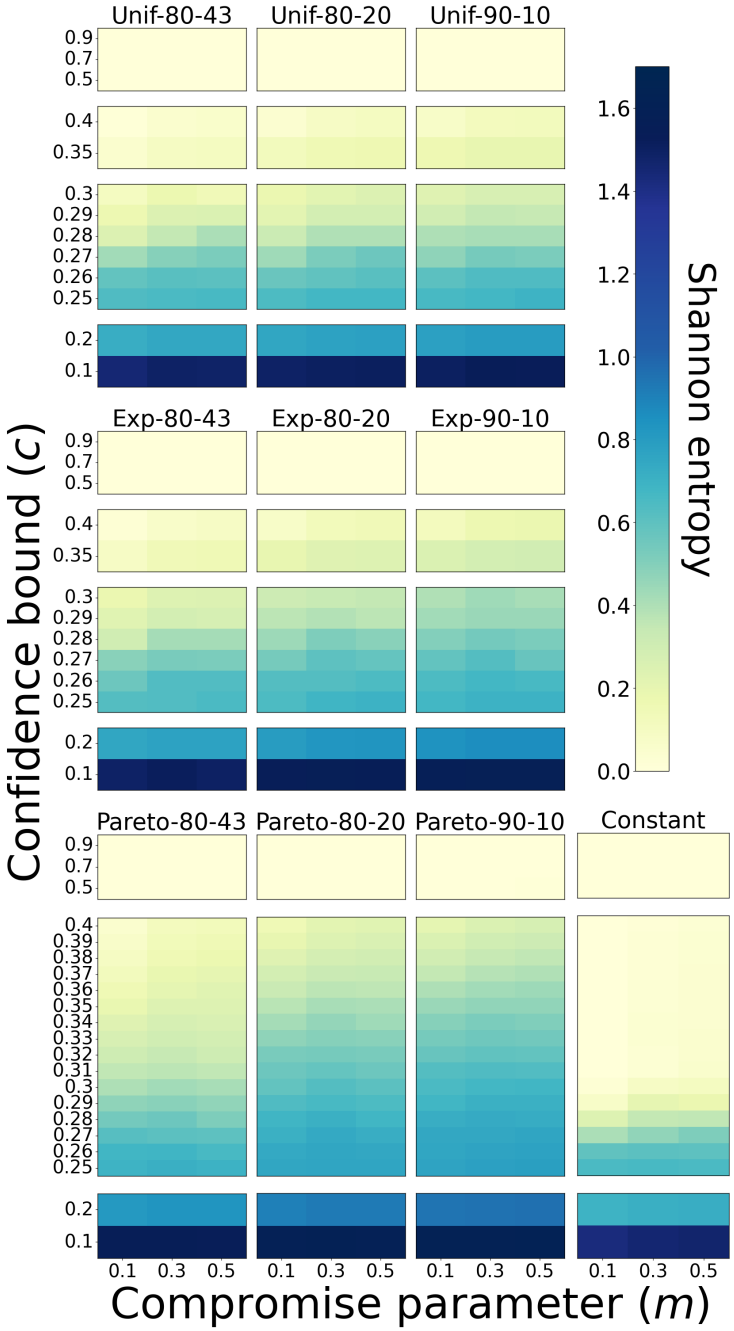

Figure 4.3: Shannon entropies of the steady-state opinion-cluster profiles in simulations of our node-weighted BCM on a 500-node complete graph with various node-weight distributions. (This figure originally appeared in [LP23].)

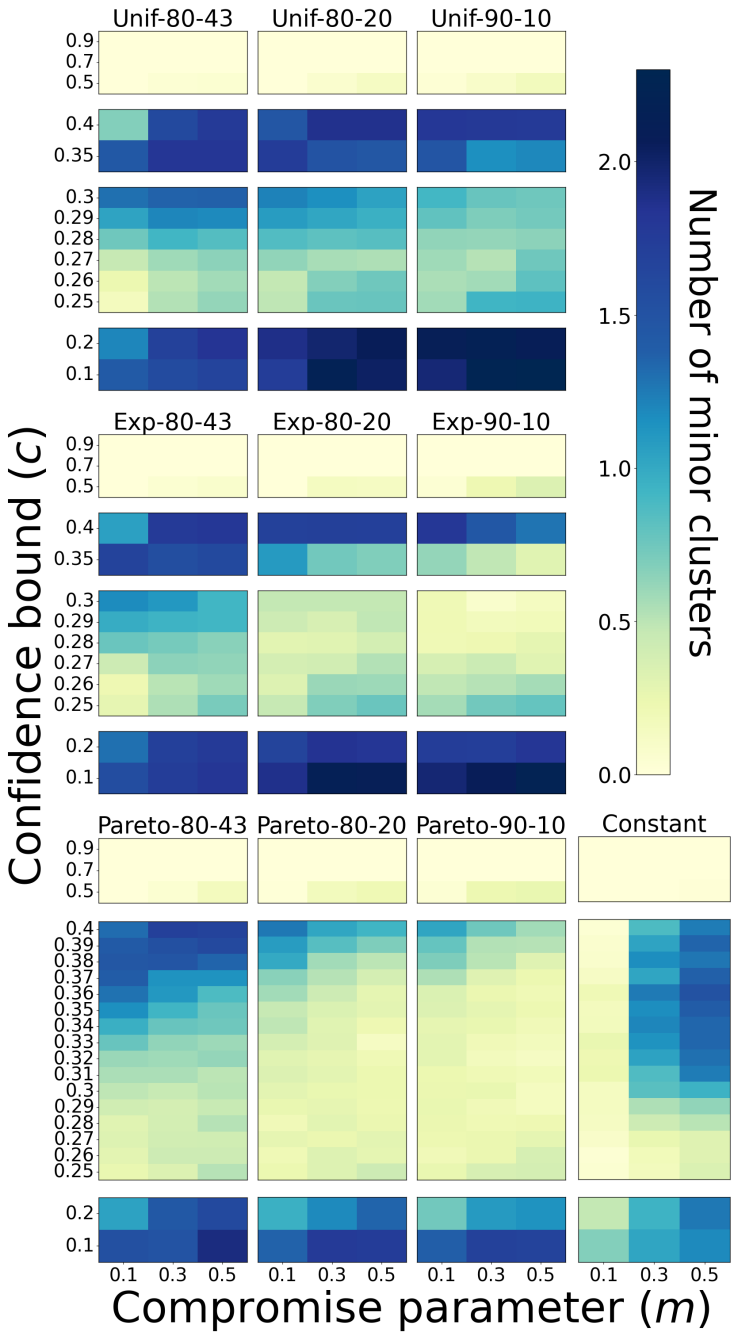

Figure 4.4: Steady-state numbers of minor opinion clusters in simulations of our node-weighted BCM on a 500-node complete graph with various node-weight distributions. We consider a cluster to be minor cluster if it has at most 2% of the nodes of a network. In this case, a minor cluster must have 10 nodes or fewer. (This figure originally appeared in [LP23].)

and some small-weight nodes are left behind to form the two minor clusters. More generally, in our simulations of our BCM on a complete graph, we observe that heterogeneity in the node weights results in large-weight nodes interacting more frequently than other nodes and quickly settling into steady-state major opinion clusters. Small-weight nodes that are not selected for opinion updates early in a simulation are left behind to form the smallest clusters in a steady-state opinion-cluster profile; this increases the amount of opinion fragmentation. When we increase the mean node weight, increase the relative proportion of large-weight nodes (by increasing the heaviness of the tail of a distribution), or decrease the value of the compromise parameter m, small-weight nodes tend to take longer to settle into opinion clusters. Node-weight heterogeneity may thereby promote both opinion fragmentation and the formation of minor opinion clusters.

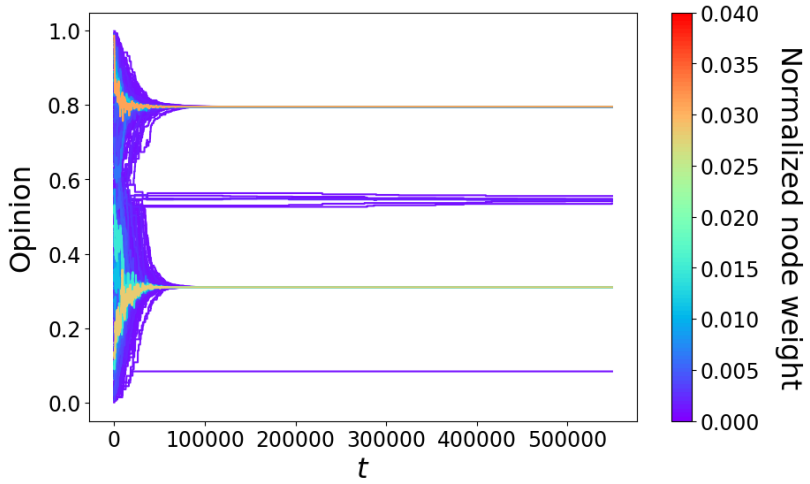

Figure 4.5: Sample trajectories of agent opinions versus time t in a single simulation of our node-weighted BCM on a 500-node complete graph with BCM parameters $c = 0.2$ and $m = 0.1$ and node weights that we draw from a Pareto-80-43 distribution. We color the trajectory of each agent by its node weight, which we normalize so that the sum of all node weights is 1. The nodes in the two minor opinion clusters are all small-weight nodes; their weights are close to 0 (and are hence in purple). (This figure originally appeared in [LP23].)

4.4.2 ER graphs

We now discuss our simulations of our BCM on $G(N, p)$ ER random graphs, where p is the homogeneous, independent probability of an edge between any pair of nodes (see Section 2.3.2). In our simulations, we examine edge probabilities $p \in \{0.1, 0.3, 0.5, 0.7\}$ and generate five graphs for each value of p. Each graph has $N = 500$ nodes. For each value of p, we observe each of the trends in Table 4.3.

In Figure 4.6, we show the steady-state Shannon entropies of our simulations for various node-weight distributions and values of p. The entropies are comparable to those that we obtained in our simulations on a 500-node complete graph (see Section 4.4.1). When $c \in [0.1, 0.4]$, for each of our three node-weight distribution families and for fixed values of p, c, and m, the 80-20 distribution tends to yield a larger Shannon entropy than the 80-43 distribution (which has a smaller mean).

For larger p, we expect the results of our simulations on $G(500, p)$ networks to be similar to those of our simulations on a 500-node complete graph. For $p \in \{0.3, 0.5, 0.7\}$ and $N = 500$, the number of major opinion clusters and the mean local receptiveness are comparable to the corresponding results for a 500-node complete graph. When $p = 0.1$ and there is opinion fragmentation, for a fixed node-weight distribution and fixed values of c and m, we usually observe fewer major opinion clusters than for larger values of p. For $p = 0.1$, a fixed node-weight distribution, and fixed $c \in [0.1, 0.4]$ and m, we also observe that the mean local receptiveness tends to be larger than it is for larger p. One possible contributing factor for this observation may be that smaller values of p yield $G(N, p)$ graphs with more small-degree nodes; these small-degree nodes have fewer available values of local receptiveness than larger-degree nodes. For example, a node with degree 2 can only have a local receptiveness of 0, 0.5, or 1. Unless a small-degree node is an isolated node in the steady-state effective graph $G_{\text{eff}}(T_f)$, its presence may help inflate the value of the steady-state mean local receptiveness.

For a fixed node-weight distribution and fixed values of c and m, decreasing p tends to

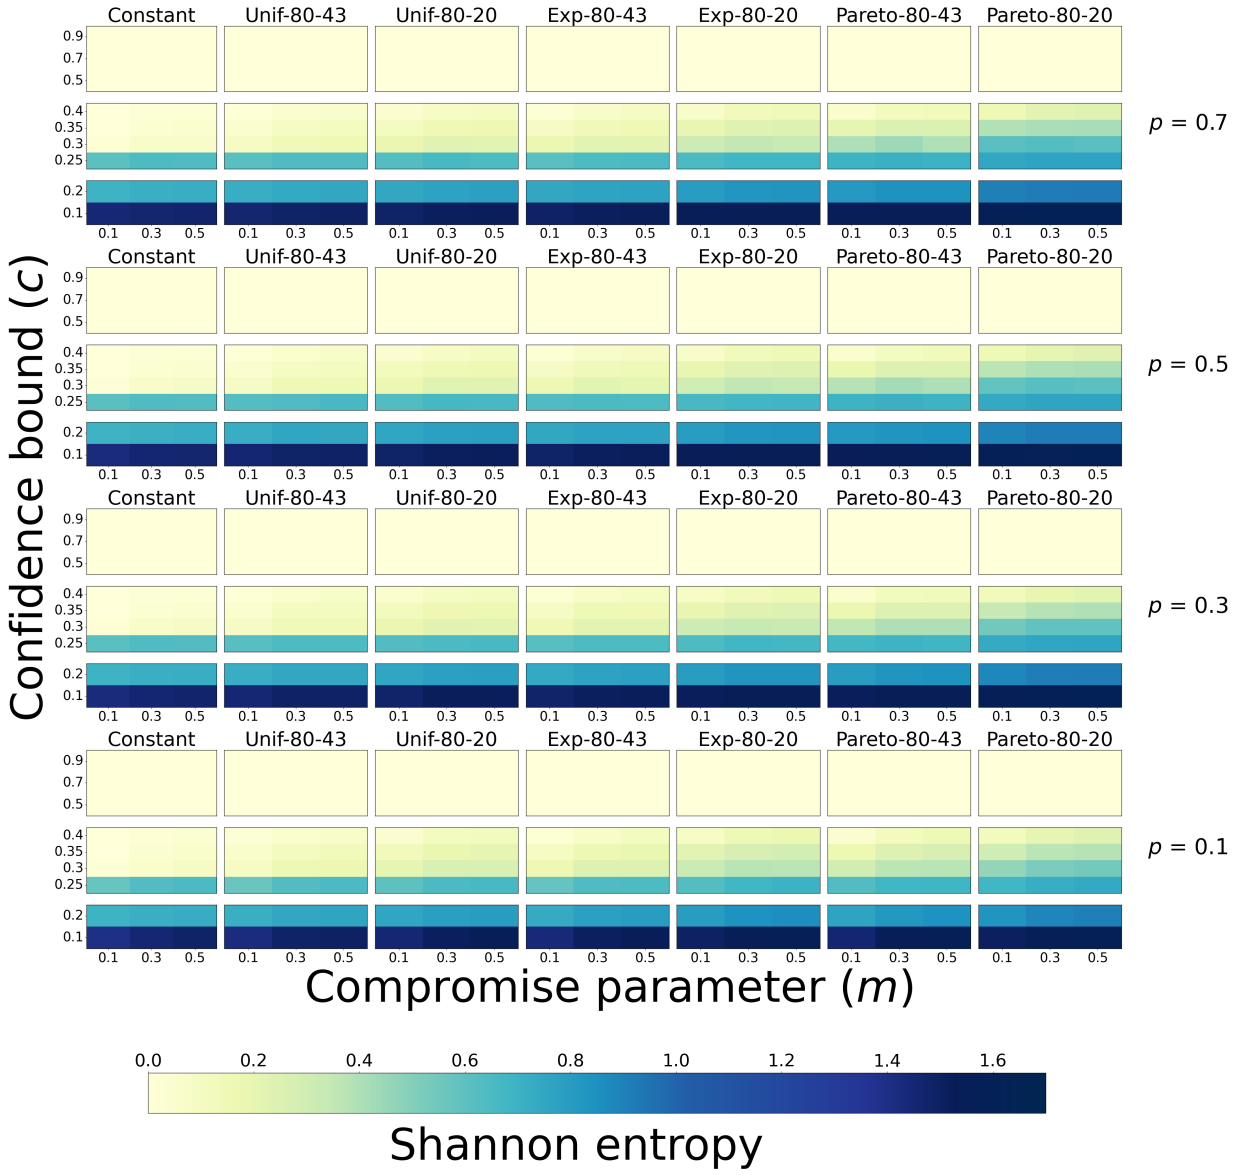

Figure 4.6: Shannon entropies of the steady-state opinion-cluster profiles in simulations of our node-weighted BCM on $G(500, p)$ ER random graphs for various node-weight distributions and several values of p. (This figure originally appeared in [LP23].)

increase the steady-state number of minor opinion clusters. For $p \in \{0.5, 0.7\}$, the steady-state numbers of minor clusters are comparable to the numbers that we obtained for a 500-node complete graph. When $p \in \{0.5, 0.7\}$, for each node-weight distribution and each value of c and m, when we take the mean of our 500 simulations, we obtain at most 3 steady-state minor clusters. For these simulations, we observe the most minor clusters when $c \in \{0.1, 0.2\}$. For $p = 0.1$, the mean number of steady-state minor clusters is at most 9; this occurs when $c \in \{0.35, 0.4\}$. It seems sensible that smaller values of p yield more minor opinion clusters. For small p, there are more small-degree nodes than for larger values of p. It is easier for small-degree nodes than for large-degree nodes to be in a minor opinion cluster, as small-degree nodes need to become unreceptive to few neighbors to end up in a minor cluster at steady state. That is, if i is a small-degree node, few neighbors j need to satisfy the inequality $|x_i - x_j| < c$.

4.4.3 SBM graphs

We now discuss our simulations of our BCM on 500-node SBM random graphs that we generate using the parameters in Table 4.1. For both the two-community and core–periphery SBM graphs, we observe the trends in Table 4.3.

For the two-community SBM graphs, the steady-state Shannon entropies and numbers of major opinion clusters are comparable to those in our simulations on a 500-node complete graph. When there is opinion fragmentation, for a fixed node-weight distribution and fixed values of c and m, the steady-state values of mean local receptiveness tend to be similar to the values for $G(500, 0.1)$ graphs and larger than the values for a complete graph. The steady-state numbers of minor opinion clusters are similar to those for the $G(500, 0.1)$ random graphs.

For the two-community SBM graphs, for each node-weight distribution and each value of c and m, when we take the mean of our 500 simulations, we obtain at most 10 steady-state minor clusters. We observe the most steady-state minor clusters when $c \in \{0.35, 0.4\}$. Recall

that we select the edge probabilities of the two-community SBM so that each of the two communities has an expected mean degree that matches that of $G(500, 0.1)$ graphs. Therefore, it is reasonable that we obtain similar results for the two-community SBM and the $G(500, 0.1)$ random graphs. In our numerical simulations, we assign the node weights randomly without considering the positions (which, in this case, is the community assignments) of the nodes of a network. When we assign weights to nodes uniformly at random, it seems that graph sparsity may be more important than community structure for determining if our BCM reaches a consensus or a fragmented state.

For a fixed node-weight distribution and fixed values of c and m, the core–periphery SBM graphs tend to have fewer major clusters than a complete graph. Additionally, both the steady-state Shannon entropy and the mean local receptiveness tend to be larger for the core–periphery SBM graphs than for a complete graph. Larger entropy and smaller local receptiveness are both indications of more opinion fragmentation. If we consider only the number of major opinion clusters, it seems that the core–periphery SBM graphs yield less opinion fragmentation than a complete graph. However, when we examine the entire opinion-cluster profile of a network and account for the cluster sizes and the minor clusters, the Shannon entropy reveals that there is more opinion fragmentation for our core–periphery SBM graphs than for a complete graph. The steady-state mean local receptiveness indicates that the nodes of a core–periphery SBM graph tend to be receptive to a larger fraction of their neighbors than the nodes of a complete graph.

We believe that Shannon entropy gives a more useful quantification of opinion fragmentation than mean local receptiveness. For networks with a large range of degrees, small-degree nodes can inflate the mean local receptiveness. (Analogously, a network's mean local clustering coefficient places more importance than its global clustering coefficient on small-degree node [New18].) In the context of our node-weighted BCM, consider a node with degree 2 and a node with degree 100, and suppose that both of them have a local receptiveness of 0.5. The larger-degree node's local receptiveness of 0.5 gives a better indication that there may

be opinion fragmentation than the smaller-degree node's local receptiveness of 0.5. However, we treat both nodes equally when we calculate the mean local receptiveness. We believe that local receptiveness is a useful quantity to calculate for individual nodes to determine how they perceive the opinions of their neighbors. However, mean local receptiveness appears to be less useful than Shannon entropy for quantifying opinion fragmentation in a network.

For a fixed node-weight distribution and fixed values of c and m, the steady-state numbers of major opinion clusters that we obtain in the core–periphery SBM graphs are comparable to the numbers for a complete graph. The steady-state numbers of minor opinion clusters tend to be larger for core–periphery SBM graphs than for two-community SBM graphs (which have more minor clusters than a complete graph). For each node-weight distribution and each value of c and m, when we take the mean of our 500 simulations, we observe at most 12 steady-state minor clusters; this occurs when $c = 0.1$. One possibility is that the core–periphery structure makes it easier to disconnect peripheral nodes of an effective graph, causing these nodes to form minor clusters. (Recall that we defined effective graphs in Section 3.3.2.) For core-periphery SBM graphs, it seems interesting to investigate the effect of assigning node weights in a way that depends on network structure. For example, if we assign all of the large weights to core nodes, will these nodes pull many peripheral nodes into their opinion clusters? If we place a large-weight node in the periphery, will it be able to pull core nodes into its opinion cluster?

4.4.4 Caltech network

We now discuss the Caltech Facebook network, which is an empirical data set in which the nodes are individuals with Caltech affiliations and the edges represent "friendships" on Facebook on one day in fall 2005 [RKM11, TMP12]. We consider the network's largest connected component, which has 762 nodes and 16,651 edges. The Caltech network has all but one of the trends that we reported in Table 4.3; the only exception is the trend in the number of minor opinion clusters. When there is opinion fragmentation, the Caltech network

has more steady-state minor clusters and larger steady-state Shannon entropies than in our synthetic networks.

In Figure 4.7, we show the steady-state numbers of minor opinion clusters in simulations of our BCM on the Caltech network. We obtain the most minor clusters when $c = 0.1$, which is the smallest value of c that we examine. Given a node-weight distribution and values of c and m, when we take the mean of our 100 simulations on the Caltech network, we often observe a large number of minor clusters (including as many as 78 of them, which is much larger than the single-digit numbers that we usually observed for our synthetic networks). Additionally, unlike in our synthetic networks, for all node-weight distributions (not just the constant weight distribution), the Caltech network tends to have more minor clusters when $m \in \{0.3, 0.5\}$ than when $m = 0.1$.

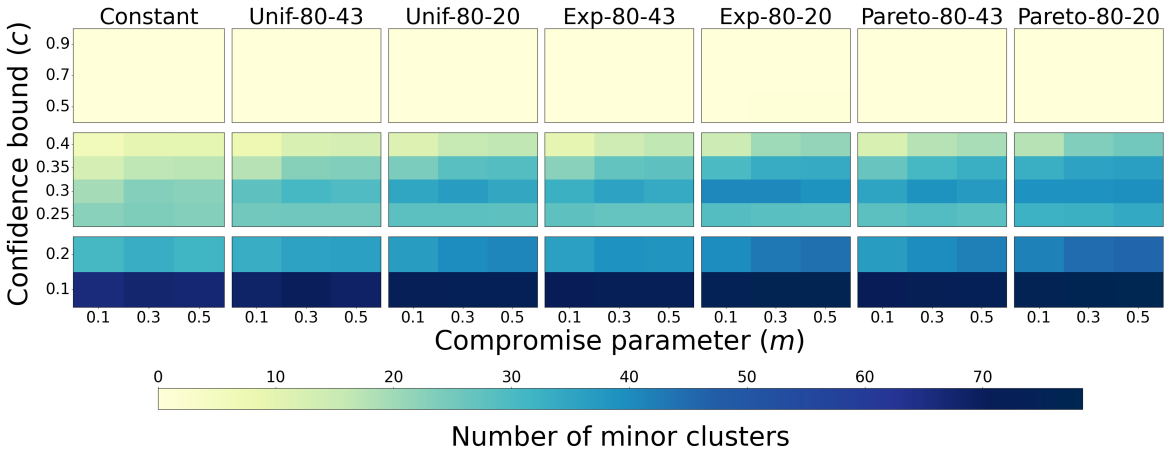

Figure 4.7: Steady-state numbers of minor opinion clusters in simulations of our node-weighted BCM on the Caltech Facebook network with various distributions of node weights. We consider an opinion cluster to be minor cluster if it has at most 2% of the nodes of a network. In this case, a minor cluster has at most 15 nodes. (This figure originally appeared in [LP23].)

In Figure 4.8, we show the steady-state numbers of major opinion clusters in simulations of our BCM on the Caltech network. For a fixed node-weight distribution and fixed values of c and m, the Caltech network tends to have fewer major opinion clusters than our synthetic

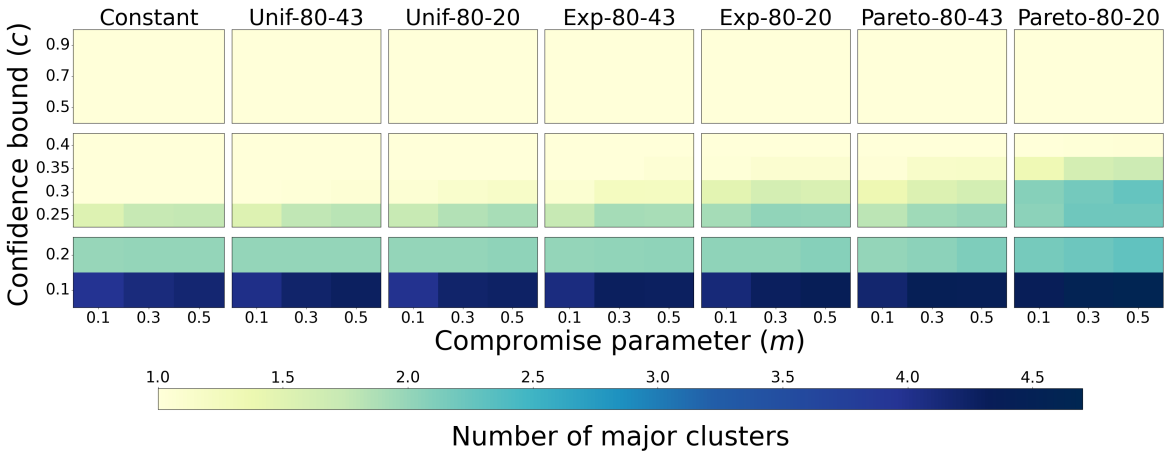

Figure 4.8: Steady-state numbers of major opinion clusters in simulations of our node-weighted BCM on the Caltech Facebook network with various distributions of node weights. We consider an opinion cluster to be major cluster if it has more than 2% of the nodes of a network. In this case, a major cluster must have at least 16 nodes. (This figure originally appeared in [LP23].)

networks. This observation is in line with our observation that the Caltech network can have many more minor clusters than our synthetic networks.

In Figure 4.9, we show the steady-state Shannon entropies for the Caltech network. For a fixed node-weight distribution and fixed values of c and m, when there is opinion fragmentation, we observe a larger entropy for the Caltech network than for our synthetic networks. This aligns with our observation that the Caltech network has many more minor opinion clusters than our synthetic networks. We show a plot of the steady-state values of mean local receptiveness for the Caltech network in our code repository. The values of the mean local receptiveness tend to be larger for the Caltech network than for a 500-node complete graph. We suspect that this arises from the presence of many small-degree nodes in the Caltech network. In Section 4.4.3, we discussed the impact of small-degree nodes on the mean local receptiveness.

The histogram of the node degrees of the Caltech network (see Figure 4.10) differs dramatically from those of our synthetic networks. Unlike in our synthetic networks, the most

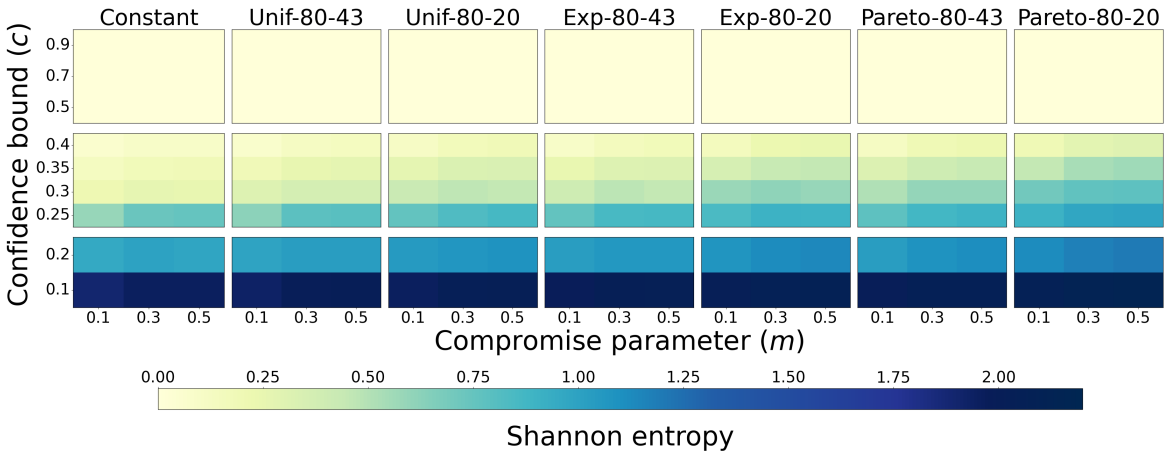

Figure 4.9: Shannon entropies of the steady-state opinion-cluster profiles in simulations of our node-weighted BCM on the Caltech Facebook network with various node-weight distributions. (This figure originally appeared in [LP23].)

common degrees in the Caltech network are among the smallest degrees. In Figure 4.10, the tallest bar in the histogram is for nodes with degrees 1–9. These abundant small-degree nodes are likely to disconnect from the largest connected components of the effective graph and form minor opinion clusters. Because we select the initial opinions uniformly at random from $[0, 1]$, when $c = 0.1$, it is possible that small-degree nodes are initially isolated nodes of the effective graph because of their initial opinions. The abundance of small-degree nodes in the Caltech network helps explain its larger steady-state numbers of minor opinion clusters and the correspondingly larger entropies than for our synthetic networks. Despite the fact that the Caltech network is structurally very different from our synthetic networks, it follows all of the trends in Table 4.3 aside from the one for the number of minor opinion clusters. Therefore, it seems that the trends that we observe in our node-weighted BCM when we assign weights to nodes uniformly at random (and hence in a way that is independent of network structure) are fairly robust to the underlying network structure.

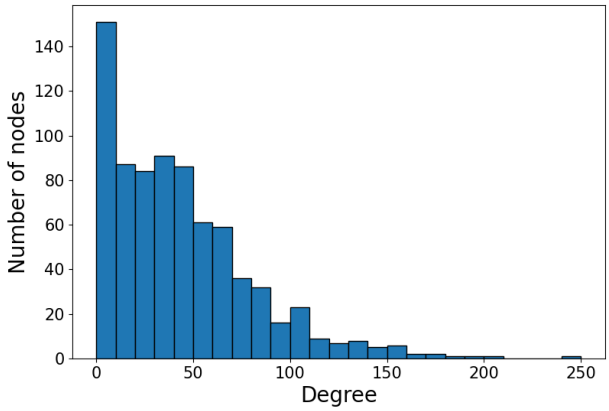

Figure 4.10: Histogram of the node degrees of the Caltech Facebook network. The bins have width 10 and originate at the left end point (i.e., the bins indicate degrees of 0–9, 10–19, and so on). (This figure originally appeared in [LP23].)

4.4.5 Finite-size effects

We now investigate finite-size effects by simulating our BCM on complete graphs of different sizes. Previously, to ensure reasonable computation times, we examined synthetic networks with 500 nodes. However, it is useful to get a sense of whether or not the trends in Table 4.3 hold for networks of different sizes. We thus simulate our BCM on complete graphs of sizes $N \in \{10, 20, 30, 45, 65, 100, 150, 200, 300, \ldots, 1000\}$. We examine $m \in \{0.3, 0.5\}$, and $c \in \{0.1, 0.3, 0.5\}$, which give regimes of opinion fragmentation, a transition between fragmentation and consensus for the constant weight distribution, and opinion consensus. We consider the constant weight distribution and the 80-43 distributions (i.e., the uniform, exponential, and Pareto distributions with a mean node weight of 2.8836). We do not examine any larger-mean distributions because they require longer computation times.

In this chapter, we show our results for convergence times and steady-state Shannon entropies. To visualize our results, we plot graph sizes on a logarithmic scale. We include plots of our simulation results at steady state for the numbers of major opinion clusters, the numbers of minor opinion clusters, and the values of mean local receptiveness in our code

repository.

In Figure 4.11, we show the convergence times of our simulations of our BCM on complete graphs of various sizes. For all node-weight distributions, the convergence times become longer as we increase the graph size. For each graph size, the convergence times for the heterogeneous weight distributions are similar to each other and are longer than those for the constant weight distribution.

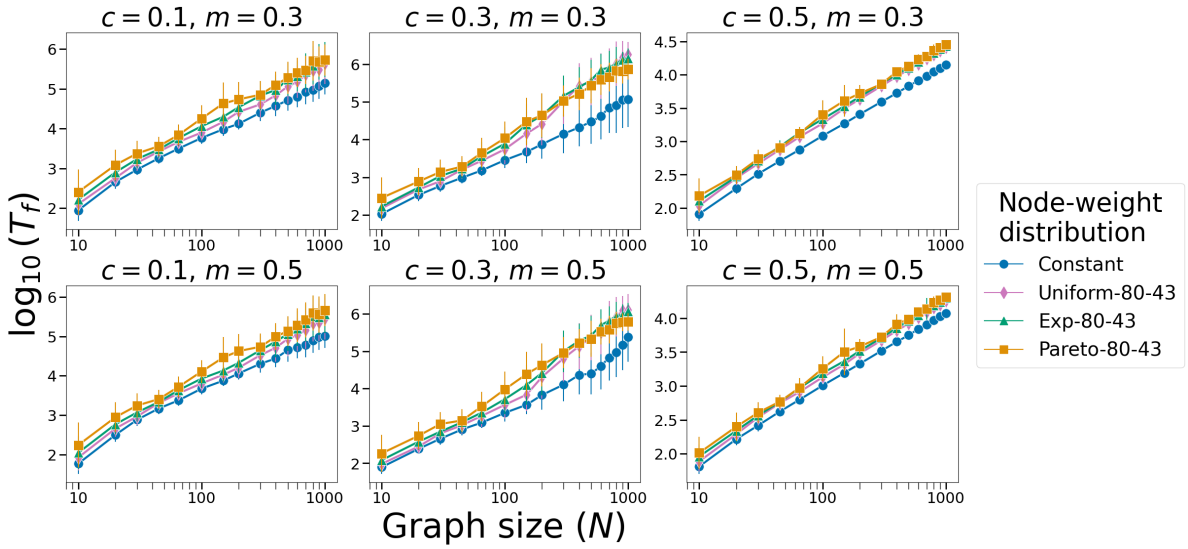

Figure 4.11: Convergence times (in terms of the number of time steps) in simulations of our node-weighted BCM on complete graphs of various sizes. We show results for various choices of c and m; the marker shape and color indicate the node-weight distribution. For this figure and subsequent figures of this type, the points are means of 100 simulations and the error bars indicate one standard deviation from the mean. The horizontal axis gives the graph size on a logarithmic scale. For clarity, the vertical axes of the plots have different scales. (This figure originally appeared in [LP23].)

