

Strategic Voting and Social Networks

by

Alan Tsang

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Computer Science

Waterloo, Ontario, Canada, 2018

© Alan Tsang 2018

Examining Committee Membership

The following served on the Examining Committee for this thesis. The decision of the Examining Committee is by majority vote.

External Examiner	Jeff Rosenschein Professor
Supervisor	Kate Larson Professor
Internal Members	Robin Cohen Professor Peter van Beek Professor
Internal-external Member	Stan Dimitrov Associate Professor

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

With the ever increasing ubiquity of social networks in our everyday lives, comes an increasing urgency for us to understand their impact on human behavior. Social networks quantify the ways in which we communicate with each other, and therefore shape the flow of information through the community. It is this same flow of information that we utilize to make sound, strategic decisions. This thesis focuses on one particular type of decisions: voting. When a community engages in voting, it is soliciting the opinions of its members, who present it in the form of a ballot. The community may then choose a course of action based on the submitted ballots. Individual voters, however, are under no obligation to submit sincere ballots that accurately reflects their opinions; they may instead submit a strategic ballot in hopes of affecting the election’s outcome to their advantage. This thesis examines the interplay between social network structure and strategic voting behavior. In particular, we will explore how social network structure affects the flow of information through a population, and thereby affects the strategic behavior of voters, and ultimately, the outcomes of elections.

We will begin by considering how network structure affects information propagation. This work builds upon the rich body of literature called opinion dynamics by proposing a model for skeptical agents — agents that distrust other agents for holding opinions that differ too wildly from their own. We show that network structure is one of several factors that affects the degree of penetration that radical opinions can achieve through the community. Next, we propose a model for strategic voting in social networks, where voters are self-interested and rational, but may only use the limited information available through their social network contacts to formulate strategic ballots. In particular, we study the “Echo Chamber Effect”, the tendency for humans to favor connections with similar people, and show that it leads to the election of less suitable candidates. We also extend this voter model by using boundedly-rational heuristics to scale up our simulations to larger populations. We propose a general framework for voting agents embedded in social networks, and show that our heuristic models can demonstrate a variation of the “Micromega Law” which relates the popularity of smaller parties to the size of the population. Finally, we examine another avenue for strategic behavior: choosing *when* to cast your vote. We propose a type of voting mechanism called “Sticker Voting”, where voters cast ballots by placing stickers on their favored alternatives, thereby publicly and irrevocably declaring their support. We present a complete analysis of several simple instances of the Sticker Voting game and discuss how our results reflect human voting behavior.

Acknowledgements

Foremost, I would like to thank my supervisor Professor Kate Larson for her support, knowledge, and guidance, on both academic and professional matters. Her positivity is contagious and makes even the most daunting challenge achievable. She is a truly inspiring role model for graduate students.

I also thank my committee members for their invaluable feedback to further improve this thesis — Professor Peter van Beek for his insightful questions, Professor Stan Dimitrov for offering a different perspective on my research problems, and Professor Robin Cohen for her attention and guidance in my career development. I am also deeply grateful for my external-examiner, Professor Jeff Rosenschein, for his continued interest in my research, and his excellent suggestions for future directions.

I appreciate the support of the wonderful administrative staff at the Cheriton School of Computer Science for assisting me with university bureaucracy, especially Margaret Towell, Neoma Power, Helen Jardine, Paula Roser, and Wendy Rush.

During my Ph.D. journey, I have crossed paths with many supportive friends, graduate students, and fellow researchers. John Doucette and Hadi Hosseini, both collaborators and academic predecessors, livened up my journey with stimulating discussions and encouraging camaraderie. Many others, I thank as well for enriching my Ph.D. life: Amirali Salehi-Abari, Vijay Menon, Michael Cromier, Curtis Bright, Ivana Kajic, Aaron Moss, Nathan Fish, Caitlin McLaren, Navid Vafaei, Cecylia Bocovich, Graham Pinhey, Noel Sardana, Sarah Kaiser, Catherine Holloway, Jane Smart, Debbie Spear, Krishna Vaidyanathan, Haotian Zhang, Ben Armstrong, and Pan Shengying. I also thank the outstanding researchers whom I have met with for their insights and discussions: David Parkes, Edith Elkind, Omer Lev, Reshef Meir, Nick Mattei, Yair Zick, John Dickerson, Joanna Drummond, Tyrone Strangway, Svetlana Obraztsova and Zinovi Rabinovich.

Finally, and most importantly, I thank my family for providing me with this opportunity. Crossing the oceans with my sister and I to start a new life in a foreign country must have seemed a daunting prospect. But now, I hope it has finally borne fruit.

Dedication

For my family, who braved the unknown so I may have this opportunity.

Table of Contents

List of Tables	xi
List of Figures	xii
1 Introduction	1
1.1 Voting Theory and Strategic Voting	2
1.2 Social Networks	3
1.3 The Intersection	3
1.4 Statement of Contributions	5
2 Background	7
2.1 Game Theory	7
2.2 Social Choice and Voting	13
2.3 Graph Theory	15
2.4 Social Networks	15
2.4.1 Opinion Dynamics	16
2.4.2 Knowledge Graphs	16
2.5 Random Graph Models	16
2.5.1 Erdős-Rényi	17
2.5.2 Barabási-Albert	17
2.5.3 Homophily	18

3	Related Work	20
3.1	Voting Models	20
3.1.1	Voting Equilibria	20
3.1.2	Poisson Games	21
3.1.3	Local Dominance	23
3.1.4	Sequential Voting	23
3.1.5	Sequential Voting and Herding	24
3.1.6	Stackelberg Voting	26
3.1.7	Iterative Voting	26
3.1.8	Modal Logic Voting Models	27
3.2	Voter Behavior	28
3.2.1	Pólya’s Urn Model	28
3.2.2	Approval Voting in Doodle Polls	28
3.2.3	Observing Human Strategic Voting Behaviors	29
3.2.4	Strategic Effects of Early and Late Ballots	32
3.3	Voting in Networks	35
3.3.1	Preference Aggregation in Social Networks	35
3.3.2	Exploring Duverger’s Law	37
3.3.3	Iterative Voting in Social Networks	38
3.3.4	Other Related works	39
4	Opinion Dynamics	41
4.1	Related Work	42
4.2	Opinion Dynamics Model	44
4.2.1	Our Model	46
4.2.2	Graph Models	47
4.2.3	A-priori Trust Models	48
4.3	Empirical Simulations	48

4.3.1	Experimental Design	49
4.3.2	Influence of Extremists	50
4.3.3	Opinion Polarization	54
4.4	Discussion	61
4.5	Conclusion	63
5	Voting in Social Networks	65
5.1	Model	66
5.1.1	Model of Voters	67
5.1.2	Response Model	68
5.1.3	Sequential vs Simultaneous Updates	69
5.2	Experimental Design	70
5.3	Results	70
5.4	Convergence	73
5.5	Discussion	77
5.6	Conclusion	79
6	Heuristic Voter Models	81
6.1	Framework	81
6.1.1	Voting Model	83
6.1.2	Fully Rational Voter	84
6.1.3	Voter Heuristics	86
6.2	Comparison via Simulations	98
6.3	Addressing the Desiderata	101
6.4	Case Study: The Micromega Rule	103
6.4.1	Results	104
6.5	Conclusion	106

7	Vote Timing	107
7.1	Introduction	107
7.2	Sticker Voting Model	108
7.2.1	Plurality Sticker Voting	109
7.2.2	Solution Concept	109
7.3	Complete Information Game	110
7.3.1	Rational Voter Behavior	115
7.4	Incomplete Information Game	115
7.4.1	Utilities for Round t	125
7.4.2	The Uncontrolled Subgames	128
7.5	Discussion & Conclusion	129
8	Conclusion	134
8.1	Summary of Main Results	135
8.2	Future Directions	136
	References	140

List of Tables

2.1	Prisoner’s Dilemma Payoff Matrix	8
2.2	Cuban Missile Crisis Game Payoff Matrix	10
5.1	Effects of Update and Tie-Breaking Methods (ER and BA graphs)	72
6.1	Adherence to the Desiderata by the Heuristics	102
6.2	Average SF-Ratios. TieH and Poisson models	105
7.10	EVO 2017 “People’s Choice” Voting Results	131

List of Figures

2.1	Cuban Missile Crisis Game.	10
2.2	Examples of ER and BA random graphs	18
4.1	A Erdős-Rényi graph with homophily	47
4.2	Opinions of moderates over the course of an experiment	51
4.3	The convergence mean opinion of moderates	52
4.4	Effects of introducing noise to the model	53
4.5	The average polarization of moderates when exposed to extremists of opposing camps	56
4.6	Evolution of opinions in moderates, with partially polarized initial opinions	57
4.7	The average polarization of moderates with initial opinions drawn from $\beta(0.5, 0.5)$	58
4.8	Evolution of opinions in moderates, on a modified ER-graph with homophily, with partially polarized initial opinions	59
4.9	Average polarization of moderates on a modified ER-graph with homophily, with partially polarized initial opinions	60
5.1	Fraction of agents strategizing (3- and 4-candidates)	74
5.2	Price of Honesty and Stability under various conditions	75
5.3	Degree of convergence to a 2-candidate system, measured as the SF ratio (3- and 4-candidates)	76
5.4	Distribution of SF Ratios show the degree with which results from each graph model conform to Duverger's Law.	77

5.5	Voters need not converge to stability.	78
6.1	Runtime in seconds to construct one voter response ($m = 5, n = 100$). . . .	98
6.2	Runtime in seconds to construct one voter response ($m = 5, n = 500$) . . .	98
6.3	Runtime in seconds to construct one voter response ($m = 6, n = 100$) . . .	99
6.4	Rate of Disagreements from the Full Voter Model ($m = 5, n = 500$). . . .	100
6.5	Conditions where Full votes sincerely in green.	101
6.6	Conditions where TieH disagrees in red or green.	101
6.7	Conditions where Poisson disagrees in red or green.	101
6.8	Average SF Ratios of graphs of different sizes	105

Chapter 1

Introduction

As social media becomes increasingly ubiquitous in recent years, the influences of social networks become increasingly subtle and far reaching. They range from the innocuous — for instance, marketing of new products, and spread of rumors or viral content — to the monumental — propagation of news and information, mobilization of political activism, and even influence over outcomes of major governmental decisions.

These networks represent complex webs of interactions between both individuals and institutions. They capture relationships and social structures that define communities both niche and vast. The relationships within these communities hold the key to how information flows within the network, and ultimately, how individuals' actions may be influenced by each other and by the institutions whom they respect.

The last decade has seen tremendous growth in the popularity of social networks in both popular media and research communities. The availability of “big data” gives scientists and researchers detailed information on how these networks and their members evolve and change over time. In this thesis, we focus our attention on studying how the structure of social networks shapes the flow of information within a community, and thereby impacts the actions of the individuals and the choices of the community as a whole. This thesis builds upon the intersection of two major concepts: **voting theory**, which studies how communities make collective decisions informed by its members; and **social networks**, which mediate the flow and availability of information within the community.

1.1 Voting Theory and Strategic Voting

Many early developments in voting theory came about as a result of the French revolution. Nicolas de Condorcet was a mathematician who would become intimately involved in the revolution, and would be responsible for the new state education system, as well as authoring a draft for the new constitution.

But his interest in voting and democratic processes predates the revolution. In his 1785 essay, Condorcet describes (amongst many other important ideas) what would become known as Condorcet’s Jury Theorem [37]. He presents a scenario where a jury is tasked with the decision to convict or acquit a defendant. Given this binary decision, he states that if each juror has at least a better than random chance of arriving at the correct conclusion, then increasing the number of juror increases the probability that a majority vote amongst the jurors will yield the correct outcome. As the number of jurors increases, it becomes almost certain that the jury will produce the correct decision.

The Jury Theorem describes a simplified case for general voting problems. In essence, a voting process (i.e. an election) is about taking the preferences of a community of agents, and aggregating them to produce a final decision that is somehow representative of all the agents’ preferences. The agents could be humans that are part of a community, or they could be virtual agents representing a sensor network or robotic swarm. The agents may be co-operative with similar goals and an objectively “correct” solution, or they could be self-interested or adversarial, each with their own preferred outcomes, possibly seeking to manipulate the result to their own advantage.

We are primarily interested in the second scenario, where individual voters are self-interested, have different subjective preferences, and are only concerned with voting in such a way to best sway the outcome of the election in their favor. As before, the voting process is used to elicit the private opinions from members of the community, so that the group as a whole may choose the most appropriate alternative. However, individual voters are under no obligation to submit a sincere ballot that reflects their true preferences. For instance, voters often face a decision between casting a sincere ballot for their favorite candidate, or casting a strategic ballot for a less preferred candidate with more promising prospects. In fact, the Gibbard-Satterthwaite theorem states that any such voting system with at least 3 candidates is either dictatorial¹ or is susceptible to this form of strategic voting [65, 118], so strategic voting may be seen as an intrinsic and unavoidable aspect of the voting process that demands closer analyses.

¹i.e There exists some voter whose favorite candidate is always the winning candidate

1.2 Social Networks

Social networks have gained tremendous popularity in both research communities and popular media in the last decade. These networks capture the relationships and social structures that define communities both large and small. While early research into social networks produced theories on how information flows through such a network [35, 122], and how influence is exerted by the popular and powerful [77], recent years have demonstrated the power of social networks, with social media playing a major role in the formation and execution of political revolutions[54].

The notion that the web of relationships and interactions one participates in, can be captured and analyzed can trace its roots back to theories of social groups of Durkheim and Tönnies from the 1890s [49, 125]. But it was in Jacob Moreno's *sociograms* from the 1930s that see a strong resemblance to the modern concept of the social network [96]. Sociograms represent social interactions using tools from graph theory, with individuals being represented as a set of vertices V , and an edge $(x, y) \in E$ connects two vertices $x, y \in V$ if the two individuals share in some sort of social interaction. Edges may also be directed to represent asymmetric power or influence relationships.

Modern social network research continues to use a similar terminology, but frequently expands the criteria necessary to qualify for a social interaction. This follows from the “weak ties” theory, which depicts the strength of seemingly tenuous connections between mere acquaintances, and emphasizes their influence in processes of information gathering, decision making and innovation[66]. The explosion of data made available through the Internet, and online social media platforms in particular, has also greatly expanded the size of social networks available for study. The size and density of modern social networks make computational methods indispensable for studying human behavior in these systems. Modeling the evolution of and behaviors on these networks is an important and evolving area of research.

1.3 The Intersection

Social networks describe how information and opinions flow within a community; Voting theory discusses how information and opinions might be aggregated from a community. It seems only natural, then, to examine both topics in conjunction with each other: How does the exchange of ideas within a social network affect the aggregation of those thoughts?

Voting theory is part of the wider field of social choice, and while this latter body of

literature offers a vast array of tools and mechanism to choose from, relatively little attention has been paid on understanding exactly how presence in, and influence within these networks affect how people vote. While several researchers [10, 63] have examined social choice problems where friends influence one’s votes in abstract ways, few have specifically incorporated the structure of the social network into the analysis [10].

The main research question I wish to consider is the impact of social networks on the behavior of strategic voters. In order for voters to strategize, they must have some knowledge of the relative competitiveness of the candidates; i.e. they must possess the information to judge when their favorite candidate has become a “lost cause” and direct their ballot in support of their second-choice candidates. In this thesis, we will use the user’s social network to represent the sources of information from which they may poll this information. While a casual definition of social network may include only friends and relatives, our definition also includes knowledgeable acquaintances, valued advisors, and trusted media institutions. Voters in the network adjust their outlook over time based on their observations, and may revise their ballots accordingly. In the next few chapters, we will explore several aspects of the interactions between social networks and strategic voting behavior in order to fulfill the following thesis statement:

Thesis Statement. *This thesis intends to advance our understanding of the interactions between social network structure and strategic voter behavior, by examining the question of how different network structures alter the flow of information through a network, and thereby alter the aggregate outcomes of independent and strategic voters over time.*

These insights will help refine explanatory and predictive models both within the social choice community and further abroad in substantive domains. In particular, they will help explain the causes of social phenomena, as well as gauge the potential success of social choice strategies, and the ultimate success of various candidates.

In Chapter 4, we answer the question of how network structure may affect the spread of information across a community of skeptical agents by drawing upon and extending existing models in opinion dynamics. In Chapter 5, we propose a model of strategic voting behavior in social networks, and show that we can replicate the “Echo Chamber Effect” by tweaking structural parameters of networks. In Chapter 6, we extend our voter model by incorporating efficient heuristics to allow our simulations to scale up to larger populations; we also propose a general framework for voting agents embedded in social networks. Finally, in Chapter 7, we isolate the temporal component of our model and examine how voters may strategize by controlling when they commit their ballots in a “Sticker Voting” election.

1.4 Statement of Contributions

This section lists the technical contributions of each chapter. We begin with Chapter 4 where we examine models of opinion dynamics and extend them to reflect skepticism between agents of differing opinions:

- We propose a model of opinion dynamics incorporating skepticism.
- We analyze the convergence behavior in social networks using random graph models, in the presence of extremists.
- We identify that voter empathy and network structure (network type, connectivity and homophily) are key factors in determining the influence of extremists.
- We analyze the conditions that allow for opinion stratification.

In Chapter 5, we propose and analyze a model for how strategic voting occurs in social networks:

- We propose a model of strategic voting based on incomplete information in social networks.
- We analyze strategic voting behavior in social networks using random graph models.
- We show that strategic behavior always improves social welfare.
- We show that strategic behavior diminished in networks with homophily, due to inability to observe opportunities, and that this leads to lower social welfare.
- We show that Duverger's Law holds in many graphs, but not in sparse networks.

In Chapter 6, we extend that model of strategic voting by proposing and analyzing heuristic models that allow simulations to scale up to larger elections:

- We propose a set of desiderata for voter models in social networks.
- We propose heuristic models for strategic voters in social networks that allow our model to scale up to larger populations and more candidates.

- We analyze these heuristic models based on desiderata, and show two heuristic models offer the best combination of desiderata: **TieH** and **Poisson**
- We demonstrate our heuristics by examining the Micromega Rule in elections with varying number of voters.

Finally, in Chapter 7, we examine the issue of how voters choose *when* to cast their ballots by proposing and analyzing the Sticker Voting model:

- We propose the Sticker Voting Mechanism and introduce the concept of strategic timing.
- We analyze the equilibrium behavior in sticker voting using simple scenarios of complete and incomplete information.
- We discussion of the applicability of our results to human voters and proposed further extensions.

Chapters 4, 5 and 7 are based on published works coauthored with my thesis advisor Kate Larson [128, 129, 130]. Chapter 6 is based on a paper that is currently under review, coauthored with Kate Larson and Amirali Salehi-Abari. The **Max-M** voter heuristic model proposed in this chapter was an idea originated from Amirali Salehi-Abari. This thesis contains text borrowed from those original publications, but were written by me.

Chapter 2

Background

This chapter provides the background on voting and social networks necessary for understanding the remainder of this thesis. We begin with a review of several basic game theory concepts, which allows us to formally define a voting process as a voting game. We will define and elaborate on different solution concepts (equilibria) useful for discussing strategic behavior for players (voters) in these voting games. Next, we will define basic concepts from graph theory and discuss how we use them to model social networks. In particular, we focus on a number of models of randomly generated graphs that we use in several parts of the thesis.

2.1 Game Theory

Economic game theory models complex decision making scenarios between multiple parties as *games*. A game begins with a set of players. Each player is faced with a choice between a number of possible actions, and the combination of actions selected by all players forms an outcome for this game. The outcome, in turn, determines the payoffs of the game to each individual player.

