A Survey-Informed Evolutionary Opinion Dynamics Model of Political Activism with an Application to the 2022 Panamanian Protests

Arturo F. Serrano Borrero

Advisor: Olivia J. Chu

Department of Mathematics, Dartmouth College

May 26, 2024

Table of Contents

0.1 Abstract

The formation of activist groups can spark social movements, coalition building, and revolutions. Although exogenous shocks to a system can often cause individuals to develop or strengthen their commitment to a cause, the mechanisms through which like-minded individuals build confidence in their views and establish social connections is not thoroughly understood. In this work, we develop a theory that begins to explain two phenomena: 1) how an activist's conviction co-evolves with their immediate social network, and 2) how so-called "politically viable networks" (PVNs), or clusters of activists, tend to arise or disappear based on the distribution of potential activists and the sociopolitical environment. We illustrate this theory by modifying the adaptive voter model (AVM) with an evolutionary game theory-inspired conviction variable, which represents both the strength with which an individual holds on to their beliefs and the comfort of holding on to such beliefs in the context of their network, encapsulating the co-evolutionary dynamics of networks and sociopolitical attitudes. In this thesis, we run systematic simulations of this conviction-moderated adaptive voter model (CMAVM). As is expected from empirical evidence in countries with stable pre-shock socio-political landscapes, we find that activists are often systematically discouraged by their exposure to disengaged individuals. However, we also identify situations that favor their survival within society as members of politically viable networks, chiefly stronger homophily preferences and higher initial activist conviction. Next, we propose three mechanisms through which to leverage survey data to introduce realistic heterogeneity using: 1) geographic location to inform graph structure and connections; 2) demographic information to endow political node attributes such as opinion and conviction; and 3) identity to differentiate the level of impact each interaction has based on the perceived similarity or difference among the individuals involved. A parameter sweep identifies that the introduction of heterogeneity magnifies or dampens the influence of certain parameters, informing our views of what "potential successful pathways" towards mobilization might look like. Finally, we build a Panama-specific model, motivated by the protests of 2022, that illustrates how the introduction of heterogeneity—informed by reliable survey data—describes several different potential evolutionary pathways that this movement may have followed. Overall, this work contributes to strengthening our understanding of social movements and other social issues where buy-in matters.

0.2 Dedication

Esta tesis de honores va dedicada a mi madre patria, Panamá, y a todos sus habitantes nobles y trabajadores, cuyas idiosincrasias, motivaciones, frustraciones y acciones motivan este proyecto.

This honors thesis is for my homeland, Panama, and all its noble and hardworking inhabitants whose idiosyncrasies, motivations, frustrations, and actions motivate this project.

0.3 Acknowledgments

First and foremost, I would like to express my deepest gratitude to my advisor, Professor Olivia Chu, whose enthusiastic and caring advice has enabled me through this thesis's research and writing processes. Words do not suffice to express my gratitude for her dedicated mentorship; I thank you for recognizing my research potential in the absence of relevant coursework, for keeping an open door and inbox to my endless flow of questions, for spending hours on curating irreplaceable simultaneously critical and encouraging feedback, for remaining flexible when we encountered significant research challenges along the way, and for helping me navigate my career by investing in my educational, professional, and personal development far beyond the wildest expectations. I will forever remember your spontaneous offer to submit my work for presentation and showcase at the 2024 Joint Mathematical Meeting's Special Session on Undergraduate Research in Mathematical Biology in San Francisco and the vote of confidence that represented. I am also deeply grateful to Mathematics Professors Daniel Rockmore, Ho-Chun Herbert Chang, and Peter Mucha for thoroughly supporting me and my project with invaluable advice and encouragement, and to Dylan Green and Dari Daniel Seo Seo '16 for invaluable input in code proofreading and optimization. I could never have asked for a better set of sponsors.

I am also grateful to the Department of Mathematics for all of the support they have provided me throughout my time at Dartmouth College. I am particularly grateful for the continuous encouragement I have received from Kyle Schubert, Lucy Knight, and Professors Mitsuo Kobayashi, Rosa Orellana, and Jan Glaubitz and for the help and advice I received from Department Administrator Tracy Moloney, Major Adviser Professor Dana Williams, and Department Chair Professor Sergi Elizalde.

This work would not have been possible without funding from the Mark C. Hansen Undergraduate Creativity Fund, the Neukom Scholar Fund, and the Rockefeller Center Senior Honors Thesis Grant from the Undergraduate Advising and Research, the Neukom Institute for Computational Science, and the Rockefeller Center for Public Policy and the Social Sciences at Dartmouth College, respectively. This project also hinges on data kindly provided by the World Justice Project, and I would like to thank their Director Dr. Alejandro Ponce and Senior Program Associate Natalia Rodríguez Cajamarca. I would also like to acknowledge the support and expertise of numerous experts at the Centro Internacional de Estudios Políticos y Sociales de Panamá (CIEPS), who are at the forefront of quantitative social science research innovation on our country, and Professor Álvaro José Riascos Villegas at the Universidad de Los Andes, Colombia, for his vital advice and mentorship on quantitative research on Latin America over the last six years. The political and sociological accuracy of this model greatly benefited from the long conversations I have had with the many Panamanian citizens throughout the country who entertained my questions, and I am eternally grateful for the time, effort, and thought of my informants. In particular, I would like to kindly thank a few experts of Panamanian public life who gave valuable input on the model: Hugo Woods, Lawyer and Professor of Political Science at the Catholic University of Saint Mary of Panama; H.E. María Fernanda Cortizo Frías, Lawyer and Ambassador of Panama at the Organization of American States; Dr. Nelva Marissa Araúz Reyes, Human Rights Researcher at CIEPS; Daniela Noriega García, President of the Panamanian Debate Association; and Gary Martin, Data Scientist and Co-Founder of the Panamanian new media conglomerate Praxis Communications.

I owe the opportunity to write a thesis like this to those who contributed to my passion for Mathematics before I came to Dartmouth. I am most grateful to Professor Ana Cecilia Tordecilla, indubitably an unsung hero of Mathematics development in the Republic of Panama, who recognized, awakened, and nurtured me as well as my coaches and students within the math team at Academia Interamericana de Panamá (AIP). I hope you see in my work the payoff of the hundreds of hours you invested in me on countless afternoons and weekends over the course of six years. Secondly, I would also like to thank AIP Professors Maruja de Sáiz, Itzy Pamela Lajón, and Renny Vassíleva Ingram for their wonderful instruction and caring support and to the Fundación Olimpiada Panameña de Matemática and the Math Department at the Universidad de Panamá's School of Natural and Exact Sciences for the opportunities to grow as a mathematician through competitions and advanced workshops during foundational years for my mathematical development. I am only one of tens of Panamanian students whose lives have been touched by heroes like them, who time and again take unpaid, additional formal and informal roles as coaches, teachers, mentors, organizers, and researchers in the hopes of improving mathematics education in Panama. I can only hope to pay this kindness forward and contribute myself to the growth of our discipline in our country.

I would like to thank my family and friends, first and foremost among them my parents, Arturo and Patricia, and my sister, Mandy, for supporting me through every step of my education and life. I am also grateful to José Velasco and his family, Xiaoyun Zhu, my grandparents Jorge, Magda, José Luis, and Beatriz, and all my aunts, uncles, and cousins who have extended their love and support from every corner of the world and beyond; this work goes to you all. Last and above all, I would like to thank the Lord my God Whose unconditional love and presence in my life has sustained me through every challenge and brought me to where I am today. In Your Name I pray, Amen.

Chapter 1

Introducing Conviction Moderation of Opinion Dynamics

> "I simply wish that, in a matter which so closely concerns the wellbeing of the human race, no decision shall be made without all the knowledge which a little analysis and calculation can provide."
>
> Daniel Bernoulli

1.1 The Mathematics of Socio-Political Activism

Activism requires the presence of activists—individuals who not only hold an opinion but hold it strongly enough to mobilize themselves and influence others. The emergence of a socio-political movement requires that these activists become embedded within activist networks, which are groups of interconnected individuals who can coordinate their efforts to effect change. The strength and connectivity of these networks are crucial for successful mobilization. However, socio-political stability often thwarts social movements by systematically discouraging activists. Stable systems tend to maintain the status quo and suppress dissent, making it challenging for activists to gain momentum. However, when there are shocks to the system—such as economic crises, political scandals, or social unrest—the frequency and significance of interactions between individuals increase. These shocks lead to more frequent social rewiring, where people seek out and connect with others who share similar views. As individuals find like-minded peers, they start building collective confidence. Over time, these networks can give rise to a politically viable network (PVN), where the growing confidence and mobilization efforts of activists reinforce each other, leading to a self-sustaining movement.

How does this theory translate into mathematics? To understand the mechanisms of association and network formation, we can use mathematical models from disciplines such as graph theory and dynamical systems. These models, collectively known as sociophysics [1, 2] or opinion dynamics models [3], allow us to study how individuals form connections, how opinions spread through networks, and how collective behavior emerges from individual interactions [4]. The models' behaviors are largely influenced by four key factors. First, the network's structure and layout, including the density of connections and the attachment patterns between individuals, determine who interacts with whom and how often, and therefore significantly influence the spread of opinions and the formation of activist networks [5]. Second, the rate at which individuals interact affects how quickly opinions can spread and mobilizations can arise; in particular, shocks to the system can temporarily increase this rate, resulting in more rapid changes in network structure leading to collective behavior [6]. Third, the initial distribution of opinions within the population determines how likely activists are to connect with other activists and how resilient they will be in the face of discouragement. Fourth, individuals' ability and willingness to change their connections based on their interactions and experiences is crucial for the dynamic formation of activist networks; this rewiring process can be modeled using stochastic rules that capture the likelihood of forming or severing ties [7].

To apply these mathematical insights to real-world scenarios, we require thorough and representative real-world data from sources such as survey responses on political opinions and social interactions. Identifying key demographic or psychological traits of individuals and populations that can drive the co-evolution of opinions and social connections is essential to understanding the mechanisms through which activism emerges. Although data-integration can prove challenging at times, this project finds computationally simple ways to leverage survey data at different stages of an opinion dynamics model to more realistically simulate the evolution of each individual's opinion over time and introduce heterogeneity at the individual level. By challenging ourselves to build realistic models from this data, we can iteratively refine our models to better capture the complexities of socio-political activism through calibration, simulation, and model iteration. Through this process, we can develop a deeper understanding of how socio-political activism emerges, spreads, and sustains itself, and the mechanisms that drive it in both general and specific settings.

1.2 Graph and Network Theory

Graph theory is an area of mathematics concerned with the representation of different problems and structures as graphs, which are collections of nodes (also vertices or points) and edges (also links or ties)

[8]. The size of a graph is given by its number of nodes. Edges can be undirected or directed to represent symmetrical or asymmetrical relationships, respectively, between the nodes. Network theory is a specific subdiscipline of graph theory where nodes or edges are endowed with attributes; in particular, social network theory aims to represent the structure of relationships between social elements [9], which are most often individuals but could also be groups, organizations, institutions, nations, or even websites. Node attributes are most often stored as vector-structured data, and can be used to track characteristics; in a social context, these are most often demographic data (e.g. age, gender, race, ethnicity, income, socio-economic status, location), psychological data (e.g. trust or talkativeness), and political data (e.g. political orientation, party affiliation, talkativeness). A key distinction can be drawn in network theory between graphs with weighted and unweighted edges; these weights are often used to designate the relative importance or capacity of an edge in a network. For this thesis, however, edges are always undirected and unweighted, representing symmetric and equal social relationships between individuals. Altogether, graphs can be analyzed according to some of their well-defined properties; those most relevant to social networks include assortativity (the extent to which neighbors share certain attributes), transitivity (the likelihood that if two nodes have a common neighbor then they too are neighbors), and propinquity (the extent to which geographical distance affects the likelihood of nodes being network neighbors). Meanwhile, nodes can also be described according to subtly different yet well standardized measures of centrality, which identify the relative importance of each node.

Another important feature of graph theory is that of induced subgraphs, which are subgraphs of an original graph composed of a certain subset of all nodes and all of the edges from the original graph which connect the nodes of that subset [10]. In social network theory, these induced subgraphs are often used to understand the social structure of specific groups of interest, such as members of a particular political party, social movement, or collective action initiative. In this thesis, politically viable networks (PVNs) are defined as induced subgraphs.

In the political science literature, social networks have been used to understand issues of social cooperation, defection, and collective action since at least 1988 [11, 12]. In particular, networks have been used extensively to describe peace and conflict issues and dynamics, both at the domestic and international levels. For instance, there are models that represent global-scale issues of peace and violence as well as domestic issues of violent, criminal, and terrorist organizations [13, 14, 15, 16]. The latter more closely resemble the problem of activism emergence, and researchers have assessed the role of networks in the recruitment and mobilization of transcendental political developments such as the rise of ethnic revolutions and terrorist organizations. Some projects have focused specifically on how social media or identity-based networks often underpin these movements [17]. Yet, there is little exploration of the role that networks play in the creation of socio-political movements in the absence of repression, war, and similar sources of political violence or when these movements form across the boundaries of particular identities. Thus, this project contributes to social network theory by discussing how networks are important in the creation of socio-political movements across identity cohorts and within political environments that have historically shown very high relative stability without cross-cohort movements or significant, violent repression.

1.3 Opinion Dynamics

Opinion dynamics refers to the study of the co-evolution of opinions and social structures. It arose as a branch of sociophysics, a field of science that employs mathematical models inspired by physics to understand how human societies behave and evolve. This tradition has yielded a number of useful "voter models" to understand how an individual's beliefs can influence those around it, how societal beliefs influence different individuals, and how beliefs become established over time. However, it is important to consider that these models do not necessitate a single voting event such as an election; instead, they describe the opinions individuals have about certain issues. Two important predecessors of the type of model described in the next two chapters of this thesis are discussed below.

1.3.1 The Classic Voter Model

One early representation of opinion dynamics hinged upon assigning "spins" (with values of ± 1) to individuals who occupied certain distinct spaces and letting them influence those who were around them. In sociophysics, each spin was mapped to an opinion (yes or no), and each space was mapped to a voter. Among the first of such models was the "voter model," proposed by Richard A. Holley and Thomas M. Liggett [18]. In this iterated discrete-time model, a voter is placed at each node of a graph and assigned an opinion. Then, the model randomly selects an individual—a random voter—whose new, updated opinion will be determined by its neighbors according to a stochastic rule. Specifically, one of its neighbors is randomly chosen according to a set probability distribution, and the neighbor's opinion is effectively transferred (or spread) to the random voter, as illustrated in Figure 1.1. Because only one opinion change happens at each discrete time step, these models are often represented in terms of coalescing Markov Chains. As such, voter models can be studied as dual and branching processes, which provide insights into their long-term behavior and the probability of different outcomes.

An important feature of these models is that they tend toward consensus; that is, information tends to flow outward from where an opinion is held, and populations reach a steady state only when one opinion becomes dominant. This phenomenon illustrates how, over time, a single opinion can come to dominate an entire population. Despite the tendency toward consensus, these models can also exhibit more interesting phenomena such as coexistence and clustering. Coexistence refers to the scenario where multiple opinions persist in different regions of the graph; this is more commonly observed in higher dimensions or for specific graph structures. Meanwhile, clustering refers to the formation of regions where a single opinion is dominant. Importantly, clusters can grow over time, eventually leading to consensus, but the path to consensus can involve complex spatial patterns and temporal dynamics. Still, the structure of the model's graph is extremely important for predicting whether a model can reach consensus and the required steps for it to reach this state. For example, a few naïve arrangements include placing voters in a line graph, or on a square lattice using Von Neumann or Moore neighborhoods [19]. These configurations influence how opinions spread: in a line graph, consensus will form very slowly and with very simple clustering. Von Neumann neighborhoods, because of their two-dimensional arrangements, allow for moderate consensus speed and they display relatively simple, two-dimensional clustering. Moore neighborhoods are very similar, but the slight change in connection rules allows for much faster consensus and complex clustering patterns, showing the importance of spatial relationships in determining a model's outcomes. For instance, the classical voter model on random graphs delivers rapid consensus, unpredictable clustering, and very short coalescence time depending entirely on connectivity. In reality, human systems must be represented by much more complex graphs, and this increases a model's complexity.

The classical voter model is helpful in simulating the spatial spread of opinions and the role of initial social structures in influencing individuals' views and shaping social configurations. However, the model's simple assumptions compromise its applicability in the socio-political context. Notably, the clearest socio-political interpretation of this model is that a naive, unconfident voter yields to a neighbor under social pressure. But social dynamics are more complex; specifically, even in high-pressure situations, individuals can choose to sever a social tie that causes them anxiety and/or establish ties to those who hold similar views. Additionally, most iterations of this simple model presume that all neighbors are equally likely to influence an individual's opinions, ignoring the nuances of weaker versus stronger ties, persuasiveness, and conviction. In light of these limitations, researchers have refined voter models to better capture the intricacies of real-world social dynamics, such as incorporating heterogeneity in influence strength and allowing for the dynamic formation and dissolution of social ties. The latter is captured by the standard adaptive voter model; the former is more of an open question that this thesis will later seek to address.

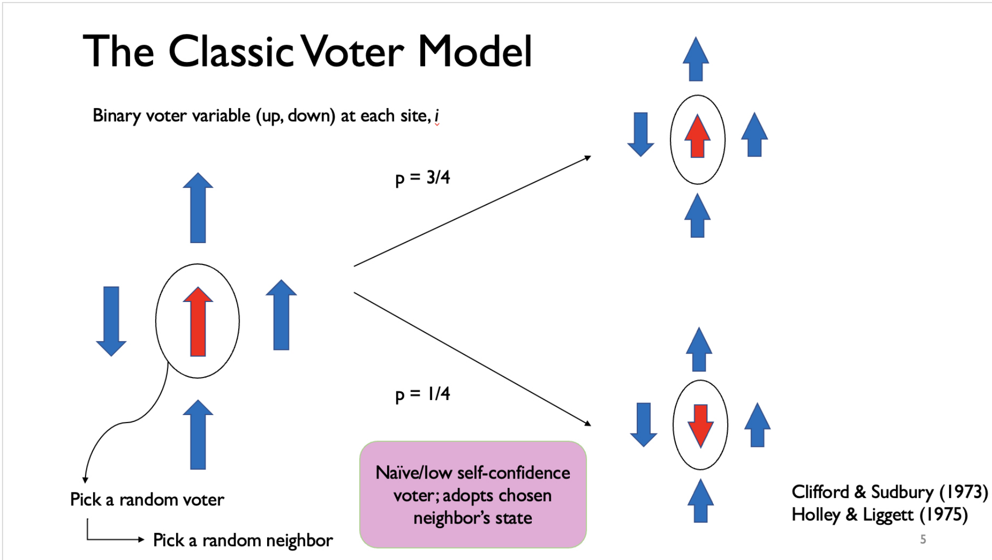

Figure 1.1: A random voter selects a random neighbor and adopts said neighbor's opinion.