In Figure 4.12, we show the steady-state Shannon entropies from our simulations of our BCM on complete graphs of various sizes. For a fixed value of c, we observe similar results when $m = 0.3$ and $m = 0.5$. When $c = 0.5$, for each node-weight distribution, our simulations always reach a consensus (i.e., there is exactly one steady-state major opinion cluster) for $N \geq 200$. Correspondingly, the steady-state entropies are close to 0. (They are not exactly

0 because the calculation of Shannon entropies includes information from minor clusters.) As we increase the network size, the error bars (which indicate one standard deviation from the mean) become progressively smaller. When $c \in \{0.1, 0.3\}$, for sufficiently large graph sizes (specifically, when $N \geq 100$), we observe that the entropy increases as we increase the heaviness of the tail of a distribution. For $c = 0.3$, the mean steady-state entropies appear to no longer change meaningfully with N when $N \geq 400$. For $c = 0.1$, this is the case when $N \geq 100$.

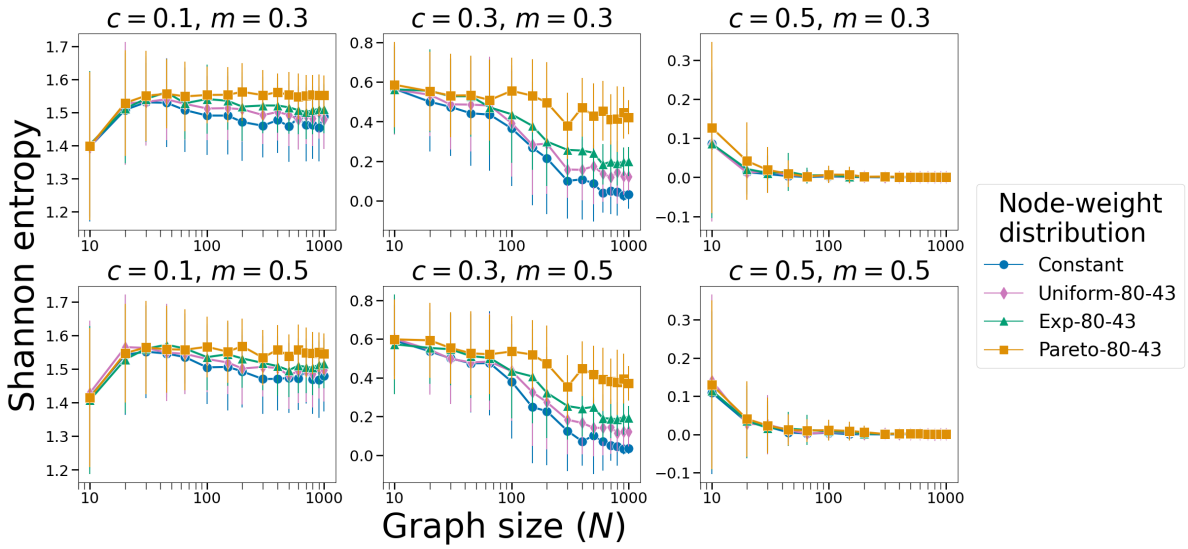

Figure 4.12: Shannon entropies of the steady-state opinion-cluster profiles in simulations of our node-weighted BCM on complete graphs of various sizes. We show results for various choices of c and m; the marker shape and color indicate the node-weight distribution. (This figure originally appeared in [LP23].)

When there is opinion fragmentation, the heterogeneous node-weight distributions tend to yield larger steady-state Shannon entropies (and hence more opinion fragmentation, if one is measuring it using entropy) than the constant weight distribution for each graph size. Additionally, for the 80-43 distributions and sufficiently large graph sizes, we obtain larger entropies as we increase the heaviness of the distribution tail. We have not explored the effect of graph size on the observed trends (see Table 4.3) when we increase the distribution mean for a fixed family of distributions. In our code repository, we include a plot of the steady-state

mean local receptiveness for complete graphs of various sizes. In that plot, we also observe that there tends to be more opinion fragmentation (in the sense of a smaller mean local receptiveness) for heterogeneous node-weight distributions with increasingly heavy tails.

We also examine the steady-state numbers of major and minor opinion clusters in simulations of our BCM on complete graphs of various sizes; we include plots of them in our code repository. For a fixed value of c, we observe similar results when $m = 0.3$ and $m = 0.5$. When $N \leq 49$, there are no minor opinion clusters, by definition, because minor clusters can include at most 2% of the nodes of a network (and even a single node constitutes more than 2% of all nodes for such small networks). When $N \geq 65$ and $c \in \{0.1, 0.3\}$, for each distribution, the number of minor clusters tends to increase as we increase N. We do not observe a clear trend in which node-weight distributions yield more minor clusters. For all N and $c = 0.5$, the mean number of minor clusters is close to 0. We now consider major opinion clusters. When $c = 0.5$ and $N \geq 200$, all simulations yield one major opinion cluster (i.e., they all reach consensus). When $c = 0.3$, for all graph sizes, there are more major opinion clusters as we increase the heaviness of the tail of a distribution. Additionally, when $c = 0.3$, for the Pareto-80-43 distribution, the number of major clusters tends to increase as we increase the graph size. For the other node-weight distributions, the number of major clusters tends to decrease as we increase the graph size. When $c = 0.1$ and $N \geq 200$, there again tends to be more major clusters as we increase the heaviness of the tail of a distribution, although the trend is not as clear as it was for $c = 0.3$.

Based on our exploration of finite-size effects, we are confident that complete graphs with $N \geq 500$ nodes have the trends in Table 4.3. For graphs with $N \geq 500$ nodes, the mean steady-state Shannon entropies for each node-weight distribution appear to no longer change meaningfully with respect to N. For each graph size, the heterogeneous 80-43 distributions have longer convergence times than the constant weight distribution. For a fixed graph size and fixed values of c and m, we observe more opinion fragmentation as we increase the heaviness of the tail of a distribution. Because of computation time, we have not examined

finite-size effects for distributions other than the 80-43 distributions. However, because the mean Shannon entropies no longer change meaningfully with respect to N for graphs with $N \geq 500$ nodes, we hypothesize that the trends in opinion fragmentation and convergence time in Table 4.3 continue to hold for our synthetic networks when there are more than 500 nodes.

4.5 Conclusions and discussion

We studied a BCM with heterogeneous node-selection probabilities, which we modeled using node weights. One can interpret these node weights as encoding phenomena such as heterogeneous agent sociabilities or activity levels. We studied our node-weighted BCM with fixed node weights that we assign in a way that disregards network structure and node opinions. We demonstrated that our node-weighted BCM has longer convergence times and more opinion fragmentation than a baseline DW BCM in which we uniformly randomly select nodes for interaction. It is straightforward to adapt our BCM to assign node weights in a way that depends on network structure and/or node opinions. See Section 4.5.2 and Section 4.5.3 for discussions.

4.5.1 Summary of our main results

We simulated our node-weighted BCM with a variety of node-weight distributions (see Table 4.2) on several random and deterministic networks (see Table 4.1). For each of these distributions and networks, we systematically investigated the convergence time and opinion fragmentation for different values of the confidence bound c and the compromise parameter m. To determine if the nodes of a network reach consensus or if there is opinion fragmentation, we calculated the steady-state number of major clusters in our simulations. To quantify the amount of opinion fragmentation, we calculated the steady-state Shannon entropy and the mean local receptiveness. For a given network, we found that entropy and mean local

receptiveness follow the same trends in which node-weight distributions have more opinion fragmentation (see Table 4.3). Additionally, based on our results, we believe that Shannon entropy is more useful than mean local receptiveness for quantifying opinion fragmentation in a network. However, calculating local receptiveness is insightful for explorations of the opinion dynamics of individual nodes.

In our simulations of our node-weighted BCM, we observed a variety of typical trends (see Table 4.3). In particular, we found that heterogeneous node-weight distributions tend to yield longer convergence times and more opinion fragmentation than the baseline DW model (which we obtain by using a constant weight distribution) in simulations of our BCM. Opinion fragmentation also tends to increase if either (1) for a fixed distribution mean, we increase the heaviness of the tail of a distribution or (2) for a fixed distribution family, we increase the distribution mean. Given a set of heterogeneous node weights, we hypothesize that large-weight nodes are selected early in a simulation with large probabilities and quickly settle into their associated steady-state major opinion clusters. Small-weight nodes that are not selected early in a simulation are left behind to form small opinion clusters, resulting in more opinion fragmentation than in the baseline DW model.

4.5.2 Relating node weights to network structure

We examined deterministic and random graphs with various structures, and we observed the trends in Table 4.3. For each of our BCM simulations, we selected node weights from a specified distribution and then assigned these weights to nodes uniformly at random. Therefore, our investigation conveys what trends to expect with fixed, heterogeneous node weights that are assigned to nodes without regard for network structure. However, our model provides a flexible framework to study the effects of node weights when they are correlated with network structure. For example, one can assign weights to nodes in a way that depends on some centrality measure (such as degree). In our BCM, we expect nodes with large degree or large weight to have more interactions than nodes with small degree

or small weight. Nodes with larger degrees have more neighbors that can select them for an interaction, and nodes with larger weights have larger probabilities of being selected for an interaction. One possible area of future work is to investigate the combined effects of node weights and node degrees on the frequencies of interactions and the distribution of steady-state opinions in our BCM. Mean-field approaches, such as the one in [FBQ21], may offer insights into these effects.

For a given set of node weights, larger-weight nodes have larger probabilities of interacting with other nodes, so their positions in a network likely influence the dynamics of BCMs and other models of opinion dynamics. One can also investigate the effect of homophily when choosing how to assign node weights. For example, in social-media platforms, very active accounts may engage with each other more frequently by sharing or commenting on each others' posts. We can incorporate such features into our BCM through a positive node-weight assortativity, such that large-weight nodes are more likely to be adjacent to each other than to other nodes.

As in the standard DW model, we assign the initial agent opinions uniformly at random in our BCM. However, in a real social network with community structure, this choice may not be realistic. One can consider a social network with communities with different mean opinion values and examine the effect of placing large-weight nodes in different communities. For example, how does placing all large-weight nodes in the same community affect opinion dynamics and steady-state opinion-cluster profiles? How does the presence of a small community of "outspoken" (i.e., large-weight) nodes influence the final opinions of nodes in other communities of a network? Will the small community quickly induce an echo chamber [FGR16], will it pull other nodes into its final opinion cluster, or will something else occur?

4.5.3 Relating node weights to node opinions

In the present chapter, we considered fixed node weights that are independent of node opinions. One can readily adapt our BCM to incorporate time-dependent node weights,

such as ones that depend on node opinions. One can allow the probability of selecting a node for interaction to depend on how extreme its opinion is [AC15] or on the similarity of its opinion to that of another node [SPG19].

Sîrbu et al. [SPG19] studied a modified DW model with heterogeneous node-selection probabilities that model algorithmic bias on social media. In their model, one first selects an agent uniformly at random. One then calculates the magnitude of the opinion difference between that agent and each of its neighbors and then selects a neighbor with a probability that is proportional to this difference. In the context of our BCM, one can represent their agent-selection mechanism using time-dependent node weights. To do this, at each time t, one assigns the same weight to all nodes when selecting a first node i. When selecting a node to interact with i, one then assigns weights to the neighbors j of node i that depend on the opinion difference $|x_i(t) - x_j(t)|$. One assigns a weight of 0 to nodes that are not adjacent to i. Simulations by Sîrbu et al. on complete graphs suggest that more algorithmic bias results in longer convergence times and more opinion clusters [SPG19]. Pansanella et al. [PRM22] observed similar trends in a study of the algorithmic-bias model of Sîrbu et al. for various random-graph models.

In our simulations of our BCM with heterogeneous node-selection probabilities, we observed similar trends of longer convergence times and more opinion clusters (and opinion fragmentation) than in our baseline DW model. When studying a BCM with heterogeneous node-selection probabilities, we see from our results that it is important to consider the baseline influence of assigning node weights uniformly at random before attributing trends such as longer convergence times and more opinion fragmentation to specific mechanisms such as algorithmic bias. Different mechanisms can yield very similar empirical observations.

4.5.4 Edge-based heterogeneous activities

In the standard DW model, at each time, one selects an edge of a network uniformly at random and the two agents that are attached to that edge interact with each other [WDA03].

Most past work on the DW model and its extensions has focused on this edge-based selection mechanism [NVT20]. In our BCM, to incorporate node weights (e.g., to encode heterogeneous sociabilities or activity levels of individuals), we instead used a node-based selection mechanism. For voter models of opinion dynamics, it is known that the choice between edge-based and node-based agent selection can substantially affect a model's qualitative behavior [KP20]. We are not aware of a comparison of edge-based and node-based agent selection in asynchronous BCMs (and, in particular, in DW models), and it seems interesting to investigate this issue.

Our BCM has node weights to encode heterogeneous activity levels of individuals. One can also examine heterogeneous pairwise activity levels to account for the fact that individuals do not interact with each of their social contacts with the same probability. To encode such heterogeneity, one can construct a variant of our BCM that incorporates edge weights. At each time step, one can select a pair of agents to interact with a probability that is proportional to the weight of the edge between them. Additionally, one can relate edge-selection mechanisms to node-selection mechanisms. We have not examined edge-based heterogeneous activity levels in a BCM, and we expect that it will be interesting to investigate them.

4.5.5 Importance of node weights

The key feature of our BCM is our incorporation of node weights into opinion dynamics. Node weights have been used in activity-driven models of temporal networks [PGP12], and activity-driven frameworks have been used to model which agents can interact with each other in models of opinion dynamics [LHM17, ZXL18]. In our BCM, the node weights determine the probabilities to select agents for interaction in a time-independent network. Alizadeh and Cioffi-Revilla [AC15], Sîrbu et al. [SPG19], and Pansanella et al. [PRM22] examined specific scenarios of heterogeneous node-selection probabilities in modified DW models. Our node-weighted BCM gives a general framework to incorporate node weights into asynchronous BCMs. Using our framework, one can consider node weights that are

fixed and assigned uniformly at random to nodes (i.e., as we investigated in this chapter), fixed and assigned according to some other probability distribution (see the discussion in Section 4.5.2), or assigned in a time-dependent way (see the discussion in Section 4.5.3). In network science, node weights have been studied far less than edge weights, and even the term "weighted network" usually refers specifically to edge-weighted networks by default. For example, it is very common to study centralities in edge-weighted networks [OAS10], but studies of centralities in node-weighted networks (e.g., see Refs. [HDZ12, SSI20]) are much less common. Heitzig et al. [HDZ12] generalized common network statistics to node-weighted networks and used node weights to represent the "sizes" of the nodes of a network. They used their framework to study brain networks with node weights that encode the areas of regions of interest, international trade networks with node weights that encode the gross domestic products (GDPs) of countries, and climate networks with node weights that encode areas in a regular grid on the Earth's surface. Singh et al. [SSI20] developed centrality measures that incorporate both edge weights and node weights, and they used them to study service-coverage problems and the spread of contagions. These studies demonstrate the usefulness of node weights for incorporating salient information in network analysis in a variety of applications.

In our node-weighted BCM, we are interested in determining which nodes of a network are (in some sense) more influential than others and thereby have a larger impact on steady-state opinion-cluster profiles. Brooks and Porter [BP20] quantified the influence of media nodes in a BCM by examining how their ideologies influence other nodes of a network. An interesting area of future work is to develop ways to quantify the influence of specific nodes in models of opinion dynamics with node weights. For example, can one determine which nodes to seed with extreme opinions to best spread such opinions? Are there nodes that make it particularly easy for communities to reach consensus and remain connected in a steady-state effective graph $G_{\text{eff}}(T_f)$? One can tailor the node weights in our BCM to examine a variety of sociological scenarios in which nodes have heterogeneous activity

levels or interaction frequencies. More generally, our model illustrates the importance of incorporating node weights into network analysis, and we encourage researchers to spend more time studying the effects of node weights on network structure and dynamics.

CHAPTER 5

Bounded-Confidence Models with Adaptive Confidence Bounds

In this chapter, we investigate BCMs in which each pair of agents has a distinct confidence bound that changes when the pair interacts. We analytically and numerically explore the limiting behaviors of these BCMs. This chapter is adapted from an original paper [LLP23] that I co-authored with Jiajie Luo and my advisor Mason A. Porter.[1] We include our code and all figures (including those that are not shown in this chapter) in our repository at `https://gitlab.com/graceli1/Adaptive-Confidence-BCM`.

5.1 Introduction and motivation

In this chapter, we formulate and study adaptive-confidence BCMs that generalize the HK and DW models by incorporating distinct, time-dependent confidence bounds for each pair of adjacent nodes (i.e., each dyad). The confidence bounds in our adaptive-confidence BCMs change after nodes interact with each other. These changes highlight the idea that the quality of an interaction between individuals can affect how much they trust each other [GAP13, LMM17, CS20]. For example, in online marketplaces, trust between users depends on their past experiences with each other and on the reported experiences of other users in reputation systems [RKZ00, RKK07, SNP13]. The word "trust" can have different

[1] My contributions to the paper [LLP23] were formulating the update rules for our BCMs jointly with Jiajie Luo, proving Theorem 5.3.4 jointly with Jiajie Luo, proving Theorem 5.3.7, and running and interpreting the numerical simulations. All co-authors wrote the paper together.

meanings in different disciplines; one interpretation is that trust represents an expectation about future behavior [SNP13]. Rather than considering trust, our BCMs use a notion of "receptiveness", which encodes the willingness of an individual to consider the future opinions of another individual. When two nodes interact with each other, their mutual receptiveness changes. See [BCF07, XLC17, NGW23] for other opinion models with interaction-influenced receptiveness.

In our adaptive-confidence BCMs, when two nodes successfully compromise their opinions in an interaction (i.e., they have a "positive interaction"), they become more receptive to each other. Likewise, when two nodes interact but do not change their opinions (i.e., they have a "negative interaction"), they become less receptive to each other. When nodes i and j interact and influence each others' opinions (i.e., their current opinion difference is smaller than their current confidence bound), we increase their confidence bound c_{ij}. When nodes i and j interact and do not influence each others' opinions (i.e., their current opinion difference is at least as large as their current confidence bound), we decrease their confidence bound c_{ij}. In our adaptive-confidence BCMs, each dyad has a distinct confidence bound and interactions are symmetric (i.e., either both nodes influence each other or neither node influences the other). One can interpret the increase of a dyadic confidence bound in our BCMs as nodes becoming more receptive to nodes with whom they compromise, and one can interpret the decrease of a dyadic confidence bound as nodes becoming less receptive to nodes with whom they do not compromise. When nodes in our BCMs have a negative interaction, they adapt their dyadic confidence bounds, but their opinions stay the same. Other researchers have considered BCMs with "repulsion", in which the opinions of interacting nodes with sufficiently different opinions move farther apart from each other [HDJ08, AC15, KF23].

5.1.1 Related work

Many researchers have generalized the HK and DW models by incorporating heterogeneity into the confidence bounds. Lorenz [Lor10] extended these BCMs so that each node has

its own confidence bound, which can result in asymmetric influence and opinion updates. Using numerical simulations, Lorenz demonstrated that these BCMs are more likely than the baseline BCMs to reach a consensus state when there are both open-minded and close-minded nodes (which have large and small confidence bounds, respectively). By analyzing the heterogeneous-confidence DW model of [Lor10] on a complete graph, Chen et al. [CSM20] proved almost-sure convergence of opinions for certain parameter values and derived sufficient conditions for the nodes of a network to eventually reach a consensus. In a related work, Chen et al. [CSD20] examined a heterogeneous HK model with "environmental noise" (e.g., from media sources) and showed that heterogeneous confidence bounds in this setting can yield larger differences in node opinions in the infinite-time limit. Su et al. [SGW17] examined the heterogeneous-confidence HK model of [Lor10] and proved that at least some nodes of a network converge to a steady-state opinion in finite time.

Researchers have also incorporated edge-based heterogeneities in the confidence bounds of BCMs. Etesami [Ete19] examined an HK model on networks with time-independent edge-heterogeneous confidence bounds and proved that their model is Lyapunov stable. Shang [Sha14] studied a DW model in which each edge has a confidence bound that takes a value from an independent and identically distributed Poisson process. They derived sufficient conditions for consensus to occur almost surely for a one-dimensional lattice graph.

Other generalizations of BCMs and related opinion models generalize the model parameters by making them time-dependent or adaptive. Weisbuch et al. [WDA02] studied a generalized DW model in which each node has a heterogeneous, time-dependent confidence bound that is proportional to the standard deviation of the opinions that that node observed in all prior interactions. They also considered a variant of their model that places more weight on the observed opinions from recent interactions. Deffuant et al. [DAW02] examined a DW model with "relative agreement". In their model, each node has an uncertainty parameter that determines (1) whether it and the node with which it interacts influence each other and (2) the amount by which they influence each other. A node changes both its opinion and

its uncertainty when it is influenced by another node. Bagnoli et al. [BCF07] considered a BCM on complete graphs in which each pair of adjacent nodes (i.e., each dyad) has an associated time-dependent affinity value (which determines whether or not they can influence each other) that depends on the magnitude of their opinion difference. Chacoma and Zanette [CZ15] examined opinion and confidence changes in a questionnaire-based experiment, and they then proposed an agent-based opinion model based on the results of their experiment. Their model is not a BCM, but it does incorporate a notion of time-dependent confidence between nodes. Bernardo, Vasca, and Iervolino [VBI21, BVI22] developed variants of the HK model in which nodes have individual, time-dependent confidence bounds[2] that depend on the opinions of neighboring nodes. In their models, nodes adapt their confidence bounds through a heterophilic mechanism (i.e., they seek neighboring nodes whose opinions differ from theirs). By contrast, in our models, nodes do not actively seek neighbors with different opinions. Instead, their mutual receptiveness increases when their opinions are sufficiently close to each other.

In this chapter, we incorporate adaptivity into the confidence bounds of BCMs, but one can instead incorporate adaptivity in the network structures of BCMs [KB08a, KB08b, DSC17, KFP23].[3] Kozma and Barrat [KB08a, KB08b] modified the DW model to allow rewiring of "discordant" edges, which occur between nodes whose opinions differ from each other by more than the confidence bound. In their model, rewired edges connect to new nodes uniformly at random. Kan et al. [KFP23] generalized this model by including both a confidence bound and an opinion-tolerance threshold, with discordant edges occurring between nodes whose opinions differ by more than that threshold. They incorporated opinion homophily into the rewiring probabilities, so nodes are more likely to rewire to nodes with more similar opinions. They observed in numerical simulations that it is often harder to achieve consensus in their adaptive DW model than in an associated baseline DW model.

[2]The confidence bounds update with time in different ways in the models in [VBI21] and [BVI22].

[3]See the reviews [SBL23, BGK23] for discussions of various notions of adaptivity in dynamical systems.

5.2 BCMs with adaptive confidence bounds

We now describe our BCMs with adaptive confidence bounds. The idea for these models was developed primarily by Jiajie Luo. Let $G = (V, E)$, where V is the set of nodes and E is the set of edges, denote a time-independent, unweighted, and undirected graph without self-edges or multi-edges. The nodes in our BCMs represent agents that have opinions that lie in the closed interval $[0, 1]$. Each node i has a time-dependent opinion $x_i(t) \in [0, 1]$. Let $\mathbf{x}(t)$ denote the vector of the opinions of all nodes at time t (i.e., the entry $[\mathbf{x}(t)]_i = x_i(t)$).

5.2.1 Our HK model with adaptive confidence bounds

In this chapter, we refer to the standard HK model with update rule (3.2.1) as the *baseline HK model*. For convenience, we repeat the opinion update rule. For the baseline HK model, at each time t, we update the opinion of each node i by calculating

$$x_i(t+1) = |I(i, x(t))|^{-1} \sum_{j \in I(i,x(t))} x_j(t)\,, \tag{5.2.1}$$

where $I(i, x(t)) = \{i\} \cup \{j \mid |x_i(t) - x_j(t)| < c \text{ and } (i, j) \in E\} \subseteq \{1, 2, \ldots, N\}$.

Our HK model with adaptive confidence bounds is similar to the baseline HK model with update rule (5.2.1), but now each edge $(i, j) \in E$ has a dyadic confidence bound $c_{ij}(t) \in [0, 1]$ that is time-dependent and changes after each interaction between the nodes in that dyad. We refer to this model as our *adaptive-confidence HK model*. Instead of a fixed confidence bound, there is an initial confidence bound $c_0 \in (0, 1)$ and we initialize all of the confidence bounds[4] to $c_{ij}(0) = c_0$ for each edge $(i, j) \in E$. There is also a confidence-increase parameter $\gamma \in [0, 1]$ and a confidence-decrease parameter $\delta \in [0, 1]$, which control how much $c_{ij}(t)$ increases and decreases, respectively, after each interaction.

[4]When $c_0 = 0$, nodes are never receptive to their neighbors (i.e., $c_{ij}(t) = 0$ for all adjacent nodes i and j at all times t). When $c_0 = 1$, all nodes are always receptive to all of their neighbors (i.e., $c_{ij}(t) = 1$ for all adjacent nodes i and j at all times t). We do not examine these values of c_0.

At each time t, we update the opinion of each node i by calculating

$$x_i(t+1) = |I(i, x(t))|^{-1} \sum_{j \in I(i,x(t))} x_j(t)\,, \tag{5.2.2}$$

where[5] $I(i, x(t)) = \{i\} \cup \{j \mid |x_i(t) - x_j(t)| < c_{ij}(t) \text{ and } (i,j) \in E\} \subseteq \{1, 2, \ldots, N\}$. Adjacent nodes i and j are receptive to each other at time t if their opinion difference is less than their dyadic confidence bound c_{ij} (i.e., $|x_i(t) - x_j(t)| < c_{ij}(t)$). At each time, we also update each confidence bound c_{ij} by calculating

$$c_{ij}(t+1) = \begin{cases} c_{ij}(t) + \gamma(1 - c_{ij}(t))\,, & \text{if } |x_i(t) - x_j(t)| < c_{ij}(t) \\ \delta c_{ij}(t)\,, & \text{if } |x_i(t) - x_j(t)| \geq c_{ij}(t)\,. \end{cases} \tag{5.2.3}$$

That is, if the opinion difference between nodes i and j is smaller than their confidence bound at time t, their associated dyadic confidence bound c_{ij} increases; otherwise, their dyadic confidence bound decreases. Larger values of γ correspond to sharper increases in the receptiveness between nodes when nodes compromise their opinions. Smaller values of δ correspond to sharper drops in the receptiveness between nodes when nodes interact but do not compromise.

Because $c_0 \in (0, 1)$ and $\gamma, \delta \in [0, 1]$, the update rule (5.2.3) preserves $c_{ij}(t) \in (0, 1)$ for each edge (i, j) and all times t. If $(\gamma, \delta) = (0, 1)$, then $c_{ij}(t) = c_0$ for all t and all edges $(i, j) \in E$. That is, the confidence bounds are homogeneous and time-independent, so our adaptive-confidence HK model reduces to the baseline HK model.

[5]Although Equation (5.2.1) and Equation (5.2.2) look the same, they use different definitions of the quantity $I(i, x(t))$. Equation (5.2.1) has a homogeneous and time-independent confidence bound, whereas Equation (5.2.2) has heterogeneous and adaptive confidence bounds.

5.2.2 Our DW model with adaptive confidence bounds

In this chapter, we refer to the standard DW model with update rule (3.1.1) as the *baseline DW model*. For convenience, we repeat the opinion update rule. For the baseline DW model, at each time t, we uniformly randomly selects an edge $(i, j) \in E$. Nodes i and j then update their opinions through the following update rule:

$$\begin{aligned} x_i(t+1) &= \begin{cases} x_i(t) + m(x_j(t) - x_i(t)), & \text{if } |x_i(t) - x_j(t)| < c \\ x_i(t), & \text{otherwise}, \end{cases} \\ x_j(t+1) &= \begin{cases} x_j(t) + m(x_i(t) - x_j(t)), & \text{if } |x_i(t) - x_j(t)| < c \\ x_j(t), & \text{otherwise}. \end{cases} \end{aligned} \tag{5.2.4}$$

We refer to our DW model with adaptive confidence bounds as our *adaptive-confidence DW model*. As in the baseline DW model, there is a compromise parameter $m \in (0, 0.5]$. As in our adaptive-confidence HK model, we initialize the confidence bounds in our adaptive-confidence DW model to be $c_{ij}(0) = c_0$, where $c_0 \in (0, 1)$ is the initial confidence bound.[6] There again is a confidence-increase parameter $\gamma \in [0, 1]$ and a confidence-decrease parameter $\delta \in [0, 1]$, which control how much $c_{ij}(t)$ increases and decreases, respectively, after each interaction.

At each time t, we select an edge $(i, j) \in E$ uniformly at random. If nodes i and j are receptive to each other (i.e., if $|x_i(t) - x_j(t)| < c_{ij}(t)$), we update the opinions of nodes i and j using the DW update rule (5.2.4). Otherwise, the opinions x_i and x_j remain the same. We also update the dyadic confidence bound c_{ij} using Equation (5.2.3). That is, if the opinions of nodes i and j differ by less than their current dyadic confidence bound at

[6] As in our adaptive-confidence HK model, when $c_0 = 0$, nodes are never receptive to any of their neighbors (i.e., $c_{ij}(t) = 0$ for all adjacent nodes i and j at all times t). When $c_0 = 1$, nodes are always receptive to all of their neighbors (i.e., $c_{ij}(t) = 1$ for all adjacent nodes i and j at all times t). We do not examine these values of c_0.

time t, the confidence bound increases; otherwise, it decreases. The update rules preserves $c_{ij}(t) \in (0, 1)$ for each edge (i, j) and all times t. All other opinions and confidence bounds remain the same. Our adaptive-confidence DW model reduces to the baseline DW model when $(\gamma, \delta) = (0, 1)$.

5.3 Theoretical results

In this section, we discuss some theoretical guarantees of our BCMs. We discuss our theoretical results for our adaptive-confidence HK and DW models in Section 5.3.1 and Section 5.3.2, respectively.

As stated in [Lor05], Theorem 3.3.1 (see Section 3.3.1) guarantees that the opinion of each node convergences to a limit opinion in the baseline HK and DW models. Because the node opinions in our adaptive-confidence HK and DW models update in the same way as in the corresponding baseline BCMs (see Equation (5.2.1) and Equation (5.2.4), respectively), the node opinions in our models also converge to a limit opinion. The limit opinion x_i^* of node i is $\lim_{t\to\infty} x_i(t)$.

5.3.1 Adaptive-confidence HK model

In this section, we discuss our theoretical results for the confidence bounds and effective graphs in our adaptive-confidence HK model.

5.3.1.1 Confidence-bound analysis

In this section, we state Lemma 5.3.1 and Theorem 5.3.2 without proof. The proofs of Lemma 5.3.1 and Theorem 5.3.2 are primarily the work of Jiajie Luo. They are in [LLP23], but we do not include them in this dissertation.

Theorem 5.3.2 is our main result about the behavior of the confidence bounds (which

update according to Equation (5.2.3)) in our adaptive-confidence HK model. Lemma 5.3.1 is a result that we use to prove Theorem 5.3.2. In Lemma 5.3.1, we consider $\gamma \in [0,1]$ and $\delta \in [0,1]$, so it also applies to the baseline HK model. We use Lemma 5.3.1 in our proof of Theorem 5.3.4.

Lemma 5.3.1. *Consider our adaptive-confidence HK model (with update rules* (5.2.2) *and* (5.2.3)*) with parameters $\gamma \in [0,1]$ and $\delta \in [0,1]$. There is a time T such that no adjacent nodes i and j in different limit opinion clusters (i.e., $x_i^* \neq x_j^*$) are receptive to each other (i.e., $|x_i(t) - x_j(t)| < c_{ij}(t)$) at any time $t \geq T$.*

Theorem 5.3.2. *In our adaptive-confidence HK model (with update rules* (5.2.2) *and* (5.2.3)*) with parameters $\gamma \in (0,1]$ and $\delta \in [0,1)$, the dyadic confidence bound $c_{ij}(t)$ of each pair of adjacent nodes, i and j, converges either to* 0 *or to* 1*. Furthermore, if i and j are in different limit opinion clusters, then $c_{ij}(t)$ converges to* 0*.*

5.3.1.2 Effective-graph analysis

In this section, we discuss the convergence of effective graphs in our adaptive-confidence HK model (see Theorem 5.3.3) and the baseline HK model (see Theorem 5.3.4).

We first state Theorem 5.3.3, which is our convergence result for effective graphs in our adaptive-confidence HK model. The proof of Theorem 5.3.3 (which is in [LLP23]) is predominantly the work of Jiajie Luo; we do not include it in this dissertation.

Theorem 5.3.3. *In our adaptive-confidence HK model (with update rules* (5.2.2) *and* (5.2.3)*) with parameters $\gamma \in (0,1]$ and $\delta \in [0,1)$, the effective graph $G_{\text{eff}}(t)$ is eventually constant with respect to time. That is, there is some time T such that $G_{\text{eff}}(t) = G_{\text{eff}}(T)$ for all times $t \geq T$. Moreover, all of the edges of the* limit effective graph $\lim_{t\to\infty} G_{\text{eff}}(t)$ *are between nodes in the same limit opinion cluster.*

Theorem 5.3.3 states that all edges of a limit effective graph are between nodes in the same limit opinion cluster. However, the edges between nodes in the same limit opinion

cluster do not have to exist in the limit effective graph. As we will discuss in Section 5.4.3 and Section 5.5, our numerical simulations suggest that our adaptive-confidence BCMs can have adjacent nodes in the same limit opinion cluster whose associated dyadic confidence bound converges to 0. The associated edge is thus not in the limit effective graph.

We now state and prove Theorem 5.3.4, which guarantees that the effective graphs in the baseline HK model converge in the limit $t \to \infty$. Theorem 5.3.4 was proved jointly with Jiajie Luo. Unlike in our adaptive-confidence HK model, all edges between nodes in the same limit opinion cluster in the baseline HK model must exist in the limit effective graph. Therefore, the limit opinion values in the baseline HK model fully determine the structure of the limit effective graph.

Theorem 5.3.4. *In the baseline HK model (with update rule Equation* (5.2.1)*), the effective graph* $G_{\text{eff}}(t) = (V, E_{\text{eff}}(t))$ *is eventually constant with respect to time. Moreover, the edge* $(i, j) \in E$ *exists in the limit effective graph if and only if this edge is between two nodes in the same limit opinion cluster (i.e.,* $x_i^* = x_j^*$*).*

Proof. We first consider adjacent nodes, i and j, that are in different limit opinion clusters (i.e., $x_i^* \neq x_j^*$). By Lemma 5.3.1, because our adaptive-confidence HK model with $\gamma = 0$ and $\delta = 1$ reduces to the baseline HK model, there exists a time T_1 such that nodes i and j are not receptive to each other (i.e., $|x_i(t) - x_j(t)| \geq c$) at any time $t \geq T_1$. Therefore, the edge $(i, j) \notin E_{\text{eff}}(t)$ at any time $t \geq T_1$.

Now consider adjacent nodes, i and j, that are in the same limit opinion cluster (i.e., $x_i^* = x_j^*$). Choose a time T_2 such that $|x_k(t) - x_k^*| < c/2$ for each node k and all times $t \geq T_2$. For all $t \geq T_2$, we then have

$$|x_i(t) - x_j(t)| \leq |x_i(t) - x_i^*| + |x_i^* - x_j^*| + |x_j^* - x_j(t)| < c/2 + 0 + c/2 = c\,.$$

Therefore, at any time $t \geq T_2$, nodes i and j are receptive to each other and the edge $(i, j) \in E_{\text{eff}}(t)$. By taking $T = \max\{T_1, T_2\}$, for any time $t \geq T$, we have that $(i, j) \notin E_{\text{eff}}(t)$

for all edges (i, j) with $x_i^* \neq x_j^*$ and that $(i, j) \in E_{\text{eff}}(t)$ for all edges (i, j) with $x_i^* = x_j^*$.