Formally, let $V = \{1, 2, \dots, n\}$ denote the set of players. Each player $i \in V$ simultaneously plays a strategy by selecting an action b_i from a set of available actions. The available actions form the action space \mathcal{B} . For simplicity, we may assume this action space is the same across all players. Let the vector $\mathbf{b} = (b_1, b_2, \dots, b_n)$ denote the collective actions chosen by the players; $\mathbf{b} \in \mathcal{B}^n$ represents the outcome of this particular game. Each player i has a utility function $u_i : \mathcal{B}^n \rightarrow \mathbb{R}$ that maps each outcome to a real number, representing

		Column Player	
		Cooperate	Defect
Row Player	Cooperate	(-1,-1)	(-10,0)
	Defect	(0,-10)	(-8,-8)

Table 2.1: Prisoner’s Dilemma Payoff Matrix

the payoff to agent i . Note that u_i is *not* generally shared across agents; i.e. agents may have different preferences over the possible outcomes. The payoff can be thought of as a numerical measurement of satisfaction that an agent gains from a particular outcome. By way of analogy, it is sometimes compared to a monetary payoff (or cost) to the agents. The game is thus defined as $G = \langle V, \mathcal{B}^n, u \rangle$. Agents take actions to maximize their personal payoffs from the final outcome.

This basic setup for games captures many interesting scenarios. For instance, the classic Prisoner’s Dilemma is constructed as a 2-player game, where two criminals have been arrested by the police and interrogated separately. Each player has two actions available to her: $\mathcal{B} = \{C, D\}$, corresponding to Cooperate (remaining silent) and Defect (ratting the other person out). This leads to four possible outcomes of the game, with four corresponding payoffs outline in Table 2.1, typically measured as “months in prison”. For this scenario, we may assume the payoffs are symmetric between the players; i.e. $u_1(b_1, b_2) = u_2(b_2, b_1)$. From the first player’s perspective, the ideal scenario is (D, C) where she rats out her accomplice, who remains silent and is fully implicated: player one walks away free ($u_1(D, C) = 0$), while player two serves the full sentence ($u_2(D, C) = -10$); conversely, this is the least ideal outcome for player two. If both players Cooperate, they are released with a slap on the wrist ($u_1(C, C) = -1$). If both players Defect, they are both convicted and charged a nearly full length sentence ($u_1(D, D) = -8$).

A natural question to ask is, what would actual players do in such a game? Are there some outcomes that are more reasonable than others? We begin with the game theory concept of a Nash Equilibrium, named after its pioneer John Nash [101]. In its simplest form, an outcome is a Nash Equilibrium if no agent has any incentive to change her strategy; that is, no agent will derive a higher payoff by playing a different strategy (assuming everyone else plays the same strategy). Formally, an outcome $\mathbf{b} = (b_1, b_2, \dots, b_n)$ is a Nash Equilibrium if for every player i , $u_i(\mathbf{b}) \geq u_i(\mathbf{b}^*)$, where \mathbf{b}^* is obtained from \mathbf{b} by altering i ’s strategy from b_i to another strategy $b' \neq b_i$.

In the Prisoner’s Dilemma outlined above, one can verify that the outcome (D, D) is a Nash Equilibrium outcome. That is, neither player profits by altering their strategy to

C. We may also repeat this process to verify that no other outcome is a Nash Equilibrium outcome. A more expedient method of verifying this fact is through the *iterated elimination of dominated strategies*. We say player i 's strategy b dominates an alternative strategy b' when i derives more utility by playing b than b' , regardless of how others players play. Formally, let \mathbf{b}_{-i} denote the strategies played by all players other than i . Moreover, we write (\mathbf{b}_{-i}, b_i) to indicate the outcome derived from \mathbf{b}_{-i} where i plays b_i . Then, for player i , we say b dominates an alternative strategy $b' \neq b$, if $u_i(\mathbf{b}_{-i}, b) \geq u_i(\mathbf{b}_{-i}, b')$, for all \mathbf{b}_{-i} . We say the dominance is *strong* if the inequality is strict. A dominated strategy is, in this sense, an unreasonable strategy for an agent to consider playing, and therefore may safely eliminate it from consideration. By repeatedly eliminating dominated strategies, we may reduce the action space to only a single outcome remains. If that is the case, the remaining outcome is a Nash Equilibrium. If we eliminated only *strongly* dominated strategies, it is a unique Nash Equilibrium.

In this basic definition of a game, we allow players to select a single action as their strategy to play. This is called a pure strategy. We may allow players to play a *mixed* strategy, which is a distribution over the action space. For instance, in the Prisoner's Dilemma, a player may choose to Cooperate with some probability p and Defect with probability $1 - p$. This allows us to distinguish Nash Equilibria between Pure Nash Equilibria (where players play only pure strategies) and Mixed Nash Equilibria (where players may play mixed strategies). While this does not change the results of the Prisoner's Dilemma analysis (i.e. (D, D) in the unique Mixed Nash Equilibrium), it is likely to change the solutions of games in general.

We may also extend the game by allowing the agents to take actions one after another rather than simultaneously. In these Sequential Games, agents that act later in the sequence can observe the actions of preceding agents (i.e. the history of play), and select strategies using that information. The sequence of actions can be represented as a rooted tree, with branches corresponding to the different actions taken by each agent. The nodes of the tree represent distinct states of the game. The game begins at the root node. Non-terminal nodes (which includes the root node, unless the game is trivial) are labelled with the active player; i.e. the player who must now make a move by playing an action. Taking the action moves the game into an appropriate child node. Terminal nodes signal the end of the game and are labelled with the payoffs for each player. This tree based representation of the game is also called the Extensive Form Game.

Figure 2.1 shows an example of a Sequential Game inspired by the political situation during the Cuban Missile Crisis in 1962. This tense situation was precipitated by the U.S. discovery that the U.S.S.R. had supplied Cuba with nuclear weapons. President Kennedy responded by ordering a naval blockade of the island to prevent additional supplies and

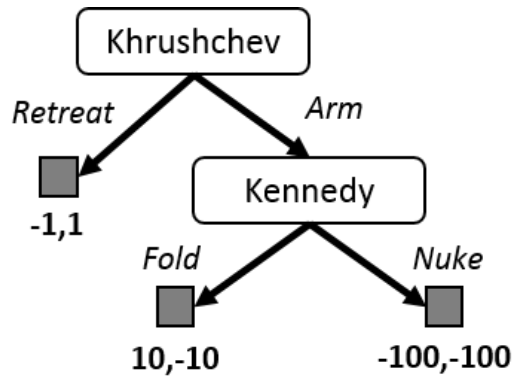


Figure 2.1: Cuban Missile Crisis Game.

		Kennedy	
		Fold	Nuke
Khrushchev	Retreat	$(-1, 1)$	$(-1, 1)$
	Arm	$(10, -10)$	$(-100, -100)$

Table 2.2: Cuban Missile Crisis Game Payoff Matrix

missiles from reaching the island. Soviet leader Khrushchev then ordered the nuclear missiles be armed and launched if the U.S. forces invaded the island. President Kennedy was then faced with a difficult decision — to withdraw the blockade and allow the Soviet supplies to reach Cuba (i.e. to “fold”), or to invade the island nation and possibly start a nuclear exchange (i.e. to “nuke”). Figure 2.1 illustrates the sequential nature of this situation — first Khrushchev decides whether to supply arms to Cuba, then Kennedy decides on a response.

Table 2.2 shows the payoffs for the players (left entry representing Khrushchev; right entry, Kennedy). Notice there are two Nash Equilibria in this game: (Retreat,Nuke) and (Arm,Fold). However, upon examining the game tree, the (Retreat,Nuke) equilibrium seems unsatisfactory because of the sequence of play. Kennedy’s threat to Nuke is a non-credible threat because once Khrushchev commits to the Arm action, Kennedy’s best course of action is to Fold.

In general, applying Nash Equilibria as a solution concept to Sequential Games leads to numerous unsatisfactory equilibria of this sort. Therefore, we refine our concept of equilibria for Sequential Games to consider the sequence of play, in the form of the Subgame Perfect Equilibrium (SPE). SPE views each subtree of the game tree as its own “subgame”,

and SPE requires that player strategies form a Nash Equilibrium for every single subgame. More formally, given an Extensive Form Game G , let us define a subgame G' rooted at node h of G as the game represented by the subtree of G rooted at h . A strategy profile s is a Subgame Perfect Equilibrium of G if s forms a Nash Equilibrium for every subgame G' rooted at *every* node h in G .

Applying this new concept to our Cuban Missile Crisis scenario, we see the the Nash Equilibrium (Retreat,Nuke) is not a Subgame Perfect Equilibrium. Kennedy's threat to Nuke is non-credible if we reach the subgame where he must act. (Arm,Fold) is the only SPE.¹

Alternatively, games may be extended to contain elements of uncertainty. Uncertainty may arise from two factors — uncertainty as to the priorities of other players, and uncertainty from the environment itself. The latter is easiest to model: we simply designate a special player as *Nature* (also called the *Chance* player in some literature). Chance players do not take actions to maximize their utility; rather they take actions based on a simple probability distribution, comparable to rolling dice. This may be treated as a mixed strategy played by the Chance player, and expected utilities for each player may then be calculated in the usual way.

To account for the former type of uncertainty, we turn to the Bayesian Game, an extension of the simultaneous (i.e. non-sequential) game. In a Bayesian game, the priorities of the players (i.e. their utility functions) are based on their *types*, which are randomly assigned at the beginning of the game. A player knows their own type, but the types of other players are unknown. The player maintains some belief over the likely types of the other players, and uses that belief to formulate a strategy based on the expected outcomes.

Formally, in addition to the set of players V and set of actions \mathcal{B} , each player also has a type t_i from a set of types T_i . The game begins with a draw ω from the set of possible states of the world Ω . Each player i acquires a type $t_i = \tau_i(\omega)$ based on a mapping function $\tau_i : \Omega \rightarrow T_i$. Let $C_i \subseteq \mathcal{B} \times T_i$ be the subset of actions available to player i due to its type t_i . The utility function $u_i : \Omega \times \mathcal{B}^n \rightarrow \mathbb{R}$ of player i is modified to map the player's type and the action profile to the payoff received by the player. Finally, each player i also has a *belief* p_i , a probability distribution representing how likely other players are to be of each type.

A pure strategy s_i maps player i 's type to an admissible action from C_i . The set S_i of

¹It is interesting to note that the outcome of the actual Cuban Missile Crisis resulted in a prolonged standoff between the two nuclear superpowers during which Kennedy maintained the Cuban blockade. After a tense seven days, Krushchev decided to withdraw Soviet supply ships from Cuban waters, ending the potentially disastrous situation.

pure strategies is thus defined as,

$$S_i = \{s_i : T_i \rightarrow \mathcal{B} \mid (s_i(t_i), t_i) \in C_i, \forall t_i\}.$$

For an action profile $\mathbf{b} = (b_1, b_2, \dots, b_n)$, the expected payoff for player i is calculated as,

$$u_i(\mathbf{s}) = \mathbb{E}_{\omega \sim p_i} u_i(\omega, s_1(\tau_1(\omega)), s_2(\tau_2(\omega)), \dots, s_n(\tau_n(\omega))).$$

A strategy profile of this Bayesian Game $G = \langle V, \Omega, \mathcal{B}, \langle T_i, C_i, u_i, p_i, \tau_i \rangle_{i \in V} \rangle$ is a Bayes-Nash Equilibrium (BNE) if it is a pure or mixed Nash Equilibrium in the game $\tilde{G} = \langle V, \tilde{A} = S_1 \times S_2 \times \dots \times S_N, u \rangle$.

Analogous to the abundance of unsatisfactory Nash Equilibria in Sequential Games, there may be many “implausible” Bayes-Nash Equilibria in a sequential Bayesian Game. One solution to this issue is a refinement called Perfect Bayesian Equilibria (PBE) that extends Bayes-Nash Equilibria to the sequential setting. Instead of beliefs being placed on the distribution of types, the beliefs are updated throughout the game, and each player seeks to maximize their expected utility according to the updated beliefs at each step of the game.

Recall in the Sequential Game, each node is labeled with the active player. We further partition the nodes of each active player into *information sets*. The nodes of each information set are indistinguishable from each other, from the active player’s perspective. The players know the game has arrived in a particular information set, but do not know which node of the information set the game is currently in. The player’s *belief system* assigns probabilities to each node in the game tree where they are the active player, such that the probabilities of nodes of each information set sum up to 1. The belief system is *consistent* if the probabilities assigned to each node reflect the probability that we reach that node through the game tree (i.e. it is updated via Bayes’ Rule).

Moreover, we say a strategy profile is *sequentially rational at a particular information set* if the active player’s chosen strategy maximizes her utility with respect to the strategies of other players (in expectation), given her belief system. The strategy profile is *sequentially rational* if it is sequentially rational at *all* information sets.

Finally, we say that a strategy profile and belief system is a Perfect Bayesian Equilibrium (PBE) if (1) the strategy profile is sequentially rational given the belief system, and (2) the belief system is consistent given the strategy profile (wherever possible). This

final “wherever possible” caveat is an important one. Information sets that are unreachable given the strategy profile cannot have probabilities computed via Bayes’ Rule. They are considered situations that are “unrealistic” and are said to be *off the equilibrium path*, and may be assigned arbitrary beliefs.

2.2 Social Choice and Voting

Social choice theory is the study of group decision making scenarios where a community of individuals must collectively agree upon a single decision². These scenarios may include the fair division of resources between members of the community, the matching of individuals needs to talents or resources, and the aggregation of the collective preferences of many diverse individuals about a number of alternatives. We are concerned with the last scenario, particularly when we are interested in finding the top ranking alternative according to the opinions of the community. This scenario is otherwise known as *voting*.

Voting may be objective or subjective in nature. In an objective setting, voters act cooperatively as noisy sensors to reveal some hidden, underlying ground truth. For instance, a panel of football judges may vote on whether or not the ball has crossed a particular line. There is an objective “best” answer, and it is in the best interest of the voters to select it. However, this thesis is primarily concerned with *subjective voting*, where voters have differing, and frequently competing, interests in the outcome of the election. Voter preferences are subjective, and voters will not necessarily vote sincerely. They may reveal preferences that are different from their true preferences in an effort to manipulate the final outcome of the election. This is called tactical or *strategic voting*, and its effects on the outcome of the election may be unpredictable and far-ranging. It is this strategic voting behavior that is one of the two focal points of this thesis.

Strategic voting may be formulated as a game. Our n players $V = \{1, 2, \dots, n\}$ will be our set of voters (or agents). They will be voting over a set of m alternatives (or candidates), $C = \{c_1, c_2, \dots, c_m\}$. The action space for each voter is denoted by the set of admissible ballots \mathcal{B} , which depends on the voting system. The voting system will also prescribe a social choice function $\mathcal{F} : \mathcal{B} \rightarrow 2^{|C|}$ that maps the set of submitted ballots $\mathbf{b} = \{b_1, b_2, \dots, b_n\}$ to a non-empty set of winners W . The winner set W is the outcome of the voting game, and each voter i derives utility $u_i(W)$ based on their utility function, which is typically private information known only to the voter herself.

²We may further specify that we operate within the field of computational social choice, where we apply techniques and concepts from computational theory and multiagent systems to the social choice domain.

While utility functions may be defined arbitrarily over the outcomes, throughout the thesis, we use a specific form of single peaked preferences. Each candidate c_j advocates a position $p(c_j)$ drawn uniformly and independently at random from some domain \mathcal{D} . Similarly, each voter i has a type p_i drawn uniformly and independently at random from \mathcal{D} . The utility i derives from the election of candidate c_j is defined as the following function on the difference between the voter’s preference and the winner’s position: $u_i(c_j) = u_i(p_i, p(c_j)) = -|p_i - p(c_j)|^2$. This particular quadratic formulation is frequently used by social choice researchers, such as Myerson and Weber [100].

For brevity, we write u_i to imply $u_i(p_i, \hat{p})$ where the position of the candidate c_i and the position favored by the agent is clear from the context. Throughout this thesis, we will refer to the social welfare of the elected outcome. If \hat{p} is the position of the elected candidate, the social welfare $SW(V)$ is the sum of the utilities for all voters for that outcome:

$$SW(V) = \sum_i u_i(p_i, \hat{p}).$$

While voting systems vary greatly, we consider only *resolute* social choice functions, which always chooses a single winner $w \in C$. Any non-resolute social choice function may be made resolute by imposing a tie-breaking scheme, typically selecting a winner from the set uniformly at random, or lexicographically. While some of our models extend to other voting systems, we will focus on the *plurality* voting rule and its variations, both because of its widespread use and for the sake of tractability. In the plurality voting system, also known as first-past-the-post, each voter is invited to declare their support for one candidate, and the candidate with the most support is the winner. In our terminology, the set of admissible ballots is the set of candidates ($\mathcal{B} = C$), and the resolute plurality social choice function $\mathcal{F} = \arg \max_c |\{b_i : b_i = c\}|$, breaking ties either lexicographically or randomly.

Under lexicographic tie-breaking, we fix some strict linear order on the candidates. Without loss of generality, we may assume the order to be (c_1, c_2, \dots, c_n) . If multiple candidates $c_{i_1}, c_{i_2}, \dots, c_{i_k}$ are co-winners (i.e. tying each other for having the most number of votes), then the winner returned by \mathcal{F} is $c_{\min(i_1, i_2, \dots, i_k)}$. Under random tie breaking, \mathcal{F} returns c_{i_z} with probability $\frac{1}{k}$ instead.

2.3 Graph Theory

A graph is a mathematical way of describing relationships between objects. A graph $G = (V, E)$ is comprised of a set of n nodes (or equivalently, vertices) $V = \{1, 2, \dots, n\}$, and a set of edges E . We denote the existence of a relationship between distinct nodes $i, j \in V$ by the edge (i, j) ; the set E contains all such relevant edges. In an undirected graph, the edges are unordered pairs (i.e. $(i, j) \in E$ implies $(j, i) \in E$); in a directed graph, the directed edges are ordered ((i, j) and (j, i) are distinct entities). Unless stated otherwise, we forbid self edges. We use the terms graph and network interchangeably.

In an undirected graph, i and j are **neighbors** if they share an edge $(i, j) \in E$; we may refer to all neighbors of i as $N_i = \{x \in V : (i, x) \in E\}$. The number of neighbors of a node i denotes its **degree**, $d_i = |N_i|$. Analogously, in a directed graph, i has out-neighbors $\overleftarrow{N}_i = \{x \in V : (i, x) \in E\}$ and in-neighbors $\overrightarrow{N}_i = \{x \in V : (x, i) \in E\}$, which differ via the orientation of edges. Note that these two sets need not be disjoint. We define out-degree d_{out} and in-degree d_{in} analogously. We adopt the convention that the direction of edges mark the *direction of admiration*; that is, the out-neighbors of i are exactly those nodes that have influence over i . In sections utilizing both directed and undirected graphs, we simplify notation by defining N_i as exactly those nodes who have influence over i , and therefore refer to the neighbors in an undirected graph, and the out-neighbors in a directed graph.

Social networks describe relationships between individuals within a community, and can be represented as graphs. The nodes V represent entities in the social network. This includes people, and may also include more abstract entities such as organization and media outlets. An edge $(i, j) \in E$ denotes a relationship between these two entities that allows for an exchange of information. This edge may be directed (implying a hierarchical, one-directional flow of information) or undirected. The choice of notation of V for both the set of vertices, and the set of voters is purposeful, as they represent the same entities in scenarios where we consider voters operating within a social network.