1.3.2 The Standard Adaptive Voter Model

A next step in opinion dynamics involves capturing the structural rearrangements of voters based on their opinions alongside the shifts in opinion based on a voter's surroundings. The Adaptive Voter Model (AVM) addresses this need by simulating the processes through which networks influence individuals' opinions and individuals modify their networks according to their connections' stances [20]. At each discrete time step of the AVM, an individual i is randomly selected. With probability ϕ, the individual will engage in social learning, meaning i will adopt the opinion of one of its neighbors. Alternatively, with probability $(1-\phi)$, i will engage in homophily by replacing its link to some j with one to another individual j' who shares i's opinion. In the AVM, ϕ is a global parameter for the population, and its exact value significantly affects the final configuration of the graph's nodes and edges. In the long time limit, the AVM typically yields a consensus state where individuals are only connected to others who share their own opinion. Holme and Newman identified in 2006 that the standard AVM's phase transition occurs at $\phi \approx 0.46$; small values of ϕ result in smaller consensus groups resembling the original opinion distributions, whereas large values of ϕ result in a single, homogeneous component where all individuals converge to the majority opinion.

Despite its simplicity, the AVM can be made significantly more realistic with certain modifications. For instance, Chu et al. propose a variation where two individuals, i and j, are randomly selected as the learner and role model, respectively [21]. In this modified model, if i and j share the same opinion, the networks and opinions remain unchanged. However, if i and j hold different opinions, with probability ϕ, i will engage in social learning, adopting j's opinion. Alternatively, with probability $(1-\phi)$, i will engage in homophily by substituting the link to j with one to another individual j' who shares i's opinion.

This modification inherently selects discordant pairs for potential changes, emphasizing the dynamics of opinion convergence and network restructuring, making it better able to capture the friction inherent in social interactions. Notably, in both the original AVM and its modified versions, ϕ represents the probability of social learning. These models, by incorporating both homophily and social learning mechanisms, provide a powerful framework for studying how social networks evolve and influence opinion dynamics, capturing the complex interplay between individual behaviors and network structures.

The AVM effectively models how opinions and social structures co-evolve. By incorporating both the likelihood of changing one's opinion and the tendency to rewire connections towards like-minded individuals, the model simulates realistic social processes. The parameter ϕ plays a crucial role in determining the model's dynamics. At high ϕ values, social learning dominates, leading to a rapid convergence towards a majority opinion and a highly interconnected social network. Conversely, at low ϕ values, homophily dominates, preserving initial opinion diversity, which results in multiple smaller, homogeneous clusters.

The AVM is not only a robust theoretical tool but also an adaptable framework that can be tailored to study various social phenomena, such as political polarization, cultural assimilation, and the spread of innovations. The AVM and its variations offer a nuanced way to understand the interplay between individual opinions and social structures, providing insights into the mechanisms driving social cohesion and division. This adaptability makes the AVM a valuable model for exploring the dynamics of opinion formation and network evolution in diverse social contexts, and this project will focus on exploring the advantages and limitations of refining the AVM in different ways.

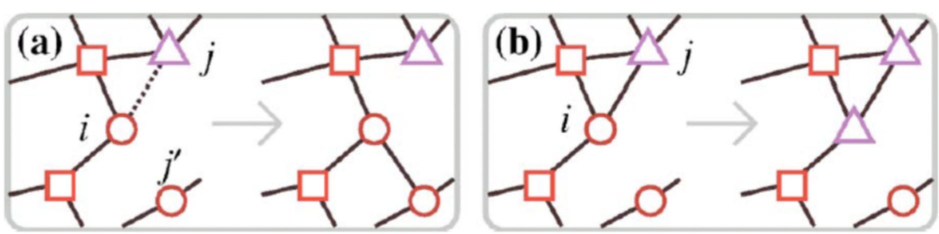

Figure 1.2: In this illustration from Holme and Newman [20], different vertex shapes are used to represent different opinions. At each time step, the system is updated according the processes shown denoted as (a) and (b). With probability $1-\phi$, homophily occurs as shown by panel (a). With probability ϕ, social learning occurs as shown by panel (b).

1.4 Game Theory

Game theory, which originated in the early 20th century, is a mathematical framework for analyzing strategic interactions between decision-makers, known as actors or players, under specific circumstances, referred to as games [22]. As each actor makes its decision, it receives a payoff, which is a numerical representation of the outcome. Though initially applied to two-person, two-strategy games, game theory has grown to have numerous applications in the modeling of a wide range of interactions and behaviors that rely on strategic decision-making.

1.4.1 Classical Game Theory

In classical game theory, actors make choices that maximize their payoffs. One of the foundational concepts in this area is the Nash equilibrium, named after John Nash. A Nash equilibrium occurs when no player can benefit by unilaterally changing their strategy, given the strategies chosen by the other players [23]. This equilibrium represents a stable state where players' expectations are aligned, and no one has an incentive to deviate.

The quintessential example of classical game theory takes the form of the Prisoner's Dilemma [24]. In this scenario, two prisoners, who are being questioned separately, are faced with the choice to either cooperate with each other by remaining silent or defect by betraying the other. The payoffs are structured so that each prisoner benefits more from betraying if the other remains silent, but if both betray, they both receive worse outcomes than if they had both cooperated. Here, the Nash equilibrium occurs when both prisoners choose to betray each other. Neither prisoner can improve their situation by changing their strategy unilaterally, making this a stable but suboptimal outcome compared to mutual cooperation.

It is important to note that the most basic models in classical game theory rely on the assumption of perfect rationality, meaning that all players are fully rational and have complete information about the game and the other players. Classical game theory assumes that players are always capable of calculating the optimal strategies they need to play to maximize their payoffs and assume that all other players are doing the same. While this assumption simplifies the analysis, it often does not hold true in real-world scenarios where information is incomplete or players have bounded rationality. For this reason, evolutionary game theory models, concepts, and approaches will be used throughout this thesis to better understand the chaotic and sometimes irrational, but unique and important, nature of human behavior.

1.4.2 Evolutionary Game Theory

In contrast to classical game theory, evolutionary game theory (EGT) emerged originally from the study of biological and later, social systems, where the focus is on the dynamics of strategy change over time rather than the static analysis of rationally-derived equilibria [25]. EGT models populations of players who interact and adapt their strategies based on their success relative to others. Instead of economic payoffs and rationality, EGT considers biological fitness and natural selection, respectively. Therefore, EGT is a way to mathematically study evolution in the Darwinian "survival of the fittest" sense. Strategies represent inherited or adopted behaviors or traits, and payoffs are interpreted as fitness or reproductive success.

Unlike classical game theory, which assumes that players consciously optimize their payoffs, EGT assumes that strategies evolve over time through processes like natural selection or learning. The Iterated Prisoner's Dilemma (IPD) is a useful example to illustrate this. In the IPD, the same players repeatedly engage in the Prisoner's Dilemma, which results in them adapting their strategies based on the outcomes of previous interactions [26]. This repetition can lead to the emergence of cooperation as a stable strategy, as players learn that mutual cooperation yields better long-term outcomes than consistent defection. Though in this example, players update their strategy based on their payoffs, in general, the individuals in EGT models need not even be aware that they are playing a game; strategies that do well simply reproduce faster.

In EGT, individuals playing their strategy interact with each other, and accumulate payoffs based on these interactions. The relative success of the strategies is determined by how well they are doing in the population (in terms of relative payoffs and fitness), leading to the prevalence of certain strategies within the population over time. As a framework, EGT analyzes the stability and dynamics of these strategies. This approach allows for the study of how complex behaviors, such as social norms and cooperation, can emerge and stabilize over time, even in the absence of perfect rationality [27, 28, 29]. Moreover, one key concept of EGT is the Evolutionarily Stable Strategy (ESS), introduced by John Maynard Smith [25]. An ESS is a strategy that, if adopted by a population, cannot be invaded by any alternative strategy because it yields a higher or equal payoff when played against itself than any mutant strategy does. When a population reaches an ESS, it means it has become a successful, consolidated population that can survive indefinitely over time.

1.5 Approach to Social Movements: Combining Opinion Dynamics and Evolutionary Game Theory

When considering how to realistically model political activism, traditional opinion dynamics seemed overly simplistic and in need of more nuanced analysis, particularly in the context of social networks and evolutionary dynamics. The application of stochastic opinion modeling to specific political problems, such as spatial polarization in Ukraine, helped us envision new paths. These models provided insights into how opinions spread and led to geographic polarization. However, striving for universality and simplicity in social modeling often comes at the expense of specificity, applicability, and realistic simulations. We discovered natural resemblances between key concepts of both evolutionary game theory (EGT) and opinion dynamics when considering political systems. Political conversations behave like games, in that political opinions resemble strategies, and the outcomes, which resemble payoffs, influence an individual's conviction. This conviction can be modeled as a fitness measure, reflecting how comfortable and confident individuals are with their opinions.

Political systems are rarely mean-fields where everyone influences everyone else equally. In reality, adaptation rates vary across individuals and even for the same individual over time. While it is true that people change their minds based on their social relations, and vice versa, we believed that understanding the circumstances in which individuals held their beliefs most strongly—especially when surrounded by a like-minded neighborhood—warranted further study. Hence, we decided to incorporate conviction moderation (motivated by EGT) into the adaptive voter model, resulting in the Conviction Moderated Adaptive Voter Model (CMAVM). This addition introduces individual-based heterogeneity into the simulation of socio-political movements. By allowing individuals' conviction levels to fluctuate based on their social interactions and personal experiences, our model better captures the complexity and variability of real-world political systems. This is informed by social science findings that have shown the importance of quantifying the communicative, collective, and combative nature of activist behavior [30].

Social movement studies might benefit from both evolutionary game theoretic and adaptive voter model study approaches. The former is relevant insofar as social movements require a particular type of "collaborator" who is more invested than the average individual while the rest function as "defectors" as they have the capacity to discourage activism. Moreover, social movements often revolve around zealots, which are individuals who have absolute adherence to their opinion [31], [32]. On the other hand, adaptive voter models show the coevolution of opinions and networks [33]. Finally, identity plays an important role in opinion formation, especially when issues are more or less prominent in people's lives because of them [34]. Our model, then, tries to capture the evolutionary nature of conviction as a fitness variable that attempts to incorporate heterogeneity into the AVM.

1.6 Data Sources

This thesis sources its data from The Latinobarómetro Corporation and The World Justice Project (WJP). This project used Latinobarómetro data chiefly to explore attitudes towards protest, likelihood likelihood of participation in both authorized and unauthorized demonstrations, and interactions between these responses and demographic features. Then, models on Panama were initialized from the WJP GPP data, taking advantage of the larger and more representative set of demographic. More details on the data sources can be found in the sections that follow, as well as the corresponding codebooks (these are available at https://www.latinobarometro.org/latContents.jsp and https://worldjusticeproject.org/2021-wjp-rule-law-index-questionnaires).

1.6.1 The Latinobarómetro Corporation

The Latinobarómetro is an annual public opinion survey conducted in Latin America. It aims to gauge the political, economic, and social attitudes of the region's inhabitants. The survey employs stratified multistage probability sampling to ensure representativeness across the population. The sample is typically designed to be representative at both national and regional levels, taking into account factors such as urban/rural distribution, gender, age, and socioeconomic status. Still, perfect representativeness is not always achieved. For its own reports, the Latinobarómetro corrects for sampling biases through weighting adjustments, which account for the probability of selection and demographic characteristics to align the sample with the actual population distribution. A sample size of 1000 respondents per country allows good representativeness of regional demographics. In the case of Panama, the country of interest in Chapter 4, this size suffices to perform some statistical analysis on attitudes based on demographics. However, it remains slightly small and biased for particularly granular demographic analysis.

The questionnaire includes a wide range of questions that cover various topics such as democracy, governance, economic conditions, social issues, and international relations. The data is available for download, allowing for independent analysis and study. Researchers and academics use the Latinobarómetro data for scholarly research and analysis. Moreover, the results of the survey are frequently cited in media reports and public discussions, contributing to informed debate and discourse on key issues facing Latin America. As such, this data plays a crucial role in informing policy decisions, academic research, and public understanding of the social, political, and economic landscape in Latin America.

Crucially, Latinobarómetro data includes information on attitudes toward protest and likelihood of participation in both authorized and unauthorized demonstrations. These questions are at the core of this project's research, so the value of these consistent measures over time is difficult to understate.

1.6.2 The World Justice Project's Rule of Law Index

Established in 2006, the WJP is an independent, international non-governmental organization that seeks to advance the rule of law around the world. It promotes a framework in which laws are clear, publicized, stable, and fair, and protect fundamental rights, ensuring that justice is delivered by competent, ethical, and independent representatives who are accessible and reflect the makeup of the communities they serve.

The WJP's Rule of Law Index (RLI) is a key initiative aimed at assessing the adherence to the rule of law in various countries. The index evaluates countries based on several factors, including the extent to which they limit government powers, ensure absence of corruption, maintain order and security, uphold fundamental rights, ensure open government, and deliver civil and criminal justice effectively. By providing detailed and comparative data on these aspects, the RLI seeks to highlight areas where governments are failing their citizens in the administration of justice and the guarantee of human rights. In so doing, the RLI promotes informed policy decisions, fostering a culture of accountability and respect for the rule of law worldwide.

The Rule of Law Index relies on rigorous and comprehensive data collection methods to ensure its accuracy and reliability. For its primary sources, the WJP conducts household surveys, also called the General Population Poll (GPP), to gather citizens' perspectives. These surveys are nationally representative, meaning they reflect a broad and diverse cross-section of the population in each country, and that the statistical inference drawn from this data serves as a reliable measurement for the country's geo-political landscape. The surveys cover a range of issues including people's experiences with crime, dispute resolution, corruption, and their perceptions of government accountability and fundamental rights. Another key aspect of the RLI are expert surveys, through which the WJP engages legal practitioners, academics, and other professionals with expertise in various aspects of the rule of law; because of their expertise, responses to the expert survey allows the WJP to delve into more technical aspects of the legal system and governance. For the purposes of this thesis, which is concerned with the everyday decisions and opinions of ordinary citizens, only data from the GPP is analyzed.

To ensure representativeness, the WJP uses random sampling techniques for the GPP and aims for a balanced demographic representation. This explicitly includes considerations for urban and rural populations, gender, age, and socioeconomic status. In turn, these generally assure a balanced representation of other variables of interest, such as religious faith or partisan affiliation. Moreover, The WJP employs robust validation methods to verify the reliability of the data collected: this includes cross-referencing survey responses and comparing the data with other reputable sources. In the case of Panama, which is the country of interest in Chapter 4 of this thesis, simple data analysis shows that the GPP very closely resembles the distributions of the aforementioned variables that are present in recent data from the 2020 Panamanian census, recent annual reports from UN agencies, and the well-respected Latinobarómetro Corporation. The sample size for Panama was 2503, a size that allows more granular demographic analysis due to its greater representation. Finally, the RLI is updated annually to reflect changes and trends in governance and rule of law within countries over time. This allows for tracking progress, identifying regressions, and making comparative analyses between different countries and regions.

The GPP includes a thorough demographic questionnaire. The GPP includes a mix of quantitative and qualitative questions designed to capture a broad array of dimensions related to the rule of law. These questions might cover personal experiences with the justice system, perceptions of law enforcement effectiveness, experiences of bribery or corruption, and overall trust in institutions.

Chapter 2

The Conviction Moderated Adaptive Voter Model (CMAVM)

"All models are wrong. Some are useful."

George Box

2.1 Methods

This project relies on systematic simulations of the conviction-moderated adaptive voter model (CMAVM), which is introduced in Chapter 2. The code that runs these simulations consists of four main components: initialization, interaction, adaptation, and evaluation. The adaptation step's code is mainly derived from Chu et al.'s baseline (geographically uninformed) model, the scripts for which are available on Github. The same source is used as the basis for an important component of evaluation, which relies on visualizing periodic "snapshots" that show the state of the system at different time steps. The rest of the code is original.

The main extension to the code from its previous iteration is the definition of politically viable networks (PVNs) as independent connected components of the induced sub graph comprised of activists at the end of a run. To thoroughly understand the impact of conviction moderation in the context of political activism, we conduct a behavior space experiment, in which we run a parameter sweep of the model, each a total of 10 times in an attempt to reduce noise.

For our first behavior space experiment, we considered four values of N (250, 500, 750, and 1000), the number of individuals; nine values of ϕ (between 0.30 and 0.70 with increments of 0.05), the probability of social learning; four values of x (0.10, 0.15, 0.20, and 0.25), the initial proportion of activists; and three values of mean α (0.45, 0.50, and 0.55), which corresponds to conviction. The second behavior space experiment looked at the interaction of N and ϕ more specifically, sweeping through the aforementioned values of N and every value of ϕ between 0.20 and 0.80 with increments of 0.025.

2.1.1 Initializing the Model

We initialize a population of N individuals, a fraction x of which is labeled as activists (opinion = 1) and the rest as others (opinion = 0). Then, each activist i is assigned a conviction α_i between 0 and 1 according to a beta distribution with mean α. This data, apart from the constant population size N, is dynamically updated as the model runs.

The network is generated as a Watts-Strogatz Small World model on N nodes, with mean degree 6, and a rewiring probability of 0.5. This code is imported according to Matlab documentation. Importantly, connections in this model are reciprocal and undirected and are easily represented by an adjacency matrix. The adjacency matrix is then transformed into a graph object for its continuous manipulation.

2.1.2 Running the Model

1. **Interaction Step**: First, an individual i is chosen randomly. If it is an activist, then it experiences two random interactions with members of its neighborhood (and if it is not an activist, we move on to the next iteration of the simulation). These interactions affect i's fitness according to a pre-determined payoff structure that depends on the similarity or difference in opinion between two interacting individuals. Namely, finding two activists will increase i's fitness by a factor of 1.75, finding one will increase it by 1.35, and finding none will decrease it by 0.90.

2. **Self-assessment Step**: The conviction value, α_i, recalculated based on i's interactions and interpreted as a fitness measure, governs the decision to adapt or not. A higher α indicates that an individual is more committed to their opinion and that they are more comfortable with their network. Then, i will either retain all of its properties (with probability α) or choose to adapt (with probability 1-α).

3. **Adaptation Step**: If i adapts, it either engages in social learning (with probability ϕ) or in homophily (with probability 1-ϕ). In the case of social learning, the social learner adopts the role model's conviction as well as their opinion.

The above process repeats itself for a total of $t = 25 * N$ time steps; this number was empirically found to be close to peak PVN presence in the system and theoretically consistent with issue permanence and interaction frequency. A flowchart describing the above steps is shown in Figures 2.1 and 2.2.