□

5.3.2 Adaptive-confidence DW model

In this section, we discuss our theoretical results for the confidence bounds and effective graphs in our adaptive-confidence DW model. Both the baseline DW model and our adaptive-confidence DW model are asynchronous and stochastic. At each discrete time, we uniformly randomly select one pair of adjacent nodes to interact. Because of the stochasticity in the baseline and adaptive-confidence DW models, our theoretical results for them are in an "almost sure" sense. By contrast, our theoretical results (see Section 5.3.1) are deterministic for the baseline and adaptive HK models.

5.3.2.1 Confidence-bound analysis

In Theorem 5.3.5, we state our main result about the behavior of the confidence bounds in our adaptive-confidence DW model. This result mirrors Theorem 5.3.2 for our adaptive-confidence HK model. The proof of Theorem 5.3.5 (which is in [LLP23]) is predominantly the work of Jiajie Luo; we do not include it in this dissertation.

Theorem 5.3.5. *In our adaptive-confidence DW model (with update rules Equation* (5.2.4) *and Equation* (5.2.3)*)with parameters* $\gamma \in (0, 1]$ *and* $\delta \in [0, 1)$*, the dyadic confidence bound* $c_{ij}(t)$ *converges either to* 0 *or to* 1 *almost surely. Moreover, if nodes* i *and* j *are in different limit opinion clusters (i.e.,* $x_i^* \neq x_j^*$*), then* $c_{ij}(t)$ *converges to* 0 *almost surely.*

5.3.2.2 Effective-graph analysis

In this section, we discuss the convergence of effective graphs in our adaptive-confidence DW model (see Theorem 5.3.6) and the baseline HK model (see Theorem 5.3.7).

We first state Theorem 5.3.3, which is our convergence result for effective graphs in our adaptive-confidence DW model. The proof of Theorem 5.3.3 (which is in [LLP23]) is predominantly the work of Jiajie Luo; we do not include it in this dissertation.

Theorem 5.3.6. *In our adaptive-confidence DW model (with update rules* (5.2.4) *and* (5.2.3)*) with parameters* $\gamma \in (0, 1]$ *and* $\delta \in [0, 1)$*, the effective graph* $G_{\text{eff}}(t)$ *almost surely eventually has edges only between nodes of the same limit opinion cluster. That is, there is almost surely some time* T *such that* $(i, j) \in E_{\text{eff}}(t)$ *implies that* $x_i^* = x_j^*$ *for all* $t \geq T$*.*

Unlike in our adaptive-confidence HK model, $\lim_{t\to\infty} G_{\text{eff}}(t)$ may not exist in our adaptive-confidence DW model. When the limit does exist, we refer to $\lim_{t\to\infty} G_{\text{eff}}(t)$ as the *limit effective graph.*

We now state Theorem 5.3.7, which guarantees the almost-sure convergence of the effective graphs as $t \to \infty$ in the baseline DW model.

Theorem 5.3.7. *Consider the baseline DW model (with update rule* (5.2.4)*). Almost surely, the effective graph* $G_{\text{eff}}(t)$ *is eventually constant with respect to time. That is, there is almost surely a time* T *such that* $G_{\text{eff}}(t) = G_{\text{eff}}(T)$ *for all times* $t \geq T$*.*

Furthermore, suppose that the limit effective graph $\lim_{t\to\infty} G_{\text{eff}}(t)$ *exists. If adjacent nodes* i *and* j *have the same limit opinion (i.e., if* $x_i^* = x_j^*$*), then the edge* (i, j) *is in the limit effective graph. Additionally, if the edge* (i, j) *is in the limit effective graph, then* $x_i^* = x_j^*$ *almost surely.*

We prove Theorem 5.3.7 by first proving Lemma 5.3.8 and Lemma 5.3.9. The proof for Lemma 5.3.9 uses ideas from a discussion with Weiqi Chu. Because the baseline DW model has a time-independent confidence bound c, our proof of Theorem 5.3.7 uses different ideas than our proof of Theorem 5.3.6.

Lemma 5.3.8. *Consider the baseline DW model (with update rule* (5.2.4)*). There is a time* T_1 *and there is almost surely a time* T_2 *such that the following statements hold for all adjacent nodes* i *and* j*.*

(1) If $|x_i^* - x_j^*| < c$, then $|x_i(t) - x_j(t)| < c$ and the edge (i, j) is in the effective graph for all times $t \geq T_1$.

(2) If $|x_i^* - x_j^*| > c$, then $|x_i(t) - x_j(t)| > c$ and the edge (i, j) is not in the effective graph at any time $t \geq T_1$.

(3) If $|x_i^* - x_j^*| = c$, then $|x_i(t) - x_j(t)| \geq c$ and the edge (i, j) is not in the effective graph at any time $t \geq T_2$.

Proof. Consider adjacent nodes i and j, and let $\Delta_{ij} = |x_i^* - x_j^*|$ denote the difference between their opinions.

We first consider the case in which $\Delta_{ij} \neq c$. Choose a time T_{ij} such that

$$|x_k(t) - x_k^*| < \frac{1}{2}|c - \Delta_{ij}| \tag{5.3.1}$$

for node $k \in \{i, j\}$ and all times $t \geq T_{ij}$.

Suppose that $\Delta_{ij} < c$. For all times $t \geq T_{ij}$, we have

$$\begin{aligned} |x_i(t) - x_j(t)| &\leq |x_i(t) - x_i^*| + |x_i^* - x_j^*| + |x_j(t) - x_j^*| \\ &< 2\left(\frac{1}{2}\right)(c - \Delta_{ij}) + \Delta_{ij} \\ &= c\,. \end{aligned}$$

Therefore, the edge (i, j) is in the effective graph for all $t \geq T_{ij}$.

Now suppose that $\Delta_{ij} > c$. Without loss of generality, let $x_i^* > x_j^*$. For all times $t \geq T_{ij}$,

we have

$$\begin{aligned} x_i(t) - x_j(t) &> \left(x_i^* - \frac{1}{2}\left|c - \Delta_{ij}\right|\right) - \left(x_j^* + \frac{1}{2}\left|c - \Delta_{ij}\right|\right) \\ &= (x_i^* - x_j^*) - |c - \Delta_{ij}| \\ &= \Delta_{ij} - \Delta_{ij} + c \\ &= c\,. \end{aligned}$$

Therefore, the edge (i, j) is not in the effective graph at any time $t \geq T_{ij}$.

If there are no adjacent nodes i and j with $|x_i^* - x_j^*| \neq c$, then let $T_1 = 0$. Otherwise, let

$$T_1 = \max_{(i,j)\in E} \{T_{ij} \text{ such that } |x_i^* - x_j^*| \neq c\}\,. \tag{5.3.2}$$

We have shown that statements (1) and (2) hold for all times $t \geq T_1$.

We now consider the case $\Delta_{ij} = c$. Without loss of generality, let $x_i^* > x_j^*$. Choose a time $\tilde{T}_{ij}$ so that

$$|x_k^* - x_k(t)| < \frac{mc}{2(1+2m)} \tag{5.3.3}$$

for node $k \in \{i, j\}$ and all times $t \geq \tilde{T}_{ij}$. We will show that, almost surely, there are a finite number of times $t \geq \tilde{T}_{ij}$ such that $|x_i(t) - x_j(t)| < c$. Suppose on the contrary that there is a sequence $t_1, t_2, \ldots$ of times such that $t_k \geq \tilde{T}_{ij}$ and $|x_i(t_k) - x_j(t_k)| < c$ for all k. At each time t, nodes i and j interact with probability $1/|E| > 0$, where $|E|$ is the number of edges in the graph. Therefore, with probability 1, nodes i and j interact at some time $t_k \geq \tilde{T}_{ij}$ with $|x_i(t_k) - x_j(t_k)| < c$. Nodes i and j compromise their opinions at time t_k, so the inequality

(5.3.3) implies that

$$\begin{aligned}
|x_i(t+1) - x_i(t)| = m|x_i(t) - x_j(t)| &\geq m\left[\left(x_i^* - \frac{mc}{2(1+2m)}\right) - \left(x_j^* + \frac{mc}{2(1+2m)}\right)\right] \\
&= m\left[c - 2\left(\frac{mc}{2(1+2m)}\right)\right] \\
&= \frac{mc[(1+2m)-m]}{1+2m} \\
&= (1+m)\frac{mc}{1+2m} \\
&> \frac{mc}{1+2m}\,. \qquad (5.3.4)
\end{aligned}$$

From the inequality (5.3.3), we have

$$|x_i(t+1) - x_i(t)| \leq |x_i(t+1) - x_i^*| + |x_i^* - x_i(t)| < 2\left(\frac{mc}{2(1+2m)}\right) = \frac{mc}{1+2m}\,,$$

which cannot hold simultaneously with inequality (5.3.4). Therefore, with probability 1, there are a finite number of times $t \geq \tilde{T}_{ij}$ such that $|x_i(t) - x_j(t)| < c$. Consequently, there almost surely exists some time $T_{ij} \geq \tilde{T}_{ij}$ such that $|x_i(t) - x_j(t)| \geq c$ and the edge (i,j) is not in the effective graph for any $t \geq T_{ij}$.

If there are no adjacent nodes i and j with $|x_i^* - x_j^*| \neq c$, then let $T_2 = 0$. Otherwise, let

$$T_2 = \max_{(i,j)\in E}\{T_{ij} \text{ such that } |x_i^* - x_j^*| = c\}\,, \qquad (5.3.5)$$

where T_2 exists almost surely because each T_{ij} exists almost surely. We have shown that statement (3) holds for all times $t \geq T_2$ if T_2 exists.

□

Lemma 5.3.9. *For adjacent nodes i and j with $|x_i^* - x_j^*| < c$, we have that $x_i^* = x_j^*$ almost surely.*

Proof. Fix adjacent nodes i and j with $|x_i^* - x_j^*| < c$, and let $\Delta_{ij} = |x_i^* - x_j^*|$ denote

the distance between their opinions. Without loss of generality, let $x_i^* > x_j^*$. We want to show that $\Delta_{ij} = 0$ almost surely. Suppose instead that $\Delta_{ij} > 0$. Fix ϵ so that $0 < \epsilon < \min\{\frac{1}{4}(c - \Delta_{ij}), \frac{\Delta_{ij}}{2(1+1/m)}\}$ and choose T_{ij} so that

$$|x_k^* - x_k(t)| < \epsilon \tag{5.3.6}$$

for each node k and all times $t \geq T_{ij}$.

By the Borel–Cantelli lemma, there is almost surely some time $t \geq T_{ij}$ at which nodes i and j interact. The inequality $\epsilon < \frac{1}{4}(c - \Delta_{ij})$ implies that

$$\begin{aligned} |x_i(t) - x_j(t)| &\leq |x_i(t) - x_i^*| + |x_i^* - x_j^*| + |x_j^* - x_j(t)| \\ &< \frac{1}{4}(c - \Delta_{ij}) + \Delta_{ij} + \frac{1}{4}(c - \Delta_{ij}) = \frac{1}{2}(\Delta_{ij} + c) \\ &< c\,, \end{aligned}$$

so nodes i and j are receptive to each other at time t. Consequently, if they interact at time t, they update their opinions and

$$\begin{aligned} x_j(t+1) &= x_j(t) + m[x_i(t) - x_j(t)] \\ &\geq x_j(t) + m[x_i^* - \epsilon - (x_j^* + \epsilon)] \\ &= x_j(t) + m(\Delta_{ij} - 2\epsilon) \\ &\geq (x_j^* - \epsilon) + m(\Delta_{ij} - 2\epsilon) \\ &> x_j^* + \epsilon\,, \end{aligned} \tag{5.3.7}$$

where the last inequality holds because $\epsilon < \frac{\Delta_{ij}}{2(1+1/m)}$, which we rearrange to obtain $2\epsilon < m(\Delta_{ij} - 2\epsilon)$. The inequality Equation (5.3.6) implies that

$$|x_j^* - x_j(t+1)| < \epsilon\,, \tag{5.3.8}$$

which cannot hold simultaneously with the inequality Equation (5.3.7). Therefore, if $0 < x_i^* - x_j^* < c$, then nodes i and j cannot interact at times $t \geq T_{ij}$. However, by the Borel–Cantelli lemma, nodes i and j almost surely interact infinitely often. Consequently, $0 < x_i^* - x_j^* < c$ with probability 0. Therefore, we almost surely have $x_i^* = x_j^*$.

□

We now use Lemma 5.3.8 and Lemma 5.3.9 to prove Theorem 5.3.7.

Proof of Theorem 5.3.7. There is a time T_1 such that statements (1) and (2) of Lemma 5.3.8 hold, and there is almost surely a time T_2 such that statement (3) of Lemma 5.3.8 holds. Therefore, there is almost surely a time $T = \max\{T_1, T_2\}$ such that all three statements (1)–(3) of Lemma 5.3.8 hold for all times $t \geq T$. Consequently, the edges of the effective graph satisfy $E_{\text{Eff}}(t) = E_{\text{Eff}}(T)$ for all $t \geq T$. The effective graph is thus eventually constant with respect to time for all $t \geq T$.

Suppose that the limit effective graph $\lim_{t\to\infty} G_{\text{eff}}(t)$ exists. For adjacent nodes i and j with the same limit opinion (i.e., $x_i^* = x_j^*$), we know that $|x_i^* - x_j^*| = 0 < c$. By statement (1) of Lemma 5.3.8, there thus exists a time T_1 such that the edge (i, j) is in the effective graph for all times $t \geq T_1$. Therefore, the edge (i, j) is in the limit effective graph.

Now suppose that the edge (i, j) is in the limit effective graph. We seek to show that $x_i^* = x_j^*$ almost surely. Because the edge (i, j) is in the limit effective graph, there exists a time $\tilde{T}$ such that $(i, j) \in E_{\text{Eff}}(t)$ for all times $t \geq \tilde{T}$. Consequently, by statement (2) of Lemma 5.3.8, it cannot be the case that $|x_i^* - x_j^*| > c$. Therefore, by statement (3) of Lemma 5.3.8, it almost surely cannot be the case that $|x_i^* - x_j^*| = c$. Consequently, we almost surely have $|x_i^* - x_j^*| < c$. By Lemma 5.3.9, it is almost surely the case that $x_i^* = x_j^*$.

□

5.4 Details of our numerical simulations

We now discuss the details of our numerical simulations of our adaptive-confidence HK and DW models.

5.4.1 Network structures

We first simulate our adaptive-confidence HK and DW models on complete graphs to better understand their behaviors. We subsequently examine how different network structures affect those behaviors. We simulate our adaptive-confidence HK model on synthetic networks that we generate using random-graph models, and we simulate both adaptive-confidence BCMs on networks from empirical data. Because of computational limitations, we consider larger networks for the adaptive-confidence HK model than for the adaptive-model DW model.

We simulate our adaptive-confidence HK model on a complete graph, $G(N, p)$ ER random graphs (see Section 2.3.2), and two-community SBM random graphs (see Section 2.3.3). In each case, we consider graphs with 1000 nodes. We consider $G(N, p)$ graphs with $p \in \{0.1, 0.5\}$ to vary the sparsity of the graphs while still yielding connected graphs for our simulations. To construct our 2×2 two-community SBMs, we partition a network into two sets of nodes; set A has 750 nodes (i.e., 75% of the network) and set B has 250 nodes (i.e., 25% of the network). For our two-community SBMs, we consider the symmetric edge-probability matrix in Equation (2.3.1). In our simulations, we take $P_{\mathrm{AA}} = P_{\mathrm{BB}} = 1$ and $P_{\mathrm{AB}} = 0.01$.

In addition to synthetic networks, we also simulate our adaptive-confidence HK model on several real-world networks. For each network, we use the largest connected component. In Table 5.1, we give the numbers of nodes and edges in the largest connected components of these networks, which are social networks from the FACEBOOK100 data set [RKM11, TMP12]. In each of these networks, the nodes are the Facebook pages of individuals at a university and the edges encode Facebook "friendships" between individuals in a one-day snapshot of the network from fall 2005 [RKM11, TMP12]. The numbers of nodes in the

largest connected components of the examined Facebook networks range from 962 to 14,917.

For our adaptive-confidence DW model, we examine a complete graph and one real-world network. We simulate our adaptive-confidence DW model on a 100-node complete graph, which is one tenth of the size of the complete graph that we consider for our adaptive-confidence HK model. We use this smaller size because of computational limitations. Our simulations of our adaptive-confidence DW model on a 100-node complete graph frequently reach our "bailout time" (see Section 5.4.2 and Table 5.5) for small initial confidence bounds. We also simulate our adaptive-confidence DW model on the LCC of the real-world NetScience network of coauthorships between researchers in network science [New06].

Table 5.1: The real-world networks on which we simulate our adaptive-confidence BCMs. For each network, we use the largest connected component and indicate the numbers of nodes and edges in that component.

Network	**Number of Nodes**	**Number of Edges**	**Model**
NetScience	379	914	DW
Reed	962	18,812	HK
Swarthmore	1657	61,049	HK
Oberlin	2920	89,912	HK
Pepperdine	3440	152,003	HK
Rice	4083	184,826	HK
UC Santa Barbara	14,917	482,215	HK

5.4.2 Simulation specifications

In Table 5.2, we indicate the values of the model parameters that we examine in our simulations of our BCMs. The BCM parameters are the confidence-increase parameter γ, the confidence-decrease parameter δ, the initial confidence bound c_0, and (for the adaptive-confidence DW model only) the compromise parameter m. For both our HK and DW models, the parameter pair $(\gamma, \delta) = (0, 1)$ corresponds to the associated baseline BCM.

Our BCM simulations include randomness from the initial opinions of the nodes and

from the specific networks in random-graph ensembles. The adaptive-confidence DW model also has randomness from the selection of nodes at each time step. We use Monte Carlo simulations to mitigate the effects of noise. For each parameter set of a random-graph model (i.e., the ER and SBM graphs), we generate five graphs. Additionally, for each graph, we generate 10 sets of initial opinions uniformly at random and reuse these sets of opinions for all BCM parameter values.

Table 5.2: The BCM parameter values that we examine in simulations of our adaptive-confidence BCMs. We consider more parameter values for complete graphs than for the other networks. We consider all of the indicated values for complete graphs, and we consider values without the asterisk (*) for the ER, SBM, and real-world networks.

Model	**BCM Parameters**
Adaptive-Confidence HK	$\gamma \in \{0, 0.0001^*, 0.0005^*, 0.001, 0.005, 0.01, 0.05, 0.1^*\}$
	$\delta \in \{0.01^*, 0.1^*, 0.5, 0.9, 0.95, 0.99, 1\}$
	$c_0 \in \{0.02, 0.03, \ldots, 0.19, 0.20, 0.30, 0.40, 0.50\}$
Adaptive-Confidence DW	$\gamma \in \{0.1, 0.3, 0.5^*\}$
	$\delta \in \{0.3^*, 0.5, 0.7^*\}$
	$c_0 \in \{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9\}$
	$m \in \{0.1, 0.3, 0.5\}$

* We consider these parameter values only for complete graphs.

In our numerical simulations in this chapter, we use the stopping criterion in Equation (3.3.8). We specify that a simulation has converged if each opinion cluster has a diameter that is less than a tolerance values. That is, for each opinion cluster S_r, we have $\max_{i,j \in S_r} |x_i - x_j| <$ tolerance. We use a tolerance value of 1×10^{-6} for our adaptive-confidence HK model. Because of computational limitations, we use a tolerance value of 0.02 for our adaptive-confidence DW model. We denote the "convergence time" (see Section 3.3.5), which is the time step in which a simulation reaches the stopping criterion, by T_f. In our simulations, the "final effective graph" is the effective graph at the convergence time T_f. We refer to the connected components of the final effective graph as the "final opinion clusters" of a simulation; they approximate the limit opinion clusters.

Our theoretical results about effective graphs inform our stopping criterion. Theorem 5.3.3 and Theorem 5.3.4 give theoretical guarantees for our adaptive-confidence HK model and the baseline HK model, respectively, that eventually the only edges of an effective graph are those between adjacent nodes in the same limit opinion cluster. Theorem 5.3.6 and Theorem 5.3.7 give similar but weaker guarantees for our adaptive-confidence DW model and the baseline DW model. Consequently, if one of our simulations runs for sufficiently many time steps, its final opinion clusters are a good approximation of the limit opinion clusters. However, instead of imposing a set number of time steps for our simulations, we use a tolerance value as a proxy to determine a "sufficient" number of time steps.

The final and limit opinion clusters in our models may not be the same, as our choice of tolerance values can lead to simulations stopping before we can determine their limit opinion clusters. Additionally, if two distinct connected components in an effective graph converge to the same limit opinion value, the nodes in those connected components are in the same limit opinion cluster. However, the sets of nodes that constitute these connected components are distinct final opinion clusters. In practice, our simulations are unlikely to have distinct opinion clusters that converge to the same opinion in the infinite-time limit. Therefore, for small tolerance values, our final opinion clusters are a good approximation of the limit opinion clusters.

To ensure that our simulations stop after a reasonable amount of time, we use a bailout time of 10^6 time steps. Our simulations of the adaptive-confidence HK model never reach this bailout time. However, our simulations of the adaptive-confidence DW model frequently reach the bailout time for small values of c_0. See Section 5.5.2.1 and Table 5.5.

5.4.3 Quantifying model behaviors

In our numerical simulations, we investigate the convergence time and characterize the final opinions. To examine the convergence time, we record the number T_f of time steps that it takes for simulations to reach our stopping criterion. To characterize node opinions in

our adaptive-confidence BCMs, we calculate the numbers of major and minor final opinion clusters (see Section 3.3.3), quantify the opinion fragmentation using the Shannon entropy (see Equation (3.3.7)) of the final opinion-cluster profiles (see Section 3.3.4), and examine the numbers of nodes and edges in each final opinion cluster.

In this chapter, we say that an opinion cluster is a "minor opinion cluster" if it has at most 1% of the nodes of a network. That is, an opinion cluster S_r is a minor opinion cluster if $|S_r| \leq 0.01N$. We say that an opinion cluster that is not a minor cluster is a "major opinion cluster". As we discussed in Section 5.4.2, the final opinion clusters approximate, but may not be exactly the same as, the limit opinion clusters. In this chapter, we say that a simulation that results in one major final opinion cluster yields a "consensus" state and that a simulation that results in at least two major final opinion clusters yields "opinion fragmentation" (i.e., a "fragmented" state).

To examine the structure of final opinion clusters, we study the properties of final effective graphs. Our theoretical results allow the possibility that some adjacent nodes converge to the same opinion without being mutually receptive. We observe this phenomenon in our numerical simulations. (See Section 5.5 for more discussion.) To quantify this behavior, we calculate a weighted average of the fractions of edges (which we call the "weighted-average edge fraction") that are in each opinion cluster of the final effective graph. In an opinion-cluster profile $\{S_r(t)\}_{r=1}^{R}$, let $E(r)$ denote the set of edges of the original graph G between nodes in opinion cluster r and let $E_{\text{eff}}(t, r)$ denote the set of edges of the effective graph $G_{\text{eff}}(t)$ that are in opinion cluster r. That is,

$$E(r) = \{(i, j) \in E \text{ such that } i, j \in S_r(t)\},$$
$$E_{\text{eff}}(t, r) = \{(i, j) \in E_{\text{eff}}(t) \text{ such that } i, j \in S_r(t)\}.$$

The weighted average of the fractions of edges (i.e., the weighted-average edge fraction) that

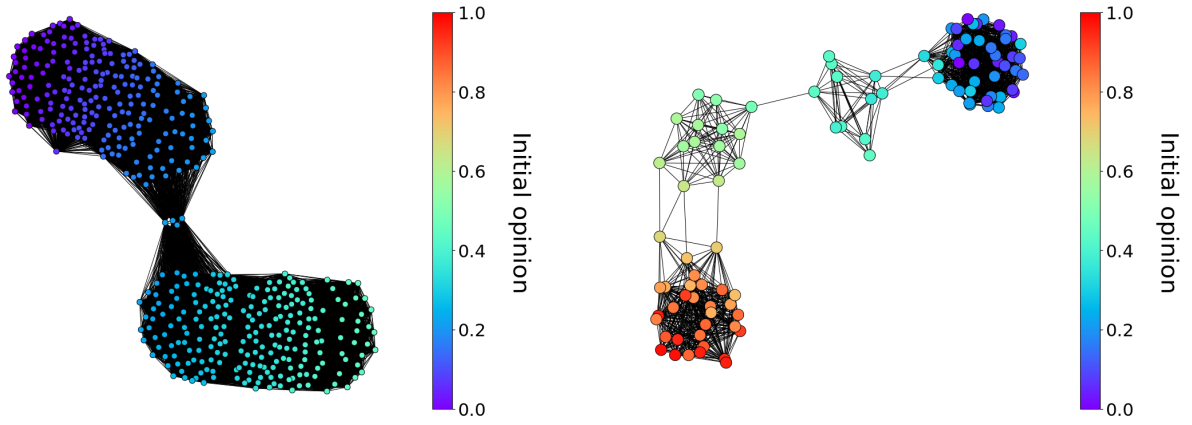

(a) One final opinion cluster in a simulation of our adaptive-confidence HK model that does not reach consensus on a 1000-node complete graph with $\gamma = 0.001$, $\delta = 0.5$, and $c_0 = 0.1$.

(b) The final effective graph in a simulation of our adaptive-confidence DW model that reaches consensus on a 100-node complete graph with $\gamma = 0.1$, $\delta = 0.5$, $c_0 = 0.1$, and $m = 0.1$.

Figure 5.1: Examples of final effective graphs with $W(T_f) < 1$. We color the nodes by their initial opinion values. (This figure originally appeared in [LLP23].)

are in the effective graph for each opinion cluster is

$$W(t) = \sum_{\substack{r=1 \\ E(r)\neq 0}}^{R} \left(\frac{|S_r(t)|}{N-\ell}\right)\left(\frac{|E_{\text{eff}}(t,r)|}{|E(r)|}\right), \tag{5.4.1}$$

where ℓ is the number of isolated nodes of the effective graph.[7] An isolated node is an opinion cluster with $E(r) = 0$. We are interested in the value of $W(t)$ at the convergence time T_f. Therefore, we calculate the weighted-average edge fraction $W(T_f)$. If each opinion cluster of an effective graph has all of its associated original edges of G, then $W = 1$. The value of W is progressively smaller when there are progressively fewer edges between nodes in the same opinion cluster of the effective graph. In Figure 5.1, we show examples of effective graphs with $W(T_f) < 1$.

[7] If every component of an effective graph is an isolated node (i.e, $N = \ell$), then one can take either $W(t) = 0$ or $W(t) = 1$. In our simulations, this situation never occurred.

5.5 Results of our numerical simulations

We now present the results of numerical simulations of our adaptive-confidence BCMs. We consider various values of the BCM parameters, which are the initial confidence bound c_0, the confidence-increase parameter γ, the confidence-decrease parameter δ, and (for the adaptive-confidence DW model only) the compromise parameter m. We use the BCM-parameter values in Table 5.2, including the values that correspond to the baseline models (i.e., $(\gamma, \delta) = (0, 1)$). Our code and plots are available in our code repository.

As we described in Section 5.4.3, for both of our adaptive-confidence BCMs, we examine the number of major clusters (which we use to determine whether a simulation reaches a consensus state or a fragmented state), the number of minor clusters, the Shannon entropy $H(T_f)$ (see equation Equation (3.3.7)), the weighted-average edge fraction $W(T_f)$ (see equation Equation (5.4.1)), and the convergence time T_f. When the Shannon entropy and the number of major clusters follow similar trends, we only show results for the number of major clusters, as it is easier to interpret than the entropy. To avoid drowning readers with too much repetition, we do not show plots for all of our numerical results; the omitted plots are available in our code repository.

Our simulation results and theoretical results about effective graphs complement each other. Theorem 5.3.2 states that all dyadic confidence bounds in our adaptive-confidence HK model converge either to 0 or to 1. Additionally, the dyadic confidence bounds for node pairs in different limit opinion clusters must converge to 0. However, we have not proven whether or not the dyadic confidence bounds for nodes in the same limit opinion cluster converge to 1, so it is possible for such confidence bounds to converge to 0. If a dyadic confidence bound convergences to 0, then the corresponding edge is absent in the limit effective graph (which is guaranteed to exist by Theorem 5.3.3). Our numerical simulations suggest that a final opinion cluster can include adjacent nodes whose dyadic confidence bound converges to 0. In particular, in many simulations, we observe that the weighted-average edge fraction

$W(T_f) < 1$, which corresponds to absent edges of the final effective graph between nodes that are in the same final opinion cluster. Our adaptive-confidence DW model has analogous theoretical results (see Theorem 5.3.5 and Theorem 5.3.6) and we again observe simulations with $W(T_f) < 1$.

5.5.1 Adaptive-confidence HK model

5.5.1.1 Summary of our simulation results

For our adaptive-confidence HK model, all of our numerical simulations reach a consensus state for $c_0 \geq 0.3$. (As we discussed in Section 5.4.3, a consensus state has exactly one major opinion cluster.) We show our simulation results for $c_0 \in \{0.02, 0.03, \ldots, 0.20\}$; we include the results for the other examined values of c_0 (see Table 5.2) in our code repository. We examine the numbers of major and minor clusters, the Shannon entropy $H(T_f)$ (see equation Equation (3.3.7)), the weighted-average edge fraction $W(T_f)$ (see equation Equation (5.4.1)), and the convergence time T_f. We plot each of these quantities versus the initial confidence bound c_0. For each value of the confidence-increase parameter γ, we generate one plot; each plot has one curve for each value of the confidence-decrease parameter δ. Each point in our plots is a mean of our numerical simulations for the associated values of the BCM parameter set (γ, δ, and c_0). We also show one standard deviation from the mean. For our simulations on a complete graph and the FACEBOOK100 networks, each point in our plots is a mean of 10 simulations (from 10 sets of initial opinions). For our simulations on $G(N, p)$ ER random graphs and SBM random graphs, each point in our plots is a mean of 50 simulations (from five random graphs that each have 10 sets of initial opinions). We include all plots, including those that we do not present in this dissertation, in our code repository. In Table 5.3, we summarize the trends that we observe in our simulations.

In all of our simulations of our adaptive-confidence HK model, we observe that $\gamma \geq 0.001$ results in fewer major clusters than in the baseline HK model. For a fixed initial confidence

Table 5.3: Summary of the observed trends in our adaptive-confidence HK model. Unless we note otherwise, we observe these trends for the complete graph, all examined random-graph models, and all examined real-world networks.

Quantity	**Trends**
Convergence Time	• For fixed values of the initial confidence bound c_0, our adaptive-confidence HK model tends to converge more slowly than the baseline HK model. • When our simulations reach a consensus state, for fixed values of c_0 and the confidence-increase parameter γ, our model converges faster when the confidence-decrease parameter $\delta = 1$ than when $\delta \leq 0.9$. For $\delta \in \{0.95, 0.99\}$, the convergence time transitions from the $\delta \leq 0.9$ behavior to the $\delta = 1$ behavior as we increase c_0.
Opinion Fragmentation[1]	• Our adaptive-confidence HK model yields consensus for $\gamma \geq 0.05$. • For fixed values of c_0, our adaptive-confidence HK model tends to yield fewer major clusters than the baseline HK model. When we fix the other BCM parameters, the number of major clusters decreases as either (1) we decrease δ or (2) we increase γ. • For our synthetic networks, we observe that the trends in Shannon entropy match the trends in the numbers of major clusters and that our adaptive-confidence HK model tends to yield less opinion fragmentation than the baseline HK model.[1]
$W(T_f)$	• When $\delta = 1$, both our adaptive-confidence HK model and the baseline HK model yield $W(T_f) = 1$. • For fixed γ and c_0, as we increase δ, the weighted-average edge fraction $W(T_f)$ tends to decrease. • For simulations that reach a consensus state, we observe two qualitative behaviors for $W(T_f)$: for $\delta \leq 0.9$, we observe that $W(T_f) < 1$ and that there is a seemingly linear relationship between $W(T_f)$ and c_0; for $\delta = 1$, we observe that $W(T_f) = 1$. Additionally, for $\delta \in \{0.95, 0.99\}$, the behavior of $W(T_f)$ transitions from the $\delta \leq 0.9$ behavior to the $\delta = 1$ behavior as we increase c_0.

[1]For the FACEBOOK100 networks, we do not observe this trend, seemingly because of the large numbers of minor clusters (which are incorporated in our calculation of Shannon entropy in Equation (3.3.7)) for these networks

bound c_0, our adaptive-confidence HK model tends to yield fewer major opinion clusters and less opinion fragmentation as either (1) we increase γ for fixed δ or (2) we decrease δ for fixed γ. Intuitively, one expects larger values of γ to encourage consensus because a larger γ entails a larger increase in a dyadic confidence bound after a positive interaction. Less intuitively, smaller values of δ, which entail a larger decrease in a dyadic confidence bound after a negative interaction, also seem to encourage consensus. In our adaptive-confidence HK model, we update opinions synchronously, with each node interacting with of all its neighboring nodes at each time. When two adjacent nodes are mutually unreceptive, their dyadic confidence bound decreases. Given the synchronous updates in our adaptive-confidence HK model, we hypothesize that small values of δ yield a faster decrease than large values of δ in the dyadic confidence bound between mutually unreceptive nodes. For small values of δ, pairs of nodes may quickly become mutually unreceptive and remain mutually unreceptive. Meanwhile, individual nodes can be receptive to (and thus average) fewer "conflicting" opinions[8], possibly aiding in reaching a consensus.

In the baseline HK model, the number of major opinion clusters tends to decrease as we increase c_0. For intermediate values of γ (e.g., $\gamma \in \{0.005, 0.01\}$ for a complete graph; see panels (E) and (F) in Section 5.5.1.2), we do not observe this trend in our adaptive-confidence HK model. Instead, as we increase c_0, we observe an increase and then a decrease in the number of major clusters; unlike for the baseline HK model, small values of c_0 tend to promote more consensus (i.e., it tends to yield fewer major clusters). For smaller values of c_0, it seems that nodes tend to be receptive to fewer nodes early in a simulation, so fewer opinions can influence them. Therefore, for these values of c_0, nodes that are mutually receptive may quickly approach a consensus when γ is sufficiently large. Our calculations of the weighted-average edge fraction $W(T_f)$ support this hypothesis. For sufficiently large γ, small values of c_0 yield small values of $W(T_f)$, indicating that final opinion clusters tend to

[8]When the neighbors to which a node is receptive have large differences in opinions with each other, we say that that node is receptive to "conflicting" opinions.

have many pairs of mutually unreceptive nodes.

5.5.1.2 A complete graph

We now discuss the simulations of our adaptive-confidence HK model on a complete graph. We summarize the observed trends in Table 5.3.