2.4 Social Networks

Social networks are mathematical descriptions of how individuals interact in a community. They capture how information from factual sources and subjective evaluations are transmitted between people, and what social structures are present in different communities. As the Internet matures as a technology, more and more information about these social

networks are captured as “big data”. At the same time, these online social structures wield ever increasing influence over our lives at all scales — from the minutiae of our day-to-day moods [79], to turnout at congressional elections [25]. It is therefore of paramount importance that we understand the mechanisms by which social networks affect decision making.

2.4.1 Opinion Dynamics

One major application of social networks is in modeling how individuals influence each others’ opinions. This field has been variously called opinion dynamics, diffusion of innovations, and information cascades in networks, depending on the emphasis of the particular mechanisms. In general, the field of opinion dynamics captures the idea that humans are fundamentally social creatures that seek to exchange information and influence each others’ opinions through repeated interaction. A simple model of this process may represent opinions by a simple value x as a real number from the interval $[0, 100]$. Each agent i has such an opinion x_i with some initial value. An interaction between i and j will cause their respective opinions to shift closer together, $x_i \leftarrow f(x_i, x_j)$ for some function f . We defer the discussion of related works in opinion dynamics to Chapter 4.

2.4.2 Knowledge Graphs

Many voting models that we will examine in Section 3.1 assume that the submitted ballots are public knowledge. However, this is an unrealistic assumption for large elections. Instead, we may consider that voters may only observe the actions of their neighbors in their social network. This idea was first formalized by Pacuit, and Parikh as the Knowledge Graph model [104], and applied to computational social choice in a follow-up paper by Chopra, Pacuit and Parikh [32]. In the latter paper, the authors propose a general framework for limiting voting knowledge, restricting each voters’ observations to their neighbors in the knowledge graph. However, they do not define any response behavior for individual voters, nor explore the aggregate behavior of the population.

2.5 Random Graph Models

Random graph models refer to probability distributions over related graphs with interesting or useful mathematical properties. Many random graph models have efficient means by

which we may sample from them. Some of these sampling methods are merely descriptive of the final product, while others are generative and shed clues on how the network may be formed. These graph models allow us to generate graphs that mimic real world networks in important ways, on demand, and with control over parameters such as network size (i.e. the number of nodes n) and edge density (i.e. the average degree of nodes).

Real world social networks also exhibit a number of important structural characteristics. Two that we will focus on are the *small-world* and *scale-free* properties.

In small-world networks, the average distance between any two vertices in the graph grows as a logarithm of the number of vertices. We expect information to travel quickly through small-world networks, which may have an effect on the aggregate strategic behavior of the population.

Real world networks are often scale-free, which means they are comprised of a handful of highly-connected hubs and many sparsely connected vertices. Highly-connected hubs may represent popular public figures or mass media outlets. They may wield considerable influence within the network. Scale-free graphs are so named because plotting the number of vertices of a particular degree on log-log scale results in a linear trend. The connectivity between highly popular, hub vertices follows a similar “scale-free” pattern as less popular vertices.

2.5.1 Erdős-Rényi

The Erdős-Rényi (ER) random graph [56] is a standard graph model used in graph theory literature. It is simple to construct and incorporates minimal assumptions. A (directed) Erdős-Rényi random graph with connectivity probability p is constructed by considering every admissible (directed) edge (i, j) , and adding it to E with fixed probability p . An ER graph exhibits the small-world property, but is not scale-free in general. Note that the resulting graph may be disconnected. If a connected graph is required, we discard the result and regenerate the graph.

2.5.2 Barabási-Albert

A Barabási-Albert (BA) random graph [9] with attachment parameter d is constructed by iteratively adding vertices, and connecting them to d existing vertices with probability proportional to their respective degrees. New vertices are preferentially drawn toward popular nodes in the network in a process of *preferential attachment*. The result degree distribution

follows a “rich get richer” scheme, resulting in a scale-free distribution of degrees. A BA network also exhibits the small-world property as well.

The original BA model assumes an undirected graph. A number of models extend the notion of preferential attachment to directed graphs. We use a simplistic extension where edges from newly added nodes are oriented toward existing nodes in the network, producing a strongly hierarchical network. Note that the attachment parameters for the directed BA network must be doubled to preserve the same average out-degree as the undirected network.

Figure 2.2 shows an undirected example from each random graph model. Both graphs have 40 vertices and are parameterized so that each node has average degree 3.

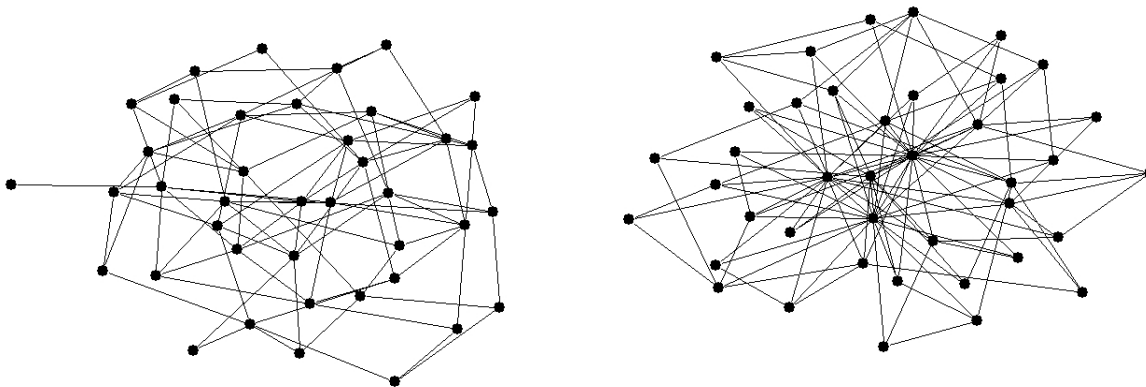


Figure 2.2: Example of an ER random graph (left) and a BA random graph (right).

2.5.3 Homophily

Another property exhibited by real world social networks is *homophily*: the tendency for people to connect and socialize with those sharing similar characteristics, beliefs and values. This concept dates as far back as Plato, who wrote in *Phaedrus* that “similarity begets friendship”. In their seminal work, McPherson, Smith-Lovin and Cook offer a survey of evidence that adults, in particular, preferentially associate with those of similar political persuasions [90]. This effect is not only limited to individuals. Hargittai, Gallo and Kane examined the link relationships between sites of top conservative and liberal bloggers discussing political issues, and found homophily to be prevalent; i.e. sites were much more likely to discuss and reference each other when they shared political views. Even more importantly, upon examining the context of links between conservative and liberal

blogs, they found that fully half of them were embedded with “straw-man” arguments that reinforced the political position of the author by distorting the opposition’s position [67]. This is especially relevant to voting because voters derive information about the election from their neighbors in the social network. A homophily of opinions can lead to the so called *Echo Chamber Effect*, where a voter is surrounded by associates that share similar beliefs, reinforcing its validity regardless of its merit. This is discussed in more detail in Chapter 5.

We modify both the ER and BA graph models to incorporate homophily. Given a homophily factor $h(i, j)$ between two nodes i and j , we adjust the likelihood of a connection between i and j by a multiplicative factor of $h(i, j)$. Formally, in an ER graph with parameter p , the probability that a (directed or undirected) edge (i, j) exists is $h(i, j)p$. In a BA graph, the probability that a new vertex i connects to an existing vertex j with degree d_j is $d_j h(i, j)/Z$, where Z is a normalizing constant. We denote the modified graphs the hER and hBA models.

Our modified ER graph model with homophily is an example of *Latent Space Models*, where individual nodes have real valued attributes and the probability of an edge between two nodes is a function of the attributes of those nodes [72]. A more recent survey on general spatial networks is available at [14]. A similar model exists in the Social Distance Graphs [22], where connection probabilities are inversely proportional to a measure of social distance (dissimilarity) between two agents.

Chapter 3

Related Work

In this chapter, we will survey the existing literature, from the multi-agent systems and wider social choice communities, on models on voting mechanisms and voter behavior, as well as the effects of social networks on information propagation and decision making.

3.1 Voting Models

In this section, we examine several models of voting. These papers examine the process of voting from a game theoretic perspective, with an emphasis on equilibrium solution concepts, winner determination, and computability. Unless stated otherwise, we presume that the plurality voting system is used to determine the winner. Indeed, the use of plurality voting scheme is integral to many of the papers mentioned in this section.

3.1.1 Voting Equilibria

Nash Equilibrium is a central concept in game theory. At Nash Equilibrium, each player is playing a strategy such that no single player will profit from deviating from that strategy. While useful in many settings, it breaks down in large games involving many players, such as voting games: any outcome where the winner exceeds the runner-up by more than 1 ballot is a Nash Equilibrium. In their seminal work, Myerson and Weber [100] extend Nash Equilibrium to voting games, creating **voting equilibria**. The concept is motivated by the importance of electoral polling prior to the voting process, and how it impacts the voting behavior of the population. Electoral polls may be used to determine the set of “viable

candidates” that have a reasonable chance of winning the election, allowing strategic voters to possibly abandon their favorite candidates in favor of a more promising second-choice.

In their model, the result of the poll is modeled as a set of *pivot-probabilities* $\mathbf{p} = p_{i,j}$ between each pair of candidates i, j . This is the probability that a single ballot supporting i will change the winner of the election from j to i . While technically a possible outcome, three-way ties are considered vastly less likely and assigned a probability of zero to make the problem tractable. Each voter has a type corresponding to a vector containing the utilities they derive for each candidate being elected; ballots (which are positional scoring) are cast strategically to maximize their expected utility given the probability of casting a pivotal vote.

We say pivot-probabilities \mathbf{p} *justifies* an electoral result μ (itself a probability distribution over candidates) if μ is in the set of possible responses of voters reacting strategically based on \mathbf{p} . We say \mathbf{p} obeys the *Ordering Condition* for μ if candidate i has a higher expected score than j in μ , then $p_{i,h} \geq p_{j,h}$ for any third candidate h .

Then an election result μ is an **Voting Equilibrium** if there exists pivot probabilities \mathbf{p} s.t. \mathbf{p} justifies μ , and obeys the Ordering Condition of μ for all $\epsilon > 0$.

Myerson and Weber establish that all election scenarios have a non-empty set of voting equilibria. More importantly, if μ is a voting equilibrium, then there exists pivot-probabilities \mathbf{q} justifying μ s.t. $q_{i,j} > 0$ only for candidates that have the top two maximal expected scores. That is, no voter will vote for a third place candidate.

While the Voting Equilibrium is a powerful solution concept for capturing plausible outcomes of an election, it gives no indication of how the voters arrive at that outcome. Throughout this thesis, we will explore various models of voter behaviors that arrive at similar solution concepts.

3.1.2 Poisson Games

Myerson also propose an alternative method for modeling elections, in the form of Large Poisson Games, first proposed in [98] and further expanded on in [99]. The Poisson Game models the voting process by assuming that the number of voters is *uncertain*, following a Poisson distribution. This allows us to decouple the distribution of voters supporting one candidate from those supporting another (called the *Independent Actions* property). The paper also proposes a mechanism for estimating the probability of pivot conditions between candidates in a 2-candidate election.

In a Poisson Game, the number of voters is given by a Poisson Distribution, parameterized by the expected number of players n . The probability that there are k voters is given by $P(k|n) = e^{-n}n^k/k!$. Each voter has a type drawn independently from T according to some distribution r . Each voter casts a ballot from the set of admissible ballots C , receiving a payoff according to utility function U based on the outcome of the vote and the voter's type. A Poisson Game (which can be generalized beyond the voting setting) is fully parametrized as the tuple (T, n, r, C, U) .

Let each voter act according to a strategy function σ , where $\sigma(c|S)$ is the probability that a voter having a type in $S \subseteq T$ will play action $c \in C$. Since each voter acts independently, the number of voters playing any action $c \in C$ is itself a Poisson random variable, with mean $n\tau(c)$, where $\tau(c)$ is the probability any given voter in the population plays action c . This is the aforementioned *Independent Actions* property.

Another interesting property exhibited by the Poisson Game is *Environmental Equivalence*. Given a Poisson Game with players distributed according to $Poisson(n)$, from the perspective of a player *within* the game, the number of *other* players in the game is also distributed according to $Poisson(n)$ (i.e. not $n - 1$). This is because the fact that the player was chosen to be part of the game in the first place, suggests that the game has many players. This exactly cancels out the effect that the number of other players must be 1 smaller to account for the perspective player.

These tools allow us to calculate the exact tie probability between two candidates. For instance, in a plurality election with two candidates 1 and 2, where a tie is broken randomly, the probability that casting a ballot for candidate $c \in \{1, 2\}$ will be pivotal (cause c to win, where c was not winning) is given by

$$v(c|n\tau) \approx \frac{e^{n(2\sqrt{\tau(1)\tau(2)} - \tau(1) - \tau(2))}}{4\sqrt{\pi n\sqrt{\tau(1)\tau(2)}}} \frac{\sqrt{\tau(1)} + \sqrt{\tau(2)}}{\sqrt{\tau(c)}}$$

where $\tau(1)$ is the probability a given voter in the population will support 1 (and likewise for $\tau(2)$), and n is the expected number of total voters. The \approx symbol denotes that the function v converges to the correct probability as $n \rightarrow \infty$.

In Chapter 6, we extend the solution of the Poisson Game to a multi-candidate election to obtain a voter heuristic that is both fast and accurate.

3.1.3 Local Dominance

Meir, Lev and Rosenschein [94] expand on the concept of Voting Equilibria pioneered by Myerson and Weber. They construct a framework for examining equilibria in various auction types as well as produce a behavioral voter model that has certain desirable properties (called the “Desiderata for Voting Models”); most notable amongst these criteria, they require their model to be predictive (able to predict a small but nonempty set of possible winners) and realistic (computationally feasible for human voters and accounting for limited knowledge; most importantly, they should not be required to compute exact probability).

In Meir’s model, voters’ observation of the electoral situation is limited to the retrieval of a set of possible winners. Each voter then iteratively revises their votes based on their private preferences. The system is at a Voting Equilibrium if no revision occurs.

The authors analyze the system using a strategic voter model where the set of possible winners are exactly those candidates that are within r votes of the current winner. If this set is empty, they default to voting truthfully. They find that when vote revision is done in a sequential manner (i.e. with an external scheduler) and starts from the truthful state, it always converges. They conjecture that it will also converge if starting from an arbitrary state. However, if voting is simultaneous, then cycles can occur.

The paper also presents a simulation to analyze the behavior of the electorate under varying conditions. The most important tuning parameter was r . They find there is a “peaked” value for r (not too high and not too low) where voters are most strategic (i.e. reporting anything other than their favorite option), and social welfare is maximized. When they repeat the simulation with heterogenous r values, they find this effect is less pronounced.

In Chapter 6, we extend Meir’s Desiderata to the social network domain, and consider other plausible and tractable heuristics for voter behavior.

3.1.4 Sequential Voting

Sequential Voting occurs in T rounds, where each voter v_i has been assigned to vote in some round t_i *a-priori*. Voting within each round occurs simultaneously amongst voters assigned that round, and each vote has full knowledge of the history of voting in prior rounds. Dekel and Piccione [44] are the first to compare the results of Sequential Voting with that of a simultaneous voting. They consider a setting with $m = 2$ alternatives, where each voter has a private, noisy signal on the utility she would gain from each alternative. They show

that any symmetric equilibrium using informational strategies in the simultaneous game is also an equilibrium in any sequential game, and therefore, information cascades in the sequential game do not affect its ability to aggregate information. However, it does leave open the question of whether the sequential game admits uniquely sequential equilibria.

Battaglini [16] further extend this work by making voting costly and allowing one to avoid this cost by abstaining. He shows that even a small cost of voting (i.e. lost productivity and time) substantially changes the equilibria in the sequential and simultaneous games, and that they are generally disjoint. When voters are allowed to abstain, simultaneous voting is uniformly more informative than sequential voting. Moreover, Battaglini, Morton and Palfrey [17] establish in both theoretical models and in laboratory experiments that early voters bear a larger cost when they choose to contribute to the information aggregation process.

Desmedt and Elkind [46] explore strategic behavior in Sequential Voting with abstention with multiple candidates. They show how the subgame perfect Nash equilibrium may be computed, and that when there are more than 3 candidates, the equilibrium behavior of voters is complex and sometimes counterintuitive. The outcome of the election is sensitive to the risk adversity of the voters, and the voter order.

3.1.5 Sequential Voting and Herding

Callander [29] proposed a voting mechanism for exploring the so-called “bandwagon” or “momentum” effects highly publicized in American primary elections. This mechanism is founded on the behavioral notion that an agent becomes invested once they commit to a particular option, and therefore derive utility for having voted for the winner (in addition to the objective utility associated with the winning candidate). Conceptually, it may be easier to frame this model via social network phenomena such as “Liking” on Facebook or “Retweeting” on Twitter, which presents a publicly visible vote as a public show of support for an idea. The model is also applicable in any setting where voting occurs sequentially and your votes become public knowledge. For example, when electing a new department head with public voting, it is not only profitable to elect your preferred candidate, but it would also be ideal if you had voted for the eventual winner.

Callander’s model is based on Sequential Voting with two alternatives A and B , where the world exists in one of two states corresponding to whether A or B is the better alternative. Voters all have the same preference for electing the correct alternative, but also derive additional utility for *having voted for the winner*. Each voter has a private, independent, noisy signal on which candidate is better. There are a countably infinite number of voters,

and the winning candidate is one that maintains a lead in the limit (or a tie, if there is no limit). Voters arrive one-by-one to cast their ballots, knowing the current standings from prior ballots. Based on this standing, voters may cast an informative ballot according to their signal, or cast an uninformative ballot, following the current leader. It is this latter case that leads to the “bandwagon” effect.

Callander observes that the bandwagon phase is inevitable as the size of the population grows large, and can form a Perfect Bayesian Equilibrium. Conversely, there are no PBE where all voters only vote informatively. Surprisingly, sometimes a voter will strategically vote against her preferred candidate, even if he is currently winning, in order to elicit more information from other voters (termed “Buyer’s Remorse”). Callander also examines the amount of informativeness of voting under these conditions. Each election proceeds from an informative phase where votes help aggregate the private opinions of voters, before entering an uninformative cascade phase. Surprisingly, even under strong preference for conformity, and a population size that converges to infinite, some degree of information can be extracted.

Alon et al. [10] examine an alternative formulation of the sequential voting model, where preferences are subjective and private. Specifically, each voter has, a-priori, a preference for either A or B (with equal probability). Each voter derives maximum utility from voting truthfully and winning, but failing that, the agent would prefer voting strategically (opposite of their preference) and having that vote match the winning candidate (whom they do not prefer).

As before, there are a countably infinite number of voters. In this version of the election, the winner is determined by establishing a leading gap of size $M > 0$. Each voter will follow a *threshold strategy*, where they will vote truthfully unless the opposing candidate has established a lead of size at least $0 < r \leq M$. They establish that the unique symmetric subgame perfect Nash equilibrium is a threshold strategy where r depends only on the agent’s prescribed utilities, and surprisingly, *does not depend on M* .