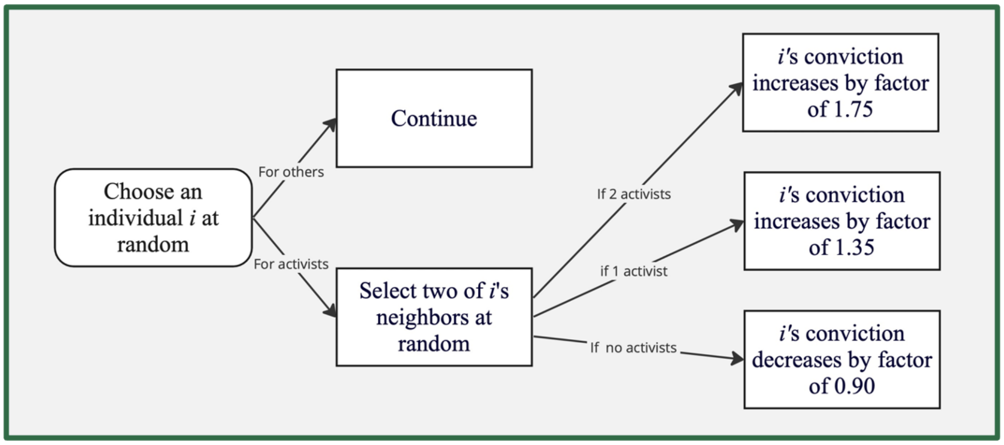

Figure 2.1: The interaction step recalculates individuals' α. The interaction step is followed by the adaptation step.

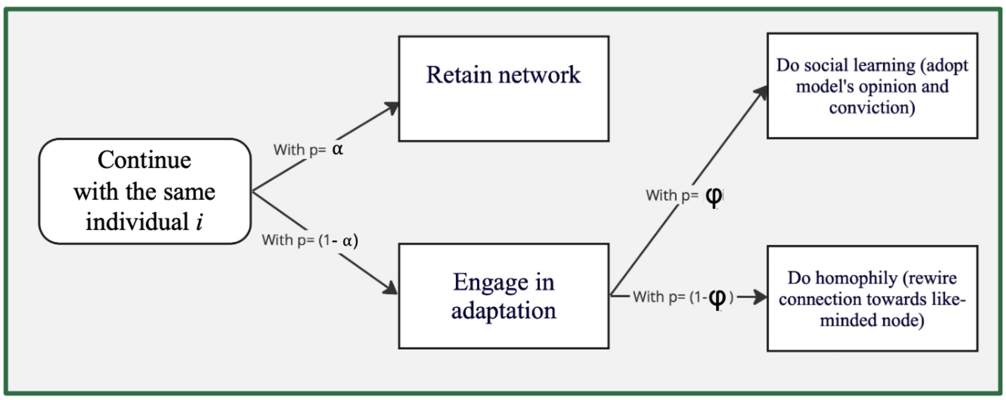

Figure 2.2: The adaptation step allows for homophily or social learning. Note that for a given time step, the randomly chosen individual i is the same individual chosen in the interaction step

2.2 Results and Model Evaluation

While "others" (i.e. those who are not activists) certainly play a significant role in this model through the systematic discouragement of activists as shown in social science research [35], model evaluation is mainly concerned with the latter. Thus, an indented sub-graph is extracted from the general graph consisting solely of activist nodes (if any). From this sub-graph and its underlying matrix, each run of the model reports the share of activists, their final mean conviction, the number of edges between them, and the number of politically viable networks (defined as the number of independent connected components of the indented sub-graph).

Moreover, we extract periodic "snapshots" of the system at predetermined time steps to evaluate the system's behavior over time. At designated intervals, the model saves the resulting graph and underlying data as outputs, which are vital to make qualitative and quantitative statements about the changes the system experiences over time.

In the limit of long time, our simulations show general tendencies towards the systematic discouragement of activists (see Figure 2.3). Thus, we run the model for $25 * N$ time steps and take a census of activists, their connections, and their PVNs. This time is interpreted as the point by which each individual has interacted with its neighborhood and at least considered adaptation a number of times over the course of the period of exogenous shock to the political system. So, the measured variables at this point meaningfully represent the effect of the shock to the system.

As should be expected from the evolutionary game theoretic dynamical updates to activists' conviction as a function of their neighborhood, the model only preserves activists in groups of similarly-minded individuals. While those nodes with high initial conviction might last for several more time steps, they eventually succumb to discouragement if they cannot surround themselves with activists too. Normally, these networks are not completely detached from the main component; rather, these PVNs tend to retain some connection through highly convicted (i.e., less vulnerable) nodes (see Figure 2.4). Moreover, PVNs are characterized by high internal connectivity and high average conviction. This is consistent with the social structure found by research on the anatomy of activist networks, both in-person and on social media [13, 36].

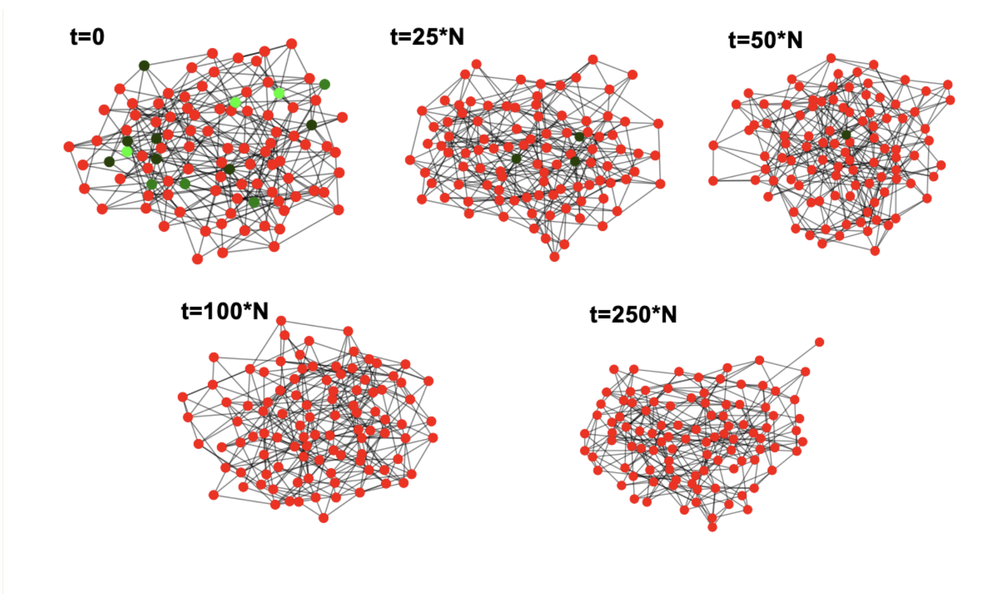

Figure 2.3: In the image, green represents activists; darker shades represent those with higher convictions. Meanwhile, red represents others. The image shows how, in most cases, activists are systematically discouraged as the limit approaches long time.

Our simulations reveal that lower ϕ (a higher preference for homophily) helps the establishment of PVNs, as it seems to help activists isolate themselves from the discouragement they feel when neighbored by others (i.e. non-activists). This is clear in Figures 2.5, 2.6, 2.7, and 2.8, which plot the outcomes of our

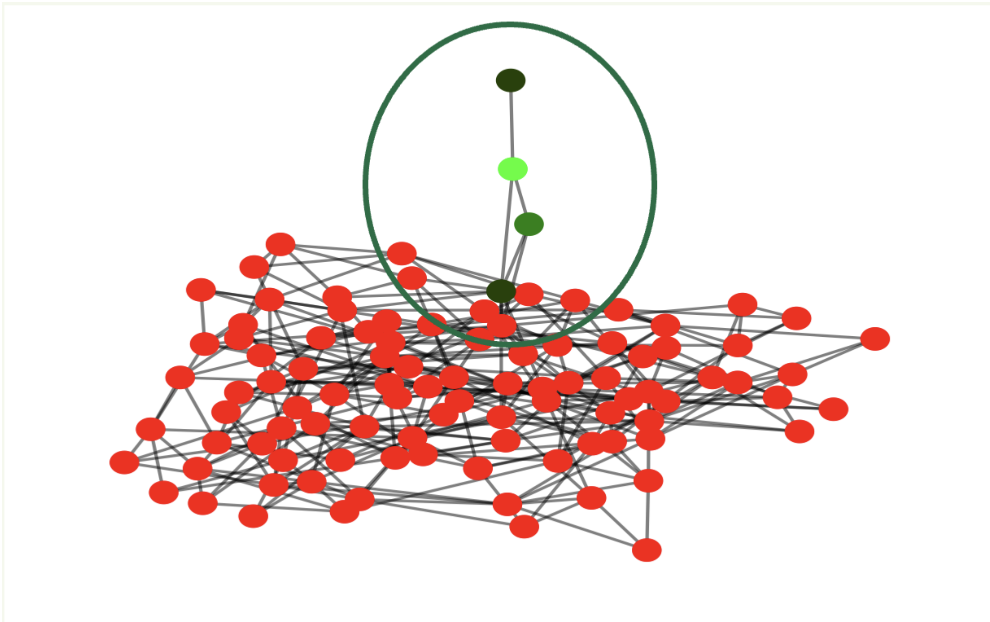

Figure 2.4: See typical isolation of a PVN for mutual reinforcement; only extremely highly convicted nodes remain connected to main component.

behavior space experiment. Additionally, while higher N leads to proportional increases in PVNs, it does not necessarily lead to higher shares of activists or mean conviction. Finally, as is expected, higher values of α and x increase the likelihood that activists are present and typically result in networks with higher shares of activists and more viable networks; this is shown by Figures 2.9 through 2.12.

The trends of our model map nicely to those of the general AVM, as there is a clear phase transition at which runs tend to lead to the consensus state. It occurs at a similar value, $\phi = 0.54$, compared to Holme and Newman's $\phi = 0.46$. So despite our modifications, including the introduction of many more auxiliary parameters to the model which allow us to understand the social dynamics involved, the most significant parameter remains ϕ, and it behaves similarly regardless of N.

Importantly, the fact that there are enduring PVNs at the final time step does not mean that activism has "succeeded", it instead just shows that activism has emerged and persisted. Emergence and persistence, however, should not be dismissed, as they are incredibly integral to the study of social movements. Specifically, the emergence of communities has been cited by the field of community psychology as an important and almost necessary pre-requisite to the emergence of modern activism in the context of the perceived antagonism and hostility towards certain forms of activism . So, the findings concerning PVN formation and survival directly engage with these fundamental research trends in social science [37]. Despite this, measuring "success" is still important. However, this is rarely attempted because it is extremely challenging to define the success or failure of activism, as it is frequently unclear which goals or objectives would represent

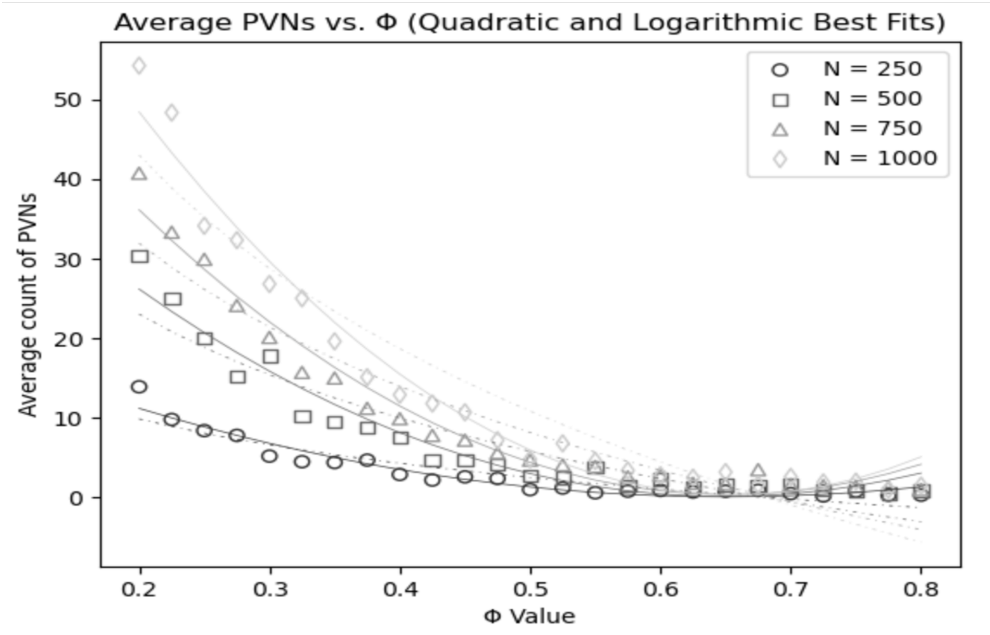

Figure 2.5: Lower ϕ (higher homophily) is associated with higher counts of PVNs; note the phase transition near 0.575 after which PVNs become exceedingly rare.

success, and it is difficult to identify when, how, and whether these goals are reached within a movement. Research in this area, though, has begun to emphasize the importance of considering activist perspectives in defining success [38]. Through the interpretation of conviction as a measure of comfort of the individual holding an opinion within their neighborhood, we contribute to this tendency of activism success metrics centered around the individual activist and their perception of change.

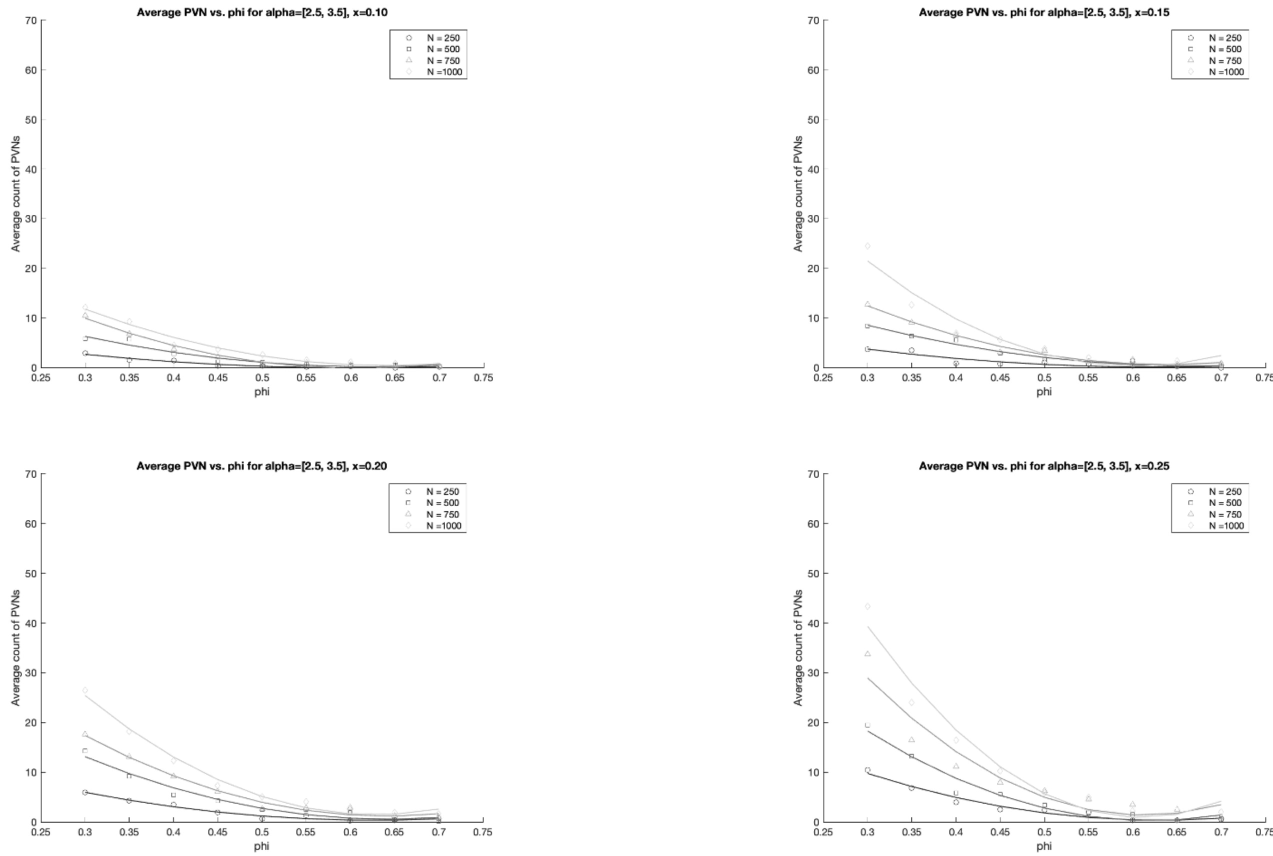

Figure 2.6: These plots, which correspond to the lowest fixed $\alpha = 0.45$, show how the highest values of x are needed for PVNs to persist.

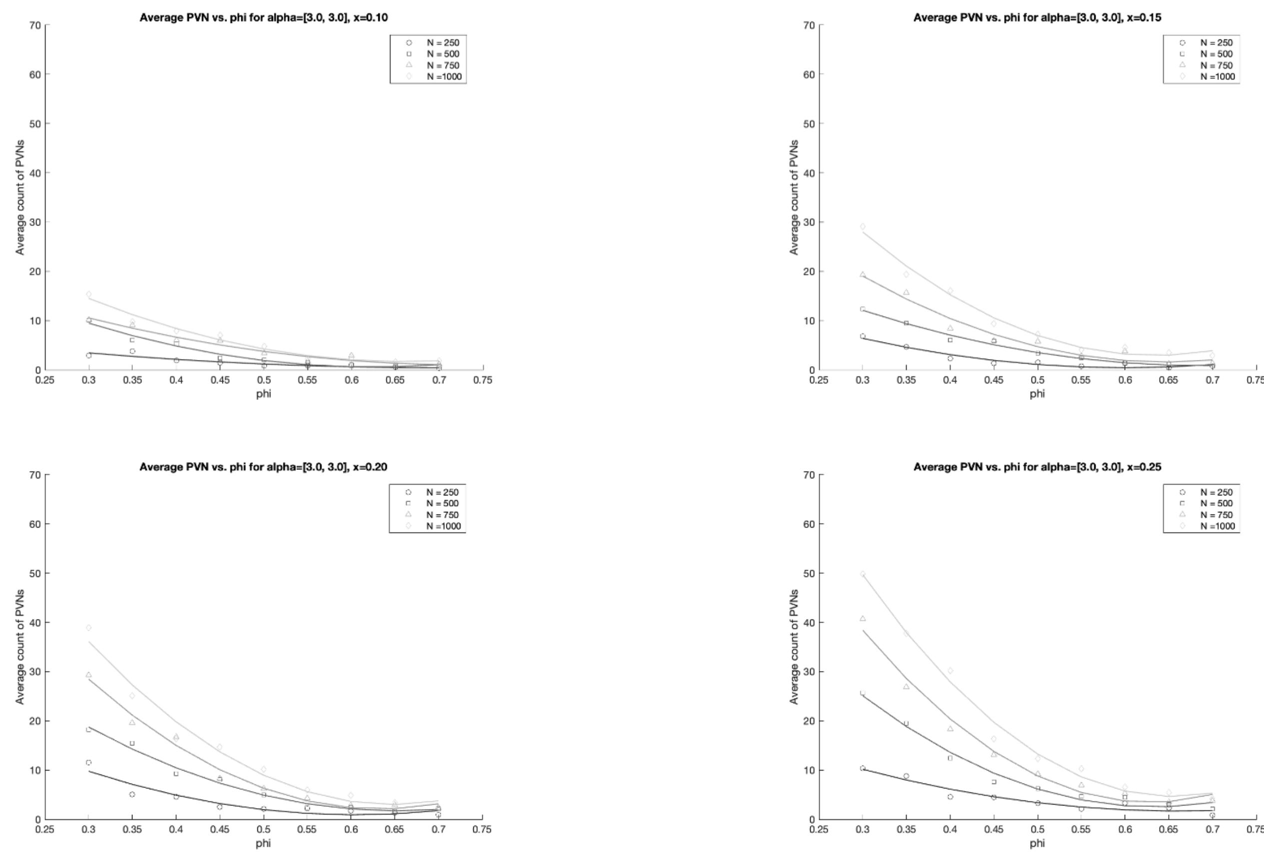

Figure 2.7: These plots, which correspond to a fixed $\alpha = 0.50$, show how sensitive the model is to the inputted values of x and α near the phase transition.