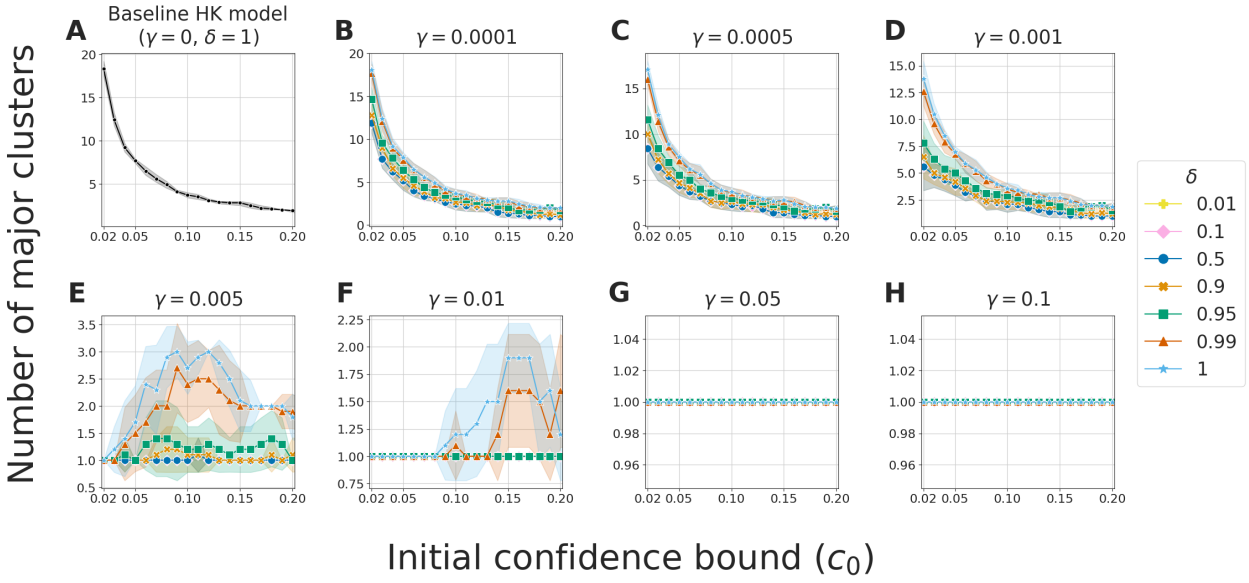

Figure 5.2: The numbers of major clusters in simulations of our adaptive-confidence HK model on a 1000-node complete graph for various combinations of the BCM parameters γ, δ, and c_0. In this and subsequent figures, we plot the mean value of our simulations for each set of BCM parameters. The bands around each curve indicate one standard deviation around the mean values. For clarity, in this figure and in subsequent figures, the vertical axes of different panels have different scales. (This figure originally appeared in [LLP23].)

In Section 5.5.1.2, we observe for a 1000-node complete graph that our adaptive-confidence HK model yields fewer major clusters (i.e., it encourages more consensus) than the baseline HK model for a wide range of BCM parameter values. Our adaptive-confidence HK model always reaches a consensus state for $\gamma \geq 0.05$. In our simulations that do not reach consensus, we tend to observe progressively more major clusters and more opinion fragmentation as either (1) we decrease γ for fixed δ and c_0 or (2) we increase δ for fixed γ and

c_0. For the baseline HK model and our adaptive-confidence HK model with small values of γ (specifically, $\gamma \in \{0.0001, 0.0005, 0.001\}$), the number of major opinion clusters tends to decrease as we increase c_0. We do not observe this trend for larger values of γ (specifically, $\gamma \in \{0.005, 0.01\}$). Instead, for these values of γ, small values of c_0 tend to promote more consensus. For example, simulations always reach a consensus state when $\gamma = 0.01$ and $c_0 \leq 0.08$. At the end of Section 5.5.1.1, we discussed our hypothesis behind this observation.

We observe very few minor clusters in our simulations of our adaptive-confidence HK model on a 1000-node complete graph. For each value of the BCM parameter set (γ, δ, c_0), the mean number of minor clusters in our 10 simulations is bounded above by 1. Because there are few minor clusters, the Shannon entropy (which accounts for both major clusters and minor clusters; see equation Equation (3.3.7)) and the number of major clusters follow similar trends for a 1000-node complete graph.

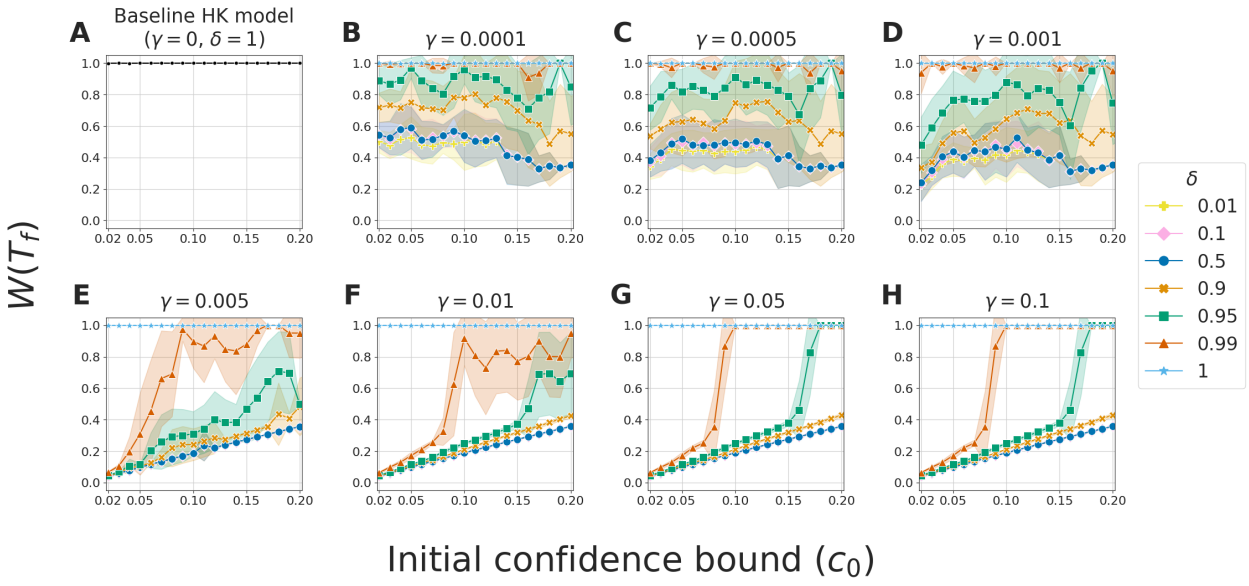

Figure 5.3: The weighted-average edge fraction $W(T_f)$ (see equation Equation (5.4.1)) in simulations of our adaptive-confidence HK model on a 1000-node complete graph for various combinations of the BCM parameters γ, δ, and c_0. (This figure originally appeared in [LLP23].)

In Section 5.5.1.2, we show $W(T_f)$ (see equation Equation (5.4.1)), which is the weighted

average of the fractions of edges in the opinion clusters of the final effective graph. When $\delta = 1$, both our adaptive-confidence HK model and the baseline HK model yield $W(T_f) = 1$. This indicates that all final opinion clusters (i.e., the connected components of the effective graph at time T_f) are complete subgraphs (i.e., cliques). By contrast, in our adaptive-confidence HK model, when $\delta < 1$ and for a wide range of the other BCM parameters, we observe that $W(T_f) < 1$. This indicates that some nodes that are adjacent in the graph G and in the same final opinion cluster do not have an edge between them in the final effective graph $G_{\text{eff}}(T_f)$. The nodes in these dyads are thus not receptive to each other (and hence do not influence each other's opinions), but they nevertheless converge to the same opinion. When δ is sufficiently small (specifically, $\delta \leq 0.9$), we observe that our adaptive-confidence HK model can reach a consensus with $W(T_f) < 1$. For sufficiently large values of γ (specifically, $\gamma \geq 0.05$), even though the nodes in some dyads are not receptive to each other, most nodes (at least 99% of them, based on our definition of major cluster) still converge to the same final opinion and hence reach a consensus.

For fixed values of γ and c_0, we observe that $W(T_f)$ tends to decrease as we decrease δ. For $\gamma \in \{0.05, 0.1\}$, our simulations always reach a consensus state. In these simulations, for each fixed δ, we observe that $W(T_f)$ appears to increase monotonically with respect to c_0. Additionally, for these values of γ, we observe a transition in $W(T_f)$ as a function of δ. For $\delta \leq 0.9$, we observe that $W(T_f) < 1$ and that there is a seemingly linear relationship between $W(T_f)$ and c_0. When $\delta = 1$, we observe that $W(T_f) = 1$. For $\delta \in \{0.95, 0.99\}$, the behavior of $W(T_f)$ transitions from the $\delta \leq 0.9$ behavior to the $\delta = 1$ behavior as we increase c_0. This transition between behaviors occurs for smaller c_0 for $\delta = 0.99$ than for $\delta = 0.95$.

In Section 5.5.1.2, for fixed c_0, we observe that our adaptive-confidence HK model takes longer to converge than the baseline HK model. For a 1000-node complete graph and fixed BCM parameters (i.e., γ, δ, and c_0), we observe that the logarithm $\log_{10}(T_f)$ of the convergence time T_f for our adaptive-confidence HK model can be up to 4 more than the logarithm

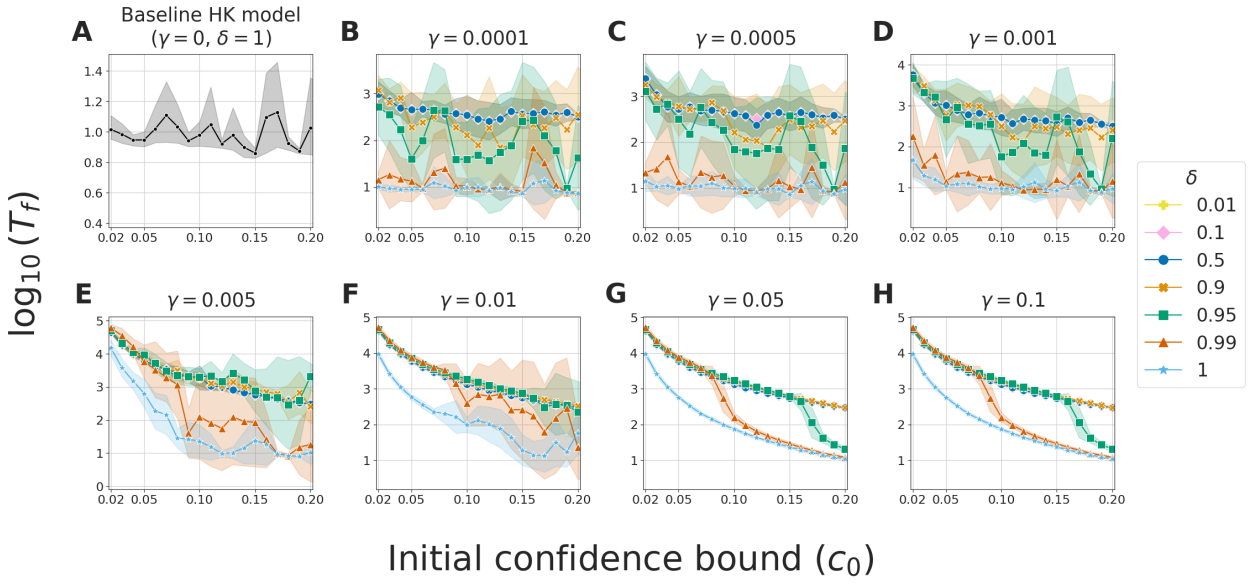

Figure 5.4: The convergence times (in terms of the number of time steps) on a logarithmic scale in simulations of our adaptive-confidence HK model on a 1000-node complete graph for various combinations of the BCM parameters γ, δ, and c_0. (This figure originally appeared in [LLP23].)

of the convergence time for the baseline HK model. That is, the convergence time can be as much as 10^4 times larger. The convergence time tends to increase as either (1) we increase γ for fixed δ and c_0 or (2) we decrease δ for fixed γ and c_0. For large values of γ (as is especially evident for $\gamma \in \{0.05, 0.1\}$), the convergence time decreases with c_0. As with $W(T_f)$, for these values of γ, we observe a transition in the convergence time as a function of δ. In Section 5.5.1.2, the curves of $\log_{10}(T_f)$ versus c_0 for $\delta \leq 0.9$ overlay each other and indicate larger convergence times than the curve for $\delta = 1$. The curves for $\delta = 0.95$ and $\delta = 0.99$ transition from the $\delta \leq 0.9$ behavior to the $\delta = 1$ behavior as we increase c_0. By contrast, for the baseline HK model and for our model with small values of γ, we observe no clear pattern between the convergence time and initial confidence bound. When our adaptive-confidence HK model reaches a consensus state, we observe from the behavior of $W(T_f)$ (see Section 5.5.1.2) and the convergence times (see Section 5.5.1.2) in our simulations that there is qualitative transition in the model behavior as we vary δ. We are not aware of previous

discussions of similar transitions in variants of the HK model.

5.5.1.3 ER graphs

We now discuss additional results of our simulations of our adaptive-confidence HK model on $G(N, p)$ ER random graphs. We generate five ER random graphs for each value of $p \in \{0.1, 0.5\}$. Each point in our plots is a mean of 50 simulations (from five random graphs that each have 10 sets of initial opinions). For fixed BCM parameters (namely, γ, δ, and c_0), our results for $G(1000, 0.5)$ graphs are more similar than those for $G(1000, 0.1)$ graphs to our results for the 1000-node complete graph.

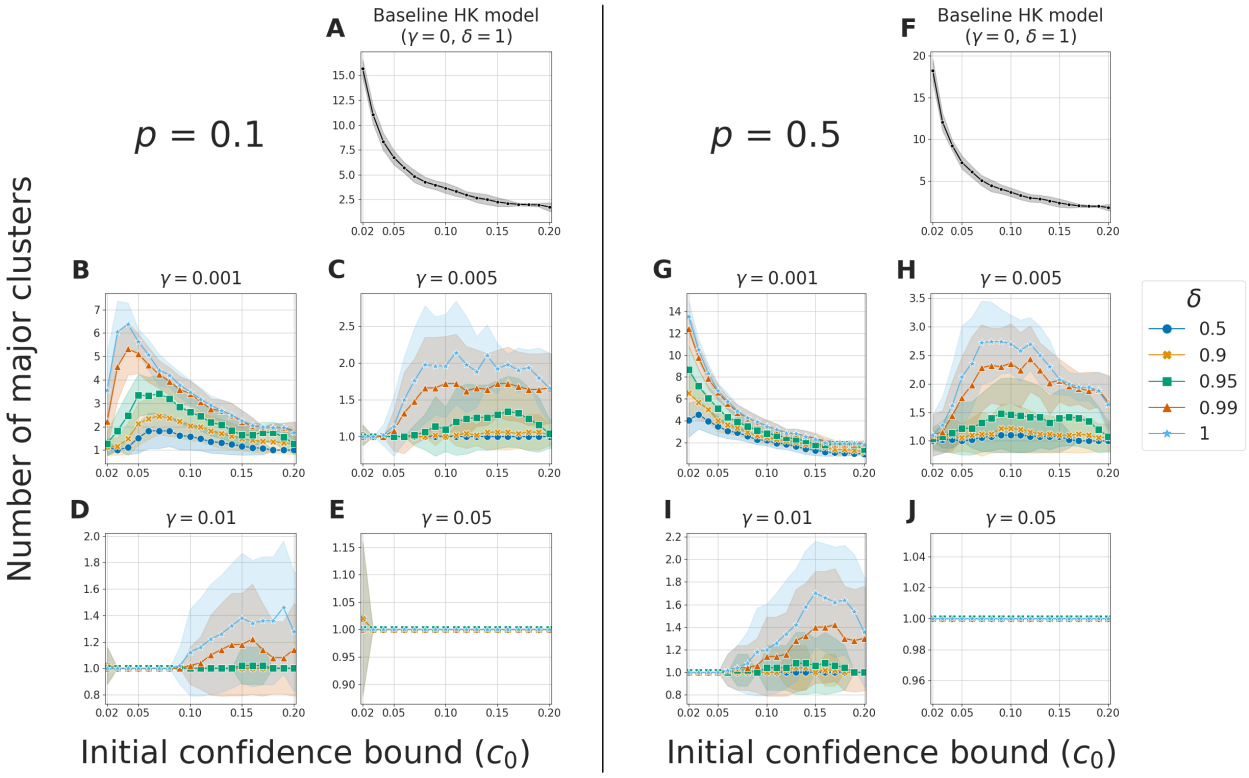

Figure 5.5: The numbers of major clusters in simulations of our adaptive-confidence HK model on $G(1000, p)$ ER random graphs with (A–E) $p = 0.1$ and (F–J) $p = 0.5$ for various combinations of the BCM parameters γ, δ, and c_0. (This figure originally appeared in [LLP23].)

In Section 5.5.1.3, we show the numbers of major clusters in our simulations of our adaptive-confidence HK model on ER graphs. For the $G(1000, 0.1)$ graphs, when $\gamma = 0.001$, small values of c_0 tend to yield few major clusters. As we increase c_0, we observe an initial increase in the number of major clusters followed by a decrease in that number. By contrast, for the 1000-node complete graph and $G(1000, 0.5)$ graphs, small values of c_0 tend to yield the most major clusters; for fixed values of γ and δ, the number of major clusters tends to decrease as we increase c_0. $G(1000, 0.1)$ graphs have more small-degree nodes than the complete graph. These small-degree nodes can easily form minor opinion clusters, especially for small values of c_0. We hypothesize that these minor clusters form quickly in a simulation and that the nodes in them quickly become unreceptive to the other nodes of a network.

For fixed BCM parameters (namely, γ, δ, and c_0), we tend to observe fewer major opinion clusters for $G(1000, 0.1)$ graphs than for the 1000-node complete graph. Additionally, for small initial confidence bounds (specifically, $c_0 \leq 0.04$), we observe more minor clusters for the $G(1000, 0.1)$ graphs than the $G(1000, 0.5)$ graphs and the 1000-node complete graph. (For $G(1000, 0.1)$ graphs, once we take the mean for each BCM parameter set, we sometimes observe as many as 20 minor clusters.) The expected mean degree of a $G(N, p)$ ER graph is $p(N - 1)$ (see Section 2.3.2). Therefore, for small probability p, we expect more nodes to have small degrees. For small initial confidence bounds, we hypothesize that many nodes with small degrees quickly disconnect to form minor opinion clusters in the effective graph.

As we just discussed, for small values of c_0, our simulations of the adaptive-confidence HK model on the $G(1000, 0.1)$ graphs yield more minor clusters than our simulations on the $G(1000, 0.5)$ graphs and the 1000-node complete graph. Nevertheless, although Shannon entropy (see equation Equation (3.3.7)) accounts for minor clusters, we still observe that it follows similar trends as the number of major clusters for both $p = 0.1$ and $p = 0.5$. Specifically, the Shannon entropy tends to increase as either (1) we decrease γ for fixed δ and c_0 or (2) we increase δ for fixed γ and c_0.

For our simulations on ER graphs with both $p = 0.1$ and $p = 0.5$, we observe the

convergence-time trends in Table 5.3. For fixed values of γ, δ, and c_0, the mean convergence time for $p = 0.1$ is at least as long as that for $p = 0.5$. Unlike for the complete graph, the ER graphs do not have a clear trend in the dependence of the convergence time either on γ (with fixed δ and c_0) or on δ (with fixed γ and c_0). As with the 1000-node complete graph, our fastest convergence times for ER graphs typically occur for $\delta = 1$. For fixed γ and c_0, we often observe that the convergence time increases as we decrease δ. However, we do not always observe this trend; for some values of γ and c_0, smaller values of δ yield faster convergence than larger values of δ.

5.5.1.4 Two-community SBM graphs

We now discuss additional results of our simulations of our adaptive-confidence HK model on two-community SBM graphs. Each of our SBM graphs consists of two complete graphs that are joined by a small number of edges (see Section 5.4.1). This yields a two-community structure. For our two-community SBM graphs, for fixed BCM parameters, the numbers of major clusters and Shannon entropies are similar to those for the complete graph. It seems that this two-community structure does not significantly impact the simulation results of our adaptive-confidence HK model.

In Figure 5.6, we show the numbers of major clusters in our simulations on SBM graphs. For fixed values of the BCM parameters (namely, γ, δ, and c_0), these simulations yield similar numbers of major clusters as in our simulations on the 1000-node complete graph (see Section 5.5.1.2) and $G(1000, 0.5)$ ER graphs (see Section 5.5.1.3). We observe few minor clusters; for each BCM parameter set, the mean number of minor clusters is bounded above by 3. Consequently, the Shannon entropy and the number of major clusters follow similar trends.

The convergence times in our simulations on SBM graphs follow the trends in Table 5.3. For fixed values of γ and c_0, we do not observe a clear trend in how the convergence time changes as we vary δ. One commonality between the SBM graphs, the ER graphs, and the

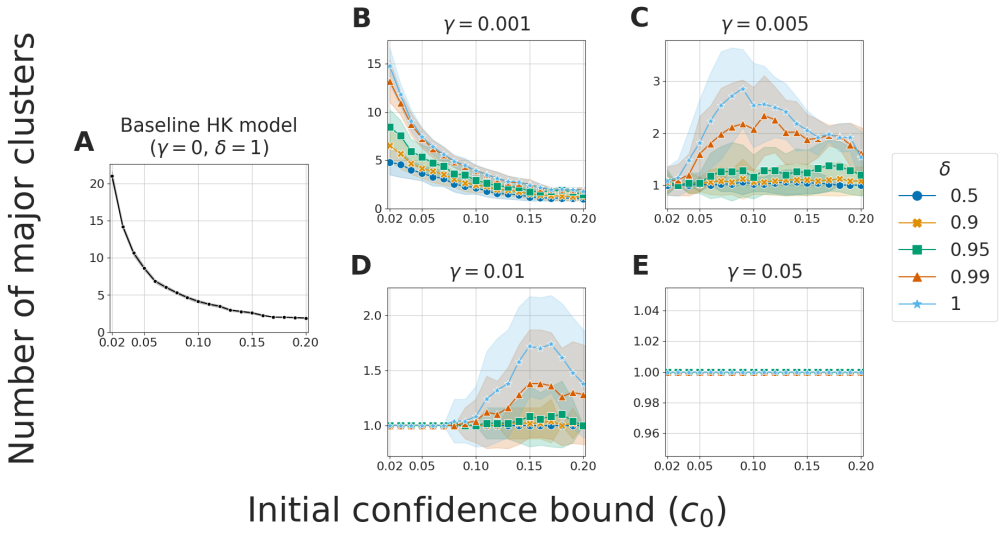

Figure 5.6: The numbers of major clusters in simulations of our adaptive-confidence HK model on 1000-node SBM random graphs with connection probabilities $P_{aa} = P_{bb} = 1$ and $P_{ab} = 0.01$ for various combinations of the BCM parameters γ, δ, and c_0. (This figure originally appeared in [LLP23].)

complete graph is that $\delta = 1$ gives the fastest convergence times. For a wide range of fixed values of γ and δ, we also observe that the convergence time tends to decreases as we increase c_0 for both our adaptive-confidence HK model and the baseline HK model.

5.5.1.5 Facebook100 university networks

We now discuss our simulations of our adaptive-confidence HK model on FACEBOOK100 networks (see Section 5.4.1) [RKM11, TMP12]. We show plots of the number of major clusters and Shannon entropy for the UC Santa Barbara network, as well as a plot of the number of major clusters for Reed College. In our code repository, we include all plots for our simulations on FACEBOOK100 networks (including plots for the other examined quantities and the other four universities in Table 5.1).

The six FACEBOOK100 networks (see Table 5.1) mostly exhibit the same trends. The Reed College network is a notable exception; we discuss it at the end of this subsubsection.

Except for the trends in Shannon entropy, we observe the same trends (see Table 5.3) for the FACEBOOK100 networks that we observed for the synthetic networks. For the FACEBOOK100 networks, most of the final opinion clusters for both our adaptive-confidence HK model and the baseline HK model are minor opinion clusters. Our simulations on the UC Santa Barbara network yield more minor clusters than our simulations on the other FACEBOOK100 networks; when $c_0 = 0.02$ and $\delta \leq 0.9$, the UC Santa Barbara network has more than 4000 minor clusters. Our calculation of Shannon entropy (see Equation (3.3.7)) includes contributions from minor opinion clusters. Therefore, because of the large numbers of minor clusters for the FACEBOOK100 networks, the Shannon entropy and numbers of major opinion clusters follow different trends. For these networks, they thus give complementary views of opinion fragmentation.

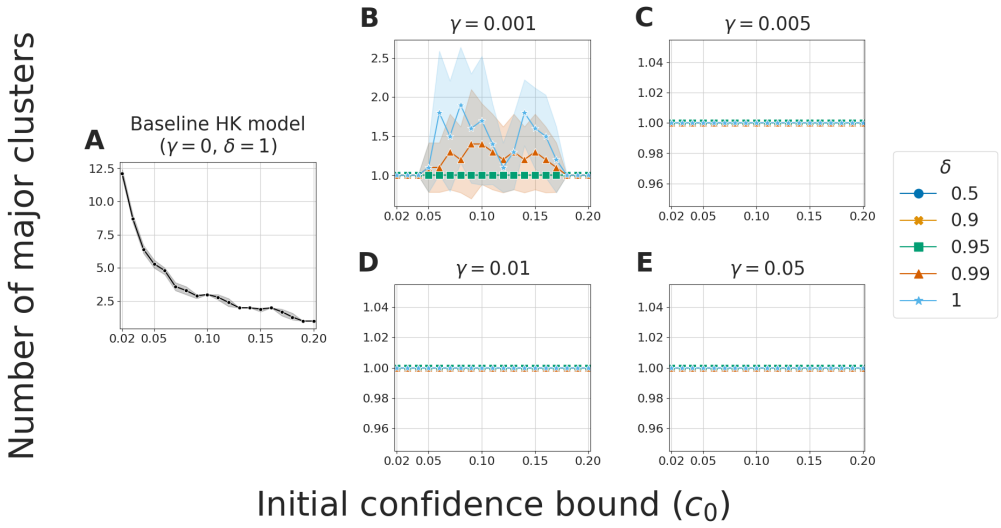

Figure 5.7: The numbers of major clusters in simulations of our adaptive-confidence HK model on the UC Santa Barbara network for various combinations of the BCM parameters γ, δ, and c_0. (This figure originally appeared in [LLP23].)

In Figure 5.7, we observe for the UC Santa Barbara network that our adaptive-confidence HK model always yields consensus (the number of major clusters is exactly 1) when $\gamma \geq 0.005$. For these values of γ, the Shannon entropy (see Figure 5.8) tends to decrease as we increase c_0 for fixed values of γ and δ. This trend occurs because the number of minor opinion

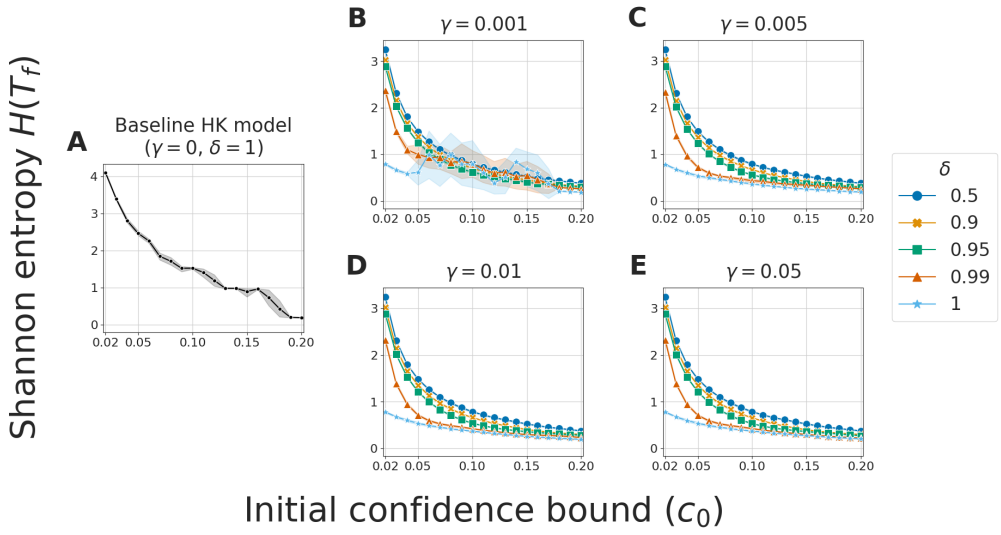

Figure 5.8: The Shannon entropies in simulations of our adaptive-confidence HK model on the UC Santa Barbara network for various combinations of the BCM parameters γ, δ, and c_0. (This figure originally appeared in [LLP23].)

clusters also tends to decrease as we increase c_0 for fixed values of γ and δ. This observation contrasts with our simulations of our adaptive-confidence HK model on synthetic networks (see Sections 5.5.1.2–5.5.1.4), for which we observed that the Shannon entropy follows similar trends as the number of major clusters as we vary one of γ, δ, or c_0 while fixing the other BCM parameters. We believe that one reason for this difference is that the FACEBOOK100 networks have many small-degree nodes, which allow more minor opinion clusters to form.

We show the numbers of major opinion clusters for the Reed College network in Figure 5.9. For very small initial confidence bounds $c_0 \leq 0.04$ and fixed values of the BCM parameters, the Reed College network tends to have more major clusters and larger Shannon entropies than the other five FACEBOOK100. This difference may arise from the small size of the Reed College network in concert with our definition of major cluster. For example, a final opinion cluster with 20 nodes is a major cluster for the Reed College network (which has 962 nodes in its largest connected component), but an opinion cluster of that size is a minor cluster for the UC Santa Barbara network (which has 14,917 nodes in its largest connected component).

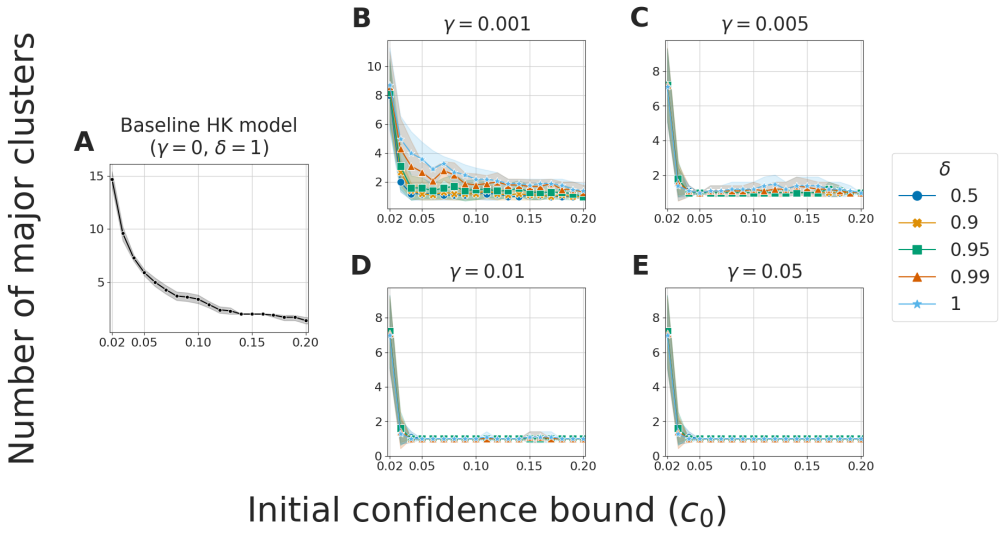

Figure 5.9: The numbers of major clusters in simulations of our adaptive-confidence HK model on the Reed College network for various combinations of the BCM parameters γ, δ, and c_0. (This figure originally appeared in [LLP23].)

5.5.2 Adaptive-confidence DW model

We now show our results of our simulations of our adaptive-confidence DW model on a 100-node complete graph and the NetScience network. In Table 5.4, we summarize the trends that we observe in these simulations. We simulate our adaptive-confidence DW model with the values of the BCM parameters (namely, the initial confidence bound c_0, the confidence-increase parameter γ, the confidence-decrease parameter δ, and the compromise parameter m) in Table 5.2. Because of the long computation times, we consider much smaller graphs and fewer BCM parameter values for our adaptive-confidence DW model than we did for our adaptive-confidence HK model. Notably, the value of the compromise parameter m (which is in the DW models but is not the HK models) affects our simulation results.

We explore the dependence of the numbers of major and minor clusters, the Shannon entropy $H(T_f)$ (see equation Equation (3.3.7)), the weighted-average edge fraction $W(T_f)$ (see equation Equation (5.4.1)), and the convergence time T_f on the initial confidence bound c_0. For each value of (γ, δ), we generate one plot; each plot has one curve for each value of the

compromise parameter m. Each point in our plots is the mean of 10 numerical simulations (from 10 sets of initial opinions) with one BCM parameter set (γ, δ, c_0, m). We also show one standard deviation from the mean. All plots, including those that we do not present in this appendix, are available in our code repository.

5.5.2.1 A complete graph

We first discuss our simulations of our adaptive-confidence DW model on a 100-node complete graph. In the present section, we show plots of the numbers of major opinion clusters (see Figure 5.10) and the weighted-average edge fractions $W(T_f)$ (see Figure 5.11).

Our adaptive-confidence DW model tends to converge more slowly than both the baseline DW model and our adaptive-confidence HK model. Our simulations of our adaptive DW model often reach the bailout time, particularly for small values of c_0 and m. In Table 5.5, we indicate the numbers of simulations that reach the bailout time. In some simulations, despite reaching the bailout time, we are still able to identify the final opinion clusters. However, the maximum difference in the opinions of the nodes in these clusters is not within our tolerance value (see Section 5.4.2) of 0.02 for our adaptive-confidence DW model. In those instances, we still use the cluster information to calculate the numbers of major and minor opinion clusters, the Shannon entropy $H(T_f)$, and the weighted-average edge fraction $W(T_f)$. For our simulations of our adaptive-confidence DW model with $(\gamma, \delta) = (0.1, 0.5)$, we run each simulation to convergence (i.e., until we reach the stopping condition that we described in Section 5.4.2). We plot the results of these simulations in Figure 5.10E and Figure 5.11B. Although some simulations reach the bailout time, the information that we are able to obtain about the opinion clusters (from both the simulations that we run to convergence and the simulations that reach the bailout time) give us confidence in the trends in Table 5.4.

In Figure 5.10, we observe for a wide range of BCM parameter values that our adaptive-confidence DW model yields fewer major clusters (i.e., it encourages more consensus) than the baseline DW model. When $c_0 \geq 0.5$, our adaptive-confidence DW model and the baseline

Table 5.4: Summary of the observed trends in our adaptive-confidence DW model.

Quantity	**Trends**
Convergence Time	• For the complete graph, for fixed values of the compromise parameter m and initial confidence bound $c_0 \leq 0.3$, our adaptive-confidence DW model tends to converge more slowly than the baseline DW model. • For the NetScience network, for fixed values of m and c_0, our adaptive-confidence DW model and the baseline DW model have similar convergence times.
Number of Major Clusters	• For the complete graph, when we fix the other BCM parameters, we (1) tend to observe fewer major clusters as we increase the confidence-increase parameter γ and (2) observe little effect on the numbers of major clusters when we vary the confidence-decrease parameter δ. • For the complete graph, for a fixed value of $c_0 \leq 0.3$, our adaptive-confidence DW model yields fewer major clusters when $m = 0.1$ than when $m \in \{0.3, 0.5\}$. The baseline DW model does not have this behavior. • For the NetScience network, for a fixed value of c_0, our adaptive-confidence DW model yields at least as many major clusters as the baseline DW model. For this network, m has little effect on the number of major clusters.
$W(T_f)$	• The baseline DW model always yields $W(T_f) = 1$. Our adaptive-confidence DW model also yields $W(T_f) = 1$ for the complete graph with $c_0 \geq 0.4$ and the NetScience network with $c_0 \in \{0.8, 0.9\}$. • When $W(T_f) < 1$, for fixed values of the parameters γ, δ, and c_0, decreasing m tends to also decrease $W(T_f)$ for both the complete graph and the NetScience network.