Finally, Gaspers, Naroditskiy, Narodytka, and Walsh [63] examine the possible and necessary winner problem in Sequential Voting (which they term “social polls”) when conducted in a social network setting. Based on examining a strict subset of ballots, a candidate is a necessary winner if no arrangement of remaining ballots will prevent the candidate from winning; a candidate is a possible winner if there exists an arrangement of remaining ballots that allow the candidate to win. They find that the possible winner problem is NP-hard to compute, but propose an efficient algorithm for finding necessary winners.

3.1.6 Stackelberg Voting

Xia and Conitzer [134] examine a variation of Sequential Voting called Stackelberg Voting, where a *finite* number of voters play the game with subjective but public preferences. It is clear that the extensive form game can be solved using backward induction, but the game space is infeasible. The paper presents a dynamic programming algorithm, utilizing *compilation functions* that solves this program in polynomial time.

They also analyze the strategic behavior of such agents, as they are able to fully determine the future actions of voters following them in the voting sequence. Their results are based on the **dominance index** $DI_r(n)$ of various voting rules — the minimum number of voters beyond half the population (i.e. $\lfloor (n/2) \rfloor$) needed to ensure a candidate’s victory. The authors find that, for voting rules with a low dominance index, the backward-induction outcome for strategic voters can result in a highly suboptimal alternative being selected. In particular, for $n - 2DI_r(n)$ voters, the winning alternative will be ranked either last or second-to-last place. Simulation results, however, show that more voters prefer the winner obtained from backward-induction, over the winner obtained by a truthful (nonstrategic) voting.

3.1.7 Iterative Voting

A more recent line of inquiry inspired by Myerson and Weber is Iterative Voting [95]. Similar to Sequential Voting, Iterative Voting proceeds in rounds, with the crucial difference that voters may choose to revise their ballots in subsequent rounds. Formally, voting begins with some initial configuration of ballots from the voters (for example, their truthful ballots). Voters have complete information on the current ballots. In each round, voters respond to this interim outcome by revising their ballots if it will yield a better outcome. This may occur simultaneously or one at a time. Iterative Voting is guaranteed to converge from a truthful state under plurality and veto [84, 95]. Branzei et al. have also investigated average utility of voters, using iterative voting under different voting rules (Plurality, Veto and Borda) [27]. They define the Dynamic Price of Anarchy (DPoA) to be the worst case ratio between the social welfare of the winner elected under truthful voting versus strategic voting. This is similar to our definition for Price of Honesty used in Chapter 5. Since their model does not operate under a social network, they are able to compute analytical bounds for DPoA under different voting rules. Similar to us, they show that strategic voting improves the elected outcome under Plurality. Other work incorporates voters who are truth biased (who prefer voting sincerely if they cannot otherwise affect

the outcome), lazy (who prefer abstaining, all else being equal) [111], or optimistic (who assume some fixed number of voters may be swayed to their cause) [102].

The voter models presented in Chapters 5 and 6 generalize Iterative Voting to a domain of incomplete information. While Iterative Voting assumes voters have complete information on the current state of the election, we assume that the social network restricts the visibility of information to only a voter’s social network neighbors. Iterative Voting, then, is a special case within our more general framework, where the voters are embedded in a complete network.

3.1.8 Modal Logic Voting Models

An alternative approach to modeling voters is the epistemic approach first formalized by Chopra, Pacuit and Parikh [32]. They highlight that in order for Gibbard-Satterthwaite’s results to apply — that any useful voting system must allow for strategic voting — the strategic voter must possess appropriate knowledge of the preferences of other voters. A voter with limited knowledge of the actions of other voters has little ability to formulate a strategic ballot. To formalize the notion of the partial knowledge available to voters, they apply techniques from modal logic. They define a logical language of Boolean formulae that captures voter preferences, and their knowledge (or lack of knowledge) of each others’ preferences. They show how the formulae may be updated as voters change their ballots. These dynamics result in a voting process that proceeds in rounds. The authors compare earlier rounds to pre-election “opinion polls”, which help inform voters’ final decisions in the last round.

Van Ditmarsch, Lang and Saffidine [131] extend this work by specifying a plausible voter responses to different states of partial information. They introduce knowledge profiles and information sets to the model. A knowledge profile corresponds to a particular configuration of ballots and epistemologies (i.e. voters’ knowledge of each others’ ballots). For each voter, the model defines a number of information sets, each of which contain a non-zero number of knowledge profiles. At any point in time, each voter knows which information set they are in, but cannot distinguish between the knowledge profiles within the set. The authors formalize several types of manipulations, and posit that voters manipulate based on a “pessimistic” (risk-averse) heuristic. Each voter associates each information set with the worst possible outcome from the knowledge profiles within that set, and casts her ballot to best mitigate this worst case outcome¹. This epistemic voting game is at an

¹We use the game theoretic term “information set” here for brevity. In the paper, the authors define states corresponding to configurations of ballots and epistemology in the voting game, and an “indistin-

equilibrium if no voter wishes to revise her ballot.

3.2 Voter Behavior

In this section, we survey literature that focuses on the behavior of individual voters, including observations from laboratory experiments or models designed to mimic human behaviors.

3.2.1 Pólya’s Urn Model

Perhaps one of the earliest attempts to model a voting process is Pólya’s Urn Process by Eggenberger and Pólya [108].² Pólya’s Urns is a sampling process that may be interpreted as the opposite of sampling without replacement. The process begins with an urn containing a finite number of balls, each painted with one of $m \geq 2$ colors. At each step, a ball is drawn from the urn independently and uniformly at random; the ball is returned to the urn *with a duplicate ball* of the same color. The resulting distribution of colored balls in the urn forms a Dirichlet-multinomial distribution, or a multivariate Pólya distribution.

While Pólya’s Urn Model has been extended in a number of interesting ways, the basic urn model can be interpreted as a Bayesian model for voter behavior: The current distribution of urns represents both the current tally of votes, and a distribution of future voter behavior. Moreover, a Bayesian update is performed after each ballot to revise future voter behaviors.

3.2.2 Approval Voting in Doodle Polls

Zou, Meir and Parkes [135] examine over 340,000 polls from the popular social polling website Doodle (restricted to the U.S. to avoid cultural confounds). Doodle allows users to conduct a quick Approval election for selecting time slots for an activity. The polls can be constructed to be open (users are expected to enter their names with their publicly revealed ballots) or closed (anonymous voting, with ballots visible only to the organizer). The paper examines the data for common trends and differences between open and closed polls, finding the following tendencies:

guishability” equivalence relations between states for each voter.

² The original article [108] is written in German. We use the definition provided at [6].

- Open polls show higher average availability.
- Both polls have monotonically increasing correlation with previous responses. Open polls have higher correlation. Correlation in closed polls may be from intrinsic popularity of some slots.
- In open polls, voters will approve a popular slot more often than in closed polls.
- In open polls, voters will approve an unpopular slot more often than in closed polls.
- In open polls, a person who approves a highly popular time slot is more likely to approve an unpopular time slot; but is also unlikely to approve a time slot of intermediate popularity.

The paper constructs a model in which voters cast their ballots in a fixed order and, in the open poll system, have a behavioral response function that examines the current state of the poll to produce a sensible ballot to maximize expected utility³. After considering several plausible behavioral functions (Random Cutoff, Random Cutoff restricted to Popular slots, and the Leader Rule), the authors propose Social Voting which attempts to also increase the number of slots approved (so the user does not appear to be “too picky”) while avoiding selecting lesser candidates; this is done by approving a number of slots of lower utility but are also Unpopular (and thus will be unlikely to win). Simulations show that Social Voting produces similar results as those observed in open polls.

Obraztsova, Polukarov, Rabinovich and Elkind [103] propose the Doodle Poll Game capturing this behavior, where users derive additional utility from appearing to be available. Reinecke et al. [113] have also examined how Doodle voting behavior may be affected by national culture and social norms. Our work on strategic timing in Chapter 7 establishes an alternative framework for examining open polls. Crucially, we are the first to consider the timing of participation to be an important element of participant strategy.

3.2.3 Observing Human Strategic Voting Behaviors

Eckel and Holt [52] examine human strategic voting behavior in an agenda-controlled committee sequential voting environment. Their experiment occurs in successive rounds.

³It is interesting to note they do not consider the question of when users choose to vote. In an open poll, it is easy for a user to visit the survey, inspect the results and decide to return at a later time of their choosing.

In each round, the voters cast a ballot in favor of two strict subsets (“options”) of remaining alternatives until only one alternative remains. The voters receive a predetermined and randomized monetary payoff, which establishes their (privately known) preference over the alternatives. A neutral third party “sets the agenda” by determining which subsets are to be voted over ahead of time. In their experiment, they consider only three alternatives, and each round eliminates one alternative from the running.

The authors consider three types of voting behavior from Plott and Levine [107]: *Sincere Voting* where the voter selects the option containing her most preferred item over one that doesn’t (and working her way down her preference list); *Avoid Worst Voting* where the voter selects the option omitting her least preferred item over one that contains it (and working her way up her preference list); and *Average Value Voting* where the voter treats all items in the options as being equally likely (and is risk neutral).

The exercise was repeated with the same group to examine longitudinal learning effects (with preferences being reshuffled after each instance of strategization⁴). This exercise series was repeated across several groups. Over the course of the experiments, certain patterns emerged. There was evidence of strategic voting in the first meeting, but significant learning effects (i.e. strategization) were present through the course of each series of meetings. Strategization was most frequent when preferences were not changed, which allowed the voters to learn each others’ preferences. Finally, once an individual behaves strategically, they are very likely to continue to do so.

Meffert and Gschwend conducted several experiments to investigate how strategic voting occurs in human participants [92, 93], focusing on coalition formation in governments with proportional representation, which they view as a more realistic scenario for strategic voting to take place. In these elections, parties win a number of seats proportional to the number of votes they receive, though often only when they have at least some threshold percentage of supporters (typically 5%). Frequently, these elections do not produce a single party with majority seats, and so the dominant party must form a coalition with a smaller party. This complicates the decision process of a voter, because they must decide between supporting the dominant, or helping elevate their favored small party to the 5% threshold.

Meffert and Gschwend initially begin by conducting economically motivated experiments [93], where they constructed fictitious election scenarios where participants played the role of a strategic voter. Each fictional party held some position in a policy space, and voters were rewarded with monetary incentives to vote in a way that elected a government closest to a given position in the policy space. A coalitional government held a position

⁴The experimenters determined that once the group adopted a strategy, they would continue to use that strategy unless the profiles were changed.

that is the average of their members (weighted by vote count). The experimenter manipulate pre-election polls and other election information to set experimental conditions. They found that voters exercised certain heuristics when determining who to vote for:

1. They avoided parties distant to themselves in the policy space.
2. They avoided isolated parties (who were unlikely to be recruited by a dominant party anyways).
3. They avoided small parties, who were unlikely to pass the threshold or contribute enough votes to a dominant party to majority.

They also investigated the role of coalition signaling, a political strategy where a party explicitly announces a proposed coalition (or explicitly rejects the possibility of one), and found that to be a significant component of voter heuristics.

The authors follow the above study with a second investigation [92]. They claim that a monetarily motivated scenario is unrealistic, and its conclusions may not generalize to real elections. In this follow-up experiment, they manipulate election information embedded in the context of two actual elections in adjacent states in Germany. Participants were asked about their party preferences and sorted into one of the states that best reflected those preferences. This allowed for manipulations that would appear “realistic” to the participants. In this study, the authors found that both political sophistication (prior knowledge of politics) and time spent examining poll information to be the main factors that led to effective strategic voting decisions; however, the two factors are substitutive – a political non-savvy person could “make up for it” by spending more time studying the polls. The authors also identify a significant number of voters who voted insincerely, but non-strategically (i.e. did not utilize polling information). The authors claim these voters “reacted passively” to signals advertized by the various parties without verifying their credibility. Only 5% of all participants identified themselves as having made strategic considerations while voting.

Poncela-Casasnovas et al. [109] investigate the modes of human behavior when they are engaged in games in a controlled “within subject” experimental setting. That is, they are concerned with how the same subject treat different variations of 2-player cooperate/defect games. They focus on four types of games: the classic conflict games of Prisoner’s Dilemma, Stag Hunt, and Snowdrift Game; in addition to the Harmony Game where both players receive the maximum award if they both coordinate.

They treat the experimental data without assumptions on player behavioral types, and perform unsupervised clustering on player strategies based on the frequency of coordination in each of the game types. They find that players can be classified with high fidelity into only one of 5 types. **Optimistic** players seek to maximize their own payoff, regardless of whether their opponents will cooperate. **Pessimistic** players maximize their minimal payoff, displaying risk-averse behavior. **Envious** players seek to maximize the payoff difference between them and their opponents. **Trusting** players always play Cooperate. And finally, 12% of players are classified as **Undefined** who appear to play the game randomly.

Tal, Meir and Gal [124] study online human voting behavior in response to poll information. They conduct experiments on Amazon Mechanical Turk where participants are given preferences (in the form of small monetary rewards) for playing in a plurality voting game. The game may be one-shot, where poll information is fictitious; or it may be a game of Iterative Voting with other participants. Aside from a small number of erratic voters (who act randomly), most voters exercise either the “default” option (a truthful ballot in the one-shot game, or maintaining the same ballot in an iterated game), or utilized a myopic best response.

Reijngoud and Endriss have also modeled how voters might respond to information from a series of polls [112]. In their paper, they define poll information functions (PIF) for summarizing the information present in the current ballot profile. For instance, the PIF may only report the current winner, or may report the current score for each candidate in an election using position scoring rule. A voter may then use this information to alter their ballot to their benefit. In their model, voters have ordinal preferences (i.e a strict ordering over candidates) and will only change their ballots when there is a guarantee that they will not be worse off. They analyze the susceptibility and immunity to manipulation of different voting rules and PIFs, with expanded results in a follow-up paper by Endriss, Obraztsova, Polukarov and Rosenschein [55]. In general, the two papers show that many voting rules (including Plurality, k -Approval with $k < m - 1$, Borda, Copeland and Maximin) are susceptible to manipulation with the PIF reveals the score vector or even only the winner. Notably, Veto is immune to manipulation when only the winner is revealed and susceptible when the full score vector is revealed. They also establish that Borda, Copeland and Maximin are susceptible when the PIF reveals the graph encoding pairwise majority winners.

3.2.4 Strategic Effects of Early and Late Ballots

Dekel and Piccione[45] explore the information asymmetry that exists when voters cast their ballots earlier or later in a sequential voting process. They consider plurality with

$N \geq 4$ voters and $m = 3$ candidates. Voter preferences are randomized and private. Voters must choose to cast their ballot in one of two periods; this choice is chosen prior to the election, and indeed *prior to the voters' own preferences are realized*. The choice of time periods is made simultaneously and then made public. Voters who chose to go in period one cast their ballots first, simultaneously; this interim result is revealed to all remaining voters, who cast their ballots simultaneously.

The authors posit, as we do in Chapter 7, that voter timing is influenced by two motivations: early voters establish the lead runners of the race, while later voters are less likely to “waste” their ballots on hopeless candidates. Under their model (and certain assumptions on voter strategies), Dekel and Piccione show that the latter effect by far dominates the former, and in most cases, all voters will choose to go in the second period, making the sequential voting outcome (equilibrium) equivalent to the simultaneous outcome (equilibrium).

To formalize their model, let $\omega = \{1, 2\}^N$ be the set of all possible timing choices made by the N voters, $H^2(\omega)$ be the set of all possible ballots collected in period 1 (i.e. all possible histories when arriving in period 2), and $H^1(\omega)$ be the empty history. Then, a mixed strategy $s^i(\omega)$ for player i maps a history $H^{t_i}(\omega)$ and the type (preferences) of i to a distribution over the candidates (probabilities for casting different plurality ballots).

A strategy s^i is *persistent* if (1) i votes for her most preferred candidate in period 1 and (2) i votes for her most preferred candidate *who can still possibly win* in period 2. Assuming voter strategies are symmetric across voters and candidates, the authors show that in a game where all voters use persistent strategies, and $N \geq 6$, the sequential voting equilibrium has all voters voting in period 2, and thus is equivalent to the simultaneous equilibrium.

In order for the sequential equilibrium to be different, voters must not be persistent. This may happen if their utility for electing their second favorite candidate is closer to their top choice. If this is true, then the voters show that there cannot exist a pure strategy equilibrium where all voters vote in the same period.

Dekel and Piccione’s model differs from our alternative model of vote timing, proposed in Chapter 7, in several important ways. The primary difference is when agents choose their timing strategy. In their model, agents must choose their voting period *a-priori*, in fact, before they even know their own preferences. We view this assumption as unrealistic. Even in structured environments such as the U.S. primaries, individual states learn enough of their preferences over time to argue for shifting the timings of primaries. We propose that the decisions of when to cast and which ballot to cast are the same decision, made in each period in a Markovian process (i.e. each history maps to a decision to either wait

or to commit a particular ballot). Note that even in a 2-period scenario, the two models are *not* identical. While the decision of whether to wait or to commit for the first period is made in the absence of any ballot information, in our model, agents already know their type. A more detailed discussion of our results and how they contrast with those obtained in several of the aforementioned works is reserved for Chapter 7.

Morton and Williams [97] examine informational effects of elections that happen in stages where some subset of voters must cast their ballots first, and subsequent voter may make use of the observed voting patterns in their own assessment of likely electoral outcomes. Their experiments are inspired by the United States primaries, which are a series of elections held in sequence through the U.S. states, where elections held later in the sequence may make use of the electoral results of earlier states to make more informed decisions. The 2000 Republican convention posited that this would lead to “better” outcomes.⁵ In political science, it is suggested that “front-loaded” primaries (i.e. primaries held close together, limiting informational effects) favored well-known candidates, while a more spread out primary allowed “dark horse” candidates to be discovered by the voters and build support [39]. Our results of Chapter 7 support this hypothesis.

The authors use monetary incentives to drive voter preferences in their laboratory experiments. Voters were assigned randomly one of three (private) types: *left*, *moderate* and *right*. They cast a single plurality ballot in favor of one of three candidates: x , y , or z . The candidates were assigned to one of three allegiances, but only one of the candidates would have their types revealed. Voters received a utility of 1 for electing their first preference, 0 for their last preference, and α for their second preference. The experimental parameter α , then controlled the risk adverseness of the voting population (the *left* and *right* voters, specifically, who may pick a *moderate* candidate if they are revealed, versus taking the lottery on the 2 remaining candidates).

In the simultaneous condition, all voters cast their ballots at the same time, having only known for certain the affiliation of one candidate. In the sequential condition, half the voters were randomly selected to cast their ballots first (after learning one of the candidates); the other half voters can then view their voting behavior, however, they did not know the identity of the candidate revealed to the first group and instead, learned the identity of a different candidate. Some groups in the sequential condition were additionally privy to the distribution of voter types in the initial group; these were called the “high information condition”.

The authors analyze equilibrium strategies in each of these conditions, using Myerson

⁵Both the Republican and Democratic conventions award bonus delegates to incentivize states that scheduled their primaries later in the sequence.

and Weber’s concept of Voting Equilibria [100]. In the simultaneous condition, voters act myopically since they have only limited information to work with. But in the sequential condition, the first group acts as in the simultaneous condition, but the second group may converge to multiple sequential voting equilibria, depending on the information revealed by the first group and their own risk-averseness.

In the laboratory experiments, the authors found that the candidate revealed first in the sequential conditions were less likely to win than those revealed in simultaneous conditions. More importantly, the moderate candidate is more likely to win in the sequential condition than the simultaneous condition. Informational effects were more difficult to tease out, but the authors found clear evidence that voting behavior in the second group in the sequential condition were noticeably different than the simultaneous condition; moreover, this effect was more pronounced in the high information condition, and when α was high (voters were risk averse).