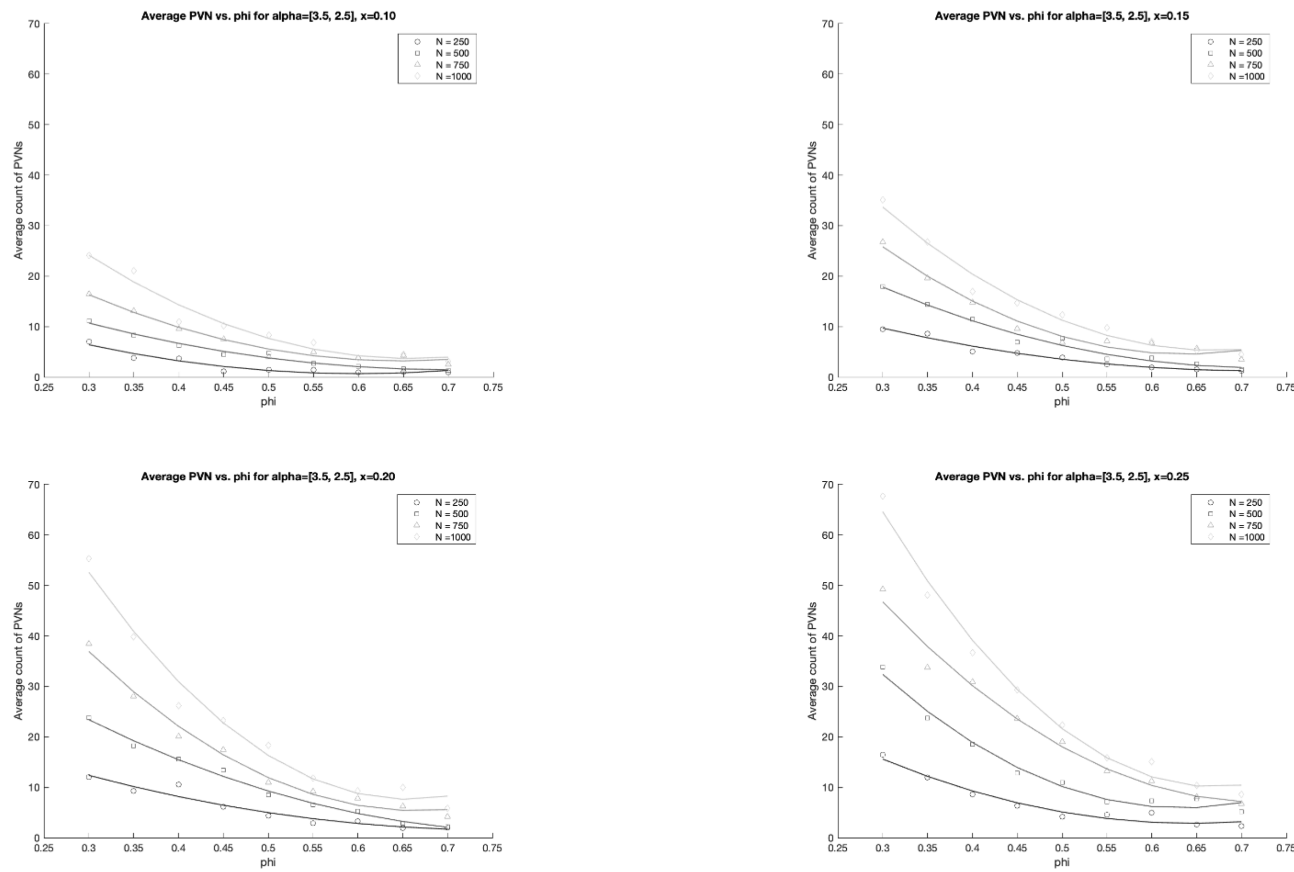

Figure 2.8: These plots, which correspond to the highest fixed α best represent how changes in x in turn increase or decrease PVN average counts. Additionally, it is worth noting that for this value of $\alpha = 0.55$, all $x > 0.15$ show persistence of PVNs.

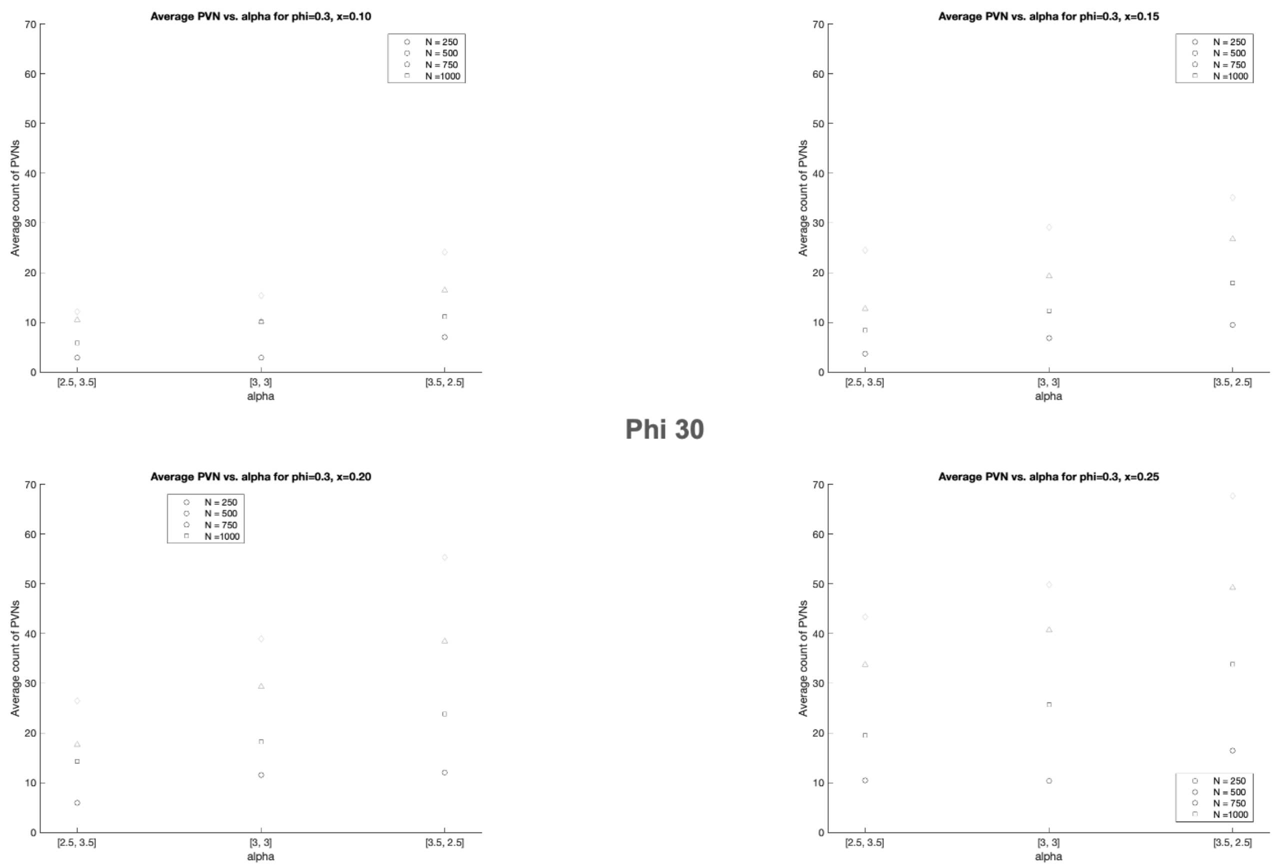

Figure 2.9: These plots represent $\phi = 0.30$, the strongest homophily preference in the parameter sweep. At this level of social preference for homophily, almost all simulations yield PVNs consistently, regardless of the initial shares of activists and initial mean convictions. The plots also show how, at this level of homophily preference, a 5% change in initial mean conviction is linked to an increase of five in the average count of PVNs.

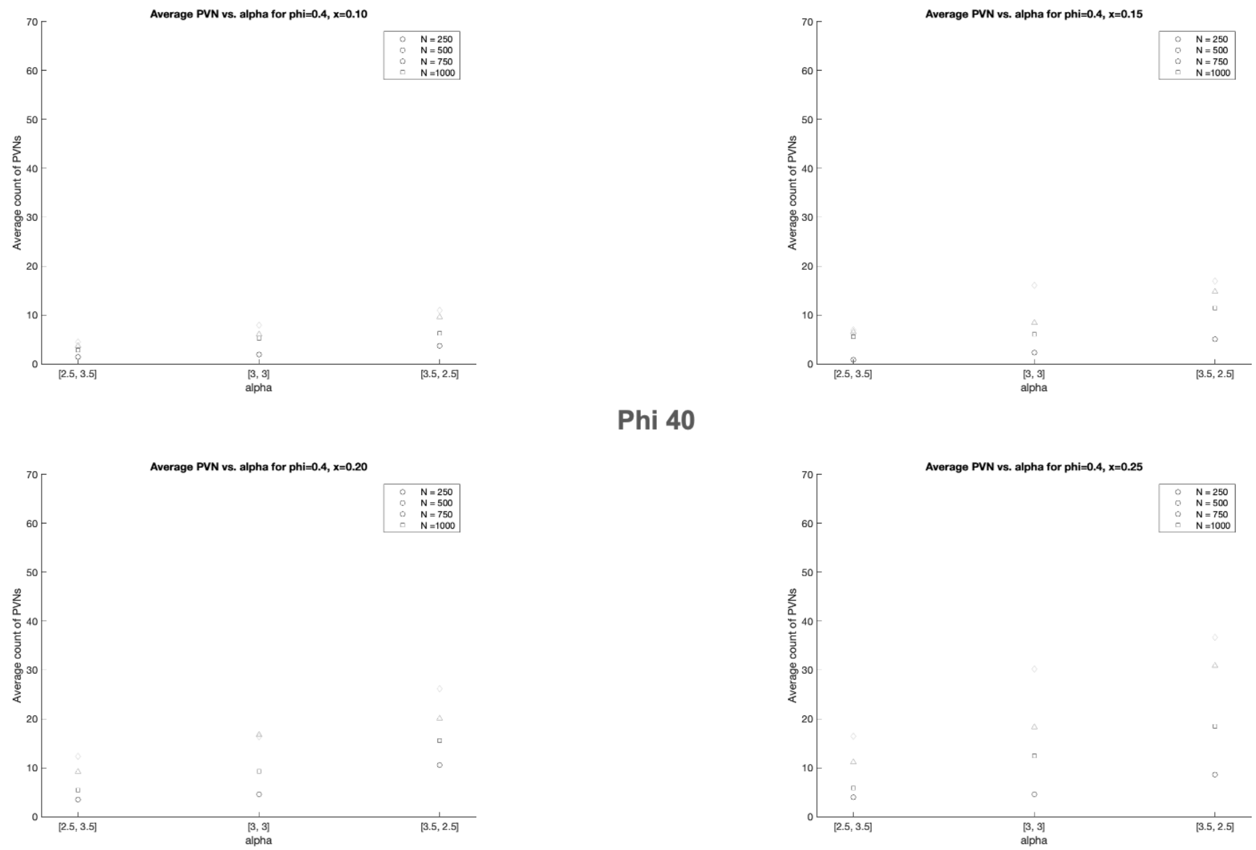

Figure 2.10: The plots for $\phi = 0.40$ best show how a 10% change in initial conviction can make the difference between a simulations consistently yielding or failing to yield PVNs, using the same value for share of activists.

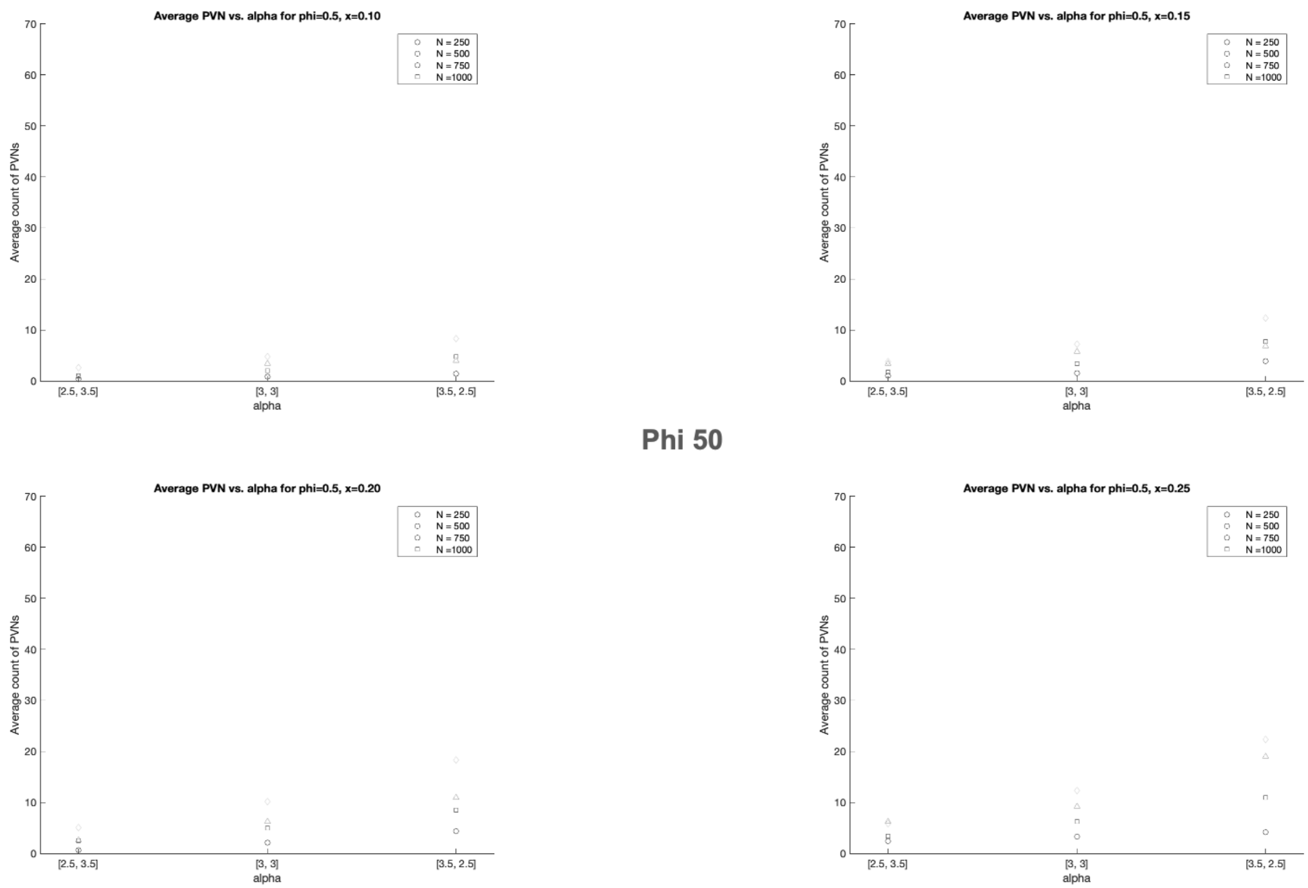

Figure 2.11: These plots represent $\phi = 0.50$, a homophily preference closely approaching the phase transition. Here we see an interesting pattern, where only simulations using relatively larger shares of activists and highest initial convictions yield PVNs consistently.

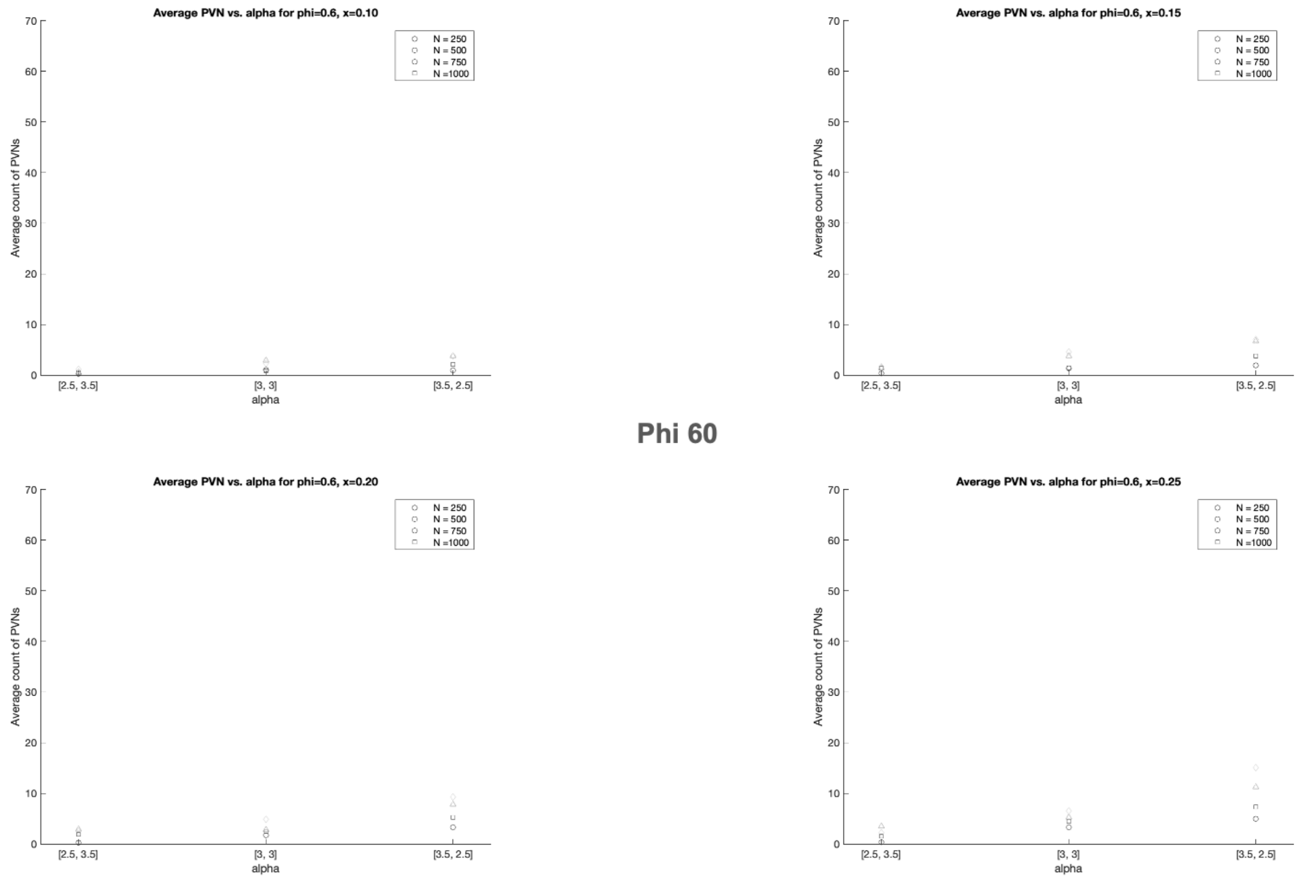

Figure 2.12: In all these plots for $\phi = 0.60$, homophily preference is so low that even simulations using the largest shares of activists and highest initial convictions failed to yield PVNs consistently.