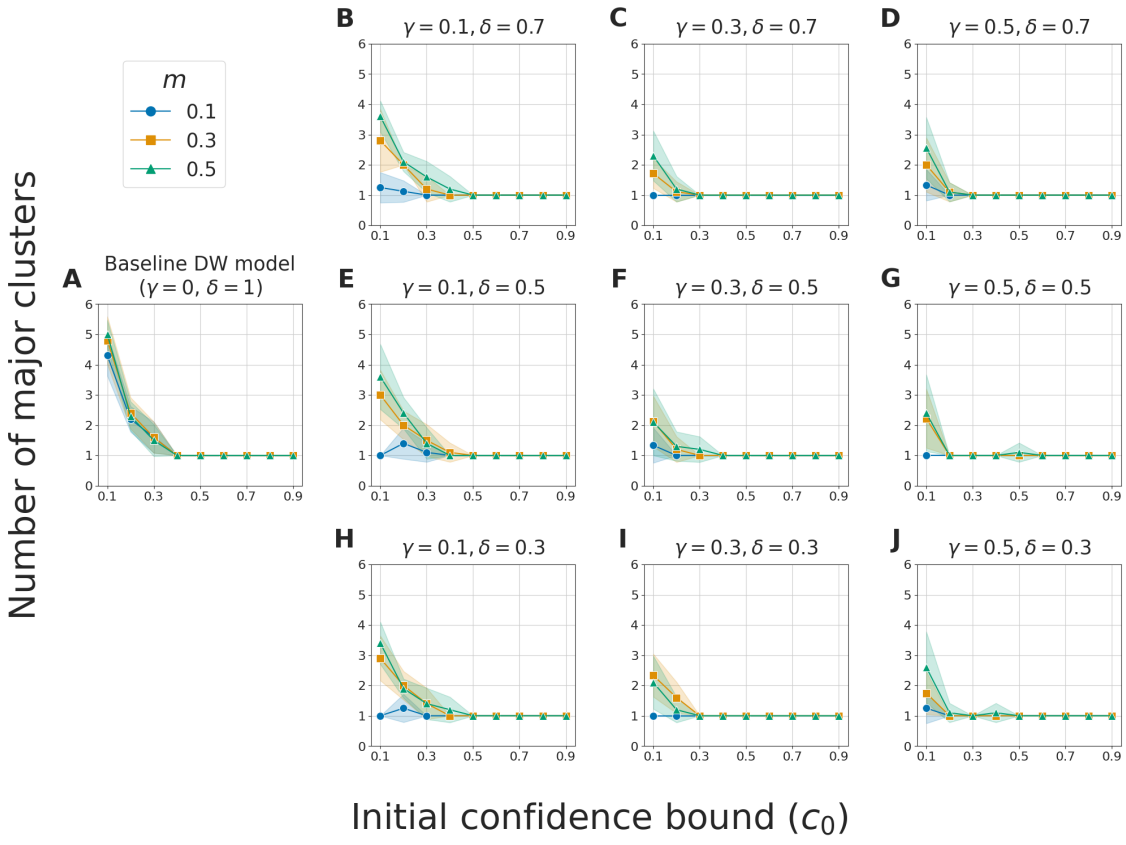

Figure 5.10: The numbers of major clusters in simulations of (A) the baseline DW model and (B–J) our adaptive-confidence DW model on a 100-node complete graph for various combinations of the BCM parameters γ, δ, c_0, and m. In this figure and subsequent figures, we do not use simulations in which we are unable to determine the final opinion clusters (see Table 5.5) to calculate the means and standard deviations. In (E), in which we show our simulations with $(\gamma, \delta) = (0.1, 0.5)$, we run all of our simulations to convergence (i.e., we ignore the bailout time) and use all of our simulations to calculate the mean numbers of major opinion clusters. (This figure originally appeared in [LLP23].)

Table 5.5: Summary of the numbers of simulations of our adaptive-confidence DW model that reach the bailout time of 10^6 time steps. For each combination of the BCM parameters (γ, δ, c_0, and m), we run 10 simulations, which each have a different set of initial opinions. In each table entry, the focal number is the number of simulations that reach the bailout time and the number in parentheses is the number of those simulations for which we are also unable to determine the final opinion clusters. We run our simulations with $(\gamma, \delta) = (0.1, 0.5)$ to convergence (i.e., without a bailout time); for those simulations, we do not track the number of opinion clusters at the bailout time.

		Number of simulations that reach bailout (number of simulations for which we are also unable to determine the final opinion clusters)					
		$m = 0.1$			$m = 0.3$		$m = 0.5$
		$c_0 = 0.1$	$c_0 = 0.2$	$c_0 = 0.3$	$c_0 = 0.1$	$c_0 = 0.2$	$c_0 = 0.1$
$\gamma = 0.1$	$\delta = 0.3$	9 (7)	2 (2)	1 (1)	0	0	0
	$\delta = 0.5$	8	1	0	1	0	0
	$\delta = 0.7$	9 (6)	2 (2)	1 (1)	2 (0)	0	0
$\gamma = 0.3$	$\delta = 0.3$	9 (5)	0	0	2 (1)	0	2 (0)
	$\delta = 0.5$	8 (7)	0	0	2 (2)	0	0
	$\delta = 0.7$	7 (4)	0	0	5 (3)	2 (1)	0
$\gamma = 0.5$	$\delta = 0.3$	9 (6)	0	0	2 (2)	1 (0)	0
	$\delta = 0.5$	8 (4)	0	0	2 (1)	0	0
	$\delta = 0.7$	6 (4)	0	0	7 (4)	0	1 (1)

DW model always reach consensus. For fixed values of γ, δ, and c_0, when our adaptive-confidence DW model does not reach consensus, decreasing the compromise parameter m tends to result in fewer major clusters. By contrast, m has little effect on the number of major clusters in the baseline DW model. Increasing γ with the other BCM parameters (i.e., δ, c_0, and m) fixed also tends to result in fewer major clusters. Changing δ with the other parameters fixed has little effect on the number of major clusters. In fact, changing δ with the other parameters fixed appears to have little effect on any of the computed quantities, so we show results only for $\delta = 0.5$ in our subsequent figures. In our code repository, we include plots for the other examined values of δ.

We observe very few minor clusters in our simulations of our adaptive-confidence DW model on the 100-node complete graph. For each BCM parameter set (γ, δ, c_0, m), the mean

number of minor clusters in our 10 simulations is bounded above by 1. Consequently, the number of major clusters and Shannon entropy follow similar trends. Overall, in our simulations on the 100-node complete graph, our adaptive-confidence DW model encourages more consensus than the baseline DW model and this difference between these two models becomes more pronounced for larger values of the confidence-increase parameter γ and smaller values of the compromise parameter m.

In Figure 5.11, we show the weighted-average edge fraction $W(T_f)$ (see Equation (5.4.1)). The baseline DW model always has $W(T_f) = 1$. By contrast, for sufficiently small initial confidence values c_0, our adaptive-confidence DW model yields $W(T_f) < 1$. For $m = 0.1$ and small c_0 (specifically, $c_0 \leq 0.3$), our adaptive-confidence DW model can reach consensus with $W(T_f) < 1$. As in our adaptive-confidence HK model (see our discussion in Section 5.5.1.2), this observation indicates that some adjacent nodes in the same final opinion cluster are not receptive to each other.

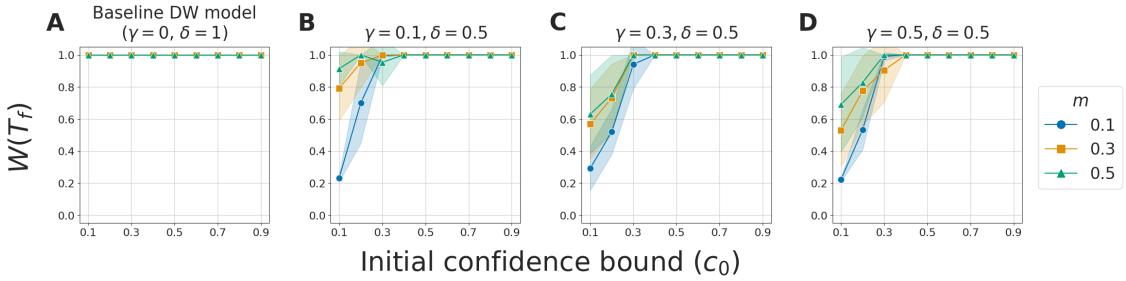

Figure 5.11: The weighted-average edge fraction $W(T_f)$ (see equation Equation (5.4.1)) in simulations of (A) the baseline DW model and (B–D) our adaptive-confidence DW model on a 100-node complete graph for various combinations of the BCM parameters γ, δ, c_0, and m. In (E), in which we show our simulations with $(\gamma, \delta) = (0.1, 0.5)$, we run all of our simulations to convergence (i.e., we ignore the bailout time) and use the resulting final opinion clusters. (This figure originally appeared in [LLP23].)

For fixed values of $c_0 \leq 0.3$ and m, our adaptive-confidence DW model tends to converge more slowly than the baseline DW model. Additionally, when we fix the other BCM parameters (i.e., γ, δ, and m), the convergence time tends to increase as we decrease c_0. As we showed in Table 5.5, for small values of c_0 (specifically, $c_0 \in \{0.1, 0.2\}$), more simulations

reach the bailout time as we decrease m. In both our adaptive-confidence DW model and the baseline DW model, $m = 0.1$ yields longer convergence times than $m \in \{0.3, 0.5\}$ for fixed values of γ, δ, and c_0.

5.5.3 Network of network-scientist coauthorships

We now discuss our simulations of our adaptive-confidence DW model on the NetScience network [New06], which is a network of network scientists with unweighted and undirected edges that encode paper coauthorships.

For the NetScience network and fixed values of c_0 and m, our adaptive-confidence DW model tends to have at least as many major opinion clusters (see Figure 5.12) and minor opinion clusters (see Figure 5.13) as the baseline DW model. In Figure 5.13, we see for $c_0 \leq 0.5$ that both our adaptive-confidence DW model and the baseline DW model yield many more minor clusters for the NetScience network than for the 100-node complete graph. For values of c_0 that are near the transition between consensus and opinion fragmentation (specifically, $c_0 \in \{0.3, 0.4, 0.5\}$), our adaptive-confidence DW model yields noticeably more major clusters and minor clusters than the baseline DW model. The transition between consensus and fragmentation appears to occur for a larger threshold in our adaptive-confidence DW model than in the baseline DW model. For the NetScience network (and unlike for the 100-node complete graph), changing the value of m with the other BCM parameters fixed appears to have little effect on the numbers of major and minor opinion clusters.

For the NetScience network and fixed values of c_0 and m, our adaptive-confidence DW model has convergence times that are similar to those of the baseline DW model. All of our simulations of our adaptive-confidence DW model on the NetScience network converge before reaching the bailout time. We obtain the longest convergence times for $c_0 = 0.3$. By contrast, for the 100-node complete graph, the convergence time increases as we decrease c_0 and many simulations reach the bailout time for $c_0 \in \{0.1, 0.2\}$. In both our adaptive-

confidence DW model and the baseline DW model, $m = 0.1$ yields longer convergence times than $m \in \{0.3, 0.5\}$ for fixed values of γ, δ, and c_0. We do not observe a clear trend in how the convergence time changes either as a function of γ (with fixed δ, c_0, and m) or as a function of δ (with fixed γ, c_0, and m).

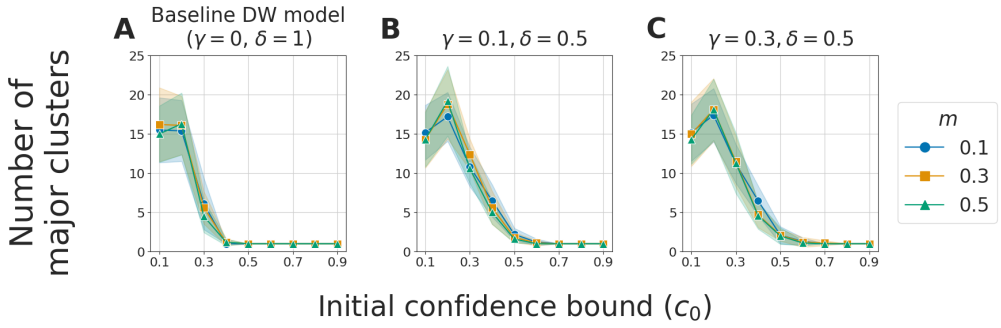

Figure 5.12: The numbers of major clusters in simulations of (A) the baseline DW model and (B, C) our adaptive-confidence DW model on the NetScience network for various combinations of the BCM parameters γ, δ, c_0, and m. (This figure originally appeared in [LLP23].)

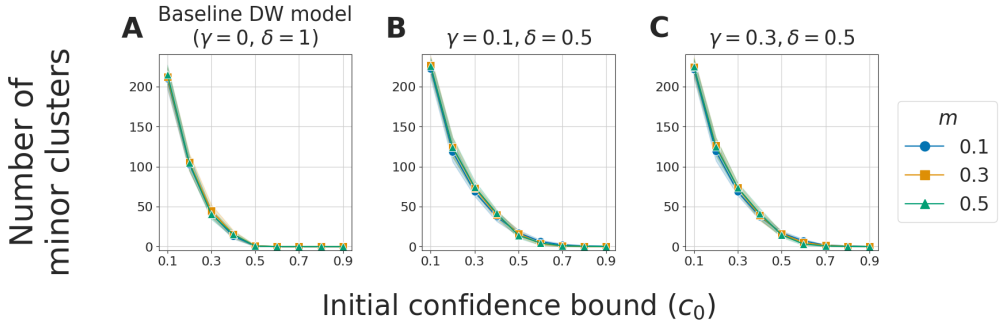

Figure 5.13: The numbers of minor clusters in simulations of (A) the baseline DW model and (B, C) our adaptive-confidence DW model on the NetScience network for various combinations of the BCM parameters γ, δ, c_0, and m. (This figure originally appeared in [LLP23].)

5.6 Conclusions and discussion

5.6.1 Summary and discussion of our results

We developed two BCMs — a synchronously-updating one that generalizes the HK model and an asynchronously-updating one that generalizes the DW model — with adaptive confidence bounds. The confidence bounds in our adaptive-confidence BCMs are distinct for each dyad of a network and change when nodes interact with each other. One can interpret the changes in confidence bounds as changes in receptiveness between nodes. We demonstrated that incorporating time-dependent, adaptive confidence bounds in our BCMs yields a variety of interesting behaviors, such as adjacent nodes that converge to the same limit opinion but are eventually unreceptive to each other.

For our adaptive-confidence BCMs[9] and the baseline BCMs, we proved results about the limiting behaviors of effective graphs, which track which nodes of a network are able to influence each other. We demonstrated using numerical simulations that our BCMs have fewer major opinion clusters and take longer to converge than the associated baseline BCMs. See Table 5.3 for a summary of the trends in our adaptive-confidence HK model, and see Table 5.4 for a summary of the trends in our adaptive-confidence DW model.

The results of our numerical simulations of our adaptive-confidence BCMs complement our theoretical results (which informed the stopping criteria in our computations). For our adaptive-confidence HK model, (1) all dyadic confidence bounds must converge either to 0 or to 1 and (2) the dyadic confidence bounds between nodes in different limit opinion clusters must converge to 0 (see Theorem 5.3.2). For our adaptive-confidence DW model, analogous results hold almost surely (see Theorem 5.3.5). However, in both of our adaptive-confidence BCMs, the dyadic confidence bounds between nodes in the same limit opinion cluster do not necessarily converge to 1, as it is possible for them to instead converge to 0. Indeed, when

[9]The proofs for the adaptive-confidence BCMs are primarily the work of Jiajie Luo.

the confidence-decrease parameter $\delta < 1$, our numerical simulations of both of our adaptive-confidence BCMs demonstrate for a wide range of the other BCM parameter values that some dyads in the same final opinion cluster have confidence bounds that converge to 0. Although the nodes in these dyads are unreceptive to each other, they still converge to the same opinion. The nodes in these dyads do not have an edge between them in the final effective graph, so the final opinion clusters (i.e., the connected components of the final effective graph) in our BCMs can have a richer structure than those in the baseline BCMs.

5.6.2 Future work

Our investigation lays groundwork and provides a point of comparison for the study of more complicated adaptive-confidence mechanisms in BCMs. Future investigations of adaptive-confidence BCMs include establishing additional theoretical guarantees, examining and validating such BCMs in sociological contexts, and generalizing these models in various ways.

It is worthwhile to analytically and numerically study the mutual receptiveness of nodes in our BCMs when they reach a consensus state. In our numerical simulations of our adaptive-confidence BCMs, when the confidence-decrease parameter $\delta < 1$, some adjacent nodes in the same final opinion cluster are eventually not receptive to each other. More specifically, our numerical simulations suggest that some adjacent nodes can converge to the same limit opinion without having an edge between them in the limit effective graph. One can explore this behavior of our BCMs and determine how the model parameters influence the existence of edges between adjacent nodes with the same limit opinion in limit effective graphs.

It is also important to consider how the behaviors of our BCMs connect to real-life social situations. One can interpret the opinion values in our models as representing outwardly expressed opinions, which may differ from internally held beliefs [Kur95]. The achievement of a "consensus" can represent agents arriving at the same outwardly expressed behavior or decision, rather than achieving an actual agreement of their internal values [Hor62]. Researchers have studied models with both internal and expressed opinions [Noo20, CLY20, HLJ21], and

one can incorporate such considerations into adaptive-confidence BCMs.

In our adaptive-confidence BCMs, adjacent agents that are unreceptive to each other's opinions can still interact with each other. Alternatively, a pair of agents can eventually stop interacting with each other — effectively changing the network structure — after repeated negative interactions. Researchers have modeled such ideas, along with network restructuring to consider new social interactions, using adaptive networks with edge rewiring [PRM22, KFP23]. A possible area of further study is the investigation of which models effectively have "mediator" nodes that assist in bringing together the opinions of agents that are unreceptive to each other or no longer interact. If there are such mediator nodes, one can examine whether or not they share common characteristics or are identifiable from network structure and initial agent opinions.

There are many possible areas to explore in the study of adaptive opinion models. In research on opinion dynamics, it is important to incorporate network adaptivity, which provides fertile ground for theoretical, computational, and empirical investigations of opinion dynamics.

CHAPTER 6

Bounded-Confidence Models with Topic-Weighted Opinion Discordance

In this chapter, we study BCMs that have multi-dimensional opinions that consist of multiple interdependent topics. This project is ongoing work in collaboration with Jiajie Luo and Weiqi Chu.[1]

6.1 Introduction and motivation

In this chapter, we generalize the DW and HK models to study agents that have vector-valued opinions on interdependent topics. Each agent in the standard BCMs (see Section 3.1 and Section 3.2) has a single scalar-valued opinion. In real life, people have opinions on topics that are related and have some interdependency. We generalize the standard BCMs to model this situation.

Political scientist Philip Converse [Con06] used the term "belief system" to describe attitudes that are bound together by some form of constraint or interdependence. Empirically, researchers have observed correlations between political positions [BL06, BL12]. In models of opinion dynamics, agents can have vector-valued opinions that encode either a combination of different issues or a single issue with multiple coordinates [Fri15]. When developing

[1]The idea for the BCMs in this project was developed primarily by Jiajie Luo. My main contributions to this project so far are deriving the region of absorption (see Section 6.3.2) and running and interpreting the numerical simulations (see Sections 6.4 and 6.5).

a discrete-time BCM with opinions in two or more dimensions, one must make choices that are not present in a model with 1D opinions. In particular, one must decide which opinion topic(s) to update at each time step and how to measure a notion of distance between opinion vectors.

Researchers have studied variants of the DW [Lor06, LX17] and HK [FLP05, Lor06, BBC13, EBN13, HK19, BP20] models with multidimensional opinions. In these papers, the authors calculated the distances between two opinion vectors using the Euclidean norm of their differences.[2] In political science, researchers often use Euclidean distance in models for mathematical convenience and tractability, but there is little empirical evidence that it is suitable for modeling opinions [BL06]. Some researchers have generalized multidimensional BCMs to have non-Euclidean distances. For example, Fortunato et al. [FLP05] studied a continuum HK model with 2-dimensional opinions in which agents are receptive to each other if and only if their opinion difference in both topics is within a confidence bound c. That is, agents with opinions $\boldsymbol{x}$ and $\boldsymbol{y}$ are receptive to each other if and only if both $|x^1 - y^1| < c$ and $|x^2 - y^2| < c$, where x^i denotes the ith entry of $\boldsymbol{x}$. Schweighofer et al. [SGS20] developed and simulated multidimensional BCMs that have a "directional" distance that measures the angle between two opinion vectors. In their models, when agents compromise their opinions, their opinion vectors rotate towards each other in opinion space.

Researchers have incorporated topic correlation or coupling parameters into models of opinion dynamics with interdependent topics. Parsegov et al. [PPT15] studied variants of the Friedkin–Johnson model, a linear model of opinion dynamics with agents that "stubbornly" factor their initial opinion into every update, with multidimensional opinions with interdependent topics. Their models have a matrix parameter C in which each entry C_{pq} determines how much topic q affects topic p. Ye et al. [YTL20] subsequently examined a continuous-time model inspired by the models in [PPT15]. Baumann et al. [BLS21] studied

[2]For two K-dimensional opinion vectors $\boldsymbol{x}$ and $\boldsymbol{y}$, the Euclidean norm of their difference is $\|\boldsymbol{x} - \boldsymbol{y}\|_2 = \sqrt{\sum_{i=1}^{K} (x^i - y^i)^2}$, where x^i and y^i denote the ith entries of $\boldsymbol{x}$ and $\boldsymbol{y}$, respectively.

a model of opinion dynamics with a "topic overlap" matrix parameter Φ. Each entry Φ_{pq} encodes the angles between topics p and q in a latent topic space and determines how much the opinions of agents on topic p affect the opinions of agents on topic q. Chen et al. [CWY21] developed and simulated a rather complex agent-based model with multidimensional opinions that involves time-dependent correlations between topics. At each time, a selected agent first updates its opinion using rules that involve the topic correlations and external interventions. It then selects another agent to interact with and updates its opinion using a modified Jager–Amblard update rule [JA05], which is a DW update (see Equation (3.1.1)) if the Euclidean distance between opinions is sufficiently small and a repulsive update if the Euclidean distance is sufficiently large.

In the present chapter, we study BCMs with multidimensional opinions and non-Euclidean distances between opinion vectors. Our topic-weighted "discordance" functions account for the magnitudes of the opinion differences in each topic. When determining whether or not two agents are receptive to each other on topic k, depending on the value of a topic-weight parameter, our topic-weighted discordance puts at least as much importance on their opinion difference on topic k than on their opinion difference on each other topic. In our models, we do not have correlation parameters that precisely control how much each pair of topics affects each other. However, inspired by [PPT15, YTL20, CWY21], this is a viable way to extend our models.

6.2 BCMs with topic-weighted opinion discordance

We now describe our BCMs with K-dimensional opinions and topic-weighted opinion distance. The idea for these models was developed primarily by Jiajie Luo.

Consider a time-independent, unweighted, undirected network $G = (V, E)$ with no self-edges or multi-edges, where V is a set of nodes and E is a set of edges between them. The nodes of the network represent agents with opinions, each of which lie in the closed interval

$[0,1]$, on K related topics. Let $\mathbf{x}_i(t) \in \mathbb{R}^K$ denote the opinion vector of node i at time t; its kth entry, which we denote by $x_i^k(t)$, is the opinion of node i on topic k.

We describe our variants of the HK model and the DW model in Section 6.2.1 and Section 6.2.2, respectively. Both of our models have a confidence bound $c \in [0,1]$ that is analogous to the confidence bound in the standard HK model (see Section 3.2) and the standard DW model (see Section 3.1).

We define topic-weighted distance functions $d_k : [0,1]^K \times [0,1]^K \mapsto [0,1]$ that we call *discordance functions*.[3] For each topic k, the discordance function d_k takes two opinion vectors as inputs and calculates a distance between them that emphasizes topic k while accounting for the other topics. For topic k, the discordance between opinion vectors $\mathbf{x}_i$ and $\mathbf{x}_j$ is

$$d_k(\mathbf{x}_i, \mathbf{x}_j; \omega) = \omega |x_i^k - x_j^k| + \frac{1-\omega}{K} \sum_{\ell=1}^{K} |x_i^\ell - x_j^\ell| \,, \tag{6.2.1}$$

where the *topic weight* $\omega \in [0,1]$ is a parameter that determines how much weight is placed on the opinion difference in topic k. In our models, when nodes i and j have opinions that satisfy $d_k(\mathbf{x}_i, \mathbf{x}_j) < c$, we say that nodes i and j are *receptive* to each other on topic k.

The *maximum discordance function* $d_{\max}$ is

$$d_{\max}(\mathbf{x}_i, \mathbf{x}_j; \omega) = \max_{1 \le k \le K} d_k(\mathbf{x}_i, \mathbf{x}_j; \omega) \,. \tag{6.2.2}$$

When nodes i and j have opinions that satisfy $d_{\max}(\mathbf{x}_i, \mathbf{x}_j) < c$, nodes i and j are receptive to each other on all topics.

6.2.1 Our HK model with topic-weighted opinion discordance

We now describe our HK model with topic-weighted opinion discordance. We refer to this model as our *topic-weighted HK model*. Our topic-weighted HK model has two parameters:

[3] Our choice to call these functions "discordance" functions was inspired by Hickok et al. [HKB22].

the confidence bound $c \in [0, 1]$ and the topic weight $\omega \in [0, 1]$.

The set[4] of nodes that a node i is receptive to on topic k at time t is

$$\mathcal{N}_i^k(t) = \{i\} \cup \{j : (i, j) \in E \text{ and } d_k(\mathbf{x}_i(t), \mathbf{x}_j(t)) < c\} \ . \tag{6.2.3}$$

At each discrete time, we update the opinions of all nodes on all topics by calculating

$$x_i^k(t+1) = \frac{1}{|\mathcal{N}_i^k(t)|} \sum_{j \in \mathcal{N}_i^k(t)} x_j^k(t) \tag{6.2.4}$$

for all nodes i and all topics k.

When the topic weight is $\omega = 1$, Equation (6.2.1) reduces to $d_k(\mathbf{x_i}, \mathbf{y_i}; \omega = 1) = |x_i^k - x_j^k|$, so the opinion discordance for topic k has no dependence on other topics. Therefore, our topic-weighted HK model with $\omega = 1$ is equivalent to K independent copies of the standard HK model (see Section 3.2), with each copy encoding the opinions for one topic. We refer to our topic-weighted HK model with $\omega = 1$ as our *baseline HK model*.

6.2.2 Our DW model with topic-weighted opinion discordance

We now describe our DW model with topic-weighted opinion discordance. We refer to this model as our *topic-weighted DW model*. Our topic-weighted DW model has three parameters: the confidence bound $c \in [0, 1]$, the compromise parameter $m \in (0, 0.5]$, and the topic weight $\omega \in [0, 1]$. The confidence bound and compromise parameter in our model are analogous to those BCM parameters in the standard DW model (see Section 3.1).

At each discrete time, we select an edge $(i, j) \in E$ uniformly at random and a topic $k \in \{1, \ldots, K\}$ uniformly at random. We then update the opinions of nodes i and j on topic

[4]The set $\mathcal{N}_i^k(t)$ is similar to the set $I(i, x(t))$ in Equation (3.2.1) for the standard HK model. Unlike in the standard HK model, in this chapter, opinions are K-dimensional, so we use the notation $\mathcal{N}_i^k(t)$ for clarity.

k by calculating

$$x_i^k(t+1) = \begin{cases} x_i^k(t) + m(x_j^k(t) - x_i^k(t))\,, & \text{if } d_k(x_i(t), x_j(t)) < c \\ x_i^k(t)\,, & \text{otherwise}\,, \end{cases}$$
$$x_j^k(t+1) = \begin{cases} x_j^k(t) + m(x_i^k(t) - x_j^k(t))\,, & \text{if } d_k(x_i(t), x_j(t)) < c \\ x_j^k(t)\,, & \text{otherwise}\,. \end{cases} \tag{6.2.5}$$

We refer to our topic-weighted DW model with $\omega = 1$ as our *baseline DW model*.[5]

6.3 Preliminary theoretical results

In this section, we discuss some of our preliminary theoretical results for our topic-weighted BCMs. In Section 6.3.1, we show that the opinions in our topic-weighted BCMs converge to limit opinions. In Section 6.3.2, we derive what we call the "region of absorption". Nodes i and j with opinion vectors $\boldsymbol{x}$ and $\boldsymbol{y}$, respectively, are receptive to each other on all topics if and only if $\boldsymbol{y}$ and $\boldsymbol{x}$ are in each other's regions of absorption. We also have results on which limit opinion states are possible and some properties of the limit effective graphs in our topic-weighted BCMs. The derivation of these results are predominantly the work of Weiqi Chu; we do not include them in this dissertation.

6.3.1 Convergence of opinions

The theorems and proofs in this subsection are joint work with Weiqi Chu and Jiajie Luo. In this subsection, we use Theorem 3.3.1 to prove that the opinions in our topic-weighted HK model (see Section 6.2.1) and topic-weighted DW model (see Section 6.2.2) converge. Let

[5]Due to the stochasticity in selecting a topic at each time in our topic-weighted DW model, our baseline DW is not equivalent to having independent copies of the standard DW model (see Section 3.1). By contrast, our topic-weighted HK model is deterministic; as we discussed at the end of Section 6.2.1, it is equivalent to K independent copies of the standard HK model.

$\mathbf{x}(t) \in [0,1]^{NK}$ denote a vector of the opinions of all nodes on all topics. We construct $\mathbf{x}(t)$ by concatenating the opinion vectors for each node (i.e., by concatenating $\mathbf{x}_1(t), \mathbf{x}_2(t), \ldots \mathbf{x}_N(t)$).

Theorem 6.3.1. *The limit* $\mathbf{x}^* = \lim_{t\to\infty} \mathbf{x}(t)$ *for our topic-weighted HK model (with update rule* (6.2.4)*) exists.*

Proof. Consider each pair of nodes i and j and each topic k. We express the update rule (6.2.4) as $\mathbf{x}(t+1) = A(\mathbf{x}(t), t)\mathbf{x}(t)$, where

$$A(\mathbf{x}(t),t)_{\alpha\beta} = \begin{cases} 1/|\mathcal{N}_i^k(t)|\,, & \text{if } \alpha = (i-1)K + k\,,\ \beta = (j-1)K + k\,, \text{ and } j \in \mathcal{N}_i^k(t) \text{ for some nodes } i \text{ and } j \text{ and some topic } k \\ 0\,, & \text{otherwise}\,. \end{cases} \tag{6.3.1}$$

It is readily checked that each $A(\mathbf{x}(t), t)$ satisfies conditions (1)–(3) in Theorem 3.3.1, which implies that $\mathbf{x}(t)$ converges to some limit opinion state $\mathbf{x}^*$. □

Theorem 6.3.2. *The limit* $\mathbf{x}^* = \lim_{t\to\infty} \mathbf{x}(t)$ *for our topic-weighted DW model (with update rule* (6.2.5)*) exists.*

Proof. We express the update rule (6.2.5) as $\mathbf{x}(t+1) = A(\mathbf{x}(t), t)\mathbf{x}(t)$. Suppose that, at time t, we select edge (i, j) and topic k to update by Equation (6.2.5). If the discordance function $d_k(\mathbf{x}_i(t), \mathbf{x}_j(t)) \geq c$, then $A(\mathbf{x}(t), t) = I_{NK}$ is the identity matrix. Otherwise, if $d_k(\mathbf{x}_i(t), \mathbf{x}_j(t)) < c$, then

$$A(\mathbf{x}(t),t)_{\alpha\beta} = \begin{cases} 1-m\,, & \text{if } \alpha = \beta = (i-1)K + k \\ 1-m\,, & \text{if } \alpha = \beta = (j-1)K + k \\ m\,, & \text{if } \alpha = (i-1)K + k\,,\ \ \beta = (j-1)K + k \\ m\,, & \text{if } \alpha = (j-1)K + k\,,\ \ \beta = (i-1)K + k \\ \delta_{\alpha\beta}\,, & \text{otherwise}\,, \end{cases} \tag{6.3.2}$$

where $\delta_{\alpha\beta}$ is the Kronecker delta function (i.e., $\delta_{\alpha\beta} = 1$ if $\alpha = \beta$ and $\delta_{\alpha\beta} = 0$ otherwise). It is readily checked that each $A(\mathbf{x}(t), t)$ satisfies conditions (1)–(3) in Theorem 3.3.1, which implies that $\mathbf{x}(t)$ converges to some limit opinion state $\mathbf{x}^*$. □

Theorems 6.3.1 and 6.3.2 give us the existence of a *limit opinion vector* $x_i^* = \lim_{t\to\infty} \mathbf{x}_i$ for each node i in our topic-weighted BCMs. For each distinct limit opinion vector, we say that the set of nodes that have that limit opinion vector is a *limit opinion cluster* (see Section 3.3.2).

6.3.2 Region of absorption

Fix the opinion vector $\boldsymbol{x}$. We want to find the region in opinion space of opinions $\boldsymbol{y}$ such that nodes with opinions $\boldsymbol{x}$ and $\boldsymbol{y}$ are receptive to each other on all topics. We call this the *region of absorption* for opinion $\boldsymbol{x}$. That is, we want to find the region of opinions

$$\mathcal{R}_{\boldsymbol{x}} = \{\boldsymbol{y} \in \mathbb{R}^K : d_{\max}(\boldsymbol{x}, \boldsymbol{y}) < c\}\,. \tag{6.3.3}$$

Theorem 6.3.3. *For our topic-weighted BCMs (with update rules* (6.2.4) *and* (6.2.5)*), the region of absorption* $\mathcal{R}_{\boldsymbol{x}}$ *(see Equation* (6.3.3)*) for an opinion vector* $\boldsymbol{x} \in \mathbb{R}^K$ *is a* K*-dimensional polytope. Furthermore, for* $s \in \{1, 2, \dots K\}$*, let*

$$\mathcal{U}_s = \left\{ \boldsymbol{u} \in \mathbb{R}^K : s \text{ entries of } \boldsymbol{u} \text{ are } \frac{c}{\omega + \frac{s}{K}(1-\omega)} \text{ and the remaining } K - s \text{ entries are } 0 \right\}.$$

The set of vertices $\hat{\mathcal{R}}_{\boldsymbol{x}}$ *for the polytope* $\mathcal{R}_{\boldsymbol{x}}$ *is*

$$\hat{\mathcal{R}}_{\boldsymbol{x}} = \left\{ \boldsymbol{y} \in \mathbb{R}^K : |\boldsymbol{x} - \boldsymbol{y}| \in \mathcal{U}_s \text{ for some } s \in \{1, 2, \dots K\} \right\}, \tag{6.3.4}$$

where $|\boldsymbol{x} - \boldsymbol{y}|$ *is the vector with* k*th entry* $|x^k - y^k|$*.*

Proof. Let $\boldsymbol{d}(\boldsymbol{x}, \boldsymbol{y}) \in \mathbb{R}^K$ denote the vector with $d_k(\boldsymbol{x}, \boldsymbol{y})$ in the kth entry. Let $\boldsymbol{u}(\boldsymbol{x}, \boldsymbol{y}) =$

$|\boldsymbol{x}-\boldsymbol{y}|$ denote an opinion-difference vector with $|x^k - y^k|$ in the kth entry. We have that

$$\boldsymbol{d}(\boldsymbol{x},\boldsymbol{y}) = \left(\omega I_K + \frac{1}{K}(1-\omega)\mathbf{1}\mathbf{1}^T\right)\boldsymbol{u}(\boldsymbol{x},\boldsymbol{y})\,, \tag{6.3.5}$$

where I_K is the $K \times K$ identity matrix. Let $B = \omega I + \frac{1}{K}(1-\omega)\mathbf{1}\mathbf{1}^T \in \mathbb{R}^{K\times K}$. That is,

$$B_{ij} = \begin{cases} \omega + \frac{1}{K}(1-\omega)\,, & \text{if } i = j \\ \frac{1}{K}(1-\omega)\,, & \text{otherwise}\,. \end{cases}$$

Fix the opinion vector $\boldsymbol{x}$. An opinion vector $\boldsymbol{y} \in \mathcal{R}_{\boldsymbol{x}}$ if and only if $d_{\max}(\boldsymbol{x},\boldsymbol{y}) < c$. By Equation (6.3.5), $d_{\max}(\boldsymbol{x},\boldsymbol{y}) < c$ if and only if $B\boldsymbol{u}(\boldsymbol{x},\boldsymbol{y}) < c\mathbf{1}$, where the vector inequality $B\boldsymbol{u}(\boldsymbol{x},\boldsymbol{y}) < c\mathbf{1}$ signifies that each entry of the vector $B\boldsymbol{u}(\boldsymbol{x},\boldsymbol{y})$ is less than c. Therefore, the region of absorption $\mathcal{R}_{\boldsymbol{x}}$ is a K-dimensional polytope containing all opinion vectors $\boldsymbol{y}$ that satisfy $B\boldsymbol{u}(\boldsymbol{x},\boldsymbol{y}) < c\mathbf{1}$. The boundary of $\mathcal{R}_{\boldsymbol{x}}$ consists of all opinion vectors $\boldsymbol{y}$ that satisfy $B\boldsymbol{u}(\boldsymbol{x},\boldsymbol{y}) = c\mathbf{1}$.