In related, non-voting literature, Sandholm and Vulkan [117] examine bargaining games in distributed systems where agents have externally imposed deadlines. Prior to their deadline, agents may negotiate with each other by making offers in continuous time. Interestingly, they find that the sequential equilibrium behavior for the agents is to wait until the deadline, at which point they will concede fully. This is due to the informational effect that accompanies making an early offer, which signals a weakness in bargaining position. Moreover, an accepted offer shows that the offerer has already conceded too much, and would have been better off by waiting.

3.3 Voting in Networks

Finally, we survey current literature that specifically examine the topic of voting in social networks.

3.3.1 Preference Aggregation in Social Networks

In their paper, Dhamal and Narahari [47] describe a method by which preferences can be aggregated over a social network efficiently by only eliciting the preferences at certain critical nodes within the network. Their method “exploits network structure and homophily of relationships”. In their tests, the method works well for networks that exhibit homophilic properties, with certain assumptions on how much deviations voters’ preferences may have from those critical nodes.

To obtain data to validate their method, the authors employ a survey spread along the social network via known “seeders”, with an option for completing the survey anonymously to address privacy concerns. A total of 26 participants completed the survey, and homophily is demonstrated within the network. Using this data, the authors are able to calculate the preference-similarity matrix between the users. They also propose a way for estimating the similarity matrix when such data is unavailable: for connected neighbors, the amount of similarity is related to the *cliqueness coefficient* (a measure of the “clusters” that both voters belong to); for unconnected neighbors, the amount of “dissimilarity” is related to the shortest path between the two nodes.

Feldman et al. [59] examine the convergence of opinions within networks that resemble large social networks. Large social networks are characterized by being *expansive* (no sparse cuts) and *sparse* (low edge density). In previous studies, Bayesian learning models favor sparsely connected graphs, since highly-connected graphs lead to the suppression of useful information. However, majority dynamics models rely on a well-connected network without highly influential individuals to perform. Thus, there is a tension between the models on the type of graphs that information propagation operate well on.

The authors suggest a model that incorporates the asynchronous element from the Bayesian learning models within a majority dynamics model. Formally, each vertex in the graph receives one of two private signals about the ground truth of the world: *blue*, and *red*. They receive the correct signal with probability $\delta > 0.5$. All vertices in the graph begin in a third *uncolored* state. At each stage, a vertex v is selected uniformly at random to be updated. v takes the color of the majority of its neighbors who have declared a color, breaking ties in favor of its private signal. Note that the update will turn an *uncolored* vertex to one of two colors.

The paper investigates the performance of their dynamics model on λ -expander graphs. Define the weighted adjacency matrix M as follows:

$$M(x, y) = \begin{cases} \frac{1}{\sqrt{d(x)d(y)}}, & \text{if } (x, y) \in G \\ 0, & \text{otherwise} \end{cases}$$

Then a λ -expander graph is a graph whose weighted adjacency matrix have all but the first eigenvalue lying within $[-\lambda, \lambda]$.⁶

The authors go on to show that the dynamics are guaranteed to stabilize over time via a potential function (on any graph). They also extend Condorcet’s Jury Theorem to any

⁶Expander graphs are generally considered to be sparse graphs that are strongly connected; i.e., they lack small subsets of edges whose removal disconnect the graph.

λ -expander graphs with fixed degree $d \geq 6\lambda$; that is, as $n \rightarrow \infty$, the probability that the network converges to the correct result approaches 1.

3.3.2 Exploring Duverger’s Law

Clough [33] examines the applicability of Duverger’s Law in social networks using agent based modeling. Eschewing the assumption of global information or common shared belief, Clough’s model allows each agent to vote strategically based only on the information present from those “nearby” in their network. The model is dynamic and agents revise their votes over time. Only plurality voting is considered.

Formally, each agent in the model has a preference drawn uniformly from $[1,100]$; each party (“alternative”) has an agenda drawn the same way, and the utility derived from an alternative being elected decreases with the square of the distance of opinions. Voting proceeds in iterations, where each agent announces their new ballot simultaneously. In order to vote strategically, each agent considers the votes of their i -neighbors (those at most distance i away in the network), and assumes they are representative as priors for a multinomial distribution for the whole electorate. Tie probabilities $T_{j,k}$ can then be computed between every pair of candidates j, k , and then Cox Prospective Ratings can be used to determine which candidate is most attractive for each voter:

$$\xi_j = \sum_{k=1}^K T_{j,k}(u_j - u_k)$$

In the first round, with no previous information to go on, agents vote truthfully. The simulation continues for 20 iterations; the author mentions that few changes occur after 20 iterations when experiments were allowed to run longer. For all experiments, agents were laid out in a 13x13 torus. Elections with 3 or 4 parties were considered, along with varying connectivity i (up to $i = 6$ for full information).

Results show that in the majority of cases, support for third and fourth parties dwindle after the third round. But even so, a small but significant number of experiments still result in 3- or 4- party systems, in contradiction to Duverger’s Law. Analysis of the SF Ratio (ratio of support between the 3rd and 2nd place candidates) reaffirms this belief. It also shows that Duverger’s Law is more difficult to uphold when there are a large number of parties, and when connectivity is sparse.

In a follow-up paper, Clough investigates the effects of homogenization of the initial distribution of agents [34]. This is motivated by the idea of homophily, that people are more likely to associate with similar ideologies in a social network⁷. This paper continues

⁷The term homophily is not mentioned in the paper, and due to the choice of graph models and

the usage of the 13x13 grid torus setup, but in the “homogenized” model, the region is subdivided into 4 quadrants. Each quadrant represents an in-group espousing a certain ideology d_l drawn uniformly from $[1,100]$. Each agent within quadrant l has a preference drawn from a Gaussian distribution with mean $\mu = d_l$ and standard deviation $\sigma = 3$. This is contrasted with the heterogenous model where agent preferences are randomly drawn from one of the four distributions (i.e. mixing the homogenous model).

The analysis of experimental results are almost identical to the earlier paper[33]: In both models, a 4-party system is more difficult to coordinate than a 3-party system, with the difficulty mitigated by connectivity. Interestingly, increased homogeneity leads to an increased diversity of opinions as agents frequently fail to coordinate to a 2-party system. And unsurprisingly, homogenous models with low connectivity exhibit similar behaviors to models with *no* connectivity at all, as most agents do not receiving meaningful information from their neighbors.

In Chapter 5, we extend Clough’s model to more realistic depictions of social networks, such as the random graph models introduced in Section 2.5. We examine the desirability of the elected candidate, by looking at computational social choice properties such as Price of Stability and Price of Honesty, and their interactions with network structure. We also examine Duverger’s Law and compare our results to those obtained by Clough. In Chapter 6, we further extend our voter model with voter heuristic to help scale up our simulations to larger elections.

3.3.3 Iterative Voting in Social Networks

Sina et al. [121] have also extended Iterative Voting more directly to the social network domain. Voters are embedded in an undirected social network. Similar to the Knowledge Graphs approach, they assume that each voter i observes the ballots of her neighbors. In addition, voters combine this local level information with a publicly available poll to derive a score for each candidate, and alters their own ballot only when it will change the winner favorably due to a pivot condition in this local score.

The paper focuses on how the outcome of the election may be altered by manipulating the social network. They show that they can make any candidate c the winner of the election by adding only a linear number of edges (in the score difference between the current winner and c), and do so in polynomial time, as long as that candidate enjoys at

experimental setup, I believe the author is not aware of more principled methods of incorporating it into the graphical models.

least some basic level of support in the population. The algorithm is successful regardless of the nature of the poll, even if it is adversarial in selecting poll participants. Finally, they propose a greedy heuristic for a manipulator who may add only a limited number of edges to the graph, and show that it is effective in simulations based on real world profiles and networks.

Sina’s model differs importantly from our model presented in Chapter 5 in how voters address the incompleteness of their information. In Sina’s model, each voter treats her social network neighbors as the *entirety* of the electorate, and only act strategically if she observes a pivot condition within her neighbors (i.e. she knows for certain she can sway the election). We argue that it is unrealistic to assume voters only wait until they observe an exact pivot condition before they act strategically, but rather, they may choose to act strategically in optimism that it may sway the election. To this end, our voters consider probable outcomes of the election, informed by their social network neighbors.

3.3.4 Other Related works

The conspicuously titled paper, *Voting in social networks* by Boldi et al. [23] actually explores a very different topic from this thesis. Motivated by the lack of voter turn out in both real elections and virtual elections such as one conducted by Facebook in 2009 (with only a 0.3% turn out of its active user base of 200 million), the authors examine a pragmatic system for voting that could be applied to social networks. In this system, users can both vote directly, or they may vote by proxy by entrusting their ballot to one of their personal contacts. This delegation is transitive; the select contacts may themselves delegate their own and the entrusted ballots further along a delegation path. The authors study the properties of such a transitive proxy voting system. They also investigate the effects of *dampening*, where successive delegation renders a ballot less effective. Their results are similar to that of the Google PageRank algorithm, with the additional limit that all vertices have out-degree 1.

Other researchers have also focused on the junction of social choice and social networks. Salehi-Abari and Boutilier [115] examine social choice in a context where individual voters are empathetic — they derive utility not only from an intrinsic preference for the outcome, but also from the happiness of their friends. Their paper focuses on the computation of the socially optimal outcome under this paradigm, and examine the Irish election data set for the presence of this for empathetic behavior.

Tosatto and van Zee [126] approach the idea from a different direction. They wish to use tools from social networks to assist in social choice by first imputing a social network

from the voters' ballots. They consider a problem where a set of m binary policies must be accepted or rejected. Each ballot b is drawn from $\{0, 1\}^m$. However, to complicate matters, certain policies are mutually incompatible, and so not all ballots / outcomes in $\{0, 1\}^m$ are admissible. They focus on the *Average Voter Rule*, where the set of admissible outcomes is restricted to the set of submitted (and therefore admissible) ballots, and the ballot closest to the average ballot is selected as the winner. The authors use a similarity measure to construct a social network based on the voters' ballots; an edge of with w connects voters i and j , where w is the number of policies that the two voters agree on (i.e. the Hamming distance). They then use a weighted extension of degree centrality to obtain the winner, and find that the result is similar to using the Average Voter Rule.

Chapter 4

Opinion Dynamics

The field of opinion dynamics draws its early roots from the study of innovation diffusion. Under these models, agents within a community individually choose whether or not to adopt a novel trait based on the actions of their neighbours, in a repeated coordination game. Early studies focused on the decision to adopt new technologies such as antibiotics [35] and hybrid corn [122]. These decisions are naturally modelled by binary variables, and the model can just as easily be applied to study operating system and social media adoption today.

While binary variables are appropriate for modelling such decisions, they lack the richness necessary to capture more gradated opinions such as political leanings, socioeconomic standings, or various fashions and fads. The field of opinion dynamics generalizes the innovation model by interpreting opinions as continuous values in the interval $[0, 1]$. Agents' opinions are swayed by each other through repeated interactions, and the opinions of the community gradually converge to an equilibrium.

The analogous problem to innovation adoption in the continuous domain is the study of the effects of extremism in a community. In the discrete model, “early adopters” are represented as agents whose opinions are fixed to a certain value. In the continuous model, agents with fixed (or merely steadfast) opinions at the ends of the spectrum are akin to extremists in a population. The pitfall is that most mathematical models in this domain tend to focus on convergence of opinions [69]. The challenge then is to devise a model that allows fractions of a population to disagree with each other, even at equilibrium.

A class of phenomena known to cognitive scientists as *cognitive bias* motivates our approach. When subjects experience cognitive bias, they arrive at skewed or irrational

conclusions based on an inaccurate and subjective reconstruction of reality [21]. One particular type of cognitive bias is *motivated cognition*, where observations are evaluated in ways most beneficial to the individual or compatible with the individual’s beliefs.¹ In one experiment [82], when asked to rate the attractiveness and personality of a confederate, participants who were led to believe they must go on a date with the confederate consistently gave more favorable ratings. In a study of a more everyday phenomenon, after observing a sports event containing a minor but questionable call, fans of the losing team were more likely to attribute the outcome to referee error over qualities of the teams, when compared to fans of the winning team; however, in games where no such a questionable call is evident, there is no such bias.

The second study is very telling. Two groups of people were exposed to the same evidence, but their opinions (on the competitive merits of the respective teams) did not converge. This seems to fly in the face of belief updates via Bayes’ rule. Jaynes provides some insight on this by allowing agents to consider the possibility that the evidence is unreliable. The further away the evidence is from an agent’s expectations, the more likely the agent is to believe that it is flawed, and therefore the less persuasive the evidence [75]. Laplace summarizes this idea nicely in his essay on probability [83], that outlandish claims “decrease rather than augment the belief which they wish to inspire; for the those recitals render very probable the error or the falsehood of their authors.”

This idea of motivated cognition is central to our model. Agents are skeptical of another agent when their opinions diverge, but are more receptive to persuasion when their opinions better align. In the rest of this chapter, we detail work on related models in opinion dynamics, then we formalize this concept of skepticism and trust² in our model of opinion dynamics, and explore its effects in simulated social networks.

4.1 Related Work

Numerous researchers in the artificial intelligence community have explored how ideas diffuse through social networks. Recent works that emphasize the convergence of opinions include a model for how language features emerge, evolve and expire [123] and how opinions can be efficiently diffused in large communities [110]; Parunak also coins the term

¹A competing theory called *cognitive dissonance* explains the same behavior through a different set of mechanisms. The specific mechanics of these behaviors are unimportant to us, as we are only concerned with the fact that these behaviors *do* occur regularly in humans and other animals.

²We use “trust” only in its plain, nontechnical sense, and not in reference to mathematical trust models within the multi-agent systems community.

“collective cognitive convergence” in his study of the phenomenon, which also includes a more comprehensive review of literature [105].

Our skeptical paradigm places emphasis on limiting interactions between agents whose opinions diverge significantly, to emulate the effects of motivated cognition. Many researchers in the 20th century have explored various linear models for opinion formation [69]. Krause [80] was the first in the field to incorporate nonlinear systems, formulating the *bounded confidence* model. In this scenario, a panel of experts must arrive at a consensus about the evaluation of a piece of work. Each begins with a private opinion and a level of confidence on the accuracy of that opinion. As they interact with each other, they allow their opinions to be swayed by only those experts who hold opinions within a certain interval of theirs. The more confident the expert, the smaller their interval. The more confident the other expert is, the larger the sway.³

Most related to our work are Deffuant’s bounded confidence models. In the basic bounded confidence models, agents may only influence each other if their opinions differ by less than a threshold parameter. In subsequent work, Deffuant [41] refines this model by incorporating a Gaussian kernel with bandwidth equal to the confidence level. This allows influence to be dropped off in a smooth, continuous manner. The initial motivation for this model was to study the emergence of “mob mentality”, where sensible individuals are driven to extreme actions when present in a crowd containing only a small fraction of radicals [43]. Interestingly, while this avalanche effect sits as a counterpoint to the skeptical behavior motivating our model, it is emergent in our experiments. In his paper, Deffuant explores the ramifications of this model on Erdős-Rényi random graphs, while a variant of his model is explored in small-world social networks [62].

The idea of skepticism arising in social networks, between agents with different opinions, has also been explored more recently by Cho, Ver Steeg and Galstyan, and verified on data from the U.S. Senate [31]. In their paper, they consider co-membership in groups as being a surrogate for trust and a driver for evolution of network structure. Salzarulo [116] also investigates a similar phenomenon based on exogenously defined “in-group” and “out-group” mentalities.

Carvalho and Larson [30] explore the role skepticism plays in expert panels. In their model, a group of experts with initially different opinions revise their evaluations, with less weight given to experts whose opinions differ greatly from their own. They show that such a panel always reaches consensus, and such a model works efficiently on real world data.

³This description is based on subsequent work in [81], as [80] is written in German. A subsequent model by Deffuant [42] also weights the amount of influence exerted by the degree of overlap between the intervals.

The concepts of trust and persuasion have also been explored from different perspectives. Fang, Zhang and Thalmann [58] proposed a model for uniting the concepts of trust and innovation diffusion by allowing trust itself to be diffused through a network; Hazon, Lin and Kraus [68] considered how group decisions may be altered by appealing to self-interested individual to change their preference ballots.

Finally, Martins [87] proposes a model bridging the continuous and discrete domains, where agents maintain an internal (continuous) probability about which of two actions is more profitable, but is only able to communicate with each other through taking (discrete) actions. His simulations on a grid lattice show that a population eventually reaches stable equilibrium configurations of actions, where certain agents can become extremely confident of their choices.

4.2 Opinion Dynamics Model

We begin first by outlining several opinion dynamics models proposed by Krause and Deffuant. In all opinion dynamics models, agents $V = \{1, 2, \dots, n\}$ each have *opinions*, labelled $\mathbf{x} = \{x_1, x_2, \dots, x_n\}$, $x_i \in [0, 1]$. An agent's opinion may be influenced by those of other agents. This process proceeds in discrete timesteps. When an agent's opinion x_i is updated to a new value $f(\mathbf{x})$, we use the notation

$$x_i \leftarrow f(\mathbf{x})$$

The simplest opinion dynamics model is the **arithmetic mean**, where an agent's opinion is updated to the average of the opinions of all agents:

$$x_i \leftarrow \sum_{j \in V} \frac{x_j}{|V|}$$

Since the right hand side does not depend on i , after one such update, all agents acquire the same opinion. We say that opinions converge in one timestep.

Krause's basic **bounded confidence model** (BC) [81] extends this basic model via the additional stipulation that agents discard any opinions that are too different from their own, acquiring the average opinion amongst the remaining agents. Formally, let $N_i = \{j \in V : |x_i - x_j| \leq \epsilon\}$, where ϵ is a model parameter representing the threshold of maximum difference in opinion that is tolerated by individuals.

$$x_i \leftarrow \sum_{j \in N} \frac{x_j}{|N|}$$

Kraus shows that even in this simplistic model, convergence is not guaranteed. Let us define a vector of opinions \mathbf{x} as an ϵ -profile if there exists an ordering $\{y_1, y_2, \dots, y_n\}$ of \mathbf{x} such that $y_1 \leq y_2 \leq \dots \leq y_n$, and $|y_i - y_{i+1}| \leq \epsilon$, $\forall 1 \leq i \leq n - 1$. Kraus shows that for $2 \leq n \leq 4$, \mathbf{x} converges if and only if it begins as an *epsilon*-profile. But for $n \geq 5$, \mathbf{x} beginning in an *epsilon*-profile is a necessary but *not* sufficient condition for convergence.

The BC model is simple, but also flawed. One major shortcoming of the model is that the strength of influence of others' opinions depends very much on ϵ . In fact, the maximal influence is attained by an agent whose opinion differs by exactly the maximal tolerated amount ϵ . This seems unintuitive and unrealistic, which prompts the extension of Deffuant, Amblard and Weisbuch, called the **smoothed bounded confidence model** (SBC) [43]. At a high level, SBC smooths the influence of other agents by applying a Gaussian kernel; agents whose opinions are similar are given higher weights, and those whose opinions are further away are given diminishingly lower weights. Each agent i has an opinion value x_i as before, and an uncertainty value c_i which controls how receptive that agent is to new and different information. As the agent updates, uncertainty may decrease, which makes the agent more discriminating and assigns less trust to agents with differing opinions.