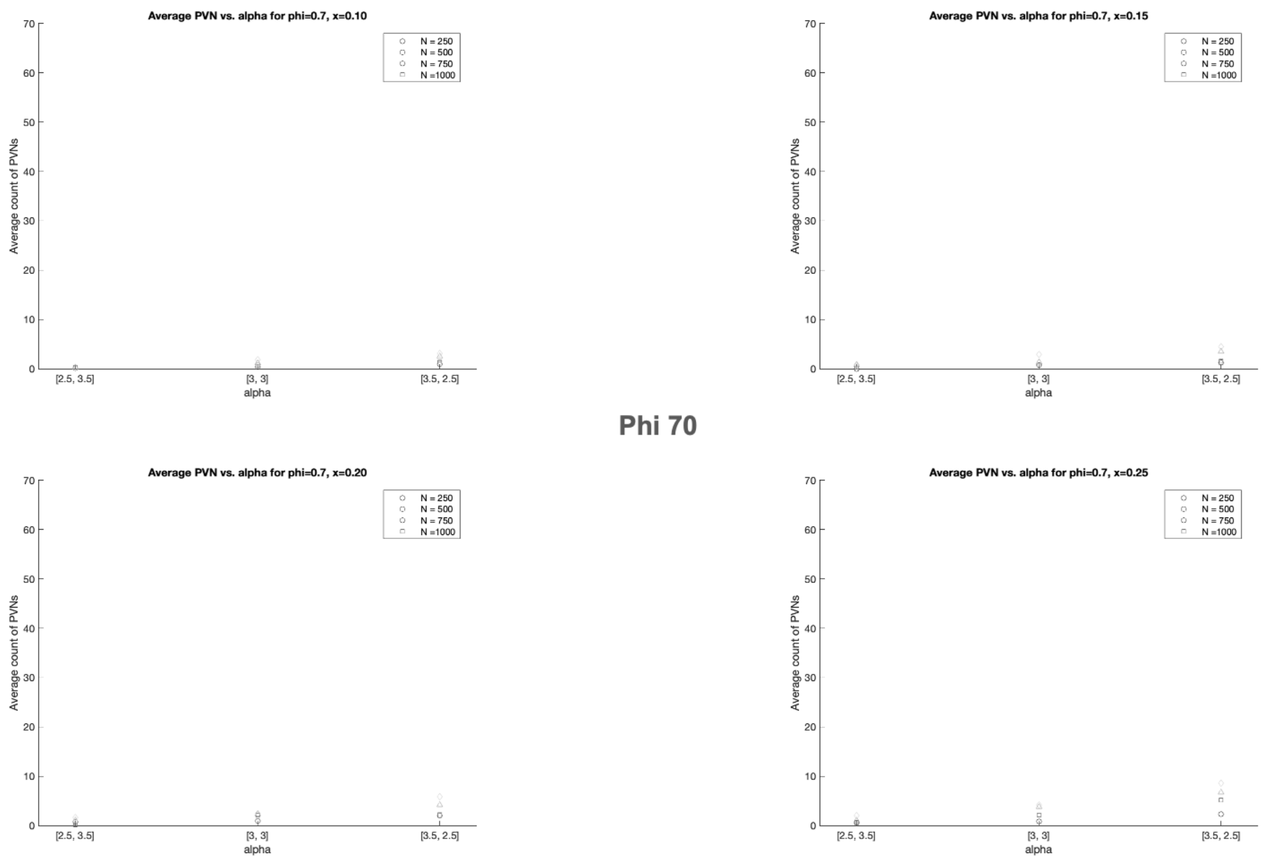

Figure 2.13: In all these plots for $\phi = 0.70$, homophily preference is so low that even simulations using the largest shares of activists and highest initial convictions failed to yield PVNs consistently.

Chapter 3

The Survey-Informed CMAVM

> "Everything should be made as simple as possible, but not simpler."
>
> Albert Einstein

The integration of real-world data into the CMAVM allows for the introduction of more realistic heterogeneity. In this project, data is leveraged in three main ways: geographic location informs graph structure and connections; demographics, political node attributes such as opinion and conviction; and identity, the level of impact each interaction has based on the perceived similarity or difference among the individuals involved. These mechanisms are inspired by research that shows how each plays an essential role in opinion formation and information flow, both in material and online communities [39, 40].

3.1 Geographical Information and Network Layout

The procedure to create the geographically-informed graph aims to model social connections in a country with population N divided into n distinct regions or communities known as $R_1, R_2, ..., R_n$ with populations $n_1, n_2, ..., n_n$ such that $n_1 + n_2 + ... + n_n = N$. Each region's social structure is generated using the Watts-Strogatz (WS) model, known for producing small-world networks characterized by high clustering and short average path lengths. These features are typical in real-world social networks, where individuals form tight-knit clusters but are still relatively close to others within the same region. Then, the regions are connected according to an algorithm based on the Stochastic Block Model such that some individuals replace an existing within-region relationship with one across-regions, to represent important and influential connections that transcend geographic or social boundaries; this is a secondary rewiring step. However, the model ensures that individuals maintain most of their social connections within their own regions by following these steps:

1. **Defining Populations**: We start by defining the population sizes of the n regions, allowing for significant population disparities if these reflect real-world scenarios. This heterogeneity allows the model to capture varying social dynamics across regions.

2. **Generating Regional Graphs**: For each region $R_1, R_2, ..., R_n$, a corresponding WS graph $G_1, G_2, ..., G_n$ is created with two parameters: a mean degree k of connection to a node's nearest (geographic) neighbors and a rewiring probability p (this model uses $k = 6$ and $p = 0.1$). Note that it is possible to assign different mean degrees and rewiring probabilities to each region to represent real-world conditions such as, for instance, the population sparsity of rural regions or extreme clustering of red-lined cities.

3. **Combining Graphs**: The individual WS graphs $G_1, G_2, ..., G_n$ are then combined into a single graph G using NetworkX's disjoint_union_all function. This step ensures that each region's graph remains distinct initially, with no inter-regional connections.

4. **Rewiring Edges Across Regions**: To introduce realistic inter-regional connections, a controlled rewiring process is applied. Each region is allowed to rewire a fixed number of its internal edges (in this case, 12 edges per region) to nodes in different regions. This step ensures that while the majority of connections remain within regions, some connections bridge across regions, reflecting occasional long-distance social ties, such as those formed through migration, work, or familial relationships. Again, note that it is possible to assign different numbers of within-region internal edges or to establish additional rules to represent real-world conditions such as, for instance, that rural communities might all be better connected to a central city than to each other.

This approach effectively creates a graph that mirrors realistic social connection patterns in a geographically divided country. The use of the WS model ensures that each region has a small-world network structure, characterized by high local clustering and short path lengths, capturing the essence of social connections within a community [41]. The controlled rewiring process introduces a small number of inter-regional connections, ensuring that while regions are predominantly internally connected, there are still some pathways for inter-regional interaction. This type of model reflects real-world social dynamics where individuals are more likely to interact with others in their immediate geographical vicinity, forming tightly-knit clusters. However, due to various factors such as travel, communication technology, and social mobility, there are still occasional long-distance connections. This balance between local clustering and occasional long-distance ties is crucial for representing key characteristics of social networks, including information flow and social integration across the entire country. Overall, this method provides a realistic simulation of social networks in a geographically divided country, highlighting the importance of both local community structures and inter-regional connections that underpin the social structure.

Moreover, the code structure for creating the graph is highly flexible, making it adaptable to various types of social divisions beyond geographical boundaries. If the primary divides in a given country are based on party lines, economic status, or other social factors, the initial communities can be configured accordingly. By defining populations for these communities, the code can generate small-world networks that reflect the social dynamics within each group. For instance, in a country where political affiliation forms the main social clusters, each Watts-Strogatz graph can represent a different political party's social network. This flexibility allows the code to simulate realistic social connection patterns in diverse contexts. Additionally, the code's parameters can be adjusted to reflect varying levels of connectivity within and between communities. By modifying the number of rewires or the rewiring probability, the model can capture different degrees of social integration. This adaptability makes the generated graphs more realistic for countries with varying strengths of social divides, providing a robust base for simulating and analyzing opinion dynamics.

3.2 Demographics and Political Attitudes

The procedure to endow each individual with a more realistic initial opinion in the graph begins by identifying a set of traits $x_1, x_2, ..., x_n$ through linear regression. These traits are identified based on their practically and statistically significant impact on the likelihood of protesting and/or supporting protests. This approach ensures that the initial opinions of individuals are assigned in a way that accurately reflects the likelihood of each individual holding a particular opinion, consistent with real-world social dynamics where traits like gender, age, and ethnicity influence political attitudes. This can be accomplished through the following steps:

1. **Extract Demographic Trait Effects**: Initially, demographic traits such as gender, age, and ethnicity are one-hot encoded to transform categorical data into a numerical format suitable for analysis. K-nearest neighbors (k-NN) matching is then applied to pair individuals with similar demographic characteristics, ensuring that each individual is matched with others who share similar traits. This regression process identifies which traits have practically and statistically significant effects on the initial views each individual possesses and helps control for confounding variables and isolates the effect of specific traits on opinions.

2. **Logistic Opinion Assignment**: Using the insights gained from the a simple linear regression on demographic covariates and the desired opinion, a simple logistic model is employed to assign initial opinions to individuals in the graph. The logistic model considers the significant traits and their effects, assigning opinions in a way that reflects the real distribution of political attitudes within the population. This probabilistic approach ensures that the assigned opinions align with the observed likelihoods based on demographic characteristics.

This procedure effectively models how demographic traits influence initial political opinions, creating a more realistic social network. By incorporating linear regression, the model accurately captures the impact of traits such as gender, age, and ethnicity on political attitudes. The logistic opinion assignment step ensures that the distribution of opinions in the network mirrors real-world scenarios. The realism introduced by this procedure is crucial for understanding the dynamics of opinion formation and change. Different demographic groups respond differently to periods of shock, such as political upheaval or social movements. By accurately reflecting these variations in initial opinions, the model provides a robust framework for studying how social networks evolve over time in response to exogenous shocks to the system. Moreover, this method's flexibility allows it to be adapted to different contexts and populations. Whether the focus is on political opinions, social behaviors, or consumer preferences, the same steps can be applied to identify key traits and assign initial states that reflect real-world distributions. This adaptability makes the procedure valuable for a wide range of applications in social science research, public policy analysis, and network theory.

3.3 Identity and Social Influence

The procedure to endow each individual with a more realistic initial opinion in the graph begins by identifying a set of traits $x_1, x_2, ..., x_n$ through covariate matching and sample average treatment effect (SATE) regression. These traits are identified based on their practically and statistically significant impact on the likelihood of protesting and/or supporting protests. This approach ensures that the initial opinions of individuals are assigned in a way that accurately reflects the likelihood of each individual holding a particular opinion, consistent with real-world social dynamics where traits like gender, age, and ethnicity influence political attitudes. This can be accomplished using the following steps:

1. **Create Covariate Matching**: Initially, demographic traits such as gender, age, and ethnicity are one-hot encoded to transform categorical data into a numerical format suitable for analysis. K-nearest neighbors (k-NN) matching is then applied to pair individuals with similar demographic characteristics, ensuring that each individual is matched with others who share similar traits. This matching process helps control for confounding variables and isolates the effect of specific traits on opinions.

2. **Calculate Difference in Attitudes**: After matching, the difference in political attitudes (likelihood of protesting/supporting protests) between matched individuals is calculated. This involves comparing the attitudes before and after a treatment or shock to the system, such as a significant political event. The difference provides insights into how specific traits influence opinion changes in response to external stimuli.

3. **SATE Regression**: The next step involves performing a linear regression analysis to assess the effect of the identified covariates on the likelihood of protesting or supporting protests. This regression helps

determine which traits have a practically and statistically significant impact on political opinions. Traits that significantly influence attitudes are identified as key factors for opinion assignment.

4. **Create Similarity Formula**: Using the relative effects of the statistically significant attributes, we design a formula as a weighted sum of traits in common or different, which is carefully encoded such that it can intake the four different types of variables most commonly present in surveys (binary, categorical, continuous, and discrete). Similarity is calculated as this pair-wise weighted sum resulting above or below threshold.

5. **Design Payoff Table**: update payoff table to show how much more or less an individual is encouraged or discouraged based on the identity of their chosen neighbor.

Payoffs									
2 S Act.	1 S/1 D Act.	2 D Act.	1 S Act./1D Non.	1 S Act./1S Non.	1 D Act./ 1 D Non.	1 D Act./1S Non.	2 D Non.	1 S/1 D Non.	2 S Non.
1.825	1.75	1.675	1.4	1.375	1.35	1.325	0.95	0.9	0.85

Table 3.1: Payoff Table, where S represents a neighbor deemed similar and D a neighbor deemed different, and Act. denotes an activist and Non. denotes a non-activist.

3.4 Methods

We conduct a parameter sweep on an informed CMAVM that incorporates the mechanism of heterogeneity introduction outlined in Section 3.3. We considered three values of N (250, 500, and 1000), the number of individuals; nine values of ϕ (between 0.30 and 0.70 with increments of 0.05), the probability of social learning; four values of x (0.10, 0.15, 0.20, and 0.25), the initial proportion of activists; and three values of mean α (0.45, 0.50, and 0.55), which corresponds to conviction.

3.4.1 Initializing the Model

We initialize a population of N individuals, a fraction x of which is labeled as activists (opinion = 1) and the rest as others (opinion = 0). Then, each activist i is assigned a conviction α_i between 0 and 1 according to a beta distribution with mean α. This data, apart from the constant population size N, is dynamically updated as the model runs.

This model also assigns four attributes randomly to each node, simulating each individual's potential responses to survey questions. According to the four variable types, each individual will have a binary, a categorical, a discrete, and a continuous attribute (respectively labeled binary, categ, discrete, and cont). These traits are fixed for each simulation.

The network is generated as a Watts-Strogatz Small World model on N nodes, with mean degree 6, and a rewiring probability of 0.5.

3.4.2 Running the Model

1. **Interaction Step**: First, an individual i is chosen randomly. If it is an activist, then it experiences two random interactions with members of its neighborhood (and if it is not an activist, we move on to the next iteration of the simulation). These interactions affect i's fitness according to a pre-determined pre-determined, expanded payoff structure that depends on the similarity or difference in opinion and in demographic traits between two interacting individuals. The result of the demographic similarity formula described in the previous section determines the payoff received according to the payoff table in Figure 3.1.

2. **Self-assessment Step**: The conviction value, α_i, recalculated based on i's interactions and interpreted as a fitness measure, governs the decision to adapt or not. A higher α indicates that an individual is more committed to their opinion and that they are more comfortable with their network. Then, i will either retain all of its properties (with probability α) or choose to adapt (with probability 1-α).

3. **Adaptation Step**: If i adapts, it either engages in social learning (with probability ϕ) or in homophily (with probability 1-ϕ). In the case of social learning, the social learner adopts the role model's conviction as well as their opinion.

The above process repeats itself for a total of $t = 25 * N$ time steps; this number was empirically found to be close to peak PVN presence in the system and theoretically consistent with issue permanence and interaction frequency. A flowchart describing the above steps is shown in Figure .

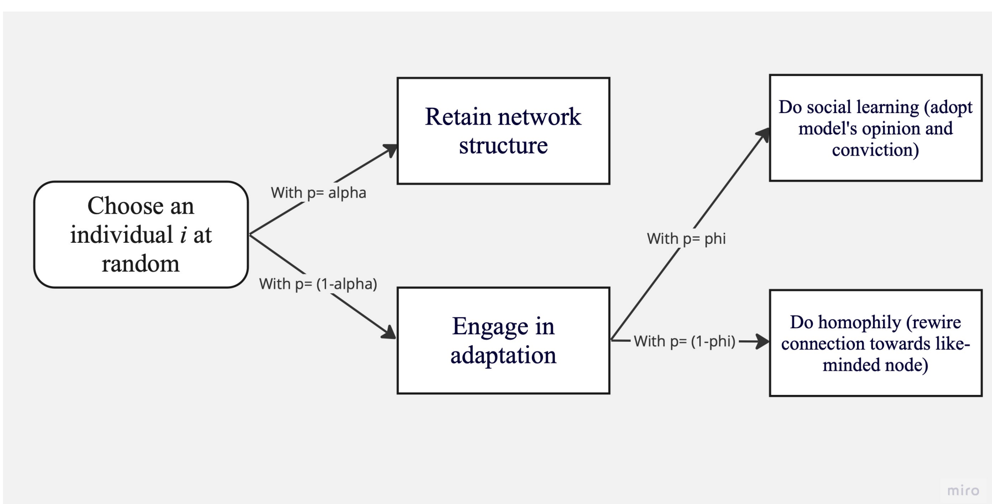

Figure 3.1: The interaction step recalculates individuals' α. The interaction step is followed by the adaptation step.