We now calculate the vertices of the polytope $\mathcal{R}_{\boldsymbol{x}}$. Suppose that $\boldsymbol{y} \in \mathcal{R}_{\boldsymbol{x}}$, so $d_{\max}(\boldsymbol{x},\boldsymbol{y}) < c$. Intuitively, if opinions $\boldsymbol{x}$ and $\boldsymbol{y}$ agree exactly on some topic k (i.e., $x^k - y^k = 0$), then the vectors $\boldsymbol{x}$ and $\boldsymbol{y}$ can have a larger difference in opinion in some other topic $\ell \neq k$ (i.e., the ℓth entry of $\boldsymbol{u}(\boldsymbol{x},\boldsymbol{y})$ can be larger) than if $x^k \neq y^k$. If $\boldsymbol{y}$ is a vertex of $\mathcal{R}_{\boldsymbol{x}}$, then $\boldsymbol{u}(\boldsymbol{x},\boldsymbol{y})$ can be 0 in as many as $K-1$ entries. Let $s \in \{1,2,\ldots,K\}$ be the number of nonzero entries of $\boldsymbol{u}(\boldsymbol{x},\boldsymbol{y})$, and let $\boldsymbol{u}_s \in \mathbb{R}^s$ denote a vector of the nonzero opinion differences of $\boldsymbol{x}$ and $\boldsymbol{y}$ in s topics. Each entry of $\boldsymbol{u}_s$ is $|x^k - y^k|$ for some topic k, and no topics repeat. Let $B_s \in \mathbb{R}^{s\times s}$ be the square matrix of the first s rows and columns of B. To find the vertices that define the polytope $\mathcal{R}_{\boldsymbol{x}}$, we want to find $\boldsymbol{u}_s$ that solves the matrix equation

$$B_s\boldsymbol{u}_s = c\mathbf{1} \tag{6.3.6}$$

for each $s \in \{1,2,\ldots,K\}$.

Fix $s \in \{1, 2, \ldots, K\}$. We calculate $B_{\boldsymbol{s}}^{-1}$ by using the Sherman–Morrison formula [Bar51, Hag89]. Let A be an invertible square matrix, and let $\boldsymbol{v}$ and $\boldsymbol{w}$ be column vectors. The Shermman–Morrison formula states that the matrix $A + \boldsymbol{v}\boldsymbol{w}^T$ is invertible if and only if $1 + \boldsymbol{w}^T A^{-1}\boldsymbol{v} \neq 0$. Furthermore, if $A + \boldsymbol{v}\boldsymbol{w}^T$ is invertible, then its inverse is $(A + \boldsymbol{v}\boldsymbol{w}^T)^{-1} = A^{-1} - \frac{A^{-1}\boldsymbol{v}\boldsymbol{w}^T A^{-1}}{1+\boldsymbol{w}^T A^{-1}\boldsymbol{v}}$.

Let $A = \omega I \in \mathbb{R}^{s\times s}$, which is invertible because $\omega \neq 0$. Let $\boldsymbol{v} = \frac{1}{K}(1-\omega)\mathbf{1} \in \mathbb{R}^s$ and $\boldsymbol{w} = \mathbf{1} \in \mathbb{R}^s$. We have that $1 + \boldsymbol{w}^T A^{-1}\boldsymbol{v} = 1/\omega \neq 0$. Therefore, by applying the Sherman–Morrisson formula, B_s is invertible and its inverse is

$$B_s^{-1} = (\omega I)^{-1} - \frac{(\omega I)^{-1}\left(\frac{1}{K}(1-\omega)\mathbf{1}\right)\mathbf{1}^T(\omega I)^{-1}}{1 + \mathbf{1}^T(\omega I)^{-1}\mathbf{1}} = \frac{1}{\omega}I - \frac{\frac{1-\omega}{\omega^2}\frac{1}{K}\mathbf{1}\mathbf{1}^T}{1 + s\frac{1}{K}\frac{1}{\omega}(1-\omega)}\,. \tag{6.3.7}$$

From Equation (6.3.6), $\boldsymbol{u}_s = B_s^{-1}c\mathbf{1}$. Using the inverse $B_{\boldsymbol{s}}^{-1}$ in Equation (6.3.7), the ith entry of $\boldsymbol{u}_s$ is

$$\begin{aligned} u_s^i = c\sum_{j=1}^{s}[B_s^{-1}]_{ij} &= c\left(\frac{1}{\omega} - s\left[\frac{\frac{1-\omega}{\omega^2}\frac{1}{K}}{1 + s\frac{1}{K}\frac{1}{\omega}(1-\omega)}\right]\right) \\ &= \frac{c}{\omega}\left(1 - \frac{\frac{s}{K}\frac{1}{\omega}(1-\omega)}{1 + \frac{s}{K}\frac{1}{\omega}(1-\omega)}\right) \\ &= \frac{c}{\omega}\left(\frac{1}{1 + \frac{s}{K}\frac{1}{\omega}(1-\omega)}\right) \\ &= \frac{c}{\omega + \frac{s}{K}(1-\omega)}\,. \end{aligned} \tag{6.3.8}$$

Therefore, the vertices of the polytope $\mathcal{R}_{\boldsymbol{x}}$ are the coordinates $\boldsymbol{y} \in \mathbb{R}^K$ that have s entries of $\boldsymbol{u}(\boldsymbol{x}, \boldsymbol{y}) = |\boldsymbol{x} - \boldsymbol{y}|$ that are equal to $\frac{c}{\omega + \frac{s}{K}(1-\omega)}$ and the remaining $K - s$ entries of $\boldsymbol{u}(\boldsymbol{x}, \boldsymbol{y})$ equal to 0.

□

We now provide some examples of what the polytope $\mathcal{R}_{\boldsymbol{x}}$ looks like. Suppose that there are $K = 2$ topics. When $s = 2$, we have $\frac{c}{\omega + \frac{s}{K}(1-\omega)} = c$, so $|\boldsymbol{x} - \boldsymbol{y}| = (c, c)$. This gives four

vertices at $\boldsymbol{x} + (a, b)$, where $a, b \in \{-c, c\}$. When $s = 1$, we have $\frac{c}{\omega + \frac{s}{K}(1-\omega)} = \frac{2c}{\omega+1}$. This gives four vertices at $\boldsymbol{x} \pm (0, \frac{2c}{\omega+1})$ and $\boldsymbol{x} \pm (\frac{2c}{\omega+1}, 0)$. When $K = 2$, the region of absorption for $\boldsymbol{x}$ consists of a polygon with eight vertices. The vertices of $\mathcal{R}_{\boldsymbol{x}}$ in this example are

$$\hat{\mathcal{R}}_{\boldsymbol{x}} = \left\{ \boldsymbol{x} + \mathbf{z} : \boldsymbol{z} \in \left\{ \pm(c, c), \pm(c, -c), \pm\left(0, \frac{2c}{\omega + 1}\right), \pm\left(\frac{2c}{\omega + 1}, 0\right) \right\} \right\}.$$

In Figure 6.1, we show a visualization of the boundary of the region of absorption in 2D for a fixed value of c and various values of ω. When $\omega = 1$ (i.e., as in the baseline BCMs), the region of absorption is a square. As we decrease ω, the area of the region of absorption increases; a node has a larger region of opinion space that a neighbor can be in for them to be mutually receptive on a particular topic.

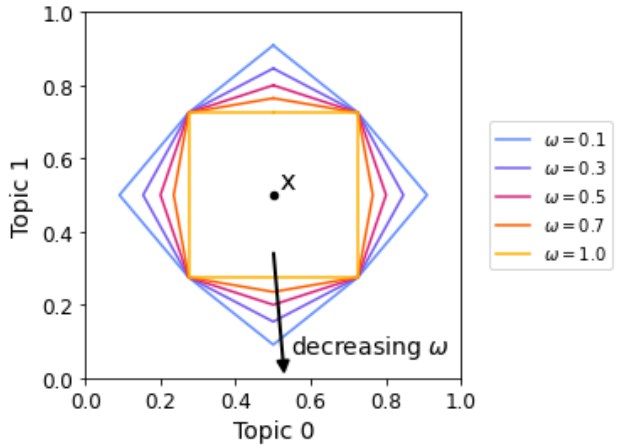

Figure 6.1: Boundaries of the region of absorption for the opinion vector $(0.5, 0.5)$ for $c = 0.225$ and various values of ω. Each colored polygon represents the region of absorption for a different value of ω. As we decrease ω, the area of the region of absorption increases.

As another example, now suppose that there are $K = 3$ topics. The region of absorption for the opinion vector $\boldsymbol{x}$ consists of a polygon with 26 vertices of the form $\boldsymbol{x} + \mathbf{z}$. For $s = 3$, there are eight vertices in which each entry of $\boldsymbol{z}$ is $\pm c$. For $s = 2$, we get 12 vertices in which one entry of z is 0 and the remaining two entries are $\pm\frac{3c}{\omega+2}$. Finally, for $s = 1$, there are six vertices in which two entries of z are 0 and the remaining entry is $\pm\frac{3c}{2\omega+1}$. The vertices of

$\mathcal{R}_{\boldsymbol{x}}$ in this example are

$$\hat{\mathcal{R}}_{\boldsymbol{x}} = \Bigg\{ \boldsymbol{x} + \mathbf{z} : \boldsymbol{z} \in \Bigg\{ \pm (c,c,c), \pm(-c,c,c), \pm(c,-c,c), \pm(c,c,-c), \\ \pm \left(0, \frac{3c}{\omega+2}, \frac{3c}{\omega+2}\right), \pm \left(0, \frac{3c}{\omega+2}, -\frac{3c}{\omega+2}\right), \\ \pm \left(\frac{3c}{\omega+2}, 0, \frac{3c}{\omega+2}\right), \pm \left(\frac{3c}{\omega+2}, 0, -\frac{3c}{\omega+2}\right), \\ \pm \left(\frac{3c}{\omega+2}, \frac{3c}{\omega+2}, 0\right), \pm \left(\frac{3c}{\omega+2}, -\frac{3c}{\omega+2}, 0\right), \\ \pm \left(\frac{3c}{2\omega+1}, 0, 0\right), \pm \left(0, \frac{3c}{2\omega+1}, 0\right), \pm \left(0, 0, \frac{3c}{2\omega+1}\right) \Bigg\} \Bigg\}.$$

6.4 Details of our numerical simulations

In this section, we discuss the setup of our numerical simulations of our topic-weighted BCMs. For our simulations, we consider opinions with $K = 2$ topics.

6.4.1 Initial opinion distributions

In this subsection, we describe the initial opinion distributions that we use in our simulations. In studies on BCMs, it is relatively uncommon to consider initial opinions that arise from distributions other than uniform distributions (see Section 3.4). In our topic-weighted BCMs, whether or not two nodes are receptive to each other on a particular topic depends on their difference in opinions on other topics. We simulate our models with various initial opinion distributions to examine the impact of introducing different dependencies of the initial opinions on the two topics.

To simulate our topic-weighted BCMs with initial opinions that have topics that are independent of each other, we examine an "independent uniform distribution". For this initial opinion distribution, we draw initial opinions for topic 1 uniformly at random from $[0, 1]$. We independently draw initial opinions for topic 2 uniformly at random from $[0, 1]$.

This initial opinion distribution is an extension to two opinion topics of the initial opinion distribution for the standard DW and HK models (see Sections 3.1 and 3.2).

To investigate the behaviors of our topic-weighted BCMs when topic 2 depends completely on topic 1, we consider what we call a "wedge distribution" of initial opinions. In this wedge distribution, we draw initial opinions for topic 1 uniformly at random from $[0, 1]$. Suppose that we select opinion x_i^1 on topic 1. We then determine the opinion x_i^2 on topic 2 by calculating

$$x_i^2 = \begin{cases} 2x_i^1\,, & \text{if } 0 \leq x_i^1 \leq 0.5 \\ 2 - 2x_i^1\,, & 0.5 < x_i^1 \leq 1\,. \end{cases} \tag{6.4.1}$$

In Figure 6.2, we show the function in Equation (6.4.1). We refer to this distribution as a wedge distribution due to the shape of the function in Figure 6.2. For our wedge distribution, the marginal probability distribution for either topic 1 or topic 2 is a uniform distribution on $[0, 1]$.

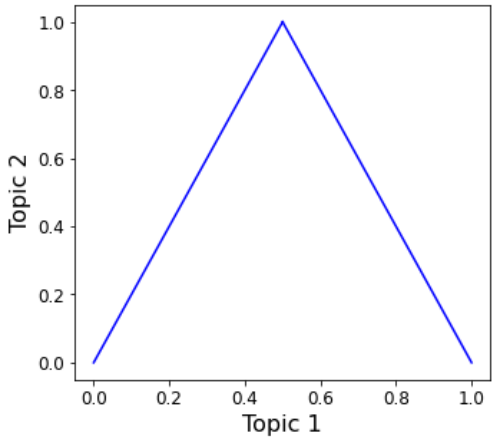

Figure 6.2: Opinion on topic 2 given the opinion on topic 1 in our wedge distribution.

To examine our topic-weighted BCMs on distributions for which we can easily control the correlation between the topics, we consider Gaussian distributions. Let Gaussian(σ, ρ)

denote a 2D Gaussian distribution with mean $(0.5, 0.5)$ and probability density function

$$f(x^1, x^2) = \frac{1}{2\pi\sigma^2\sqrt{1-\rho^2}} \times \exp\left(-\frac{1}{2\sigma^2(1-\rho^2)}\left[(x^1-0.5)^2 - 2\rho(x^1-0.5)(x^2-0.5) + (x^2-0.5)^2\right]\right), \tag{6.4.2}$$

where x^1 and x^2 are the opinions on topics 1 and 2, respectively, σ is the standard deviation of the opinions on topics 1 and 2, and ρ is the Pearson correlation between the opinions on topics 1 and 2.

For our numerical simulations, we consider a Gaussian($\sigma = 0.22, \rho = 0$) initial opinion distribution, which is a symmetric Gaussian distribution with no correlation between topics 1 and 2. We also consider the skewed Gaussian distribution Gaussian($\sigma = 0.22, \rho = 0.8$), which incorporates a correlation between topics 1 and 2. For $\sigma = 0.22$, we expect that about 95% of samples drawn from either of our Gaussian distributions will lie in our opinion space $[0, 1] \times [0, 1]$. In our numerical simulations, when we draw initial opinions from a Gaussian($\sigma = 0.22, \rho = 0$) or Gaussian($\sigma = 0.22, \rho = 0.8$) distribution, if the initial opinion does not lie in $[0, 1] \times [0, 1]$, we reject it[6] and draw a new initial opinion.

6.4.2 Simulation specifications

We simulate our topic-weighted BCMs on complete graphs. As a next step, we plan to simulate our models on additional networks (e.g., such as ER graphs and real-world networks). We simulate our topic-weighted HK model on a 2000-node complete graph. Our topic-weighted DW model is computationally more expensive than our topic-weighted HK

[6]For our simulations, to generate $40,000$ initial opinions from a Gaussian($\sigma = 0.22, \rho = 0$) distribution, we drew a total of $41,934$ samples (i.e., we rejected about 4.84% of samples). To generate $40,000$ initial opinions from a Gaussian($\sigma = 0.22, \rho = 0.8$), we we drew a total of $41,960$ samples (i.e., we rejected about 4.90% of samples).

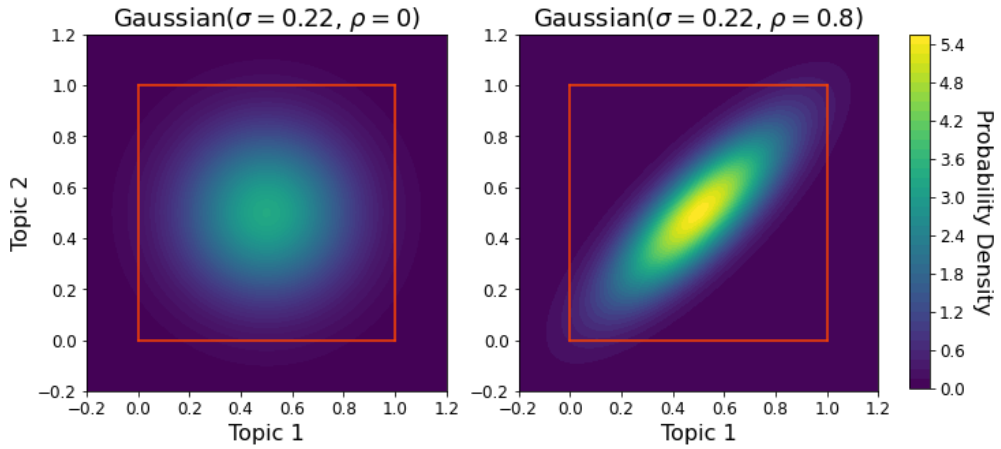

Figure 6.3: The probability density functions for the Gaussian($\sigma = 0.22, \rho = 0$) and Gaussian($\sigma = 0.22, \rho = 0.8$) distributions of initial opinions in our simulations of our topic-weighted BCMs. The red box in each panel indicates the region $[0, 1] \times [0, 1]$. When we draw an initial opinion outside this region, we reject it and resample from the Gaussian distribution.

model; we simulate it on a small 500-node complete graph.

In our numerical simulations of our topic-weighted BCMs, nodes have opinions in the 2D opinion space $[0, 1] \times [0, 1]$. Accordingly, we suppose that the confidence bound $c \in (0, 1)$. For our topic-weighted HK model, we use $c \in [0.025, 0.3]$ in our simulations. For our topic-weighted DW model, we use $c \in [0.05, 0.4]$ in our simulations. In our simulations of our topic-weighted DW model, we use the compromise parameter $m = 0.5$. Nodes average their opinions on a topic when they interact and are receptive to each other on that topic.

In our topic-weighted BCMs, the topic weight is $\omega \in [0, 1]$. We simulate our topic-weighted BCMs with values of $\omega \in [0.1, 1]$. As we discussed in Section 6.2, when $\omega = 1$, the discordance function (see Equation (6.2.1)) for topic k depends only on the opinions on topic k. Our topic-weighted BCMs with $\omega = 1$ are our baseline BCMs. When $\omega = 0$, all discordance functions are the same[7]; in our simulations, we do not consider this case.

[7]For $\omega = 0$, when adjacent nodes i and j interact on any topic, they compromise their opinions on that topic if their mean opinion difference $\frac{1}{K}\sum_{\ell=1}^{K} |x_i^\ell - x_j^\ell|$ is less than c.

In our simulations of our topic-weighted BCMs, the generation of sets of initial opinions and (for the topic-weighted DW model only) the selection of pairs of nodes to interact and the topic to interact on each time step are stochastic. We use Monte Carlo simulations to reduce these sources of noise. For our simulations of our topic-weighted HK model on a 2000-node complete graph, we randomly generate 20 sets of initial opinions for each initial opinion distribution. We reuse those sets of initial opinions in our simulations with different values of the BCM parameters (namely, the confidence bound c and the topic weight ω). Similarly, for our topic-weighted DW model, we randomly generate 20 sets of initial opinions for each initial opinion distribution.

In our topic-weighted BCMs, as in the standard HK and DW models, one can determine the limit opinion clusters in finite time (see the last paragraph of Section 3.3.2). If every pair of opinion clusters, S_a and S_b, has $d_{\max}(\boldsymbol{x}_i, \boldsymbol{x}_j) \geq c$ for all $i \in S_a$ and all $j \in S_b$ at some time $\tilde{T}$, then (because c is fixed) no node in S_a can influence the opinion of a node in S_b (and vice versa) for all $t \geq \tilde{T}$. Meanwhile, nodes in each opinion cluster continue to compromise their opinions on each topic with each other. In particular, if $d_{\max}(\boldsymbol{x}_i, \boldsymbol{x}_j) < c$ for all nodes i and j in an opinion cluster at some time $t \geq \tilde{T}$, then (1) all pairs of adjacent nodes in that opinion cluster are receptive to each other on all topics and (2) all nodes in that opinion cluster eventually converge to the same limit opinion. Therefore, for our topic-weighted BCMs, distinct limit opinion clusters form in finite time.

For our topic-weighted BCMs, if two opinion vectors $\boldsymbol{x}_i$ and $\boldsymbol{x}_j$ satisfy $|x_i^k - x_j^k| < c$ for each topic k, then $d_{\max}(\boldsymbol{x}_i, \boldsymbol{x}_j) < c$. In our numerical simulations of these BCMs, we use a modification of the stopping criterion in Equation (3.3.8). We specify that a simulation has converged if the maximum difference in opinions within each opinion cluster for each topic is less than a tolerance value. That is,

$$\max\{|x_i^k(t) - x_j^k(t)| \text{ such that } i, j \in S_r(t) \text{ for some } r \text{ and } k \in \{1, 2\}\} < \text{tolerance}\,. \quad (6.4.3)$$

For our topic-weighted HK model, we use a tolerance value of 1×10^{-6}; for our topic-weighted DW model, we use a tolerance value of 0.01. We denote the time step in which a simulation reaches the stopping criterion by T_f.

For our topic-weighted DW model, the selection of a pair of nodes to interact and a topic on which to interact at each time step is stochastic. For each simulation of our topic-weighted DW model with $\omega \in (0, 1)$, we simultaneously run a "control" simulation. In a control simulation, we use a value of $\omega = 1$ (to give the baseline DW model) and we use the same confidence bound c, set of initial opinions, and sequence of edges and topics for interactions as in the corresponding simulation of our topic-weighted DW model. If a topic-weighted DW simulation reaches the stopping criterion (see Equation (6.4.3)) before the corresponding control simulation, we stop the topic-weighted DW simulation and continue the corresponding control simulation. Similarly, if a control simulation reaches the stopping criterion before the topic-weighted DW simulation, we continue the corresponding topic-weighted DW simulation.

To characterize node opinions in our topic-weighted BCM, we calculate the numbers of major and minor limit opinion clusters (see Section 3.3.3) and a variant of the limit order parameter $Q(T_f)$ (see Equation (3.3.5)) for opinion-cluster profiles (see Section 3.3.4). In this chapter, we say that an opinion cluster is a "minor opinion cluster" if it has at most 2% of the nodes of a network. That is, an opinion cluster S_r is a minor opinion cluster if $|S_r| \leq 0.02N$. We say that an opinion cluster that is not a minor cluster is a "major opinion cluster". We calculate the order parameter using our maximum discordance function (see Equation (6.2.2)), instead of the absolute value of a 1D opinion difference. In this setting, the order parameter is thus

$$Q(t) = \frac{1}{N^2} \sum_{i=1}^{N} \sum_{j=1}^{N} \mathbf{1}_{d_{\max}(\boldsymbol{x}_i(t), \boldsymbol{x}_j(t)) < c} \,, \tag{6.4.4}$$

where $\mathbf{1}_{d_{\max}(\boldsymbol{x}_i(t), \boldsymbol{x}_j(t)) < c}$ is an indicator function that equals 1 when $d_{\max}(\boldsymbol{x}_i(t), \boldsymbol{x}_j(t)) < c$ and

equals 0 otherwise. With our definition of limit opinion clusters (see the last paragraph of Section 6.3.1) as sets of nodes that converge to the same opinion in all topics, Equation (6.4.4) is analogous to Equation (3.3.6).

6.5 Preliminary results of our numerical simulations

6.5.1 Topic-weighted HK model

We now discuss the simulations of our topic-weighted HK model on a 2000-node complete graph with various initial opinion distributions (see Section 6.4.1) and values of the BCM parameters (namely, the confidence bound c and topic weight ω). As we discussed in Section 6.4.2, for each initial opinion distribution, we use 20 distinct sets of initial opinions in Monte Carlo simulations of our topic-weighted HK model.

In Figure 6.4, we show the numbers of major limit opinion clusters in our simulations of our topic-weighted HK model for various initial opinion distributions. For a fixed initial opinion distribution, as we decrease ω, the transition between a consensus state (i.e., exactly one major opinion cluster) and a fragmented state (i.e., two or more major clusters) tends to occur at smaller values of c. For a fixed initial opinion distribution and fixed ω, there tend to be more major opinion clusters (and correspondingly less consensus) as we decrease c. However, $c = 0.025$ often yields fewer major opinion clusters than $c = 0.05$ for the independent uniform and Gaussian($\sigma = 0.22, \rho = 0$) initial opinion distributions. This arises from our ad hoc notion of a major opinion cluster having more than 2% of the nodes of a network. In Figure 6.5, we show the numbers of minor limit opinion clusters in our simulations of our topic-weighted HK model for various initial opinion distributions. When $c = 0.025$, our simulations with the independent uniform and Gaussian($\sigma = 0.22, \rho = 0$) initial opinion distributions have many minor opinion clusters. Rather than focusing on the numbers of major and minor opinion clusters, we calculate the order parameter $Q(T_f)$ to examine opinion fragmentation in our topic-weighted HK model.

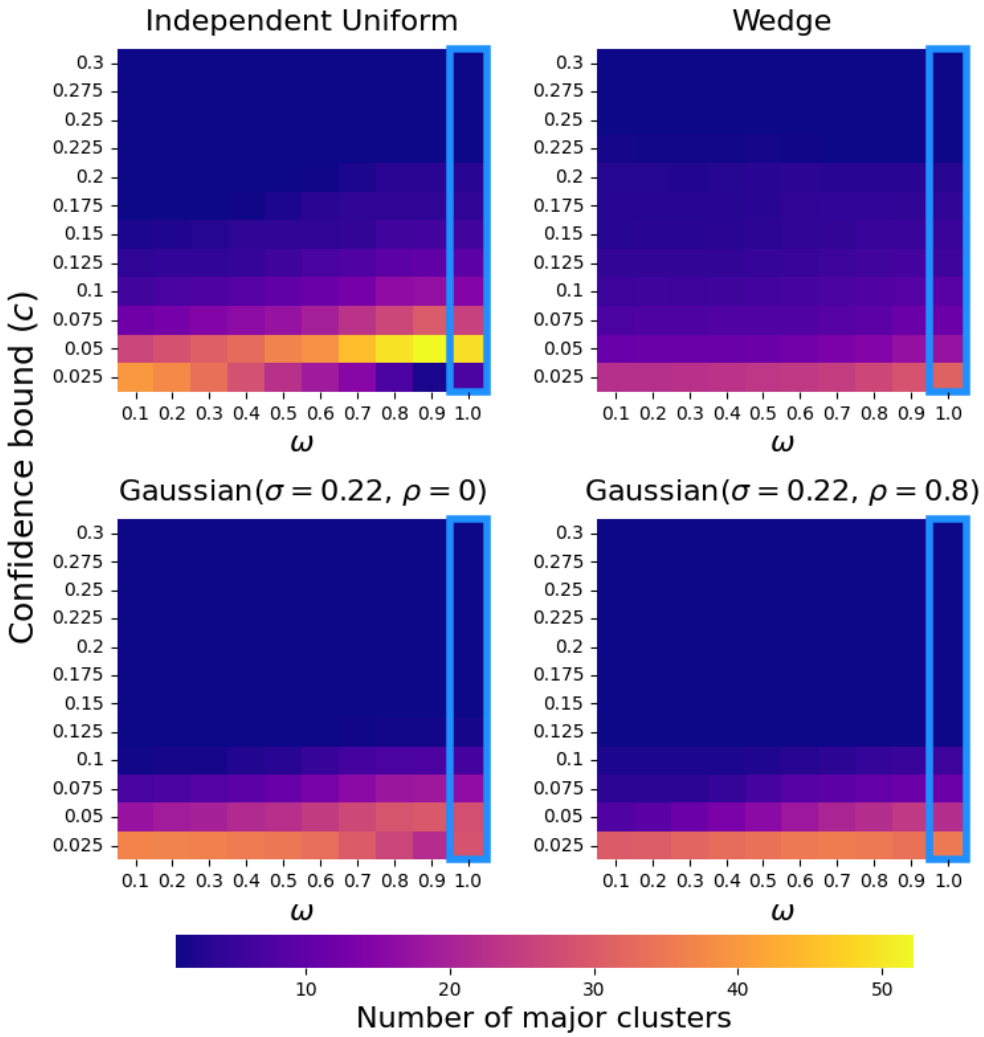

Figure 6.4: Numbers of major limit opinion clusters in simulations of our topic-weighted HK model on a 2000-node complete graph with various initial opinion distributions and various values of the BCM parameters c and ω. We consider a cluster to be major cluster if it has more than 2% of the nodes of a network. (In this case, a major cluster must have at least 41 nodes.) For this heat map and all subsequent heat maps, the depicted values are means of 20 simulations of our topic-weighted BCM for each value of the BCM parameter pair (c, ω). In each heat map, the column on the far right depicts the results of our simulations of the baseline HK model with $\omega = 1$; it is highlighted with a blue border.

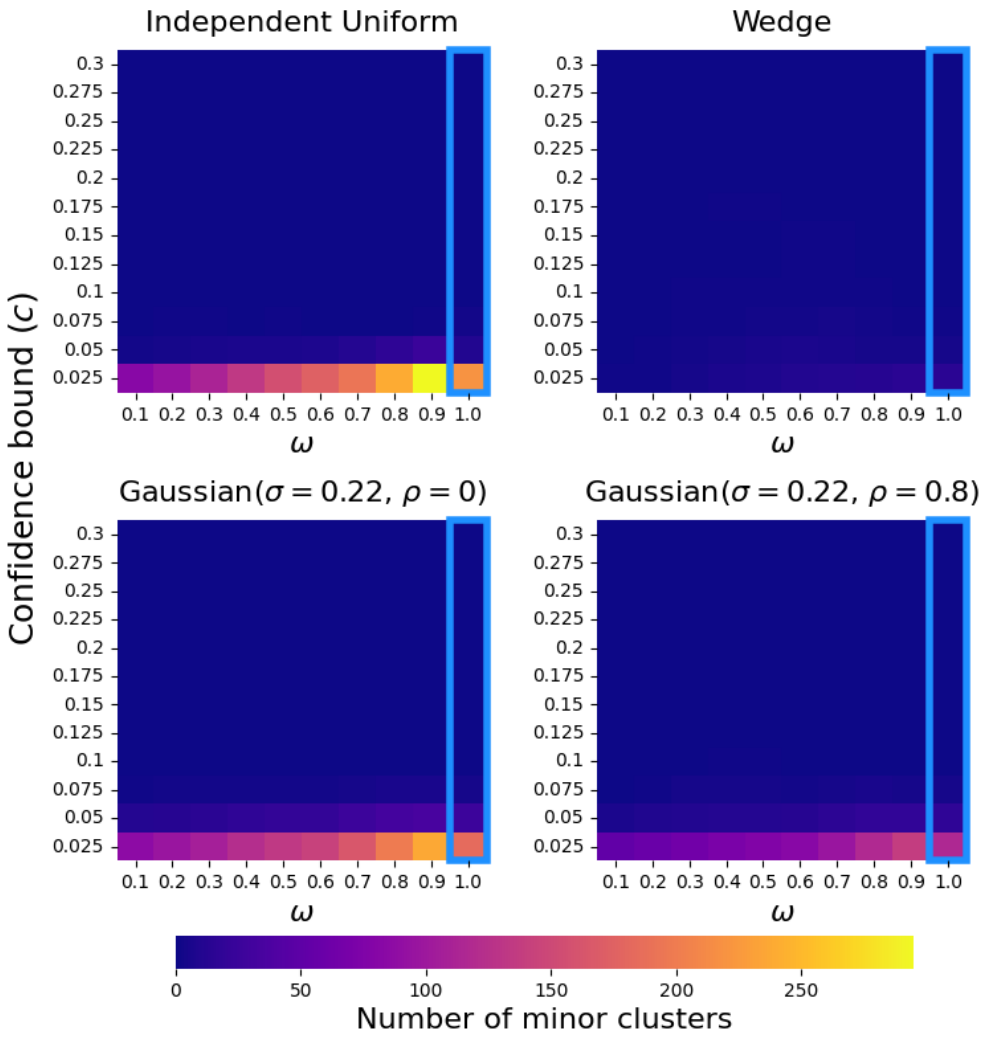

Figure 6.5: Numbers of minor limit opinion clusters in simulations of our topic-weighted HK model on a 2000-node complete graph with various initial opinion distributions and various values of the BCM parameters c and ω. We consider a cluster to be minor cluster if it has at most 2% of the nodes of a network. (In this case, a minor cluster must have 40 nodes or fewer.)

In Figure 6.6, we show the limit order parameters $Q(T_f)$ (see Equation (6.4.4)) in our simulations of our topic-weighted HK model for various initial opinion distributions. When $Q(T_f) = 1$, all nodes converge to the same limit opinion vectors and reach a consensus state. Smaller values of $Q(T_f)$ indicate more opinion fragmentation than larger values of $Q(T_f)$. For fixed ω and a fixed initial opinion distribution, we observe smaller values of $Q(T_f)$ (and correspondingly more opinion fragmentation) as we decrease c. For fixed c and a fixed initial opinion distribution, we tend to observe more opinion fragmentation as we increase ω. This is especially noticeable near the transition between consensus and fragmentation.

In our simulations, we observe that the choice of initial opinion distribution has a large effect on whether our topic-weighted HK model reaches a consensus state or a fragmented state. For fixed ω, the transition between consensus and fragmentation occurs at smaller values of c for the two Gaussian initial opinion distributions than for the independent uniform and wedge initial opinion distributions. Intuitively, it makes sense that the two Gaussian initial opinion distributions promote more consensus than the independent uniform and wedge distributions. This is because the initial opinions for the Gaussian distributions are more concentrated near the center (see Figure 6.3) of the opinion space and there is small probability of initial opinions near the boundaries of the opinion space $[0, 1] \times [0, 1]$. For the independent uniform and Gaussian($\sigma = 0.22, \rho = 0$) initial opinion distributions, as we decrease ω, the value of c at which the transition between consensus and fragmentation occurs also decreases. For these distributions and $c = 0.2$, our topic-weighted HK model with $\omega \leq 0.4$ reaches consensus, while the baseline HK model has opinion fragmentation. By contrast, for the wedge and Gaussian($\sigma = 0.22, \rho = 0.8$) initial opinion distributions, ω has little effect on the value of c at which there is a transition between consensus and fragmentation. For the wedge distribution and fixed c, varying ω typically has little effect on the limit order parameter $Q(T_f)$. However, for the wedge distribution, when $c = 0.225$ (near the transition between consensus and fragmentation), $Q(T_f)$ tends to increase as we decrease ω.