Formally, we first define the Gaussian kernel for i , which depends on the agent's uncertainty c_i , and opinions x and x' :

$$g_i(x, x') = \exp\left(-\frac{(x - x')^2}{c_i}\right)$$

Then, at each timestep, agent i pairs up with a randomly selected agent j , updating its opinion x_i and uncertainty c_i as follows (agent j performs a symmetric and simultaneous update):

$$x_i \leftarrow \frac{x_i + g_i(x_i, x_j)x_j}{1 + g_i(x_i, x_j)}$$

$$c_i \leftarrow \frac{c_i + g_i(x_i, x_j)c_j}{1 + g_i(x_i, x_j)}$$

One limitation of this model is that the influence between an agent and her neighbors is governed by a single parameter c_i , and this parameter affects all neighbors in the same

way. We propose an alternative model where each edge is weighted by how much one agent trusts another, and that these trust values are allowed to evolve independently of each other. This allows agents to maintain close ties to some individuals, while being skeptical of others.

4.2.1 Our Model

In our model, agents $\{1, 2, \dots, n\}$ are embedded in a social network represented by a simple, undirected graph $G = (V, E)$. Each agent i has an *opinion* $x_i \in [0, 1]$ and is influenced by neighbours $N(i) = \{v \in V | \{i, v\} \in E\}$. For each neighbour j , i maintains a *trust* value $w_{i,j} > 0$ representing the weight i gives to j 's opinions.

We define a trust function T , based on the distance between any particular opinions x and x' via the Gaussian kernel, described in Equation (4.1). The bandwidth parameter b represents the *empathy* of the population; a higher empathy reflects a population more willing to be persuaded by someone with a more different opinion.

$$T(x, x') = \exp\left(-\frac{(x - x')^2}{b}\right) \quad (4.1)$$

Equations (4.2) and (4.3) describe the opinion and trust updates performed at each time step: each agent i updates its opinion x_i and trust values $w_{i,j}$ via a weighted average. A lower $w_{i,j}$ indicates i is more skeptical of j , and therefore, less influenced by j 's opinions. We include a parameter $w_{i,i}$ as the inertia of i 's trust and opinions, to be weighted against those of its neighbours. A high value of $w_{i,i}$ means i changes its opinion slowly, while a low value means i is easily swayed by the opinions of others. The value of $w_{i,i}$ does not change throughout the simulation. We also define a parameter r representing the *learning rate* of the population; a higher learning rate reflects a more judgemental population that more quickly distrusts someone with a different opinion. Note also that the opinion update (4.2) is performed before the trust update (4.3) in each iteration.

$$x_i \leftarrow \frac{w_{i,i}x_i + \sum_{j \in N(i)} w_{i,j}x_j}{w_{i,i} + \sum_{j \in N(i)} w_{i,j}} \quad (4.2)$$

$$w_{i,j} \leftarrow \frac{w_{i,j} + r T(x_i, x_j)}{1 + r} \quad (4.3)$$

The majority of nodes in each network will represent *moderate* agents, with randomly chosen initial opinions which update as described above. The remainder of the vertices will represent *extremists*. Extremists have polarized opinions x_i fixed at one extreme of the spectrum (either 0 or 1), which is equivalent to setting their empathy to 0; i.e. they do not update according to equations (4.2) and (4.3).

The use of the Gaussian kernel is reminiscent of the smoothed bounded confidence (SBC) model in [43]. Our model differs by replacing each agent’s personal confidence value, with dynamically updated trust values between every pair of agents. This allows agents to remain receptive to some of their neighbours while becoming more skeptical of others, and also for trust to be gradually lost or recovered over time. The notion of equating confidence with persuasiveness is appropriate in a cooperative setting such as an expert panel, but seems less suitable in a setting where agents are skeptical in their interactions.

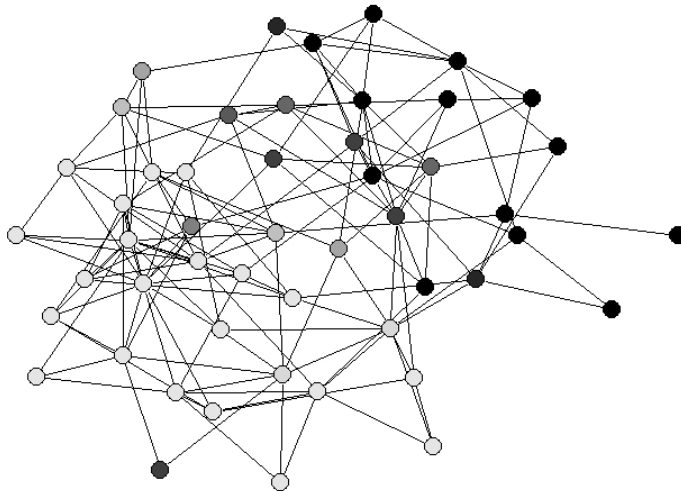


Figure 4.1: A Erdős-Rényi graph with homophily. Node colors indicate initial opinions, with progression from white (0) to black (1).

4.2.2 Graph Models

We consider two types of random graph models in our experiments: the classic Barabási-Albert random graph, and a homophily model based on Erdős-Rényi random graphs similar to that presented in [133].

A Barabási-Albert random graph with attachment parameter m is constructed by iteratively adding vertices, connecting them to m existing vertices with probability proportional to their respective degrees. It is often used to model the scale-free property of social networks where a relatively few number of vertices (“hubs”) cover most of the edges.

A directed Erdős-Rényi random graph with connectivity probability p is constructed by considering every pair of vertices i and j , and connecting them with fixed probability p . We incorporate homophily in this model by reweighting the connection probability between i and j as $(1-d)p$, where $d = |x_i - x_j|$. This causes vertices with similar opinions to be joined with higher probability than those with disparate opinions. As with the classic Erdős-Rényi model, the resulting graph may be disconnected. If this is the case, we simply discard and regenerate the graph. A typical modified Erdős-Rényi graph on 50 vertices and $p = 0.2$ is shown in Figure 4.1; agent opinions were drawn from the distribution $Beta(0.5, 0.5)$.

4.2.3 A-priori Trust Models

The initial trust between the agents represent how much the agents trust each other prior to the start of the experiment. We utilize three different trust models:

First, we have the *uniform trust* model, where $w_{i,j} = 1, \forall \{i, j\} \in E$. We define $w_{i,i} = d_i$, where d_i is the degree of vertex i , which is consistent with a degree-based voter model where the interactions between an agent and its neighbours are modeled as a series of pairwise interactions. This model makes the fewest assumptions about how trust has been established.

Next, we have the *degree based trust* model, where more initial trust given to the opinions of well-connected (“popular”) members of the community: $w_{i,j} = \frac{d_j}{d_i}, \forall \{i, j\} \in E$. Similarly by the logic above, we define $w_{i,i} = 1$.

Finally, we have the *kernel based trust* model. Here, we assume the vertices have interacted previously and their trust value have converged to equilibrium values specified by equation (4.1); that is, $w_{i,j} = T(x_i, x_j)$ and $w_{i,i} = 1$

4.3 Empirical Simulations

In this section, we describe two sets of experiments that explore the behavior of agents in our model. The first set of experiments operate only on Barabási-Albert random graphs, and aims to explore the ability of extremists to influence the moderate population on

typical (i.e. scale-free) social networks. That is, we will measure the average opinions of the moderate population at the end of each trial. By seeing how much they deviate from their initial opinions, we will know how much influence extremists have exerted on them. We will look at how variations in graph structure, edge density and agent empathy affect the magnitude of this influence. Both higher edge density and higher agent empathy are traits that facilitate the spread of opinions in networks, and so we hypothesize that they also allow extremists to have a larger impact on the population, even in our Skeptical model. Differences in graph structure are a more complicated topic. The dominant feature of Barabási-Albert graphs is the presence of highly connected hubs. We hypothesize that the conversion of these hubs will allow extremists to exert more influence in these graphs than in Erdős-Rényi random graphs.

In the second set of experiments, we explore the ability for extremists at both ends of the spectrum to polarize the moderate population, with the ultimate goal of finding necessary conditions for the opinions of the moderates to stratify and stabilize at multiple, non-polarized levels. As before, we measure the change in average opinion of the moderate population to show the ability of the extremists to polarize the population. Moreover, we will measure the number of distinct clusters of final opinions to show whether or not the entire population collapsed to a single final opinion, or become stratified with multiple clusters of opinions. We hypothesize that traits that are conducive to extremist influence in the first set of experiments will similarly facilitate polarization, and that most graph conditions will result in the convergence to one or two opinion clusters.

4.3.1 Experimental Design

For each experiment, we initialize the social network G with 200 nodes using the appropriate graph model, with varying parameters for graph construction and agent empathy. In the first set of experiments, 10% (20 nodes) of the population is chosen uniformly at random to be 1-extremists; we call this the 1-pole model. In the second set, the population contains 10% 0-extremists and 10% 1-extremists⁴, also chosen uniformly at random; we call this the 2-pole model.

The remainder of the population comprise the moderates. They begin with opinions initialized to random values: either sampled uniformly from the interval $[0, 1)$, or from the partially polarized distribution $Beta(0.5, 0.5)$. The uniform distribution serves as a

⁴We selected 10% as an arbitrary percentage that still leaves a large segment of the population as moderates to be influenced. The 2-pole experiments of Deffuant, Amblard and Weisbuch [43] used 5% extremists of each type. We do not expect the difference to have a qualitative impact on the results.

baseline incorporating minimal assumptions, while the *Beta* distribution is set so that initial opinions will already be bimodal and drawn toward the extremes; this is thought to increase the likelihood that opinions will stratify. Initial trust between them is set according to one of the models outlined in Section 4.2.3.

Once the instance is initialized, the variables are updated according to equations (4.1)-(4.3). We set $r = 1.5$ for all experiments, as preliminary tests did not find varying $r > 0$ changed our qualitative results.

The experiment terminates when no opinions changed by more than a small value ϵ , or a maximum number of iterations t_{max} has been reached. In our experiments, we set $\epsilon = 0.001$ and $t_{max} = 500$; t_{max} was rarely reached in practice. This model was implemented using Python 3.3.2. All results are averaged over 25 replicated trials.

4.3.2 Influence of Extremists

We begin by investigating the ability of extremists to affect the opinions of the moderates, and how that impact varies with graph structure and parameters of the agents. Figure 4.2 shows the evolution of opinions over the course of an experiment. Opinions of the moderates are bucketed in intervals of size 0.05 on the y-axis, and the timesteps is marked on the x-axis. Each column of the heat map shows the distribution of opinions in that particular timestep. We see that opinions are distributed randomly at the start, and begin to converge toward the mean before shifting dramatically toward 1.0. To quantify the impact of extremists, we measure the mean opinion of the moderates at the end of each experiment. If the moderates were completely unaffected by the extremists, the mean would hover around 0.5. If the extremists were completely successful at persuading the moderates, the mean would near 1.0.

Figure 4.3 shows how the average opinion at convergence changes as we adjust the empathy bandwidth parameter b , and the attachment parameter m . As expected, increasing empathy increases the impact of the extremists on the population. However, aside from the special case when $m = 1$, which imposes a tree structure on the network, increasing connectivity does not significantly impact the mean at convergence. This is likely due to the small-world property of these graphs, allowing influence to propagate quickly through the network.

Figure 4.3 also contrasts the effects of initializing using uniform trust (top) and degree-based trust (bottom). Adopting initial degree based trust introduces more degrees of freedom in the experiment, in the form of the portion of hubs becoming extremists. This accounts for the higher variability in our results.

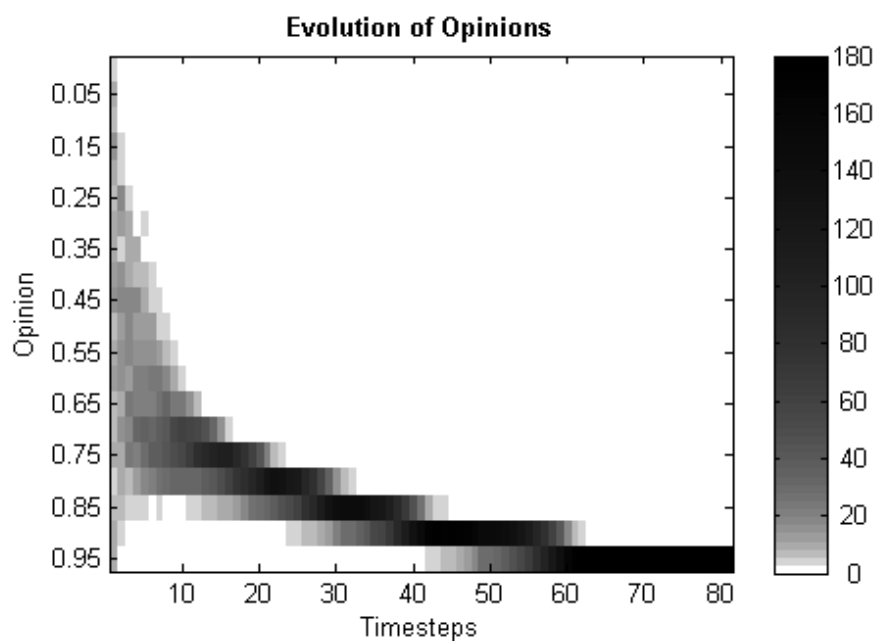


Figure 4.2: Opinions of moderates over the course of an experiment. Note the color scale is logarithmic, and the y-axis is decreasing.

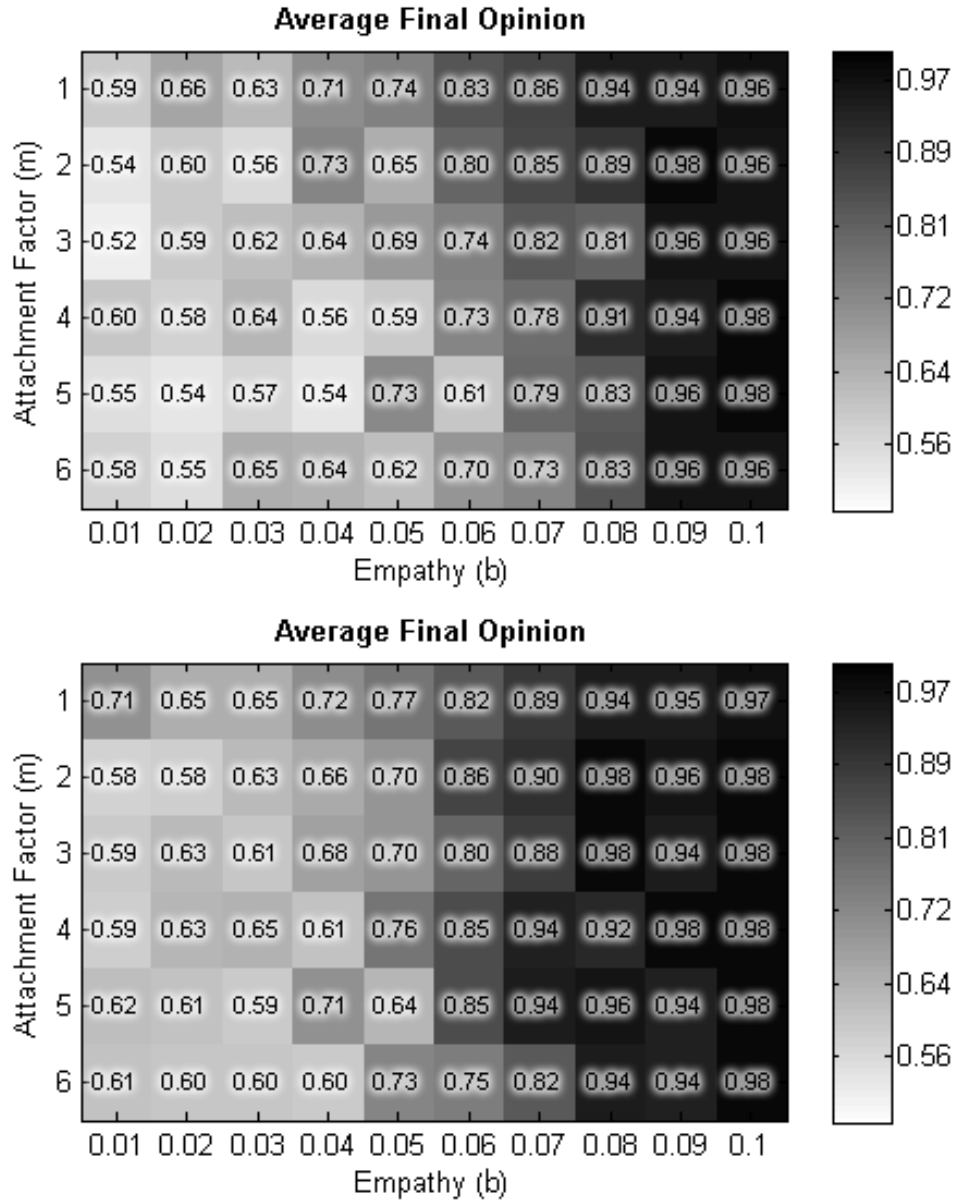


Figure 4.3: The convergence mean opinion of moderates, in the presence of 10% 1-extremists. The model at the top is initialized using uniform trust (95% confidence interval within ± 0.11 for all sets), and the bottom, using degree based trust (95% C.I. within ± 0.10).

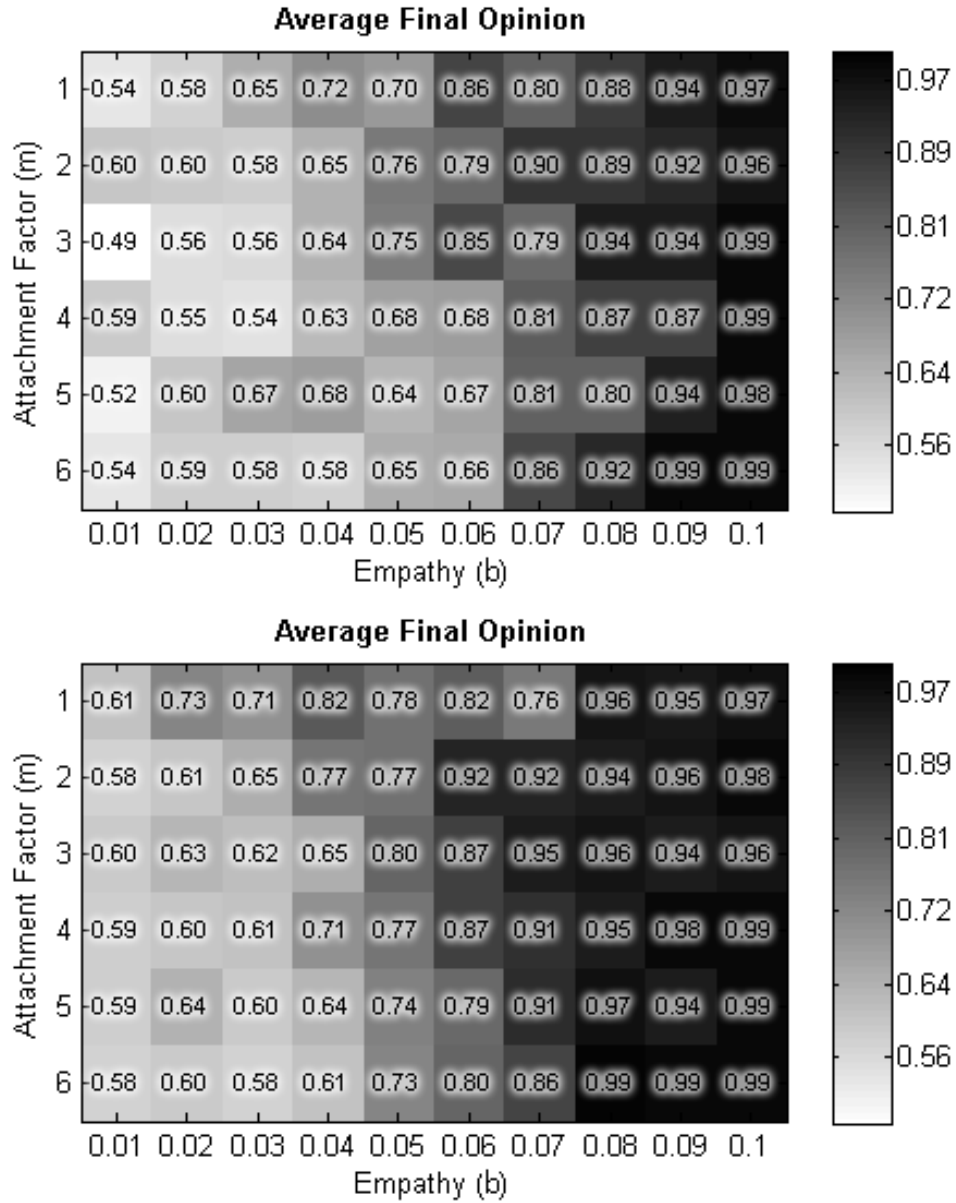


Figure 4.4: Effects of introducing noise to the model of Figure 4.3. Uniform trust (top, 95% C.I. within ± 0.11) and degree based trust (bottom, 95% C.I. within ± 0.10).