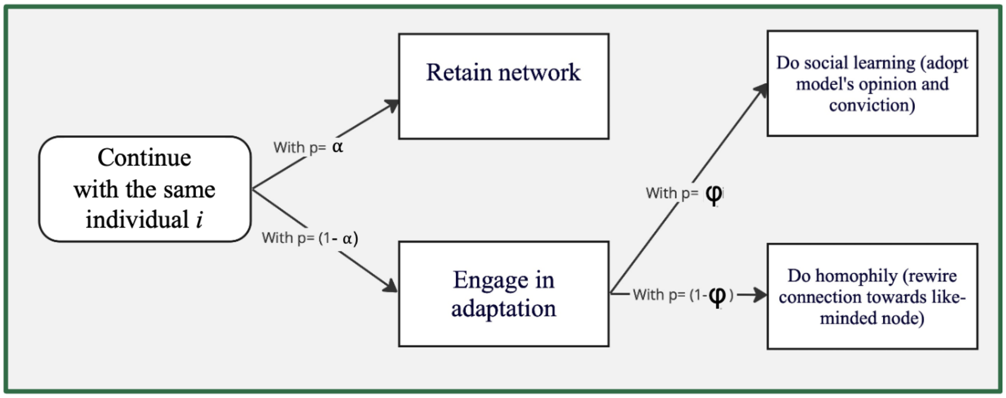

Figure 3.2: The adaptation step allows for homophily or social learning.

3.5 Results

Our simulations yield largely similar results to those of the CMAVM. Lower ϕ (a higher preference for homophily) helps the establishment of PVNs. While higher N leads to proportional increases in PVNs, it does not necessarily lead to higher shares of activists or mean conviction. Higher values of α and x increase the likelihood that activists are present and typically result in networks with higher shares of activists and more viable networks. The trends of our informed CMAVM still map nicely to those of the general AVM, including its phase transition at which runs tend to lead to the consensus state.

There seems to be one important consequence of factoring identity into the strength of social influence of pairwise interactions. Most remarkably, this modification mildly magnifies or dampens the the effect of the parameter x, meaning that it mildly decreases the likelihood of persistence for smaller initial shares of activists and mildly increases the likelihood of persistence for smaller initial shares of activists. This might be because depending on the size of the initial sets of activists, each activist may have an easier more difficult time finding another activist who is similar to them. Besides that small difference, the plots below are almost indistinguishable from those corresponding to the uninformed CMAVM shown in Chapter 2. However, the lack of huge differences should not be interpreted as evidence of an unsuccessful model; rather, it shows that the CMAVM is mostly resilient even after the introduction of heterogeneity.

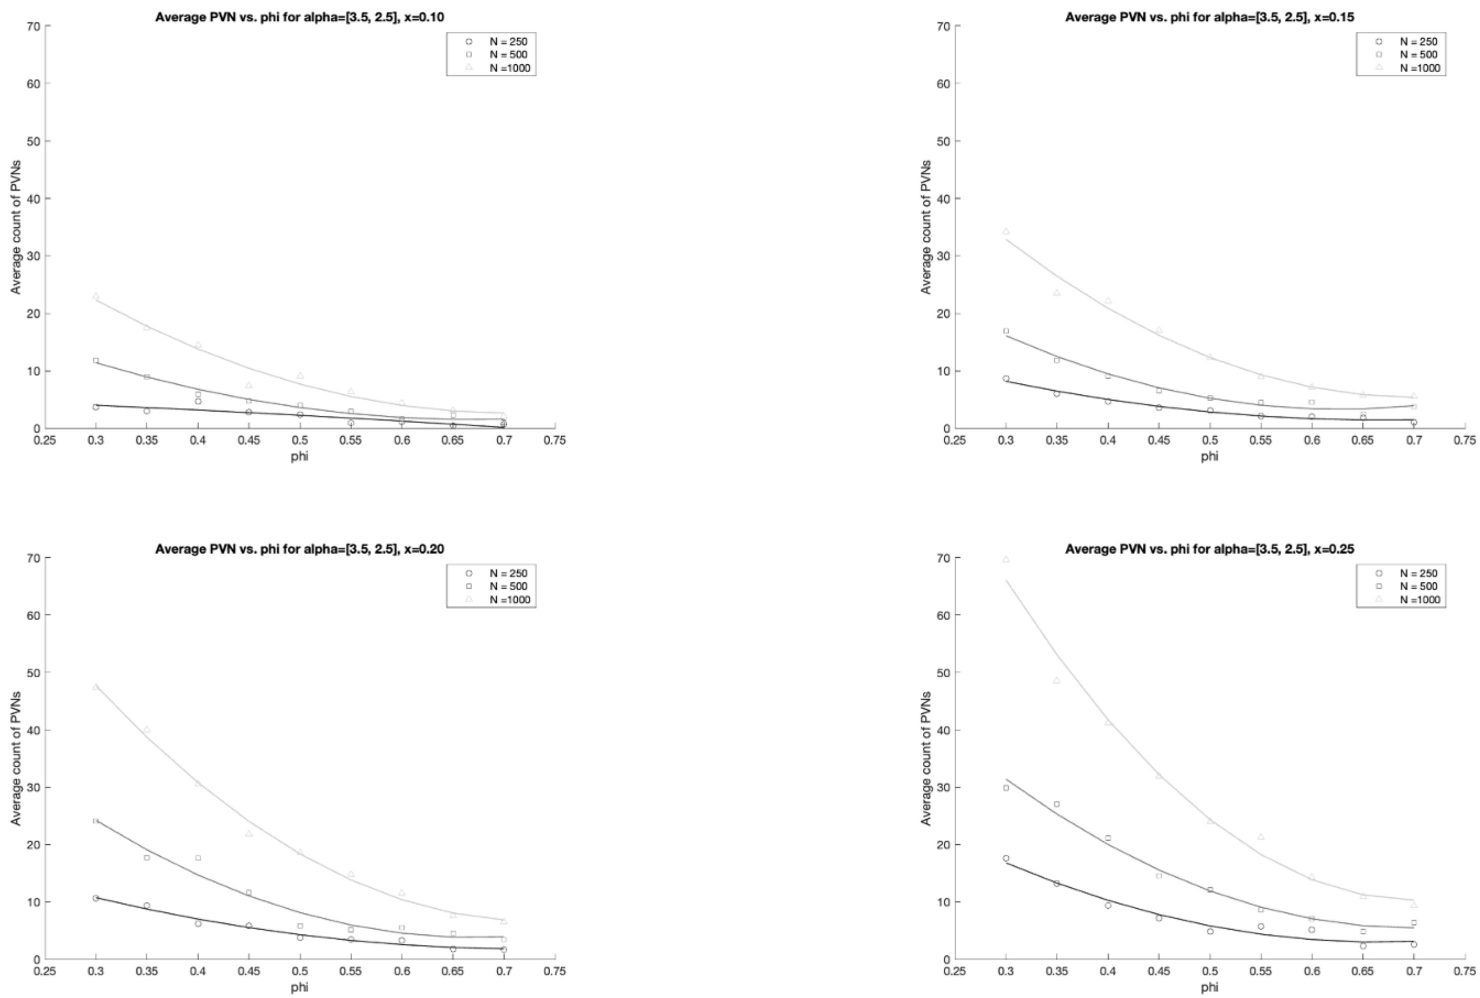

Figure 3.3: These plots are almost indistinguishable from those in Figure 2.6, yet they have some subtle differences in the panels corresponding to the lowest and highest values of x.

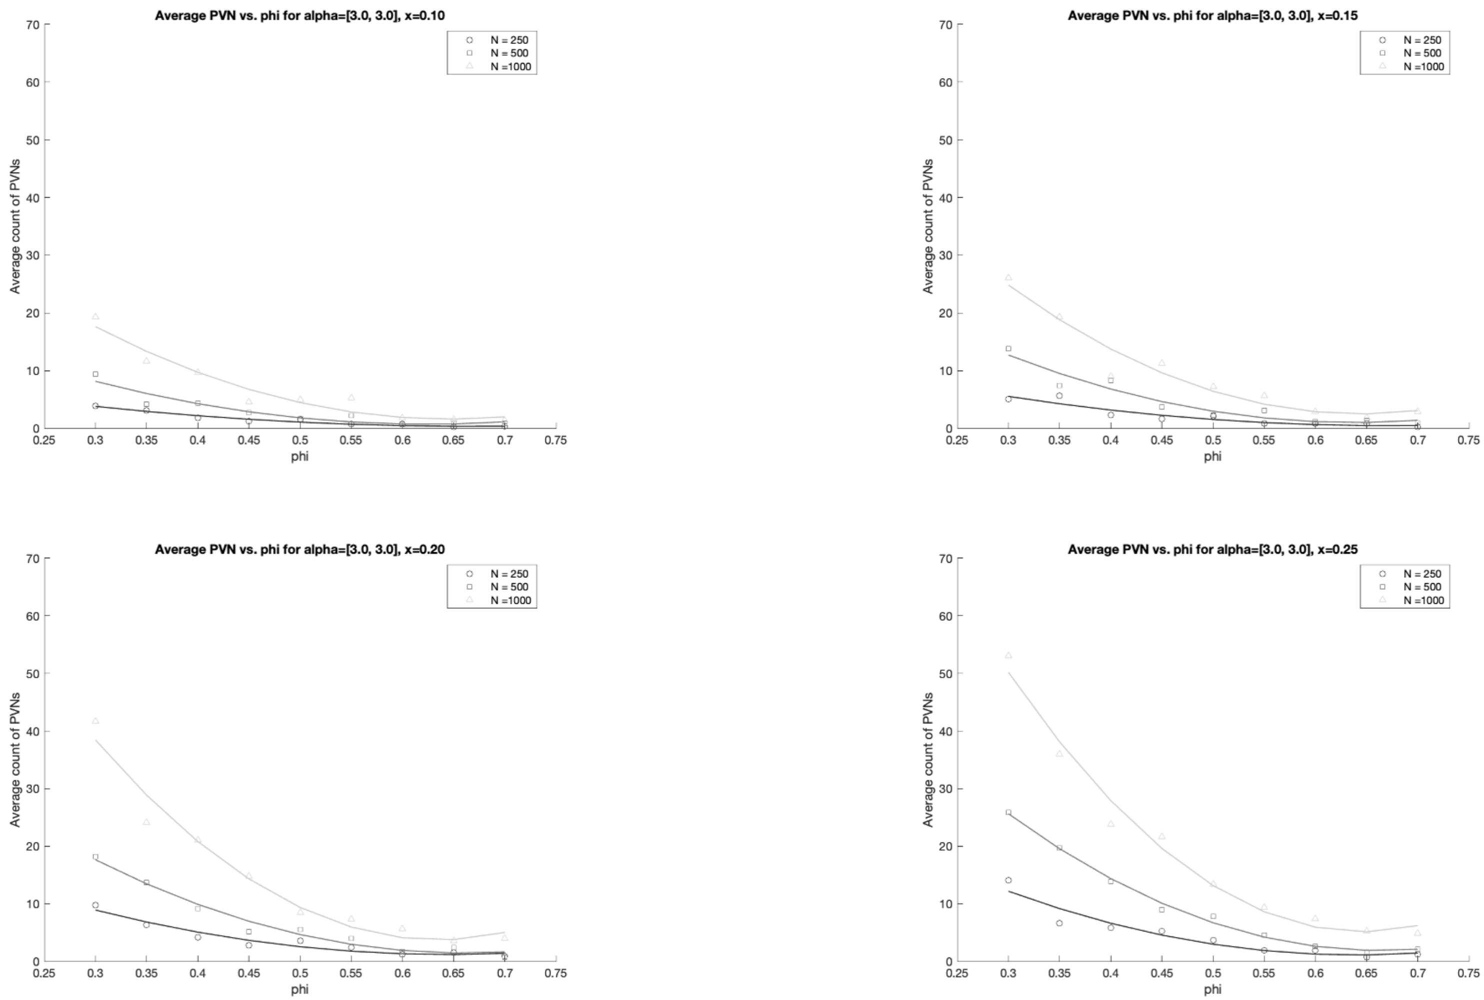

Figure 3.4: These plots are almost indistinguishable from those in Figure 2.7, yet they have some subtle differences in the panels corresponding to the lowest and highest values of x.

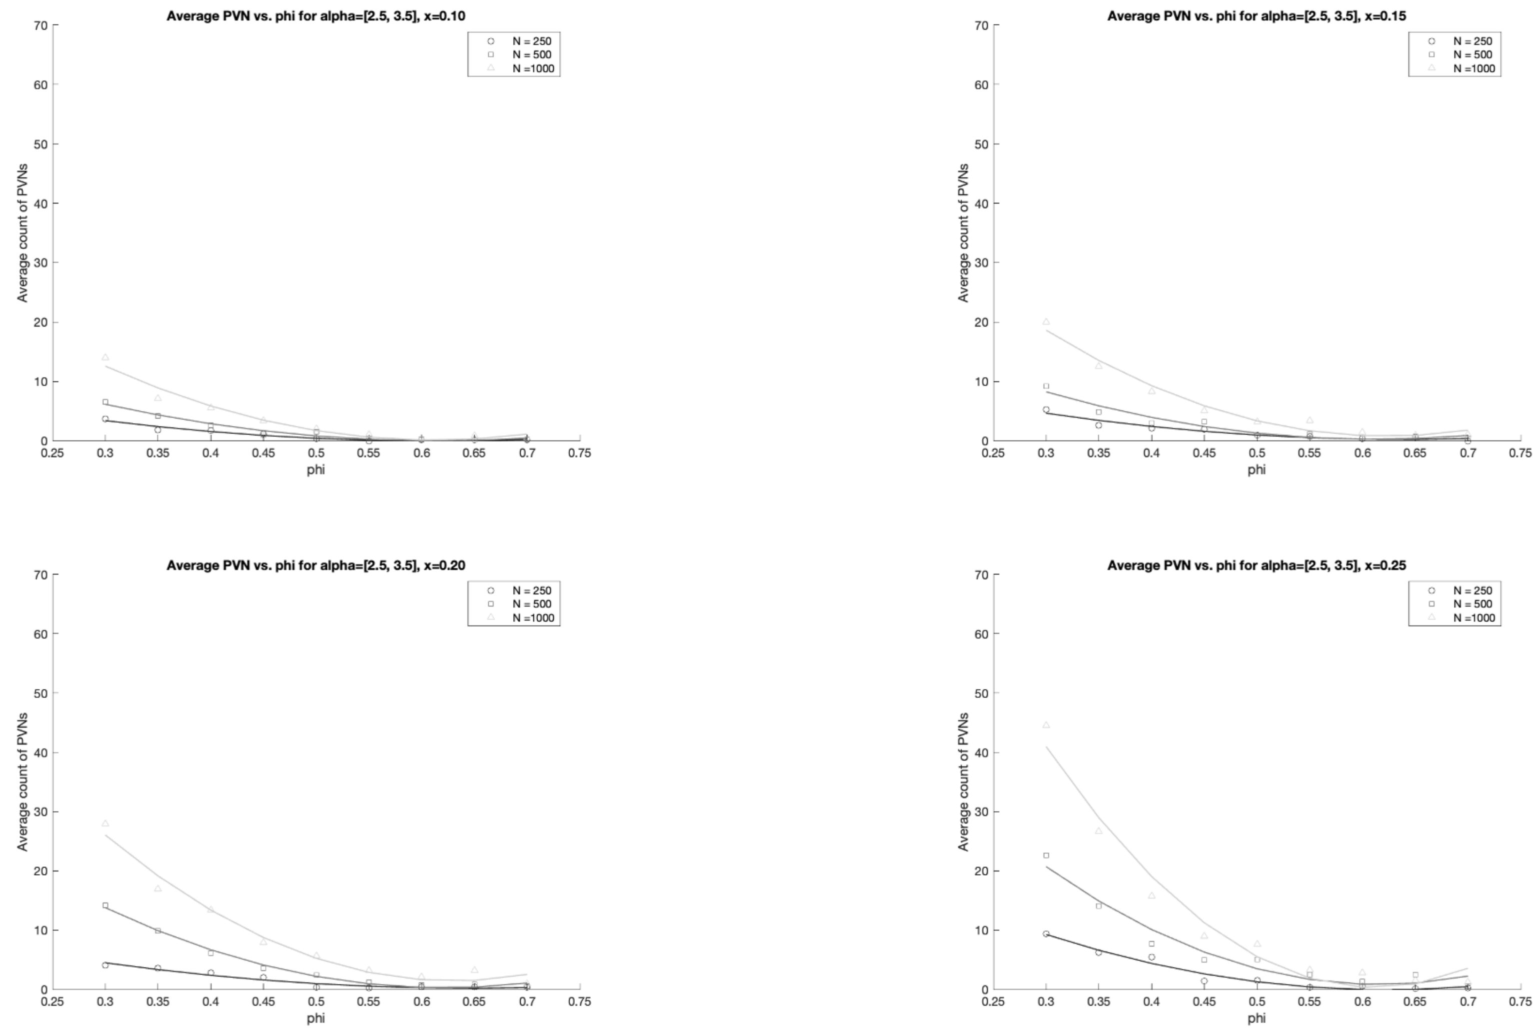

Figure 3.5: These plots are almost indistinguishable from those in Figure 2.7, yet they have some subtle differences in the panels corresponding to the lowest and highest values of x.

Chapter 4

An Application: The Panamanian Protests of July 2022

"La pelea es peleando." / "The fight requires fighting."

Victoriano Lorenzo

4.1 Background and Model Introduction

Periods of exogenous shock are linked to opinion evolution [42]. These processes of opinion evolution interact with social rewiring and exemplify the conditions that can lead to the formation of social movements [43]. This model simulates the process through which significant social re-wiring might have helped Panamanians hold more intensely onto a negative view of the country's outlook and find likeminded individuals across cohort boundaries. In doing so, we can better understand how cross-cohort activism happened in this highly unique and interesting socio-political landscape.

The Republic of Panama enjoyed a prolonged socio-political stability between its return to democracy in 1989 through 2022 [44]. However, that year's summer witnessed unprecedented levels of cross-cohort mobilization and motivated this project [45]. This was influenced by an increased frequency of political discussion during the COVID-19 pandemic, a period of shock to the Panamanian political system. During this period, severe lockdown measures and increased official communication politicized the discussion of current events. In particular, Panamanians began seeking ways to hold the government accountable for its management of the pandemic. Discomfort grew along three main issues: cost of living, public sector corruption, and neglect of indigenous rights.

None of these issues were new to Panamanians in 2022, and it is difficult to argue that they had worsened so significantly and suddenly as to awaken such a response. Historically, Panamanian unions, most notably public school teachers and construction workers, protested rising costs of living with some regularity and relative success. On the issue of corruption, movements comprising primarily younger, more affluent, and more technologically-connected Panamanians have denounced different corruption scandals and holding small demonstrations since Panama's return to democracy and very consistently since at least the early years of the Varela government (2014-2019). Lastly, indigenous and rural interest groups have a lengthy and successful tradition of demonstration, both against the government and corporations [46, 47].