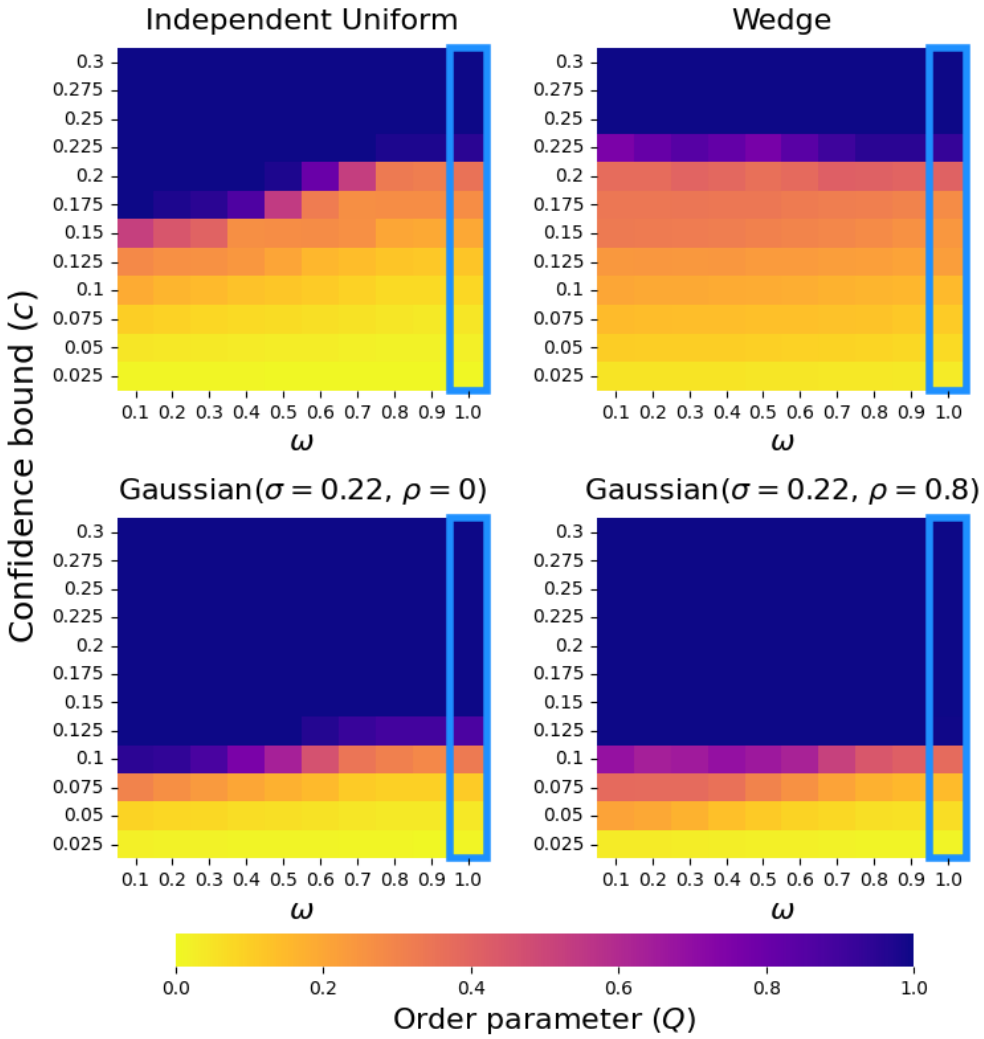

Figure 6.6: Limit order parameters $Q(T_f)$ in simulations of our topic-weighted HK model on a 2000-node complete graph with various initial opinion distributions and various values of the BCM parameters c and ω.

To investigate the impact of the distribution of initial opinions on the distribution of limit opinions, we plot the limit opinion clusters in opinion space for our simulations with confidence bound c near the transition between consensus and fragmentation. In Figure 6.7, we show an example of the limit opinion clusters for our topic-weighted HK model and the baseline HK model with $c = 0.1$ for Gaussian($\sigma = 0.22, \rho = 0$) and Gaussian($\sigma = 0.22, \rho = 0.8$) initial opinion distributions. For the simulations that we show for our topic-weighted HK model with $\omega = 0.1$ (see panels A and C of Figure 6.7), we observe a limit opinion cluster near $(0.5, 0.5)$ with almost all of the nodes for the Gaussian($\sigma = 0.22, \rho = 0$) distribution and most of the nodes for the Gaussian($\sigma = 0.22, \rho = 0.8$) distribution. The simulation with the Gaussian($\sigma = 0.22, \rho = 0.8$) distribution also has a second major limit opinion cluster with about 17% of the nodes of the network offset from the limit opinion cluster near $(0.5, 0.5)$ and lying near the central axis of the skewed Gaussian distribution (see Figure 6.3). We hypothesize that for small ω, early in a simulation, nodes that have initial opinions near $(0.5, 0.5)$ have more nodes in their regions of absorption for a Gaussian($\sigma = 0.22, \rho = 0$) initial opinion distribution than for a Gaussian($\sigma = 0.22, \rho = 0.8$) initial opinion distribution. In Figure 6.7, for our topic-weighted HK model, more limit opinions are near $(0.5, 0.5)$ and we observe more consensus for the Gaussian($\sigma = 0.22, \rho = 0$) distribution than for the Gaussian($\sigma = 0.22, \rho = 0.8$) distribution. When we compare the limit opinion clusters of our topic-weighted HK model (see panels A and C of Figure 6.7) with those of the baseline HK model (see panels B and D of Figure 6.7), we see that the baseline HK model has more small limit opinion clusters, resulting in smaller order parameters $Q(T_f)$ and less consensus than in our topic-weighted HK model. For our topic-weighted HK model, as we decrease the topic weight ω, the areas of the regions of absorption increases. Therefore, for smaller values of ω, a node can potentially have more nodes to which it is receptive, thereby promoting consensus.

In Figure 6.8, we show an example of the limit opinion clusters for our topic-weighted HK model and the baseline HK model with $c = 0.2$ for independent uniform and wedge initial

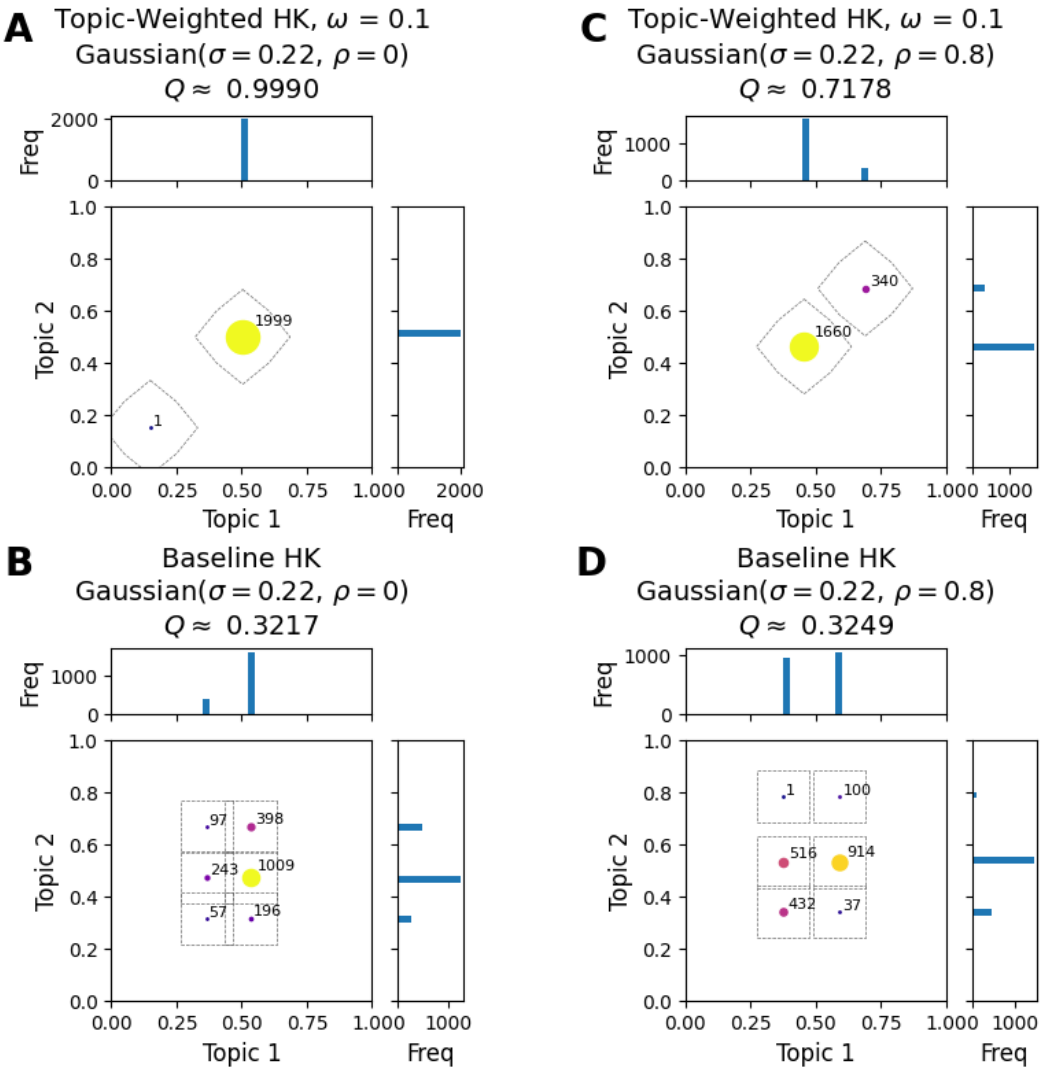

Figure 6.7: The limit opinion clusters in simulations of our topic-weighted HK model and the baseline HK model with $c = 0.1$ for the Gaussian initial opinion distributions. In (A, B), we show simulations of our topic-weighted HK model with $\omega = 0.1$. In (C, D), we show corresponding simulations of the baseline HK model (with $\omega = 1$) with the same initial opinion distributions. For this figure and the subsequent figure of this type, each panel shows the limit opinion clusters for a single simulation. For each panel, in the square plot, each point represents the location of a single limit opinion cluster. The color, size, and annotation of each point indicate the number of nodes in that limit opinion cluster. Around each limit opinion cluster, we show a dotted polygon to indicate its region of absorption (see Equation (6.3.3)). Above and to the right of each square plot, we show the marginal distributions of the limit opinions in topic 1 and topic 2, respectively.

opinion distributions. The simulation that we show of our topic-weighted HK model with $\omega = 0.5$ and an independent uniform initial opinion distribution (see panel A of Figure 6.8) reaches a consensus state; there is a single limit opinion cluster near the center $(0.5, 0.5)$ of the opinion space. For the simulations that we show of the baseline HK model (see panels B and D of Figure 6.8), there is no limit opinion cluster at the center of the opinion space. Instead, there are 2–4 major limit opinion clusters located near $(0.5, 0.5)$; their regions of absorption surround the center of the opinion space at $(0.5, 0.5)$ with little overlap. The baseline model has smaller regions of absorption than our topic-weighted HK model with $\omega < 1$. Because of the smaller regions of absorption, it seems that the limit opinions are somewhat away from the center of the opinion space, promoting opinion fragmentation. The simulation that we show of our topic-weighted HK model with $\omega = 0.5$ and a wedge initial opinion distribution (see panel C of Figure 6.8) has three major limit opinion clusters. These three limit opinion clusters are located near the wedge distribution function (see Equation (6.4.1)) in opinion space. For fixed c, our topic-weighted HK model with $\omega = 0.5$ has larger regions of absorption than the baseline model. However, because of the shape of the wedge distribution, it seems that nodes on one side of the opinion space (e.g., near $(0.25, 0.3)$) are never receptive to nodes on the other side of the opinion space (e.g., near $(0.75, 0.3)$). Consequently, in Figure 6.8C, our simulation with the wedge distribution results in a fragmented state with three major limit opinion clusters. The largest major limit opinion cluster is near $(0.5, 0.75)$ and is near the apex of the wedge (i.e., near $(0.5, 1)$). This major limit opinion cluster is likely the largest due to it containing nodes from both "legs" of the wedge distribution (i.e., nodes i with initial opinion $x_i^1(0) < 0.5$ and nodes i with $x_i^1(0) > 0.5$).

6.5.2 Topic-weighted DW model

We now briefly discuss the simulations of our topic-weighted DW model on a 500-node complete graph with the independent uniform and wedge initial opinion distributions and various values of the BCM parameters (namely, the confidence bound c and topic weight ω).

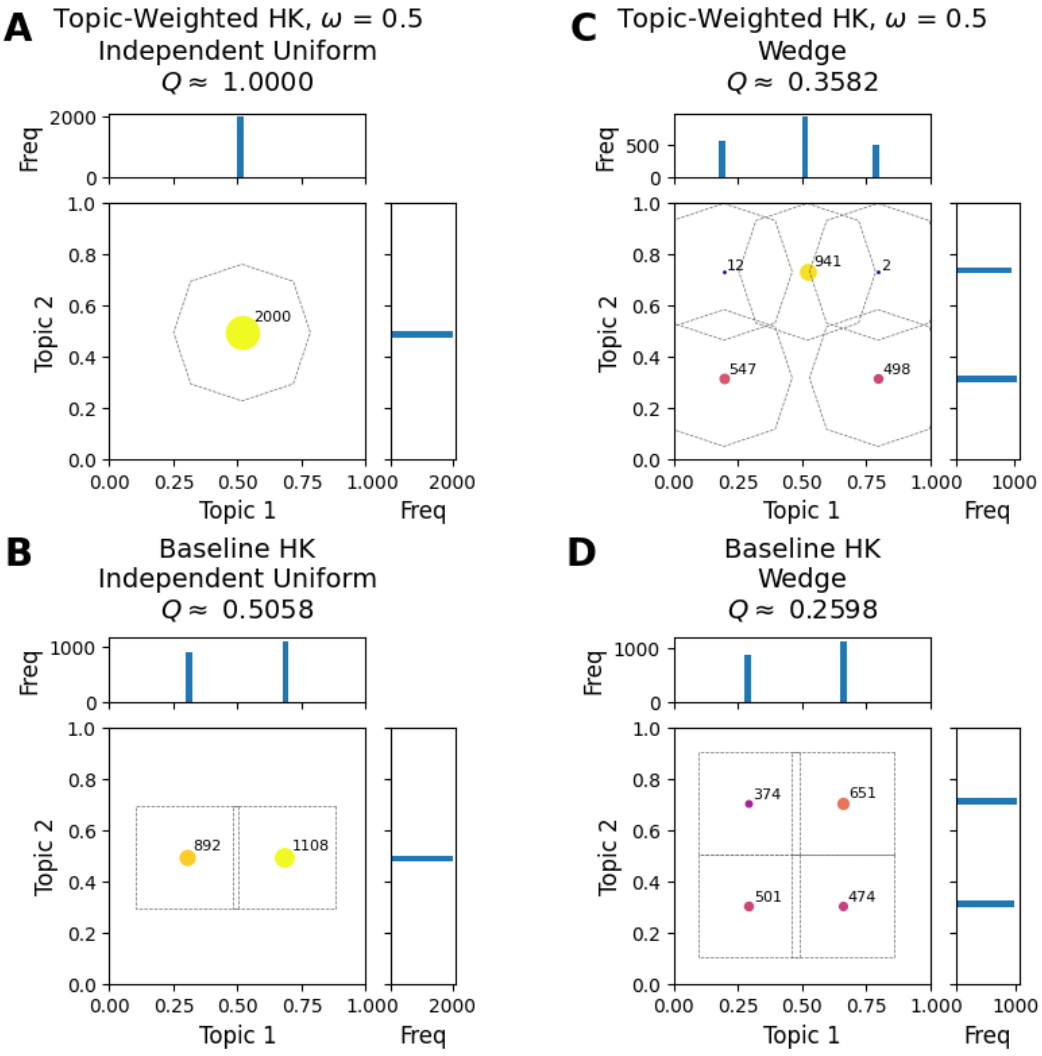

Figure 6.8: The limit opinion clusters in simulations of our topic-weighted HK model and the baseline HK model with $c = 0.2$ for the independent uniform and wedge initial opinion distributions. In (A, B), we show simulations of our topic-weighted HK model with $\omega = 0.5$. In (C, D), we show corresponding simulations of the baseline HK model (with $\omega = 1$) with the same initial opinion distributions.

We have not yet completed our simulations of our topic-weighted DW model with Gaussian initial opinion distributions and plan to do so as part of our next steps for this project. As we discussed in Section 6.4.2, for each initial opinion distribution, we use 20 distinct sets of initial opinions in Monte Carlo simulations of our topic-weighted DW model.

In Figure 6.9, we show the limit order parameters $Q(T_f)$ (see Equation (6.4.4)) in our simulations of our topic-weighted DW model for various initial opinion distributions. For fixed ω and both the independent uniform and wedge initial opinion distributions, we observe smaller values of $Q(T_f)$ (and correspondingly more opinion fragmentation) as we decrease c. For an independent uniform initial opinion distribution, we observe similar trends for our topic-weighted DW model as we did for our topic-weighted HK model. One trend is that, for fixed c and a fixed initial opinion distribution, we tend to observe more opinion fragmentation as we increase ω. Another trend is that, as we decrease ω, we observe a decrease in the value of c at which the transition between consensus and fragmentation occurs. For fixed c and the wedge initial opinion distribution, we typically observe little effect on the order parameter when we vary ω. However, for $c = 0.25$, there appears to be less opinion fragmentation for $0.4 \leq \omega \leq 0.6$ than for other values of ω. We are not sure why this seems to occur, so we plan to run additional simulations to determine if this observation is due to noise.

6.6 Discussion and next steps

We developed BCMs that generalize the HK and DW models to K-dimensional opinions with topic-weighted opinion discordance functions (see Equation (6.2.1)). As we decrease the topic weight ω in our BCMs, when two nodes interact on topic k, we more heavily weight their opinion differences on topics other than k when we determine if they are receptive to each other on topic k. As we decrease the topic weight ω, the regions of absorption (see Equation (6.3.3)) increase in area. Therefore, given a node i with opinion $\boldsymbol{x_i}$, a neighboring node j has a larger region of opinion values for nodes i and j to be receptive to each other

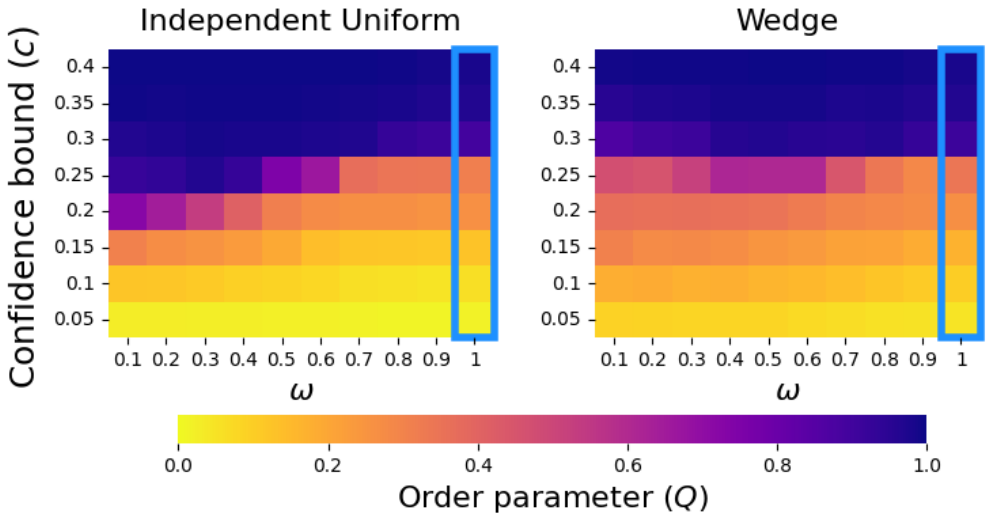

Figure 6.9: Limit order parameters $Q(T_f)$ (see Equation (6.4.4)) in simulations of our topic-weighted DW model with various values of the BCM parameters c and ω for (left) an independent uniform and (right) a wedge initial opinion distribution.

on all topics.

In our preliminary numerical simulations, we demonstrated for our topic-weighted BCMs that the choice of initial opinion distribution has a large effect on (1) whether our models reaches a consensus state or a fragmented state and (2) the amount of opinion fragmentation. For a fixed confidence bound c and a fixed initial opinion distribution, we tended to observe less opinion fragmentation (i.e., more consensus) for smaller values of ω. Smaller values of ω correspond to larger regions of absorption, so perhaps more pairs of nodes are receptive to each other on all topics, thereby promoting more consensus. However, for a fixed confidence bound c and the wedge initial opinion distribution, smaller ω did not promote more consensus; instead, we observed little effect on the order parameter Q (see Equation (6.4.4)) when varying ω. We hypothesize that this may be due to the shape of the wedge initial opinion distribution. Our simulations highlight the large effect of initial opinion distribution on the limit opinions of our topic-weighted BCMs.

We plan to run additional numerical simulations of our topic-weighted BCMs. For our

topic-weighted DW model, we will run simulations with Gaussian initial opinion distributions. To investigate the effect of network structure on our topic-weighted HK and DW models, we will simulate them on additional networks, such as synthetic networks (e.g., ER graphs and SBM graphs), and real-world networks (such as the ones that we considered in Chapter 5).

CHAPTER 7

Conclusions

In this dissertation, we developed and studied three generalizations of bounded-confidence models (BCMs) of opinion dynamics.

In Chapter 4, we generalized the Deffaunt–Weisbuch (DW) BCM by using node weights to model agents that have heterogeneous activation probabilities. Using numerical simulations, we systematically investigated (using a variety of network structures and node-weight distributions) the effects of node weights, which we assign uniformly at random to the nodes. In our node-weighted BCM, we found that heterogeneous node-weight distributions tend to yield longer convergence times and more opinion fragmentation than the associated baseline BCM. Additionally, we tended to observe more opinion fragmentation as we increased the inequality in node-activation probabilities by increasing either the heaviness of the tail or the mean of the node-weight distribution. One can tailor the node weights in our BCM to examine a variety of sociological scenarios in which nodes have heterogeneous activity levels. For example, one can adapt it to assign node weights in a way that depends on network structure and/or node opinions. Our model illustrates the importance of incorporating node weights into network analysis, and we encourage researchers to spend more time studying the effects of node weights on network structure and dynamics.

In Chapter 5, we generalized the DW and Hegselmann–Krause (HK) BCMs so that each pair of agents has a distinct confidence bound that changes when the pair interacts. We analytically and numerically explored the limiting behaviors of our adaptive-confidence BCMs. We gave convergence properties for the dyadic confidence bounds and the limiting behaviors

of effective graphs. For a variety of networks and a wide range of values of the parameters that control the increase and decrease of confidence bounds, we demonstrated numerically that our adaptive-confidence BCMs promote consensus and have longer convergence times than the associated baseline BCMs. By examining the effective graphs in our numerical simulations, we showed that our adaptive-confidence BCMs can have neighboring nodes that converge to the same opinion but are not receptive to each other. This qualitative behavior does not occur in the associated baseline BCMs.

In Chapter 6, we generalized the DW and HK models to multi-dimensional opinions that consist of multiple interdependent topics. In our preliminary numerical simulations, we demonstrated for these topic-weighted BCMs that the choice of initial opinion distribution has a large effect on (1) whether our models reaches a consensus state or a fragmented state and (2) the amount of opinion fragmentation. In studies of BCMs, it is uncommon to examine initial opinion distributions that are not uniformly random (see Section 3.4). We encourage researchers to consider the effects of initial opinion distributions in studies of BCMs.

Each of our three generalizations of BCMs incorporates a different modification to make them more realistic, while maintaining tractability. However, our models have not been validated empirically. Our BCMs rely on many assumptions that make them tractable, but these assumptions are generally not true in the complex real-life ways that people change opinions. Even with these drawbacks, we believe that our mechanistic models are insightful for future work and provide frameworks to explore opinion dynamics. We encourage researchers to consider exploring effective graphs, incorporating node weights and adaptive parameters and network structures, and examining different initial opinion distributions in models of opinion dynamics.

APPENDIX A

A Linear Analysis of Shock–Grain Interaction for Inertial Confinement Fusion (ICF) Ablators

In this appendix, we apply a linear theory of shock interaction with density non-uniformity [VWL07] to study shock interaction for a model of high-density carbon grains. This project is a collaboration with Seth Davidovits of Lawrence Livermore National Laboratory.[1] This project has been submitted as an original paper [LD24] that I co-authored with Seth Davidovits.[2]

Ablator materials, such as high-density carbon (HDC), that are used for inertial confinement fusion (ICF) have grain structure that can lead to small-scale density non-uniformity and the generation of perturbations when the materials are shocked and compressed. One challenge in simulating the effects of grains and other material microstructures on shock propagation is the required resolution and computational costs. We investigate the effects of HDC grains on shock propagation by combining the asymptotic analysis that was developed by Velikovich et al. [VWL07] for the linearized governing partial differential equations (i.e., the compressible Euler equations) with a Fourier decomposition of a model of homogeneous 2D grains. We also compare the results of the linear analysis to numerical simulations of the compressible Euler equations without the linear approximation. We assess the applica-

[1]This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

[2]My contributions to the paper [LD24] were formulating our model of 2D grains and applying the linear theory of [VWL07] to a Fourier decomposition of our model for 2D grains. This was done under the mentorship of Seth Davidovits. Seth Davidovits and I wrote the paper together.

bility of the linear analysis, examine key features of shock–grain interactions, consider the effects of the size and aspect ratio of grains, and study the problem of approximating (i.e., "de-resolving") grains in simulations.

We model HDC with grains as a non-uniform density field consisting of a periodic grid of homogeneous grains that are separated by interstitial spaces with ramped transitions between the grains and interstitial regions. We apply the linear theory (with one minor addition) that was developed by Velikovich et al. [VWL07] for shock interaction with a pre-shock density perturbation that consists of a single Fourier mode. To apply this linear analysis to our model of 2D grains, we calculate a Fourier expansion of the pre-shock density perturbations. For each Fourier mode, we calculate the post-shock perturbations following the aforementioned linear analysis. Taking the sum of these perturbations for each Fourier mode gives a total perturbation.

Through comparison to simulations, we see that shock–grain interactions are nonlinear, except in a regime in which the density contrast between the grains and interstitial spaces is sufficiently small. Nonetheless, the linear theory shows some key features of the shock–grain interactions. We focus on examining the post-shock perturbed kinetic energy, which includes contributions from density and velocity perturbations. Our calculations show that the post-shock perturbations are composed of sonic reflections off of grain boundaries and vorticity deposition along them, with the latter dominating the perturbed energy content. We find that the mean post-shock perturbed kinetic energy decreases with increasing grain size. From the perspective of the post-shock perturbed energy, the linear theory largely supports a method of Davidovits et al. [DWC22] for de-resolving the grains that treats the grains statistically.

APPENDIX B

Detection of Overlapping Communities in Undirected Graphs with Personalized PageRank

In this appendix, we investigate a method to detect overlapping communities in undirected graphs. This project is ongoing work in collaboration[1] with Trevor Steil at Lawrence Livermore National Laboratory (LLNL).[2]

Many of the methods to detect communities[3] (see Section 2.2) in graphs partition the set of nodes of a graph; they find a set of communities in which each node belongs to exactly one community [FH16]. For many real-life graphs, it can be appropriate to consider overlapping[4] communities [YL15]; a node can belong to multiple communities and communities can have overlapping node membership. For example, one can consider overlapping friend groups in a social network or overlapping conference participation of groups of researchers. Building from the work of Andersen et al. [ACL06] and Whang et al. [WGD13], we study a method based on personalized PageRank (PPR) to detect overlapping communities in unweighted, undirected graphs. For computational efficiency and to allow us to handle large graphs, we implement this method using YGM [SRP23], an asynchronous distributed computing library

[1] I am leading this project with mentorship from Trevor Steil.

[2] This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344 (LLNL TR-864751). Funding from LLNL LDRD project 21-ERD-020 was used in this work.

[3] Researchers sometimes refer to communities in a graph as "clusters" [FH16].

[4] Overlapping communities are also sometimes called "soft communities". Additionally, when nodes have different strengths of membership in different communities, overlapping communities are sometimes called "fuzzy communities" [FH16].

that was developed at LLNL.

The PPR vector for a set of seed nodes in a graph indicates the stationary distribution of a random walk (see Section 2.1) on the graph in which a random walker can "teleport" to a seed node. Consider a connected, unweighted, undirected graph $G = (V, E)$, where V is the set of nodes and E is the set of edges. Let W denote the lazy-random-walk transition matrix $W = \frac{1}{2}(I + D^{-1}A)$, where A is the adjacency matrix of the graph and D is a diagonal matrix of node degrees. Let $\boldsymbol{s}$ be a vector representing a seed set S with size (i.e., number of nodes) N_s; the vector $\boldsymbol{s}$ has entry $1/N_S$ in seed-node positions and 0 elsewhere. The PPR vector $\boldsymbol{\pi}_s$ for this seed vector $\boldsymbol{s}$ satisfies

$$\boldsymbol{\pi}_s = \alpha \boldsymbol{s} + (1 - \alpha)\boldsymbol{\pi}_s W \,, \tag{B.0.1}$$

where $\alpha \in (0, 1]$ is the teleportation parameter, which controls the probability that a random walker "teleports" to a seed node. The ith entry of $\boldsymbol{\pi}_s$ is the PPR of node i for the seed set S.

Andersen et al. [ACL06] developed a method to efficiently calculate PPR vectors and use them to find local communities for seed nodes in an undirected graph. They developed a "pushing" algorithm that efficiently approximates PPR with a runtime that is independent of the size of the graph. They also developed a "conductance-sweep" algorithm, which uses the PPR vector to find a local community with small conductance (see Equation (2.2.1)) that contains the seed node for the PPR vector. A set of nodes with smaller conductance can be considered a higher-quality community than a set of nodes with larger conductance (see Section 2.2). Whang et al. [WGD13] developed a method that uses the algorithm in [ACL06] to find overlapping communities in a graph. At each step of their method, a seed-selection procedure selects a seed node that does not yet belong to a community and then uses PPR to find a community for that seed node. Their method repeats until it reaches the maximum number of seed nodes that one wants to consider or all nodes in the graph belong

to at least one community. The resulting set of overlapping communities is a "community profile". In [WGD13], Whang et al. investigated a variety of seed-selection procedures.

We use the method of Whang et al. [WGD13] with some small modifications, and we implement it in a distributed-computing setting. Let R denote the set of all nodes that have not yet been assigned to any community. When Whang et al. [WGD13] selected a seed node in R, they took that seed node and all of its neighbors to be the seed set for their PPR calculation. By contrast, we consider seed sets that consist of a single seed node. Additionally, when we select a seed node, before calculating PPR, we first check if we can add it to an existing community without increasing the conductance of that community.

The *local conductance* [GS12] of a node i in a graph is the conductance (see Equation (2.2.1)) of the set that consists of i and its neighbors. A node that has a local conductance that is less than or equal to the local conductance of its neighbors has *minimal local conductance*. Gleich et al. [GS12] found that nodes with minimal local conductance are good seed nodes for PPR-based community detection. Whang et al. [WGD13] examined a seed-selection procedure that finds seed nodes with minimal local conductance.

In our preliminary investigation, we examine two[5] seed-selection procedures, which we call our "largest-degree" and "smallest-local-conductance" seed-selection procedures. For our largest-degree seed-selection procedure, we select a seed node in R that has the largest degree. For our smallest-local-conductance seed-selection procedure, we select a seed node in R that has the smallest local conductance. We test our method with these seed-selection procedures on a variety of graphs from the Stanford Large Network Dataset Collection [LK14].

We observe that largest-degree seed selection tends to yield fewer and larger communities than smallest-local-conductance seed selection. Despite the differences in the numbers

[5]We also consider variants of these seed-selection procedures in which only the nodes in R that have minimal local conductance are potential seed nodes. If there are no such nodes, we consider any node in R as a potential seed. We observe little difference in our results when we restrict seed nodes to have minimal local conductance when possible.

and sizes of the resulting communities, both seed-selection procedures yield community profiles with similar mean conductance. Additionally, seed nodes that we select later in our community-detection method do not necessarily yield worse communities. That is, for $k < \ell$, the ℓth community does not necessarily have larger (and therefore worse) conductance than the kth community. In our preliminary investigation, we find that the choice of seed-selection procedure can affect the number, sizes, and qualities (which we quantify with conductance) of the communities. An area for future work is to develop methods to assess and compare community profiles. Such methods can be application-dependent and will hopefully allow researchers to determine which community profiles (and which corresponding seed-selection procedures) are suitable for their application.

REFERENCES

[AC15] Meysam Alizadeh and Claudio Cioffi-Revilla. "Activation Regimes in Opinion Dynamics: Comparing Asynchronous Updating Schemes." *Journal of Artificial Societies and Social Simulation*, **18**(3): 8, 2015.

[ACL06] Reid Andersen, Fan Chung, and Kevin Lang. "Local Graph Partitioning Using PageRank Vectors." In *2006 47th Annual IEEE Symposium on Foundations of Computer Science*, FOCS'06, 475–486, Berkeley, CA, USA, 2006. Institute of Electrical and Electronics Engineers (IEEE), Piscataway, NJ, USA.

[AMS19] Alessia Antelmi, Delfina Malandrino, and Vittorio Scarano. "Characterizing the Behavioral Evolution of Twitter Users and The Truth Behind the 90-9-1 Rule." In *Companion Proceedings of The 2019 World Wide Web Conference*, WWW '19, 1035–1038, San Francisco, CA, USA, 2019. Association for Computing Machinery, New York, NY, USA.

[AWA22] Johnathan A. Adams, Gentry White, and Robyn P. Araujo. "Mathematical Measures of Societal Polarisation." *PLoS ONE*, **17**(10): e0275283, 2022.

[BAP24] Carmela Bernardo, Claudio Altafini, Anton Proskurnikov, and Francesco Vasca. "Bounded Confidence Opinion Dynamics: A Survey." *Automatica*, **159**: 111302, 2024.

[Bar51] M. S. Bartlett. "An Inverse Matrix Adjustment Arising in Discriminant Analysis." *The Annals of Mathematical Statistics*, **22**(1): 107–111, 1951.

[BB03] Albert-László Barabási and Eric Bonabeau. "Scale-Free Networks." *Scientific American*, **288**(5): 60–69, 2003.

[BBC13] Arnab Bhattacharyya, Mark Braverman, Bernard Chazelle, and Huy L. Nguyen. "On the Convergence of the Hegselmann–Krause System." In *Proceedings of the 4th Conference on Innovations in Theoretical Computer Science*, ITCS '13, 61–66, Berkeley, CA, USA, 2013. Association for Computing Machinery, New York, NY, USA.

[BC19] Anna D. Broido and Aaron Clauset. "Scale-Free Networks are Rare." *Nature Communications*, **10**: 1017, 2019.

[BCF07] Franco Bagnoli, Timoteo Carletti, Duccio Fanelli, Alessio Guarino, and Andrea Guazzini. "Dynamical Affinity in Opinion Dynamics Modeling." *Physical Review E*, **76**: 066105, 2007.

[BCP11] Andrea Baronchelli, Claudio Castellano, and Romualdo Pastor-Satorras. "Voter Models on Weighted Networks." *Physical Review E*, **83**: 066117, 2011.

[BGK23] Rico Berner, Thilo Gross, Christian Kuehn, Jürgen Kurths, and Serhiy Yanchuk. "Adaptive Dynamical Networks." *Physics Reports*, **1031**: 1–59, 2023.