One critique of Deffuant’s SBC model is its sensitivity to noise [53]. We introduce a similar level of noise to our model, allowing each moderate agent to change their opinion by a small value drawn from a Gaussian distribution, with small probability at each update. More formally, each agent at each iteration has a 0.01 probability of using the following equation in place of equation (4.2) for their opinion update. Note that the resulting opinion is bounded within $[0, 1]$.

$$x_i \leftarrow \frac{w_{i,i}x_i + \sum_{j \in N(i)} w_{i,j}x_j}{w_{i,i} + \sum_{j \in N(i)} w_{i,j}} + \Delta, \Delta \sim \mathcal{N}(0, 0.15) \quad (4.4)$$

Figure 4.4 shows the effect of introducing this degree of noise into the update process. Aside from the $m = 1$ case, there is little qualitative difference compared to Figure 4.3. Equation (4.3) controls the trust dynamics within our network and enables agents to react to sudden deviations in opinions. This gives our model the robustness to absorb noisy signals.

Examining the evolution of opinions in Figure 4.2, we see that in the initial stage of the experiment, the moderates rapidly converge toward a common opinion. Even agents near the pole are drawn in due to the initial trust conditions. This effect is amplified by the small-worlds property of these graphs. Once an early consensus is reached, the moderate opinion may slowly migrate to the extreme through gradual influence from extremists (as the case in Figure 4.2), or may successfully insulate the extremists from influencing the general opinion.

4.3.3 Opinion Polarization

In our second set of experiments, we incorporate two sets of extremists competing for the opinions of the moderate population. We initialize a randomly selected 10% of the population as 1-extremists and another 10% as 0-extremists. Deffuant [41] characterized 4 types of convergence in these two-pole situations: the moderates may converge to a single opinion that is either (I) moderate or (II) polarized, (III) the population may split in two with a portion converging at each pole, or (IV) the population may fragment, with fractions that retain non-extreme opinions.

The polarization of each agent’s opinion is the absolute difference of their final opinion from the middle ground of 0.5. The higher the polarization, the more influence is felt from the extremists. This allows us to differentiate non-polarized outcomes (types I and

IV) from polarized outcomes (types II and III). To detect whether or not moderates have stratified opinions, we examine the final distribution of their opinions to see if they are unimodal (types I and II), or multimodal (types III and IV).

To eliminate false positives due to noise, we use the following procedure for identifying multimodality. First, we form a histogram of opinions, dividing the $[0, 1)$ interval into 20 buckets b_1, \dots, b_{20} , each of width 0.05. A distribution is multimodal if there exist three buckets b_i, b_j, b_k ($i < j < k$) such that $b_j < \min(b_i, b_k)/2$, and $\min(b_i, b_k) \geq T$. We arbitrarily choose the threshold $T = 20$, which represents 10% of the agents.

Figure 4.5 shows the average polarization for our experiments. As before, higher empathy b is correlated with increased influence from extremists. However, now network structure plays a role as well: There is less impact from extremists in more highly connected networks, represented by lower average polarization. We examine the final opinions and find that, in all cases with $m > 1$, the moderates converge to a unimodal distribution that drifts toward one of the extremes, reaching a type I or type II convergence. As before, we verify that these results are robust against noise (data not shown).

One might wonder whether a population that is initially divided can produce type III or type IV convergences. We investigate this possibility by drawing initial opinions x_i from $Beta(0.5, 0.5)$. Figure 4.6 shows a run under these parameters. We observe a behavior similar to that of Figure 4.2 – the population converges toward an early consensus, and gradually shifts to a unimodal distribution near one of the extremes. This behavior is consistent across all trials, with no multimodal distributions arising when $m > 1$.

Figure 4.7 shows the average polarization of the general population using the two initial trust models. On the top, we observe that uniform initial trust allows polarization to occur rapidly, regardless of network structure, with nearly complete polarization occurring at $b > 0.04$. This is a surprising result, since in order for moderates to polarize at one extreme, a large portion of the population must be converted from their initial opinions set on the other end of the spectrum. We also observe this trend when the network is initialized using degree-based trust (Figure 4.7, bottom), but it is not as obvious as with uniform initial trust.

Thus, there appear to be two main factors preventing opinions from stratifying. The *initial trust* given to agents of significantly different opinions, and the *lack of homophily* in the graph structure. To remedy the first issue, we implement the kernel trust model, where agents are inoculated with skepticism right from the start, modeling a situation where agents have previously interacted and trust dynamics have reached an equilibrium between them. To combat the second issue, we define the modified Erdős-Rényi random graph to capture the homophily property of social networks.

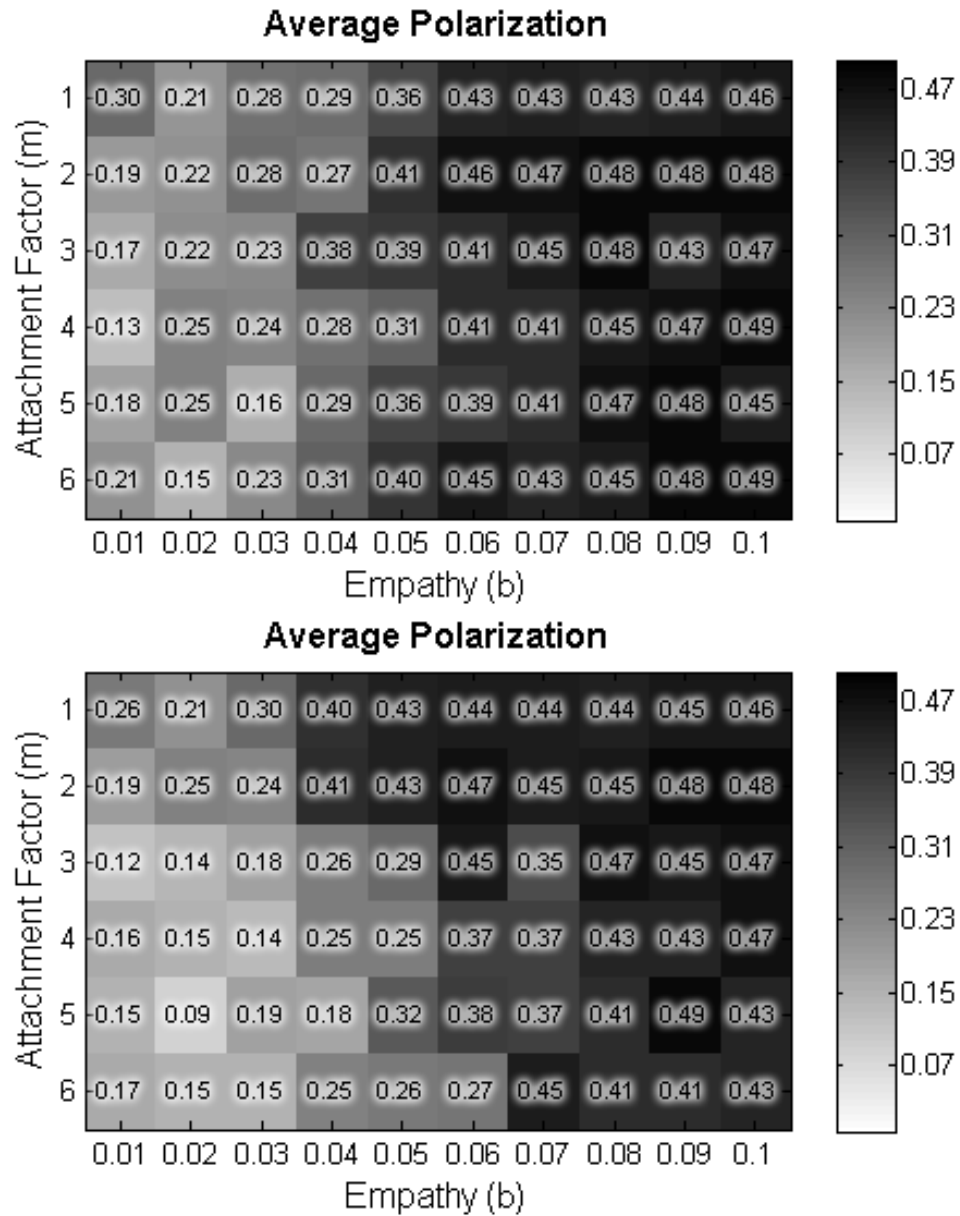


Figure 4.5: The average polarization of moderates when exposed to extremists of opposing camps. The model at the top is initialized using uniform trust (95% C.I. within ± 0.09), and the bottom, using degree based trust (95% C.I. within ± 0.09).

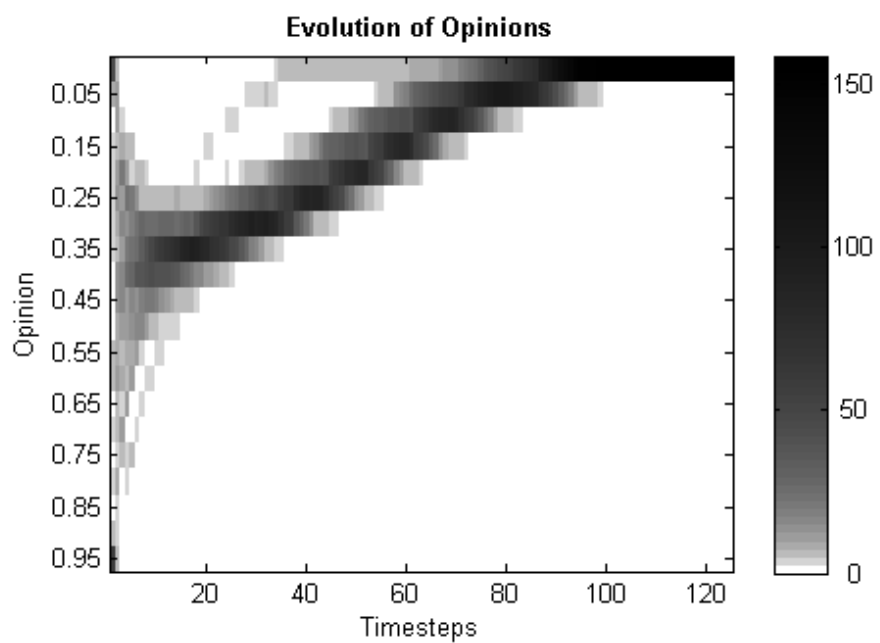


Figure 4.6: Evolution of opinions in moderates, with partially polarized initial opinions. Note the color scale is logarithmic.

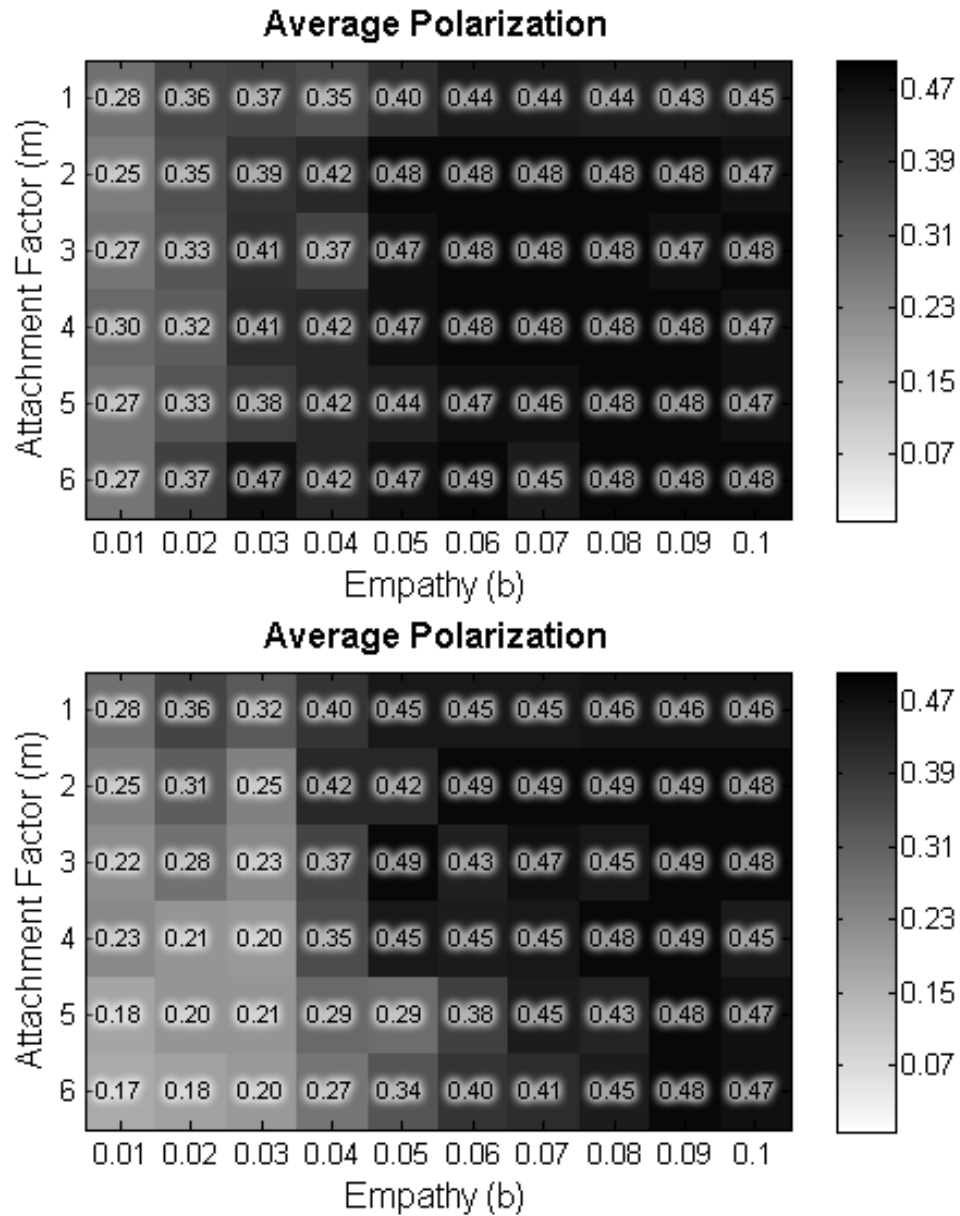


Figure 4.7: The average polarization of moderates with initial opinions drawn from $\beta(0.5, 0.5)$. The model at the top is initialized using uniform trust (95% C.I. within ± 0.07), and the bottom, using degree based trust (95% C.I. within ± 0.09).

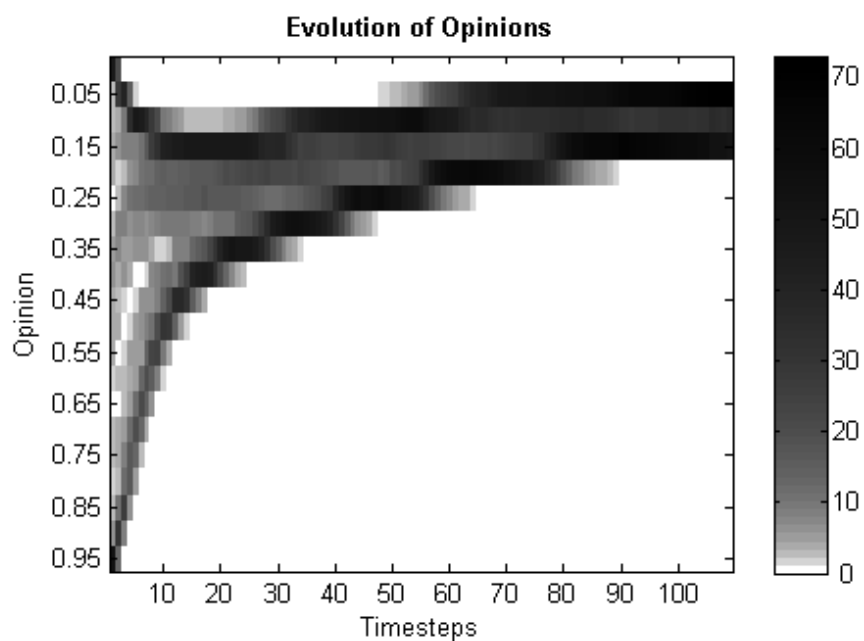


Figure 4.8: Evolution of opinions in moderates, on a modified ER-graph with homophily, with partially polarized initial opinions. Note the color scale is logarithmic.