The 2022 Panamanian protests, though different from the preceding single-sector movements, featured many of these facets of the traditions of Panamanian activism quite prominently. For instance, the protest's first movers included teachers and construction workers [48]. Crucially, though, they began collaborating with many other unions almost immediately. Younger people, including students in the country's public high schools and universities, were more involved [49]; when asked, many cited great outrage at the discovery of corruption scandals [50]. Panama's seven indigenous peoples also joined early, employing mainly the same mechanisms of protest that had delivered policy victories in the past: road blocks [51]. And these shock waves crossed the country in every direction, with very similar multi-cohort protests occurring even in Panama's isolated rural regions (see Figure 4.1). It seems that technology was at least partly responsible in bridging the historical gap between metropolitan Panama and the countryside; specifically, the pandemic increased rural access to and consumption of technology and media, putting the countryside in closer contact with urban areas. Of course, there are still major challenges to technological access in rural Panama [52], but the argument is that some technology did connect the public sentiments of the city and countryside.

A common narrative in Panamanian media and among local scholars is that the public perception of protest changed radically during the pandemic, reaching support thresholds that both allowed and required action. This was frequently cited in direct conversation with individuals from multiple regions when questioned about how they thought the protests had emerged. The interviewed Panamanian faculty and public servants also agreed with this characterization. However, though Latinobarómetro data does show a small but sizable shift towards public support of protest and towards increased participation in unauthorized demonstrations, the shift is nowhere near the drastic change of the collective mindset that Panamanians think prompted the protest (refer to Figures 4.2, 4.3, and 4.4). In other words, the protests are likely not the direct result of the small increase in support for protest and willingness to participate. Our model explores whether the mechanisms highlighted by the CMAVM could have been the main cause of these developments.

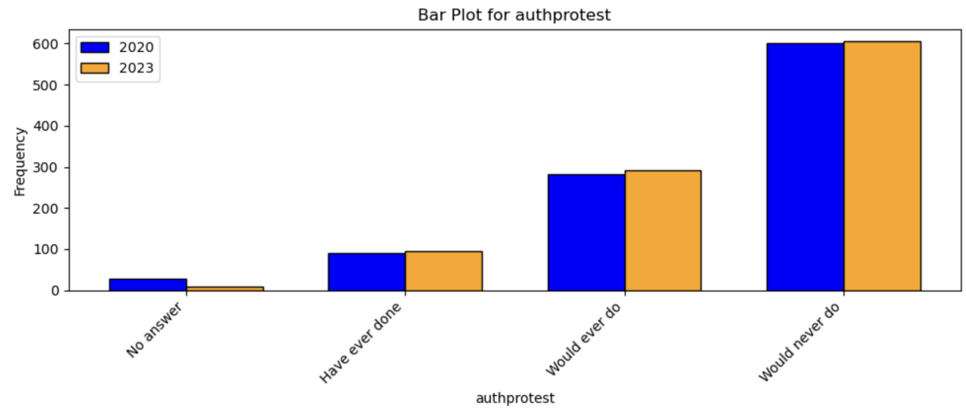

Figure 4.1: The COVID-19 pandemic did not result in meaningful increase or decrease in the likelihood of participation in authorized demonstrations and protests; however, more Panamanians gave a definitive answer to this question.

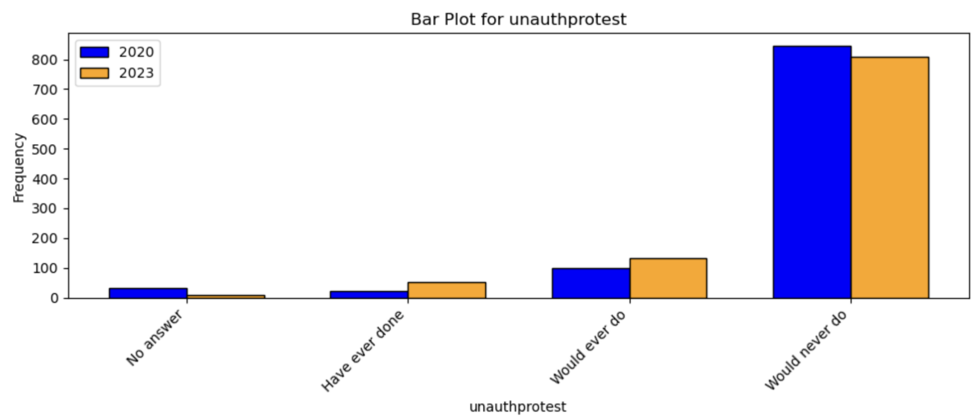

Figure 4.2: The COVID-19 pandemic resulted in meaningful increase in the likelihood of participation in unauthorized demonstrations protest, and the number of surveyed individuals who had participated in protest nearly tripled. Still, demonstration participants remained a relatively small fraction of the population.

4.2 Informing the Model

4.2.1 Geographic Creation of Network Layout

We use WJP survey data to assign each individual's region and obtain the population size of each. The choice to group by region as opposed to province was made in light of the tremendously large population ratios between provinces. We start by creating one Small-World Watts-Strogatz graph for each one of Panama's regions according to their size. Then, we add all regions together as a disjoint union graph. Finally, we use the stochastic block model to rewire edges from within-region to cross-region. This ensures that most connections stay within a Panamanian's own region, but that there are some that cross these boundaries.

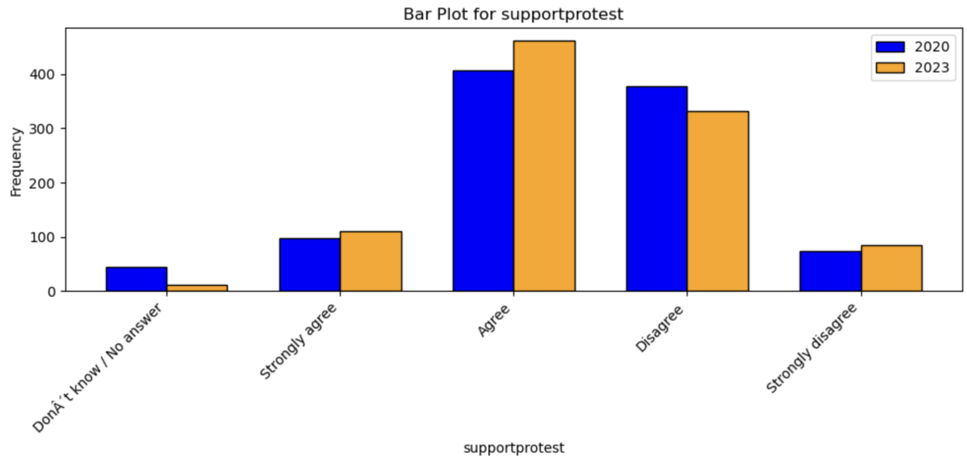

Figure 4.3: The COVID-19 pandemic resulted in meaningful increase in the support for other groups' protests among Panamanians.

4.2.2 Demographically-Informed Opinion Assignment

Our model does not endow opinions randomly; instead, we leverage data to assign activists in a more realistic fashion. We use Latinobarómetro survey data from 2020 and 2023, both of which contain the three measurements of support for protest. Then, we explored which demographic groups were more likely to be activists; these are young adults and individuals at either extreme of the financial spectrum. We use this insight to influence the assignment of opinions: our logistic equations make it so that individuals with these characteristics are more likely to be labeled as activists as the probabilistic endowment happens. The code ensures that the proportion of activists overall remains exactly as desired and that there are activists who do not belong to these identified groups

4.2.3 Identity Factored into Social Influence

The similarity formula for Panama is built as follows, taking in nine demographic factors stored as node attributes. These are gender, age, nationality, financial status, highest level of education, level of employment, geographical region, party affiliation, and ethnicity (respectively encoded as gend, age, nation, fin, edu, emp, region, paff, and ethni). Each one of these factors has an associated weight; in this model, all factors were weighted as 0.10, except for level of education and financial status which received 0.15. Our code establishes logical rules to compare the two individuals on each dimension and return a factor by which to multiply the weight. These factors can be absolute or relative, based on the demographic dimension in question. For example, the logical conditions for gender returns 1 if same or 0 if different, while age returns a float value between 0 and 1 according to how close in age the individuals are. Once the pairwise similarity score is computed, the neighbor is classified as similar if the weighted sum score is greater than or equal to 0.50 and different if below. This result then affects which payoff from the informed CMAVM table the activist receives.

We use the traditional understanding of all of these variables except for partisan affiliation, where we define a binary so that an individual either belongs or does not belong to any political party. This is because, as mentioned before, specific partisan lines were not relevant to the 2022 Panamanian protests; instead, the main division separated independents from members of all parties. Distancing oneself from partisan affiliation became a widely used mechanism to signal discomfort with the political establishment and the country's status quo.

It is worth saying that the role of race in Panamanian politics is complicated and that identity-based politics are not as common [53]. Panama does follow a different paradigm of socialization and politicization around race than countries which routinely incorporate these into social models, such as the United States and the United Kingdom [54]. However, we include ethnicity within the similarity formula as a moderately influential factor because psychology and sociology researchers have emphasized that it can play a role, either consciously or subconsciously, in social comparison and political interaction [55].

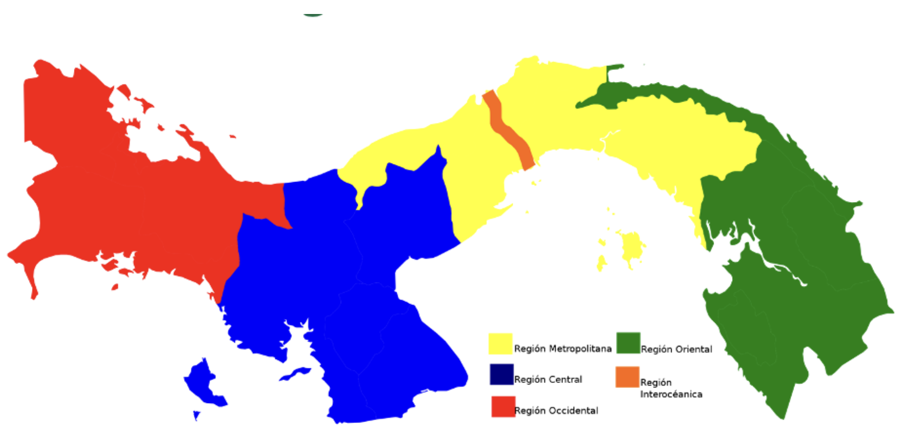

Figure 4.4: Panama's 10 provinces and 6 indigenous comarcas are divided into four regions: Western Panama (comprising Bocas del Toro, Chiriquí, Ngäbe Buglé, and Naso Tjerdi), Central Panama (comprising Veraguas, Coclé, Herrera, and Los Santos), Metropolitan Panama (comprising Panamá, Colón, and Madungandí), and Eastern Panama (comprising Darién, Emberá Wounaan, Guna Yala, and Madungandí.

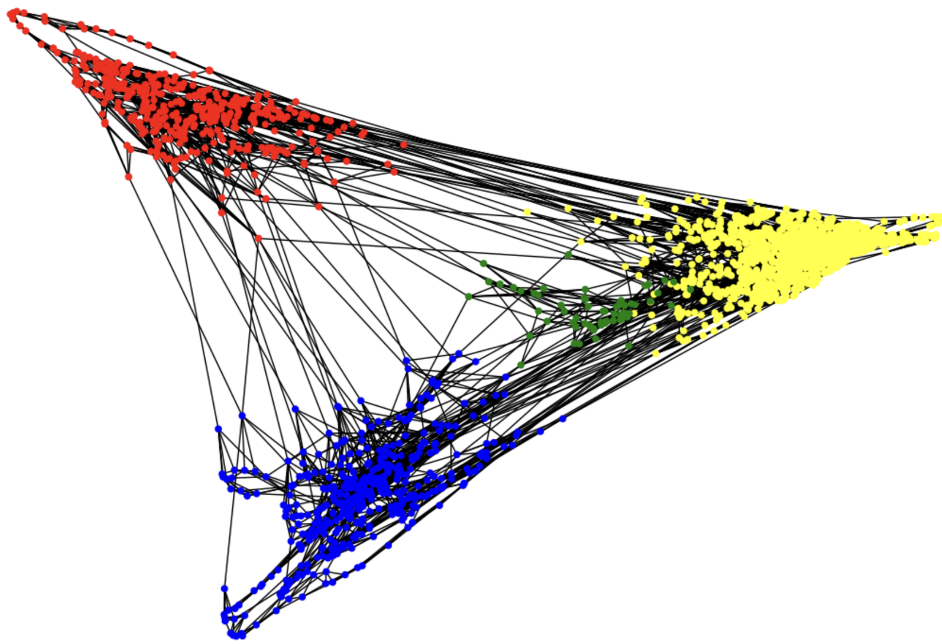

Figure 4.5: This graph shows $N = 2303$ nodes distributed according to each individual's region as assigned in the data. The spectral graph representation shows the connections between regions but overemphasizes distance.

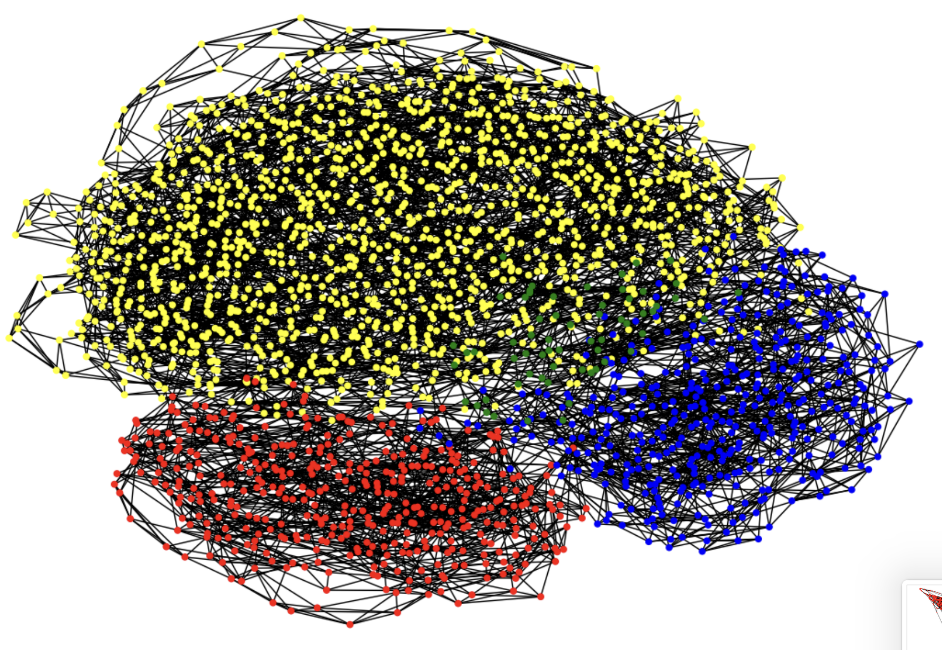

Figure 4.6: TThis graph shows $N = 2303$ nodes distributed according to each individual's region as assigned in the data. The Kawai-Kamada graph shows the clustering and de-emphasizes distance across regions.

4.3 Methods

4.3.1 Initializing the Model

We initialize a population of 2303 individuals represented by a graph created according to the procedure outlined in section 4.1.1. Each of these nodes corresponds to an individual survey respondent in the cleaned WJP dataset. A fraction x of these nodes is labeled as activists (opinion = 1) and the rest as others (opinion = 0), according to the rules described in subsection 4.1.2. Then, each activist i is assigned a conviction α_i between 0 and 1 according to a beta distribution with mean α. This data, apart from the constant population size N, is dynamically updated.

4.3.2 Running the Model

1. **Interaction Step**: First, an individual i is chosen randomly. If it is an activist, then it experiences two random interactions with members of its neighborhood (and if it is not an activist, we move on to the next iteration of the simulation). These interactions affect i's fitness according to a pre-determined payoff structure that depends on the similarity or difference in opinion between two interacting individuals; as defined in section 4.1.3.

2. **Self-assessment Step**: The conviction value, α_i, recalculated based on i's interactions and interpreted as a fitness measure, governs the decision to adapt or not. A higher α indicates that an individual is more committed to their opinion and that they are more comfortable with their network. Then, i will either retain all of its properties (with probability α) or choose to adapt (with probability 1-α).

3. **Adaptation Step**: If i adapts, it either engages in social learning (with probability ϕ) or in homophily (with probability 1-ϕ). In the case of social learning, the social learner adopts the role model's conviction as well as their opinion.

The above process repeats itself for a total of $t = 25 * N$ time steps; this number was empirically found to be close to peak PVN presence in the system and theoretically consistent with issue permanence and interaction frequency. A flowchart describing the above steps is shown in Figures 4.3.1 and 4.3.2.

4.4 Results

The Panama informed CMAVM's results show a more likely mechanization of the process that led to the 2022 Panamanian protests. We theorize that the rise of protest depended, rather, on the existing activists followed a specific process of reconfiguration catalyzed by the COVID-19 pandemic. First, Panama's existing activists began experiencing higher levels of conviction through more frequent communication and

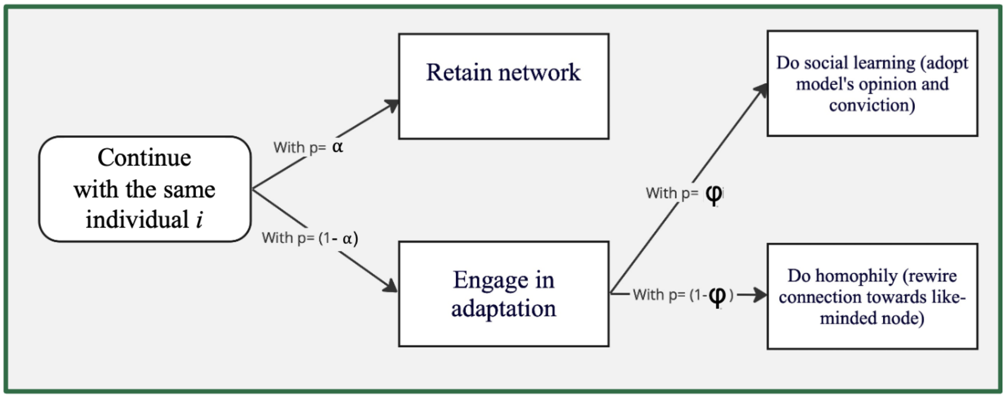

Figure 4.7: The interaction step recalculates individuals' α. The interaction step is followed by the adaptation step.