[BGS16] Aaron Bramson, Patrick Grim, Daniel J. Singer, Steven Fisher, William Berger, Graham Sack, and Carissa Flocken. "Disambiguation of Social Polarization Concepts and Measures." *The Journal of Mathematical Sociology*, **40**(2): 80–111, 2016.

[BHT10] Vincent D. Blondel, Julien M. Hendrickx, and John N. Tsitsiklis. "Continuous-Time Average-Preserving Opinion Dynamics with Opinion-Dependent Communications." *SIAM Journal on Control and Optimization*, **48**(8): 5214–5240, 2010.

[BKR03] E. Ben-Naim, P.L. Krapivsky, and S. Redner. "Bifurcations and Patterns in Compromise Processes." *Physica D: Nonlinear Phenomena*, **183**(3): 190–204, 2003.

[BL06] Kenneth Benoit and Michael Laver. *Party Policy in Modern Democracies*. Routledge, London, United Kingdom, 2006.

[BL12] Kenneth Benoit and Michael Laver. "The Dimensionality of Political Space: Epistemological and Methodological Considerations." *European Union Politics*, **13**(2): 194–218, 2012.

[BLS21] Fabian Baumann, Philipp Lorenz-Spreen, Igor M. Sokolov, and Michele Starnini. "Emergence of Polarized Ideological Opinions in Multidimensional Topic Spaces." *Physical Review X*, **11**: 011012, 2021.

[BP20] Heather Z. Brooks and Mason A. Porter. "A Model for the Influence of Media on the Ideology of Content in Online Social Networks." *Physical Review Research*, **2**: 023041, 2020.

[Bre13] Alain Bretto. *Hypergraph Theory: An Introduction*. Springer International Publishing, Cham, Switzerland, 2013.

[BVI22] Carmela Bernardo, Francesco Vasca, and Raffaele Iervolino. "Heterogeneous Opinion Dynamics with Confidence Thresholds Adaptation." *IEEE Transactions on Control of Network Systems*, **9**(3): 1068–1079, 2022.

[CCG14] Bradley Carron-Arthur, John A. Cunningham, and Kathleen M. Griffiths. "Describing the Distribution of Engagement in an Internet Support Group by Post Frequency: A comparison of the 90-9-1 Principle and Zipf's Law." *Internet Interventions*, **1**(4): 165–168, 2014.

[CDM17] Tanmoy Chakraborty, Ayushi Dalmia, Animesh Mukherjee, and Niloy Ganguly. "Metrics for Community Analysis: A Survey." *ACM Computing Surveys*, **50**(4): 54, 2017.

[CFL09] Claudio Castellano, Santo Fortunato, and Vittorio Loreto. "Statistical Physics of Social Dynamics." *Reviews of Modern Physics*, **81**: 591–646, 2009.

[CH56] Dorwin Cartwright and Frank Harary. "Structural Balance: A Generalization of Heider's Theory." *Psychological Review*, **63**(5): 277–293, 1956.

[CLY20] Chun Cheng, Yun Luo, and Changbin Yu. "Consensus for Expressed and Private Opinions Under Self-Persuasion." *IFAC-PapersOnLine*, **53**(2): 2483–2488, 2020. 21st IFAC World Congress.

[CM11] Daniel Chandler and Rod Munday. *A Dictionary of Media and Communication*. Oxford University Press, Oxford, United Kingdom, 2011.

[Con06] Philip E. Converse. "The Nature of Belief Systems in Mass Publics (1964)." *Critical Review*, **18**(1–3): 1–74, 2006.

[CP23] Weiqi Chu and Mason A. Porter. "A Density Description of a Bounded-Confidence Model of Opinion Dynamics on Hypergraphs." *SIAM Journal on Applied Mathematics*, **83**(6): 2310–2328, 2023.

[CPP21] Young-Pil Choi, Alessandro Paolucci, and Cristina Pignotti. "Consensus of the Hegselmann–Krause Opinion Formation Model with Time Delay." *Mathematical Methods in the Applied Sciences*, **44**(6): 4560–4579, 2021.

[CQ22] Dino Carpentras and Michael Quayle. "Propagation of Measurement Error in Opinion Dynamics Models: The Case of the Deffuant Model." *Physica A: Statistical Mechanics and its Applications*, **606**: 127993, 2022.

[CS20] Ginny Seung Choi and Virgil Henry Storr. "Market Interactions, Trust and Reciprocity." *PLoS ONE*, **15**(5): e0232704, 2020.

[CSD20] Ge Chen, Wei Su, Songyuan Ding, and Yiguang Hong. "Heterogeneous Hegselmann–Krause Dynamics with Environment and Communication Noise." *IEEE Transactions on Automatic Control*, **65**(8): 3409–3424, 2020.

[CSM20] Ge Chen, Wei Su, Wenjun Mei, and Francesco Bullo. "Convergence Properties of the Heterogeneous Deffuant–Weisbuch Model." *Automatica*, **114**: 108825, 2020.

[CTS13] Adrián Carro, Raúl Toral, and Maxi San Miguel. "The Role of Noise and Initial Conditions in the Asymptotic Solution of a Bounded Confidence, Continuous-Opinion Model." *Journal of Statistical Physics*, **151**(1): 131–149, 2013.

[CW17] Bernard Chazelle and Chu Wang. "Inertial Hegselmann–Krause Systems." *IEEE Transactions on Automatic Control*, **62**(8): 3905–3913, 2017.

[CW22] Johan Ugander Christopher Musco, Indu Ramesh and R. Teal Witter. "How to Quantify Polarization in Models of Opinion Dynamics." In *Proceedings of the 17th International Workshop on Mining and Learning with Graphs (MLG)*, Washington DC, USA, 2022. Association for Computing Machinery, New York, NY, USA.

[CWY21] Tinggui Chen, Yulong Wang, Jianjun Yang, and Guodong Cong. "Modeling Multidimensional Public Opinion Polarization Process under the Context of Derived Topics." *International Journal of Environmental Research and Public Health*, **18**(2): 472, 2021.

[CZ15] Andrés Chacoma and Damián H. Zanette. "Opinion Formation by Social Influence: From Experiments to Modeling." *PLOS ONE*, **10**(10): e0140406, 2015.

[DAW02] Guillaume Deffuant, Frédéric Amblard, and Gérard Weisbuch. "How Can Extremism Prevail? A Study Based on the Relative Agreement Interaction Model." *Journal of Artificial Societies and Social Simulation*, **5**(4): 1, 2002.

[DFB23] Alina Dubovskaya, Susan C. Fennell, Kevin Burke, James P. Gleeson, and Doireann O'Kiely. "Analysis of Mean-Field Approximation for Deffuant Opinion Dynamics on Networks." *SIAM Journal on Applied Mathematics*, **83**(2): 436–459, 2023.

[Dit01] Jan Christian Dittmer. "Consensus Formation Under Bounded Confidence." *Nonlinear Analysis: Theory, Methods & Applications*, **47**(7): 4615–4621, 2001.

[DNA00] Guillaume Deffuant, David Neau, Frederic Amblard, and Gérard Weisbuch. "Mixing Beliefs Among Interacting Agents." *Advances in Complex Systems*, **3**(01n04): 87–98, 2000.

[DSC17] Michela Del Vicario, Antonio Scala, Guido Caldarelli, H. Eugene Stanley, and Walter Quattrociocchi. "Modeling Confirmation Bias and Polarization." *Scientific Reports*, **7**(1): 40391, 2017.

[DWC22] S. Davidovits, C. Weber, and D. Clark. "Modeling Ablator Grain Structure in ICF Implosions." *Physics of Plasmas*, **29**(11): 112708, 2022.

[EBN13] Seyed Rasoul Etesami, Tamer Başar, Angelia Nedić, and Behrouz Touri. "Termination Time of Multidimensional Hegselmann–Krause Opinion Dynamics." In *2013 American Control Conference*, 1255–1260, Washington, DC, USA, 2013. Institute of Electrical and Electronics Engineers (IEEE), Piscataway, NJ, USA.

[ER59] P. Erdős and A. Rényi. "On Random Graphs I." *Publicationes Mathematicae Debrecen*, **6**: 290–297, 1959.

[ER60] P. Erdős and A. Rényi. "On the Evolution of Random Graphs." *Publications of the Mathematical Institute of the Hungarian Academy of Sciences*, **5**: 17–60, 1960.

[Ete19] S. Rasoul Etesami. "A Simple Framework for Stability Analysis of State-Dependent Networks of Heterogeneous Agents." *SIAM Journal on Control and Optimization*, **57**(3): 1757–1782, 2019.

[FBQ21] Susan C. Fennell, Kevin Burke, Michael Quayle, and James P. Gleeson. "Generalized Mean-Field Approximation for the Deffuant Opinion Dynamics Model on Networks." *Physical Review E*, **103**: 012314, 2021.

[FGR16] Seth Flaxman, Sharad Goel, and Justin M. Rao. "Filter Bubbles, Echo Chambers, and Online News Consumption." *Public Opinion Quarterly*, **80**(S1): 298–320, 2016.

[FH16] Santo Fortunato and Darko Hric. "Community Detection in Networks: A User Guide." *Physics Reports*, **659**: 1–44, 2016.

[FLP05] Santo Fortunato, Vito Latora, Alessandro Pluchino, and Andrea Rapisarda. "Vector Opinion Dynamics in a Bounded Confidence Consensus Model." *International Journal of Modern Physics C*, **16**(10): 1535–1551, 2005.

[For05] Santo Fortunato. "On the Consensus Threshold for the Opinion Dynamics of Krause–Hegselmann." *International Journal of Modern Physics C*, **16**(2): 259–270, 2005.

[Fri15] Noah E. Friedkin. "The Problem of Social Control and Coordination of Complex Systems in Sociology: A Look at the Community Cleavage Problem." *IEEE Control Systems Magazine*, **35**(3): 40–51, 2015.

[GAP13] Jennifer L. Glanville, Matthew A. Andersson, and Pamela Paxton. "Do Social Connections Create Trust? An Examination Using New Longitudinal Data." *Social Forces*, **92**(2): 545–562, 2013.

[GCB20] Mattia Gasparini, Robert Clarisó, Marco Brambilla, and Jordi Cabot. "Participation Inequality and the 90-9-1 Principle in Open Source." In *OpenSym 2020: Proceedings of the 16th International Symposium on Open Collaboration*, OpenSym 2020, Virtual Conference, Spain, 2020. Association for Computing Machinery, New York, NY, USA.

[GGL12] Javier Gómez-Serrano, Carl Graham, and Jean-Yves Le Boudec. "The Bounded Confidence Model of Opinion Dynamics." *Mathematical Models and Methods in Applied Sciences*, **22**(2): 1150007, 2012.

[Gil59] E. N. Gilbert. "Random Graphs." *The Annals of Mathematical Statistics*, **30**(4): 1141–1144, 1959.

[GOD21] Mirta Galesic, Henrik Olsson, Jonas Dalege, Tamara van der Does, and Daniel L. Stein. "Integrating Social and Cognitive Aspects of Belief Dynamics: Towards a Unifying Framework." *Journal of The Royal Society Interface*, **18**: 20200857, 2021.

[GS12] David F. Gleich and C. Seshadhri. "Vertex Neighborhoods, Low Conductance Cuts, and Good Seeds for Local Community Methods." In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '12, 597–605, Beijing, China, 2012. Association for Computing Machinery, New York, NY, USA.

[GTC09] Lei Guo, Enhua Tan, Songqing Chen, Xiaodong Zhang, and Yihong (Eric) Zhao. "Analyzing Patterns of User Content Generation in Online Social Networks." In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '09, 369–378, Paris, France, 2009. Association for Computing Machinery, New York, NY, USA.

[Hag89] William W. Hager. "Updating the Inverse of a Matrix." *SIAM Review*, **31**(2): 221–239, 1989.

[HDH18] Changwei Huang, Qionglin Dai, Wenchen Han, Yuee Feng, Hongyan Cheng, and Haihong Li. "Effects of Heterogeneous Convergence Rate on Consensus in Opinion Dynamics." *Physica A: Statistical Mechanics and its Applications*, **499**: 428–435, 2018.

[HDJ08] S. Huet, G. Deffuant, and W. Jager. "A Rejection Mechanism in 2D Bounded Confidence Provides More Conformity." *Advances in Complex Systems*, **11**(04): 529–549, 2008.

[HDZ12] J. Heitzig, J. F. Donges, Y. Zou, N. Marwan, and J. Kurths. "Node-Weighted Measures for Complex Networks with Spatially Embedded, Sampled, or Differently Sized Nodes." *The European Physical Journal B*, **85**(1): 38, 2012.

[Hei46] Fritz Heider. "Attitudes and Cognitive Organization." *The Journal of Psychology*, **21**(1): 107–112, 1946.

[HFQ20] Wenchen Han, Yuee Feng, Xiaolan Qian, Qihui Yang, and Changwei Huang. "Clusters and the Entropy in Opinion Dynamics on Complex Networks." *Physica A: Statistical Mechanics and its Applications*, **559**: 125033, 2020.

[HK02] Rainer Hegselmann and Ulrich Krause. "Opinion Dynamics and Bounded Confidence: Models, Analysis and Simulation." *Journal of Artificial Societies and Social Simulation*, **5**(3): 2, 2002.

[HK19] Rainer Hegselmann and Ulrich Krause. "Consensus and Fragmentation of Opinions with a Focus on Bounded Confidence." *The American Mathematical Monthly*, **126**(8): 700–716, 2019.

[HKB22] Abigail Hickok, Yacoub Kureh, Heather Z. Brooks, Michelle Feng, and Mason A. Porter. "A Bounded-Confidence Model of Opinion Dynamics on Hypergraphs." *SIAM Journal on Applied Dynamical Systems*, **21**(1): 1–32, 2022.

[HL15] Petter Holme and Fredrik Liljeros. "Mechanistic Models in Computational Social Science." *Frontiers in Physics*, **3**: 78, 2015.

[HLJ21] Jian Hou, Wenshan Li, and Mingyue Jiang. "Opinion Dynamics in Modified Expressed and Private Model with Bounded Confidence." *Physica A: Statistical Mechanics and its Applications*, **574**: 125968, 2021.

[HLL83] Paul W. Holland, Kathryn Blackmond Laskey, and Samuel Leinhardt. "Stochastic Blockmodels: First Steps." *Social Networks*, **5**(2): 109–137, 1983.

[Hor62] Irving Louis Horowitz. "Consensus, Conflict and Cooperation: A Sociological Inventory." *Social Forces*, **41**(2): 177–188, 1962.

[HS12] Petter Holme and Jari Saramäki. "Temporal Networks." *Physics Reports*, **519**(3): 97–125, 2012.

[JA05] Wander Jager and Frédéric Amblard. "Uniformity, Bipolarization and Pluriformity Captured as Generic Stylized Behavior with an Agent-Based Simulation Model of Attitude Change." *Computational & Mathematical Organization Theory*, **10**(4): 295–303, 2005.

[Jac06] Dirk Jacobmeier. "Focusing of Opinions in the Deffaunt Model: First Impression Counts." *International Journal of Modern Physics C*, **17**(12): 1801–1808, 2006.

[KB08a] Balazs Kozma and Alain Barrat. "Consensus Formation on Adaptive Networks." *Physical Review E*, **77**(1): 016102, 2008.

[KB08b] Balazs Kozma and Alain Barrat. "Consensus Formation on Coevolving Networks: Groups' Formation and Structure." *Journal of Physics A: Mathematical and Theoretical*, **41**(22): 224020, 2008.

[KF23] Claudia Kann and Michelle Feng. "Repulsive Bounded-Confidence Model of Opinion Dynamics in Polarized Communities." e-print arXiv:2301.02210, 2023.

[KFP23] Unchitta Kan, Michelle Feng, and Mason A. Porter. "An Adaptive Bounded-Confidence Model of Opinion Dynamics on Networks." *Journal of Complex Networks*, **11**: cnac055, 2023.

[KO23] Fariba Karimi and Marcos Oliveira. "On the Inadequacy of Nominal Assortativity for Assessing Homophily in Networks." *Scientific Reports*, **13**(1): 21053, 2023.

[Koz22] Ivan V. Kozitsin. "Formal Models of Opinion Formation and their Application to Real Data: Evidence from Online Social Networks." *The Journal of Mathematical Sociology*, **46**(2): 120–147, 2022.

[Koz23] Ivan V. Kozitsin. "Opinion Dynamics of Online Social Network Users: A Micro-Level Analysis." *The Journal of Mathematical Sociology*, **47**(1): 1–41, 2023.

[KP20] Yacoub H. Kureh and Mason A. Porter. "Fitting in and Breaking up: A Nonlinear Version of Coevolving Voter Models." *Physical Review E*, **101**: 062303, 2020.

[Kra00] Ulrich Krause. "A Discrete Nonlinear and Non-Autonomous Model of Consensus." In Saber N. Elaydi, Jerry Popenda, and Jerry Rakowski, editors, *Communications in Difference Equations: Proceedings of the Fourth International Conference on Difference Equations*, 227–236. CRC Press, Boca Raton, Florida, 2000.

[Kur95] Timur Kuran. *Private Truths, Public Lies*. Harvard University Press, Cambridge, MA, USA, 1995.

[KZP12] Gang Kou, Yiyi Zhao, Yi Peng, and Yong Shi. "Multi-Level Opinion Dynamics Under Bounded Confidence." *PLOS ONE*, **7**(9): e43507, 2012.

[LA18] S. Lehmann and Y.-Y. Ahn. *Complex Spreading Phenomena in Social Systems: Influence and Contagion in Real-World Social Networks*. Springer International Publishing, Cham, Switzerland, 2018.

[LAZ04] M. F. Laguna, Guillermo Abramson, and Damián H. Zanette. "Minorities in a Model for Opinion Formation." *Complexity*, **9**(4): 31–36, 2004.

[LD24] Grace J. Li and Seth Davidovits. "Microphysics of Shock-Grain Interaction for ICF Ablators in a Fluid Approach." (submitted), 2024.

[LHK10] Jure Leskovec, Daniel Huttenlocher, and Jon Kleinberg. "Signed Networks in Social Media." In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '10, 1361–1370, Atlanta, GA, USA, 2010. Association for Computing Machinery, New York, NY, USA.

[LHM17] Dandan Li, Dun Han, Jing Ma, Mei Sun, Lixin Tian, Timothy Khouw, and H. Eugene Stanley. "Opinion Dynamics in Activity-Driven Networks." *Europhysics Letters*, **120**(2): 28002, 2017.

[LK14] Jure Leskovec and Andrej Krevl. "SNAP Datasets: Stanford Large Network Dataset Collection." `http://snap.stanford.edu/data`, June 2014. Stanford University. Last Accessed: 1 May 2024.

[LLP23] Grace J. Li, Jiajie Luo, and Mason A. Porter. "Bounded-Confidence Models of Opinion Dynamics with Adaptive Confidence Bounds." e-print arXiv:2303.07563, 2023.

[LMM17] Diandian Li, Liang Meng, and Qingguo Ma. "Who Deserves my Trust? Cue-Elicited Feedback Negativity Tracks Reputation Learning in Repeated Social Interactions." *Frontiers in Human Neuroscience*, **11**: 307, 2017.

[Lor05] Jan Lorenz. "A Stabilization Theorem for Dynamics of Continuous Opinions." *Physica A: Statistical Mechanics and its Applications*, **355**(1): 217–223, 2005.

[Lor06] Jan Lorenz. "Continuous Opinion Dynamics of Multidimensional Allocation Problems under Bounded Confidence. More Dimensions Lead to Better Chances for Consensus." *European Journal of Economic and Social Systems*, **19**(2): 213–227, 2006.

[Lor07] Jan Lorenz. "Continuous Opinion Dynamics Under Bounded Confidence: A Survey." *International Journal of Modern Physics C*, **18**(12): 1819–1838, 2007.

[Lor10] Jan Lorenz. "Heterogeneous Bounds of Confidence: Meet, Discuss and find Consensus!" *Complexity*, **15**(4): 43–52, 2010.

[LP23] Grace J. Li and Mason A. Porter. "Bounded-Confidence Model of Opinion Dynamics with Heterogeneous Node-Activity Levels." *Physical Review Research*, **5**: 023179, 2023.

[LVC14] Carlos Lozares, Joan Miquel Verd, Irene Cruz, and Oriol Barranco. "Homophily and Heterophily in Personal Networks. From Mutual Acquaintance to Relationship Intensity." *Quality & Quantity*, **48**(5): 2657–2670, 2014.

[LX17] Jin Li and Renbin Xiao. "Agent-Based Modelling Approach for Multidimensional Opinion Polarization in Collective Behaviour." *Journal of Artificial Societies and Social Simulation*, **20**(2): 4, 2017.

[LYW16] Kristina Lerman, Xiaoran Yan, and Xin-Zeng Wu. "The "Majority Illusion" in Social Networks." *PLOS ONE*, **11**(2): e0147617, 2016.

[Mas19] Michael Mäs. "Challenges to Simulation Validation in the Social Sciences. A Critical Rationalist Perspective." In Claus Beisbart and Nicole J. Saam, editors, *Computer Simulation Validation: Fundamental Concepts, Methodological Frameworks, and Philosophical Perspectives*, 857–879. Springer International Publishing, Cham, Switzerland, 2019.

[MDB20] Corrado Monti, Gianmarco De Francisci Morales, and Francesco Bonchi. "Learning Opinion Dynamics from Social Traces." In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD

'20, 764–773, Virtual Event, CA, USA, 2020. Association for Computing Machinery, New York, NY, USA.

[MGR10] Naoki Masuda, N. Gibert, and S. Redner. "Heterogeneous Voter Models." *Physical Review E*, **82**: 010103, 2010.

[Mie14] Trevor van Mierlo. "The 1% Rule in Four Digital Health Social Networks: An Observational Study." *Journal of Medical Internet Research*, **16**(2): e33, 2014.

[MSC01] Miller McPherson, Lynn Smith-Lovin, and James M Cook. "Birds of a Feather: Homophily in Social Networks." *Annual Reviews of Sociology*, **27**(1): 415–444, 2001.

[MVP18] X. Flora Meng, Robert A. Van Gorder, and Mason A. Porter. "Opinion Formation and Distribution in a Bounded-Confidence Model on Various Networks." *Physical Review E*, **97**(2): 022312, 2018.

[New03] M. E. J. Newman. "Mixing Patterns in Networks." *Physical Review E*, **67**: 026126, 2003.

[New04] M. E. J. Newman. "Analysis of Weighted Networks." *Physical Review E*, **70**: 056131, 2004.

[New05] M. E. J. Newman. "Power Laws, Pareto Distributions and Zipf's law." *Contemporary Physics*, **46**(5): 323–351, 2005.

[New06] M. E. J. Newman. "Finding Community Structure in Networks Using the Eigenvectors of Matrices." *Physical Review E*, **74**: 036104, 2006.

[New18] Mark Newman. *Networks*. Oxford University Press, Oxford, United Kingdom, 2nd edition, 2018.

[NGW23] Andrew Nugent, Susana N. Gomes, and Marie-Therese Wolfram. "On Evolving Network Models and their Influence on Opinion Formation." *Physica D: Nonlinear Phenomena*, **456**: 133914, 2023.

[Nie06] Jakob Nielsen. "The 90-9-1 Rule for Participation Inequality in Social Media and Online Communities." `https://www.nngroup.com/articles/participation-inequality/`, Oct 2006. Nielsen Norman Group. Last Accessed: 7 Jan 2022.

[Noo20] Hossein Noorazar. "Recent Advances in Opinion Propagation Dynamics: A 2020 Survey." *The European Physical Journal Plus*, **135**(6): 521, 2020.

[NVT20] Hossein Noorazar, Kevin R. Vixie, Arghavan Talebanpour, and Yunfeng Hu. "From Classical to Modern Opinion Dynamics." *International Journal of Modern Physics C*, **31**(7): 2050101, 2020.

[OAS10] Tore Opsahl, Filip Agneessens, and John Skvoretz. "Node Centrality in Weighted Networks: Generalizing Degree and Shortest Paths." *Social Networks*, **32**(3): 245–251, 2010.

[PFT18] Rohit Parasnis, Massimo Franceschetti, and Behrouz Touri. "Hegselmann–Krause Dynamics with Limited Connectivity." In *2018 IEEE Conference on Decision and Control (CDC)*, 5364–5369, Miami, FL, USA, 2018. Institute of Electrical and Electronics Engineers (IEEE), Piscataway, NJ, USA.

[PFT22] Rohit Yashodhar Parasnis, Massimo Franceschetti, and Behrouz Touri. "On the Convergence Properties of Social Hegselmann–Krause Dynamics." *IEEE Transactions on Automatic Control*, **67**(2): 589–604, 2022.

[PGP12] N. Perra, B. Gonçalves, R. Pastor-Satorras, and A. Vespignani. "Activity Driven Modeling of Time Varying Networks." *Scientific Reports*, **2**(1): 469, 2012.

[PKI22] Antonio F. Peralta, János Kertész, and Gerardo Iñiguez. "Opinion Dynamics in Social Networks: From Models to Data." e-print arXiv:2201.01322; to appear as a chapter in Taha Yasseri, editor, *Handbook of Computational Social Science*, Cheltenham, UK, 2024, Edward Elgar Publishing Ltd., 2022.

[PPT15] Sergey E. Parsegov, Anton V. Proskurnikov, Roberto Tempo, and Noah E. Friedkin. "A New Model of Opinion Dynamics for Social Actors with Multiple Interdependent Attitudes and Prejudices." In *2015 54th IEEE Conference on Decision and Control (CDC)*, 3475–3480, Osaka, Japan, 2015. Institute of Electrical and Electronics Engineers (IEEE), Piscataway, NJ, USA.

[PRM22] Valentina Pansanella, Giulio Rossetti, and Letizia Milli. "Modeling Algorithmic Bias: Simplicial Complexes and Evolving Network Topologies." *Applied Network Science*, **7**(1): 57, 2022.

[RKK07] Sini Ruohomaa, Lea Kutvonen, and Eleni Koutrouli. "Reputation Management Survey." In *The Second International Conference on Availability, Reliability and Security*, ARES'07, 103–111, Vienna, Austria, 2007. Institute of Electrical and Electronics Engineers (IEEE), Piscataway, NJ, USA.

[RKM11] Veronica Red, Eric D. Kelsic, Peter J. Mucha, and Mason A. Porter. "Comparing Community Structure to Characteristics in Online Collegiate Social Networks." *SIAM Review*, **53**(3): 526–543, 2011.

[RKZ00] Paul Resnick, Ko Kuwabara, Richard Zeckhauser, and Eric Friedman. "Reputation Systems." *Communications of the ACM*, **43**(12): 45–48, 2000.

[SBL23] Jakub Sawicki, Rico Berner, Sarah A. M. Loos, Mehrnaz Anvari, Rolf Bader, Wolfram Barfuss, Nicola Botta, Nuria Brede, Igor Franović, Daniel J. Gauthier,

Sebastian Goldt, Aida Hajizadeh, Philipp Hövel, Omer Karin, Philipp Lorenz-Spreen, Christoph Miehl, Jan Mölter, Simona Olmi, Eckehard Schöll, Alireza Seif, Peter A. Tass, Giovanni Volpe, Serhiy Yanchuk, and Jürgen Kurths. "Perspectives on Adaptive Dynamical Systems." *Chaos*, **33**(7): 071501, 2023.

[SFH21] Hendrik Schawe, Sylvain Fontaine, and Laura Hernández. "When Network Bridges Foster Consensus. Bounded Confidence Models in Networked Societies." *Physical Review Research*, **3**: 023208, 2021.

[SGS20] Simon Schweighofer, David Garcia, and Frank Schweitzer. "An Agent-Based Model of Multi-Dimensional Opinion Dynamics and Opinion Alignment." *Chaos: An Interdisciplinary Journal of Nonlinear Science*, **30**(9): 093139, 2020.

[SGW17] Wei Su, YaJuan Gu, Sha Wang, and YongGuang Yu. "Partial Convergence of Heterogeneous Hegselmann–Krause Opinion Dynamics." *Science China Technological Sciences*, **60**(9): 1433–1438, 2017.

[Sha13] Yilun Shang. "Deffuant Model with General Opinion Distributions: First Impression and Critical Confidence Bound." *Complexity*, **19**(2): 38–49, 2013.

[Sha14] Yilun Shang. "An Agent Based Model for Opinion Dynamics with Random Confidence Threshold." *Communications in Nonlinear Science and Numerical Simulation*, **19**(10): 3766–3777, 2014.

[SLL14] Jiongming Su, Baohong Liu, Qi Li, and Hongxu Ma. "Coevolution of Opinions and Directed Adaptive Networks in a Social Group." *Journal of Artificial Societies and Social Simulation*, **17**(2): 4, 2014.

[SLS17] Alina Sîrbu, Vittorio Loreto, Vito D. P. Servedio, and Francesca Tria. "Opinion Dynamics: Models, Extensions and External Effects." In Vittorio Loreto, Muki Haklay, Andreas Hotho, Vito D. P. Servedio, Gerd Stumme, Jan Theunis, and Francesca Tria, editors, *Participatory Sensing, Opinions and Collective Awareness*, 363–401. Springer International Publishing, Cham, Switzerland, 2017.

[SM00] Jianbo Shi and Jitendra Malik. "Normalized Cuts and Image Segmentation." *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **22**(8): 888–905, 2000.

[SNP13] Wanita Sherchan, Surya Nepal, and Cecile Paris. "A Survey of Trust in Social Networks." *ACM Computing Surveys*, **45**(4): 47, 2013.

[Sob15] Pawel Sobkowicz. "Extremism Without Extremists: Deffuant Model with Emotions." *Frontiers in Physics*, **3**: 17, 2015.

[SPG19] Alina Sîrbu, Dino Pedreschi, Fosca Giannotti, and János Kertész. "Algorithmic Bias Amplifies Opinion Fragmentation and Polarization: A Bounded Confidence Model." *PLOS ONE*, **14**(3): e0213246, 2019.

[Spo17] Dominic Spohr. "Fake News and Ideological Polarization: Filter Bubbles and Selective Exposure on Social Media." *Business Information Review*, **34**(3): 150–160, 2017.

[SR51] Ray Solomonoff and Anatol Rapoport. "Connectivity of Random Nets." *The Bulletin of Mathematical Biophysics*, **13**(2): 107–117, 1951.

[SRP23] Trevor Steil, Tahsin Reza, Benjamin Priest, and Roger Pearce. "Embracing Irregular Parallelism in HPC with YGM." In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '23, Denver, CO, USA, 2023. Association for Computing Machinery, New York, NY, USA.

[SSI20] Anuj Singh, Rishi Ranjan Singh, and S. R. S. Iyengar. "Node-Weighted Centrality: A New Way of Centrality Hybridization." *Computational Social Networks*, **7**: 6, 2020.

[TFM16] Károly Takács, Andreas Flache, and Michael Mäs. "Discrepancy and Disliking Do Not Induce Negative Opinion Shifts." *PLOS ONE*, **11**(6): e0157948, 2016.

[TMP12] Amanda L. Traud, Peter J. Mucha, and Mason A. Porter. "Social Structure of Facebook Networks." *Physica A: Statistical Mechanics and its Applications*, **391**(16): 4165–4180, 2012.

[Vaz22] Federico Vazquez. "Modeling and Analysis of Social Phenomena: Challenges and Possible Research Directions." *Entropy*, **24**(4): 491, 2022.

[VBI21] Francesco Vasca, Carmela Bernardo, and Raffaele Iervolino. "Practical Consensus in Bounded Confidence Opinion Dynamics." *Automatica*, **129**: 109683, 2021.

[VMG16] Corentin Vande Kerckhove, Samuel Martin, Pascal Gend, Peter J. Rentfrow, Julien M. Hendrickx, and Vincent D. Blondel. "Modelling Influence and Opinion Evolution in Online Collective Behaviour." *PLOS ONE*, **11**(6): e0157685, 2016.

[VWL07] A. L. Velikovich, J. G. Wouchuk, C. Huete Ruiz de Lira, N. Metzler, S. Zalesak, and A. J. Schmitt. "Shock Front Distortion and Richtmyer–Meshkov-Type Growth Caused by a Small Preshock Nonuniformity." *Physics of Plasmas*, **14**(7): 072706, 2007.

[WDA02] Gérard Weisbuch, Guillaume Deffuant, Frédéric Amblard, and Jean-Pierre Nadal. "Meet, Discuss, and Segregate!" *Complexity*, **7**(3): 55–63, 2002.

[WDA03] Gérard Weisbuch, Guillaume Deffuant, Frederic Amblard, and Jean Pierre Nadal. "Interacting Agents and Continuous Opinions Dynamics." In Robin Cowan and Nicolas Jonard, editors, *Heterogeneous Agents, Interactions and Economic Performance*, 225–242. Springer-Verlag, Heidelberg, Germany, 2003.

[WGD13] Joyce Jiyoung Whang, David F. Gleich, and Inderjit S. Dhillon. "Overlapping Community Detection Using Seed Set Expansion." In *Proceedings of the 22nd ACM International Conference on Information & Knowledge Management*, CIKM '13, 2099–2108, San Francisco, CA, USA, 2013. Association for Computing Machinery, New York, NY, USA.

[WH19] Stefan Wojcik and Adam Hughes. "Sizing Up Twitter Users." `https://www.pewresearch.org/internet/2019/04/24/sizing-up-twitter-users/`, 2019. Pew Research Center. Last Accessed: 31 May 2021.

[WLE17] Chu Wang, Qianxiao Li, Weinan E, and Bernard Chazelle. "Noisy Hegselmann–Krause Systems: Phase Transition and the 2R-Conjecture." *Journal of Statistical Physics*, **166**(5): 1209–1225, 2017.

[XL14] Fei Xiong and Yun Liu. "Opinion formation on Social Media: An Empirical Approach." *Chaos: An Interdisciplinary Journal of Nonlinear Science*, **24**(1): 013130, 2014.

[XLC17] Fei Xiong, Yun Liu, and Junjun Cheng. "Modeling and Predicting Opinion Formation with Trust Propagation in Online Social Networks." *Communications in Nonlinear Science and Numerical Simulation*, **44**: 513–524, 2017.

[YDH14] Yuecheng Yang, Dimos V. Dimarogonas, and Xiaoming Hu. "Opinion Consensus of Modified Hegselmann–Krause Models." *Automatica*, **50**(2): 622–627, 2014.

[YL15] Jaewon Yang and Jure Leskovec. "Defining and Evaluating Network Communities Based on Ground-Truth." *Knowledge and Information Systems*, **42**(1): 181–213, 2015.

[YTL20] Mengbin Ye, Minh Hoang Trinh, Young-Hun Lim, Brian D.O. Anderson, and Hyo-Sung Ahn. "Continuous-Time Opinion Dynamics on Multiple Interdependent Topics." *Automatica*, **115**: 108884, 2020.

[ZC15] Jiangbo Zhang and Ge Chen. "Convergence Rate of the Asymmetric Deffuant–Weisbuch Dynamics." *Journal of Systems Science and Complexity*, **28**(4): 773–787, 2015.

[Zha14] Jiangbo Zhang. "Convergence Analysis for Asymmetric Deffuant–Weisbuch Model." *Kybernetika*, **50**(1): 32–45, 2014.

[ZXL18] Jun Zhang, Haoxiang Xia, and Peng Li. "Dynamics of Deffuant Model in Activity-Driven Online Social Network." In Jian Chen, Yuji Yamada, Mina Ryoke, and Xijin Tang, editors, *19th International Symposium Knowledge and Systems Sciences*, KSS 2018, 215–224, Tokyo, Japan, 2018. Communications in Computer and Information Science, **949**, Springer, Singapore.