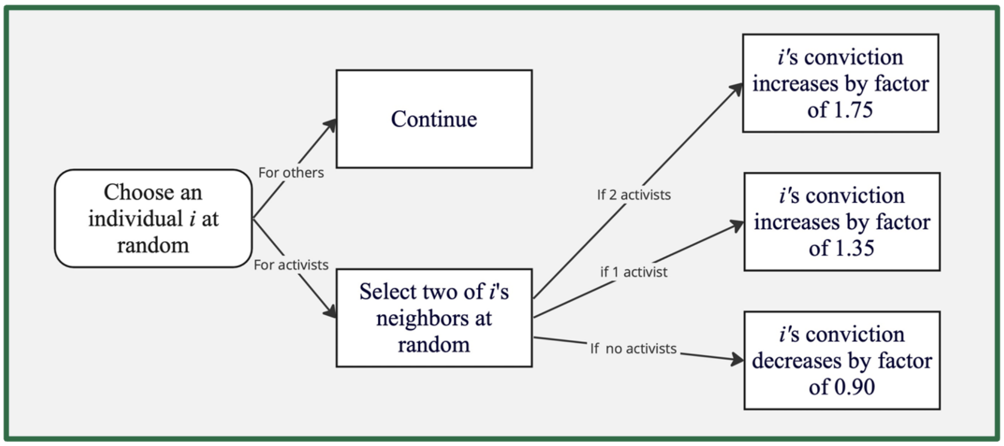

Figure 4.8: The adaptation step allows for homophily or social learning.

the increased reach of new media, while politicization of society as a whole increased the preference for homophily. These two conditions, increased initial conviction and homophily preferences, allowed for significant social reconfiguration. Activists sought and found other activists beyond their immediate circle and isolated themselves, at least in terms of political influence, from non-activists around them.

Our model's results demonstrate how the change in initial conditions induced by COVID-19, namely the higher levels of society's homophily preferences and activists' conviction, likely caused the formation of politically viable networks. This is shown through the model, where we see how PVNs can emerge across demographic boundaries without the need to assign a large initial activist share. Namely, we see that higher ϕ and α are sufficient to allow PVN formation even with the low values of x identified by Latinobarómetro data. As such, we provide a more realistic, mechanistic, and accurate representation of the social dynamics that underpinned the 2022 Panamanian protests.

Chapter 5

Discussion, Conclusion, and Future Research

"Sometimes, attaining the deepest familiarity with a question is our best substitute for actually having an answer."

Brian Greene

This thesis proposes and builds the survey-informed conviction-moderated adaptive voter model (CMAVM), which improves our understanding of the mechanisms that underpin social movement emergence through the formation of activist networks during periods of exogenous shock. Most notably, we show how protest emergence is more than a numbers game. While the initial share of activists does affect their likelihood of persistence within the system, several other factors, including the system's homophily preferences, the activists' initial conviction, and the activists' clustering also contribute significantly to the likelihood of persistence. These mathematical findings have very important social scientific implications. That homophily preferences are the strongest predictor of PVN formation implies that societies which demonstrate a preference for preserving one's own beliefs over interpersonal relationships might witness more activism. The results also show that PVNs evolve mostly as relatively stable clique-like structures in the graph's periphery. This means that even subtle degrees of insulation from non-activists significantly help activists build collective conviction over time through mutual reassurance. In other words, insulation allows for stable persistence of activism as the activists influence each other. Finally, the model shows how small changes to the activists' initial mean conviction can substantially increase the likelihood of activism persistence, even compensating for lower initial share of activists. This means that, if the first activists believe strongly in a movement, activism can emerge from an even smaller group or around a less widespread point of view.

Moreover, our model's inclusion of meaningful political interactions emphasizes the centrality of social interaction in opinion formation and network building processes, creating a much more realistic picture of the co-evolution of opinions and social ties. Even if the outcome, activism, is the same, we emphasize the mechanisms of association and confidence building, and the different patterns that society might follow based on different sets of initial conditions. This is a more realistic representation of an inherently human and social process, which is full of heterogeneity. Our model captures more than just an individual's opinion — we can measure, represent, and interpret the strength with which they hold onto it and the relative comfort that they have holding it within their current networks. Future research can use the time evolution of individual-level conviction to evaluate individual trajectories of individuals' opinions and convictions to understand how exogenous shocks and periods of intensified social reconfiguration might affect individuals. In a more interconnected and interdependent world so prone to spontaneous shocks and changes, it is essential to understand how these might affect individual persons and their networks.

The survey-informed CMAVM is novel in that it allows demographic and geographical data, inherently heterogeneous attributes of nodes and graphs, respectively, to drive social processes. The model is also extremely flexible, as the mechanisms of leveraging data for the introduction of heterogeneity are derived from simple and well-studied computational and statistical methods. By outlining specific processes to leverage several different types of key and commonly collected survey data, we hope to bridge the gap between novel theoretical models and real-world applications. We remain hopeful that these possibilities inspire others to explore wider sets of interesting, real-world phenomena and develop even more mechanisms linking computation and application.

The Panama model and this project as a whole contribute to raising awareness around the modeling of opinion dynamics in different societies, specifically Panama and Latin America, where partisan and political polarization is not necessarily the basis of some of our most recent and important developments of activism and protest. While Latin America does experience polarization, the degree to which it does and its consequences are very different from the polarization dynamics of countries like the United States and the United Kingdom upon which most polarization models are built. As we witness multi-cohort protests in previously socio-politically stable systems that do not espouse specific political ideologies, including many movements in 2022, researchers must modify their understanding of activism emergence. The factors underpinning movement emergence include exogenous shocks (like COVID-19), intensifying political conversations among immediate neighbors, and rapid social reconfiguration, and our models need to be refined to capture these features. This thesis provides new theoretical lenses and new computational methods to begin this undertaking.

If we interpret the CMAVM's conviction feature as buy-in, the model is also applicable to multiple other social dynamics that require or depend on significant buy-in, such as volunteering. In this sense, the relational nature of conviction building is an important finding, suggesting that volunteer organizations may benefit from taking advantage of existing networks of like-minded individuals or prioritize the creation of these for the preservation of their volunteer base in the long term. And while this is consistent with research in the field, our model shows a way to quantify and simulate the importance of relationships to the degree of commitment to a cause.

The model could always become more complex: incorporating a more granular understanding of geographical location, gender, race, age, and partisanship are natural extensions of this model. Integrating node attributes based on identity and allowing an individual to weigh the impact of an interaction differently based of the other party's identity could help us better understand questions surrounding the importance of these heterogeneously distributed identity features on the emergence of specific and diverse types of activism. In fact, implementing some of these extensions would enable this model to not only provide explanations for relevant social phenomena where identity-based polarization is common, such as views on different social, economic, gender, and cultural policies, but also represent the intensity of those views and the mechanisms that control those changes over time. This is all dependent on breaking the mean field assumptions that are so prevalent in network science, opinion dynamics, and evolutionary game theoretic studies of activism.

Of course, the models introduced in this work are not without limitations. As we focused on the short and medium-term aftermath of exogenous shock, we do not allow for the evolution of shock intensity over time. Realistically, exogenous shocks can worsen or improve over the period following it, and multiple shocks can coexist. These two features of shock response, spontaneous change and coexistence, are extremely important from both a mathematical and a social scientific perspective and have been the object of study of projects in evolutionary game theory and opinion dynamics. For instance, research has identified the impact of spontaneous change on the unpredictable nature of consumer choice [56] and the relevance of contagion synergy in the co-diffusion of social contagions [57]. While these few aspects could be integrated stochastically into the model by, for instance, allowing for the spontaneous addition or attrition of activists or the introduction of a new shock, this would complicate model interpretability. In all applied mathematics, there exist tradeoffs between verisimilitude and interpretability, especially in the early study of new models. Future research should consider ways to study these properties of exogenous shock in light of our model. Furthermore, the models assume, for simplicity of model execution and interpretation, that non-activists are all equally as convicted of their non-participation. However, one can imagine that a significant mechanism for non-activists to become activists might include slow socialization by activists in the non-activist's neighborhood over time, the effect which is not captured by the current version which does not allow for the evolution of non-activists' conviction. Future research can explore this extension for the potential discovery of important trends or even more potential pathways towards activism persistence.

Similarly, there are aspects of individual-level behavior that are heterogeneous but not considered in this model. For instance, several of an individual's personality traits, including sociability, political talkativeness, and trust in others, most certainly influence central aspects of the model, such as how many connections an individual can maintain, how frequent their interactions might become, and how much weight they give to each interaction. Social science research has thoroughly discussed the traits of an "activist personality" [58]. And importantly, increasingly more attitudinal surveys are asking these important personality questions, opening doors for scientific enquiry. Future research can explore these by extending this model in a few ways. For example, to understand how people with different degrees of talkativeness might experience a shock differently, the selection of a chosen individual could be probabilistic and based on the talkativeness. Another reasonable extension includes assigning individual ϕ values based on how trusting each individual is of others in their circle, as this might be a proxy for how receptive or unreceptive they will be of differing ideas. Another limitation of all AVMs is the lack of data on social connections. While this thesis uses multiple methods to produce verisimilar patterns of social connections, they are not rooted in hard data. It is difficult to imagine survey data capturing this in the near future. However, research could look at alternative data sources that can provide empirically based connection patterns, such as account interactions on social media.

Lastly, the Panama model also suffers from the need to utilize two very different data sources. Although the project was designed around the thorough WJP GPP questionnaire and survey history, multi-year data from the WJP was never made available. Additionally, behavioral and attitudinal questions surrounding protest were not asked in Panama after 2019, rendering the comparison of Panamanian society's behavior and attitude towards protest before and after the identified shock, the COVID-19 pandemic, nearly impossible. These challenges were mitigated by resorting to the Latinobarómetro as a data source, but the model would benefit from depending entirely on a single dataset for its computation, interpretation, and verification.

Bibliography

[1] Matjaž Perc. The social physics collective. *Scientific Reports*, 2019.

[2] Serge Galam. *Sociophysics: A Physicist's Modeling of Psycho-political Phenomena*. Springer, 2012.

[3] Haoxiang Xia, Huili Wang, and Zhaoguo Xuan. Opinion dynamics: A multidisciplinary review and perspective on future research. *Int. J. Knowl. Syst. Sci.*, 2(4):72–91, oct 2011.

[4] Nicholas T. Ouellette and Deborah M. Gordon. Goals and limitations of modeling collective behavior in biological systems. *Frontiers in Physics*, 9, 2021.

[5] Julian Kates-Harbeck and Michael Desai. Social network structure and the spread of complex contagions from a population genetics perspective. 08 2022.

[6] William R. Thompson and Thomas J. Volgy. *Shocks and Political Change: A Comparative Perspective on Foreign Policy Analysis*. Springer, 2023.

[7] Abdulla F. Ally and Ning Zhang. Effects of rewiring strategies on information spreading in complex dynamic networks. *Communications in Nonlinear Science and Numerical Simulation*, 57:97–110, 2018.

[8] J.A. Bondy and U.S.R. Murty. *Grpah Theory*. Springer, 2010.

[9] Jens Krause, Darren Croft, and Richard James. Social network theory in the behavioural sciences: Potential applications. *Behavioral Ecology and Sociobiology*, 62:15–27, 10 2007.

[10] Reinhard Diestel. *Graph Theory*. Springer, 2005.

[11] Jennifer M Larson. Networks of conflict and cooperation. *Annual review of political science*, 24(1):89–107, 2021.

[12] Roger V. Gould. Multiple networks and mobilization in the paris commune, 1871. *American sociological review*, 56(6):716–729, 1991.

[13] Pablo Barberá, Ning Wang, Richard Bonneau, John T. Jost, Jonathan Nagler, Joshua Tucker, and Sandra González-Bailón. The critical periphery in the growth of social protests. *PloS one*, 10(11):e0143611–e0143611, 2015.

[14] Cassy Dorff, Max Gallop, and Shahryar Minhas. Networks of violence: Predicting conflict in nigeria. *The Journal of politics*, 82(2):476–493, 2020.

[15] Jennifer M. Larson and Janet I. Lewis. Rumors, kinship networks, and rebel group formation. *International organization*, 72(4):871–903, 2018.

[16] Jennifer M. Larson, Janet I. Lewis, and Pedro L. Rodriguez. From chatter to action: How social networks inform and motivate in rural uganda. *British journal of political science*, 52(4):1769–1789, 2022.

[17] Pablo Barberá, Ning Wang, Richard Bonneau, John T. Jost, Jonathan Nagler, Joshua Tucker, and Sandra González-Bailón. The critical periphery in the growth of social protests. *PloS one*, 10(11):e0143611–e0143611, 2015.

[18] Richard A. Holley and Thomas M. Liggett. Ergodic Theorems for Weakly Interacting Infinite Systems and the Voter Model. *The Annals of Probability*, 3(4):643 – 663, 1975.

[19] Philippe Giabbanelli, Cole Freeman, Joshua Devita, Nicholas Rosso, and Zabrina Brumme. Mechanisms for cell-to-cell and cell-free spread of hiv-1 in cellular automata models. pages 103–114, 05 2019.

[20] Petter Holme and M E J Newman. Nonequilibrium phase transition in the coevolution of networks and opinions. *Physical review. E, Statistical, nonlinear, and soft matter physics*, 74(5 Pt 2):056108–056108, 2006.

[21] Olivia J Chu, Jonathan F Donges, Graeme B Robertson, and Grigore Pop-Eleches. The microdynamics of spatial polarization: A model and an application to survey data from ukraine. *Proceedings of the National Academy of Sciences*, 118(50):e2104194118, 2021.

[22] John Von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ, USA, 1944.

[23] J.F. Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286–295, 1951.

[24] Diego Ríos and Eleonora Cresto. Prisoner's dilemma, one shot and iterated. pages 930–937, 2015.

[25] J. Maynard Smith and G. Randall Price. The logic of animal conflict. *Nature*, 246:15–18, 1973.

[26] William Press and Freeman Dyson. Iterated prisoners dilemma contains strategies that dominate any evolutionary opponent. *Proceedings of the National Academy of Sciences of the United States of America*, 109:10409–13, 05 2012.

[27] Martin Nowak. Five rules for the evolution of cooperation. *Science (New York, N.Y.)*, 314:1560–3, 01 2007.

[28] Erez Lieberman, Christoph Hauert, and Martin A Nowak. Evolutionary dynamics on graphs. *Nature*, 433:312–316, 2005.

[29] Corina E Tarnita, Tibor Antal, Hisashi Ohtsuki, and Martin A Nowak. Evolutionary dynamics in set structured populations. *Proceedings of the National Academy of Sciences USA*, 106:8601–8604, 2009.

[30] Brooke W. McKeever, Robert McKeever, Minhee Choi, and Shudan Huang. From advocacy to activism: A multi-dimensional scale of communicative, collective, and combative behaviors. *Journalism & Mass Communication Quarterly*, 100(3):569–594, 2023.

[31] Jacek M Kowalski and Andrzej Pekalski. Role of zealots in a network model of cooperative political associations. *Physica A: Statistical Mechanics and its Applications*, 549:123915, 2020.

[32] Yun Luo, Chun Cheng, Yuke Li, and Changbin Yu. Opinion formation with zealots on temporal network. *Communications in Nonlinear Science and Numerical Simulation*, 98:105772, 2021.

[33] David Lazer, Brian Rubineau, Carol Chetkovich, Nancy Katz, and Michael Neblo. The coevolution of networks and political attitudes. *Political communication*, 27(3):248–274, 2010.

[34] Fergus G Neville, David Novelli, John Drury, and Stephen D Reicher. Shared social identity transforms social relations in imaginary crowds. *Group Processes & Intergroup Relations*, 25(1):158–173, 2022.

[35] Kai Epstude Maja Kutlaca, Martijn van Zomeren. Friends or foes? How activists and non-activists perceive and evaluate each other. *PLOS One*, 15(4):1–59, 2020.

[36] Sandra González-Bailón and Ning Wang. Networked discontent: The anatomy of protest campaigns in social media. *Social Networks*, 44:95–104, 2016.

[37] Flora Cornish, Catherine Campbell, and Cristián Montenegro. Activism in changing times: Reinvigorating community psychology – introduction to the special thematic section. *Journal of social and political psychology*, 6(2):526–542, 2018.

[38] D Rye. Activists and activism success: towards a grounded conceptualisation. *Interest Groups and Advocacy*, 2024.

[39] Matthew T. Ballew, Matthew H. Goldberg, Seth A. Rosenthal, Matthew J. Cutler, and Anthony Leiserowitz. Climate change activism among latino and white americans. *Frontiers in communication*, 3, 2019.

[40] AmaÇ HerdaĞdelen, Wenyun Zuo, Alexander Gard-Murray, and Yaneer Bar-Yam. An exploration of social identity: The geography and politics of news-sharing communities in twitter. *Complexity (New York, N.Y.)*, 19(2):10–20, 2013.

[41] Qawi K. Telesford, Karen E. Joyce, Satoru Hayasaka, Jonathan H. Burdette, and Paul J. Laurienti. The ubiquity of small-world networks. *Brain connectivity*, 1(5):367–375, 2011.

[42] Kimmo Eriksson, Irina Vartanova, and Pontus Strimling. How does political discussion frequency impact political moral opinions? the moral argument theory of opinion dynamics. *Frontiers in Psychology*, 13:915252, 2022.

[43] Nella Van Dyke and Bryan Amos. Social movement coalitions: Formation, longevity, and success. *Sociology Compass*, 11(7):e12489, 2017.

[44] James Loxton. The puzzle of panamanian exceptionalism. *Journal of Democracy*, 33(1):85–99, 2022.

[45] Naomi Hossain and Jeffrey Hallock. Food, energy & cost of living protests, 2022. *NY: Friedrich Ebert Stiftung NY. https://ny. fes. de/article/food-energy-cost-of-living-protests-2022*, 2022.

[46] Martinelli revokes law 8 in response to indigenous protests - panama. *Business News Americas*, 2011.

[47] Nike ditches shoe design after panama's indigenous guna protest, 2019.

[48] CE Noticias Financieras English. Crisis in panama: construction workers and educators resumed protests due to "non-compliance" with agreements with the government, 2022.

[49] CE Noticias Financieras English. Young people protest against changes in state aid and corruption in panama, 2021.

[50] CE Noticias Financieras English. Anti-corruption protests, power outages and empty shelves in panama, 2022.

[51] Indigenous groups block roads in panama in protest over land titles: Panama protests, 2022.

[52] Traben Pleasant and Shaozeng Zhang. Capabilities and aspirations: A multi-theory approach to information and communications technology for development in rural panama. *The Electronic journal of information systems in developing countries*, 87(4), 2021.

[53] Marcus A Johnson. Racial-ized democracy: The electoral politics of race in panama, 2017.

[54] Edward Telles and Stanley Bailey. Understanding latin american beliefs about racial inequality. *The American journal of sociology*, 118(6):1559–1595, 2013.

[55] Cheryl Pritlove, Clara Juando-Prats, Kari Ala-leppilampi, and Janet A Parsons. The good, the bad, and the ugly of implicit bias. *The Lancet (British edition)*, 393(10171):502–504, 2019.

[56] John Kemp. Spontaneous Change, Unpredictability and Consumption Externalities: a Dynamic Approach to Consumer Choice. *Journal of Artificial Societies and Social Simulation*, 2(3):1–1, 1999.

[57] Ho Chun Herbert Chang. Co-Diffusion of Social Contagions. pages 1–59, 2018.

[58] Jan Matti Dollbaum and Graeme B. Robertson. The activist personality: Extraversion, agreeableness, and opposition activism in authoritarian regimes. *Comparative political studies*, 56(11):1695–1723, 2023.