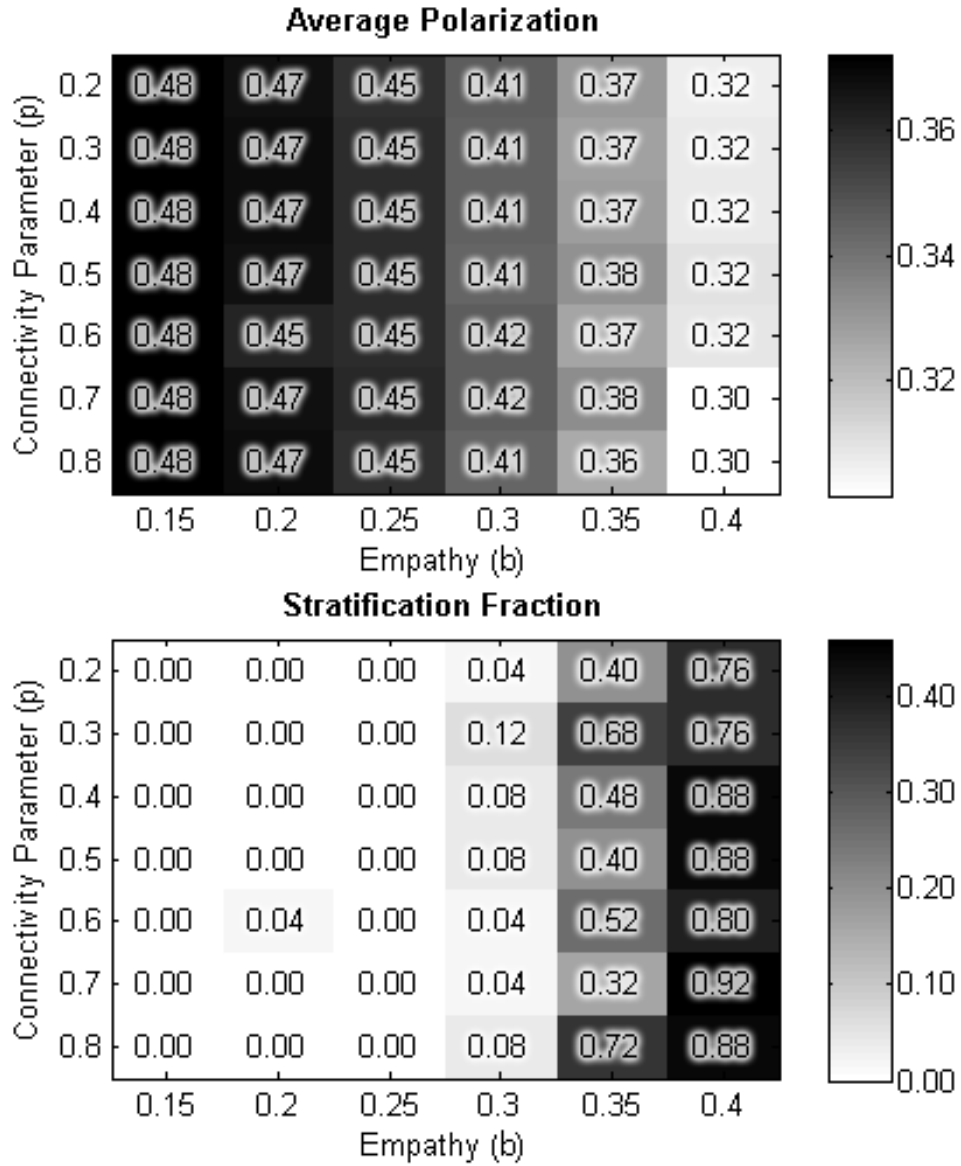


Figure 4.9: Average polarization of moderates on a modified ER-graph with homophily, with partially polarized initial opinions (top, 95% C.I. within ± 0.03), and the frequency of stratification (bottom).

Interestingly, in our simulations of this new model, stratification does not occur until empathy $b > 0.3$, far above the point at which opinions normally become polarized. Figure 4.8 shows the evolution of opinions in a run that ends in a Type IV convergence. Notice the concentration of opinions migrate gradually from the poles, but do not converge. As shown in Figure 4.9 (top), the amount of polarization actually *decreases* as empathy increases beyond 0.3. Figure 4.9 (bottom) shows the fraction of runs that converge to multimodal distributions. As empathy exceeds 0.3, the likelihood of a type III or type IV outcome increases, and the presence of type IV convergences necessarily lowers the average polarization.

The notion that agents with higher empathy, and therefore those who “listen” to, and are influenced by, a wider range of opinions, is a necessary ingredient for opinions to stratify is very surprising. We hypothesize that this is because agents with such high empathy values are simultaneously affected by extremists from both poles, stabilizing their opinions in a bimodal configuration. Similar stratification is not observed in the Barabási-Albert or the unmodified Erdős-Rényi models, even when employing kernel trust; nor is it observed in the modified Erdős-Rényi model without kernel trust (data not shown). And thus we hypothesize that these conditions represent a narrow band of conditions that are necessary for opinion stratification.

4.4 Discussion

One natural question to ask is how the conversion of half the population from one end of the opinion spectrum to the other occurs in Barabási-Albert graphs. The answer may be found by approximating the amount of influence that can be exerted on a densely connected community, even when they have already reached a unified opinion (this is a best case scenario that lower bounds the amount of influence that can be exerted on it). To do this, we extend the concept of *cluster densities* from innovation diffusion. We define a *cluster of density p* as a set of nodes in G such that no node in the cluster has more than fraction p of its neighbours outside the cluster [51].

Now, suppose A is a cluster of density p , $B = G \setminus A$, and all agents in A have opinion x , while all agents in B have opinion $x + \Delta$.

Consider a node i in A with degree d . According to Equation (4.2), x_i will be updated according to

$$\begin{aligned}
x_i &\leftarrow \frac{dx_i + \sum_{j \in N(i)} w_{i,j} x_j}{d + \sum_{j \in N(i)} w_{i,j}} \\
&= \frac{dx_i + \sum_{j \in N(i) \cap A} x_i + \sum_{j \in N(i) \cap B} w_{i,j} x_j}{d + |N(i) \cap A| + \sum_{j \in N(i) \cap B} w_{i,j}} \\
&= \frac{d(1+p)x_i + \sum_{j \in N(i) \cap B} w_{i,j} x_j}{d(1+p) + \sum_{j \in N(i) \cap B} w_{i,j}}.
\end{aligned}$$

Now, recall the trust function $T(x, x + \Delta)$, which denotes the weight given to an opinion that differs from an agent's by Δ . Let us write this as $T(\Delta)$ for brevity. If we approximate the weights $w_{i,j}$ with the target trust function $T(\Delta)$,

$$\begin{aligned}
x_i &\leftarrow \frac{d(1+p)x_i + d(1-p) T(\Delta)(x_i + \Delta)}{d(1+p) + d(1-p) T(\Delta)} \\
&= x_i + \frac{((1-p) T(\Delta))\Delta}{(1+p) + (1-p) T(\Delta)}.
\end{aligned}$$

Finally, if the difference in opinions is sufficiently large, and the clusters sufficiently dense, then we may assume $(1+p) \gg (1-p) T(\Delta)$. Then,

$$x_i \leftarrow x_i + \frac{1-p}{1+p} T(\Delta)\Delta.$$

And we substitute our trust function to get

$$x_i \leftarrow x_i + \frac{1-p}{1+p} \exp\left(-\frac{\Delta^2}{b}\right)\Delta. \quad (4.5)$$

If the right hand side of this expression is bounded within ϵ of x_i , then the simulation will terminate. By comparison, let us modify the above setup by allowing x_i to have a very small fraction p'' of its neighbours that are bridge vertices, with an intermediate opinion $x_i + \Delta/2$. x_i still has fraction p of its neighbours in A , and p' of its neighbours in B with opinion $x_i + \Delta$ ($p + p' + p'' = 1$). By a similar analysis, approximating the weights $w_{i,j}$ with T yields:

$$\begin{aligned} x_i &\leftarrow \frac{(1+p)x_i + p' T(\Delta)(x_i + \Delta) + p'' T(\frac{\Delta}{2})(x_i + \frac{\Delta}{2})}{(1+p) + p' T(\Delta) + p'' T(\frac{\Delta}{2})} \\ &= x_i + \frac{(p' T(\Delta)\Delta + p'' T(\frac{\Delta}{2})\frac{\Delta}{2})}{(1+p) + p' T(\Delta) + p'' T(\frac{\Delta}{2})}. \end{aligned}$$

And if we assume $(1+p) \gg p' T(\Delta) + p'' T(\frac{\Delta}{2})$, then,

$$\begin{aligned} &\cong x_i + \frac{1}{1+p} \left[p' T(\Delta)\Delta + p'' T(\frac{\Delta}{2})\frac{\Delta}{2} \right] \\ &= x_i + \frac{1}{1+p} \left[\exp(-\frac{\Delta^2}{b})\Delta + \exp(-\frac{\Delta^2}{4b})\frac{\Delta}{2} \right]. \end{aligned} \tag{4.6}$$

By comparing equation (4.5) with (4.6), we see that the amount of influence effected on x_i is greater in the presence of bridge vertices if $T(\Delta) < 2/3$, which is certainly true if we expect the simulation to halt in the bridgeless case.

Thus, when there is a large gulf in opinions, influence is quite limited and skepticism is high. However, the presence of even a handful of unpolarized intermediaries will serve as a siphon through which influence will flow, starting an avalanche effect where the two clusters' opinions begin to converge with increasing speed. This is reminiscent of the ‘‘mob mentality’’ that inspired Deffuant’s SBC model.

4.5 Conclusion

We have introduced a robust model of opinion dynamics that captures trust and skepticism between agents that changes over time based on the difference in opinions between the

agents. We show that agents operating in a preferential attachment, small-world network will quickly converge to an early, loose consensus before taking coordinated action to migrate the collective opinion to the equilibrium. This equilibrium may be moderate or polar, with agent empathy being the primary factor influencing the final outcome. A secondary factor is connectivity, which has a significant moderating effect, but only in the two-pole model.

Only by utilizing a homophilic graph model and inoculating our agents with an equilibrium amount of skepticism for other agents, can we cause opinions to stratify away from extreme values. We hypothesize that this stratification can only exist when individual opinions are stabilized by more extreme opinions from both ends of the spectrum.

Future work include further exploration of the properties of homophilic graph models. The model could be adjusted to include heterogeneous empathy and learning rates within the population. It could also be extended to a dynamic population, with agents entering and leaving the community over time, and with their opinions growing more confident (higher skepticism) as they interact with the community. Finally, our model could be adapted to the discrete action domain, where each agent possesses a private continuous opinion, takes discrete actions based on that private opinion, and only observes the discrete actions of their neighbours.

Chapter 5

Voting in Social Networks

Voting is a method of social choice where a community elicits the personal preferences of individuals to conduct collective decision making. A major concern in voting systems is manipulation via strategic voting. This happens when voters benefit from casting a ballot that does not reflect their true preferences; while this may be beneficial for the voter, it misinforms the community on the needs of its constituents. In order for voters to manipulate successfully, they must have some knowledge regarding the outcome of the election. One reasonable model is to view the election as a series of rounds, where voters put forth tentative ballots that may be continually revised; this is called Iterative Voting, which assumes voters have complete information on the ballots of all other voters [95]. Subsequent work incorporates voters who are truth biased (who prefer voting sincerely if they cannot otherwise affect the outcome) or lazy (who prefer abstaining, all else being equal) [111]. In a social network, however, voters are restricted to observing only the actions of their neighbors. Each voter must form a model of the likely outcome of the election based on this incomplete information, and use this model to inform their actions. This assumption may appear unrealistic at first glance. Since, after all, one does not simply make decisions based on a sampling of opinions from Facebook friends. However, our use of the term social network extends beyond relationships in online social media platforms, and also include experts and associates, media outlets, and any other source of opinion and information that may contribute to the decision making process.

In this chapter, we present a behavioral model of voters embedded in a social network. Voting occurs in successive rounds during which voters may alter their ballots. Voters can observe only the ballots of their friends in the social network. Each voter assumes her friends are representative of the wider population, and will vote strategically to maximize her own expected utility. We explore the behavior of this model on a variety of random

graph networks, including ones that exhibit homophily. We focus on using plurality as our voting system. Strategization is a major concern in plurality voting. Voters are incentivized to submit ballots that do not truthfully represent their true, underlying preferences. If voters are not being truthful, the voting mechanism may produce an outcome that is inappropriate or damaging to the population. However, we show that in our model of self-interested strategization, social welfare is *improved* via strategization, when compared to truthful voting. We also conjecture that network homophily affects strategic voting behavior, and contributes to low percentage of voters that strategize in many real world elections (for example, in [20] and [61]). In particular, when voters are surrounded by like-minded individuals, they are more likely to vote truthfully, echoing the opinions of those around them. This is called the “Echo Chamber Effect” in popular media, and can be observed in our simulations. Crucially, it lowers social welfare by decreasing the amount of strategization that occurs, hurting the overall community. Finally, while our model converges quickly in practice, we show a counterexample where voters never converge to a stable state.

5.1 Model

Let $V = \{1, 2, \dots, n\}$ represent our set of voters. They are embedded in a social network, represented as a simple, directed graph $G = (V, E)$. We adopt the convention that a directed edge $(i, j) \in E$ denotes that voter i observes voter j and as such, j 's actions may influence i . An edge may represent communication between friends, a leader's influence on followers, or patronage of media and news platforms. Let $\mathcal{N}(i)$ denote the set of voters observed by i ; i.e. $\mathcal{N}(i)$ is the out-neighbors of i .

Let $C = \{c_1, c_2, \dots, c_m\}$ represent the set of available candidates. Let \mathcal{F} be the voting function used to aggregate those ballots to choose a single winner; it may or may not be deterministic. The choice of \mathcal{F} will define a set of valid ballots that can be submitted by voters; let us denote this set as \mathcal{B} . Then, \mathcal{F} maps \mathcal{B}^n to a winner $\hat{c} \in C$.

The voting process proceeds in rounds. In round t , each voter $i \in V$ submits a ballot $b_i^{(t)} \in \mathcal{B}$. The voter formulates this ballot as a response $R_i : \mathcal{B}^{|\mathcal{N}(i)|} \rightarrow \mathcal{B}$ based on her observations of her friends – i.e. the previous ballots of her out-neighbors. These rounds may represent a series of preliminary polls leading up to the final election. We assume all voters begin with the truthful ballot.¹ Voting continues until no agents choose to revise their ballots, whereupon the winner is decided by the voting function \mathcal{F} . When no voters

¹Or a truthful ballot, depending on the voting system

wish to deviate from their current ballot, the system has converged to an equilibrium. If it reaches this state, we say the system is stable.

5.1.1 Model of Voters

Models of voters in multiagent systems literature are divided between those utilizing ordinal preferences (where only the ordering of outcomes matter) and cardinal preferences (where outcomes are associated with utility values).² While each model has its own merits, we choose the latter model because our voters infer and weight the probabilities of the different outcomes, and act rationally to maximize expected utility.

Formally, voters derive utility based on the candidate that is elected by \mathcal{F} . Each candidate $c_i \in C$ advocates a position $p(c_i)$ in some domain \mathcal{D} that is common knowledge. Each voter i favors a position $p_i \in \mathcal{D}$ known only to herself. If \hat{c} is the winning candidate elected by \mathcal{F} , then a utility function $u_i(p_i, p(\hat{c})) : \mathcal{D} \times \mathcal{D} \rightarrow \mathbb{R}$ determines the value of this outcome.

For the purpose of this chapter, \mathcal{D} are the integers from 0 to 100 (inclusive), and preferences are single-peaked. This allows us to benchmark our result to previous work (e.g. [33, 34]); this one dimensional scale is also commonly used in political science literature to represent the left-right political spectrum [19, 71]. We assume the utility a voter derives from the outcome decreases with the square of the distance between her favored position p_i and the winner’s advocated position \hat{p} :

$$u_i(p_i, \hat{p}) = -|p_i - \hat{p}|^2.$$

For brevity, we write u_i to imply $u_i(p_i, \hat{p})$ where the position of the candidate c_i and the position favored by the agent is clear from the context. Throughout this chapter, we will refer to the social welfare of the elected outcome. If \hat{p} is the position of the elected candidate, the social welfare $SW(V)$ is the sum of the utilities for all voters for that outcome:

$$SW(V) = \sum_i u_i(p_i, \hat{p}).$$

²Cardinal utility models are used commonly in the literature, for example in Random Utility Theory [12].

5.1.2 Response Model

First, let us consider a rational voter v in a complete information setting, where a voter observes all of the ballots from the rest of the population. Let $\mathbf{n}_{-i} = \{n_i\}$ be the vector enumerating the ballots supporting each candidate $i \in C$. If b represents the ballot our voter submits, let $\mathcal{F}(\mathbf{n}_{-i}, b)$ denote the winner of the resulting election. Note that in a vast majority of cases, it does not matter what v votes at all! For instance, let us assume we break ties randomly. If two candidates i and j are locked in a close race for first place, with $|n_i - n_j| \leq 1$, then v has the opportunity to cast a *pivotal* vote to break the tie in favor of the preferred option. This condition is called a *pivot* condition, and it is easy to see that such scenarios are very uncommon. Outside of a pivot condition, it does not actually matter who v votes for at all! It is, however, unrealistic to assume voters act only when they are pivotal. Indeed, turnout of elections would be very low if this is the case. To improve this model, we do away with the assumption of complete information. This section focuses on modeling how a rational voter processes *partial* information, and formulates a strategic ballot in response. Moreover, the voter implements a local response model that strategizes using only locally available information in her social network.

We begin by making a natural assumption that each voter assumes her friends are representative of the wider population. If a ballot b is observed in a fraction f of her friends, then she assumes any voter within the network will submit ballot b with probability f .

We formally specify the response model for plurality voting for simplicity, but it can be adapted to any voting system with finite $|\mathcal{B}|$. This means each ballot is an individual candidate, and $\mathcal{B} = C$. Let (s_1, s_2, \dots, s_m) represent the number of voters in $\mathcal{N}(i)$ voting for candidates (c_1, c_2, \dots, c_m) . Voter i will then assume each voter (other than herself) in the network will support candidate c_x with probability $\frac{s_x+1}{S}$, where S is a normalizing constant to make the probabilities sum to 1. The $+1$ is a Laplace smoothing, and is necessary to ensure that all ballots remain possible. This means the ballots from the rest of the electorate follow a multinomial distribution with support $\mathbf{s} = (\frac{s_1+1}{S}, \frac{s_2+1}{S}, \dots, \frac{s_m+1}{S})$, $S = |\mathcal{N}(i)| + m$.

We can calculate the probability of any outcome of the election by using the multinomial distribution. Let the vector $\mathbf{b} = (b_1, b_2, \dots, b_m)$ denote the outcome where the remaining $n - 1$ voters in the network contribute b_i ballots supporting candidate c_i . The probability of this outcome is calculated as follows:

$$Pr(\mathbf{b}; n - 1; \mathbf{s}) = \frac{(n - 1)!}{b_1! b_2! \dots b_m!} \frac{\prod_{i=1}^m (s_i + 1)^{b_i}}{S^{n-1}}$$

With complete information, a rational voter only profits from casting a ballot when it is pivotal. With incomplete information, however, our voter must calculate the probability of each winning tie, and cast a ballot that, in expectation, will break ties to maximize her utility. For simplicity, we assume that winning ties between 3 or more candidates are such remote possibilities that they functionally have probability zero. Then, let $T(y, x)$ be the probability of a winning tie between candidates x and y , calculated by enumerating all possible such ties and summing their probabilities. Additionally, we also consider all near-ties, where the addition of one vote to candidate x will cause a winning tie with y ; let $\tilde{T}(y, x)$ be the probability of this outcome.

Finally, voter i revises her ballot to support the candidate x with the maximal marginal gain in expected utility \mathcal{C}_x , calculated below. If a voter observes no other ballots (i.e. $\mathcal{N}(i) = \emptyset$), her ballot remains fixed.

We consider two tie-breaking rules: probabilistic and lexicographic tie-breaking. Below is a modification of prospective ratings introduced by Myerson and Weber [100], for unbiased probabilistic tie-breaking and risk-neutral voters:

$$\mathcal{C}_x = \sum_{y=1}^m \left(\frac{1}{2} T(y, x)(u_x - u_y) + \frac{1}{2} \tilde{T}(y, x)(u_x - u_y) \right)$$

The first term in the summation calculates the expected utility where the voter breaks a winning tie between x and y in favor of x . Since we break ties randomly, without the additional vote, x would win half the time; the additional ballot changes the outcome so that x wins with certainty. This is the reason behind the $\frac{1}{2}$ coefficient for this term. Similarly, the second term in the summation calculates the expected utility where the voter casts a ballot that causes a winning tie between x and y , which allows x to win half the time; without this ballot, y would win with certainty.

An analogous modification exists for lexicographic tie-breaking, where $\mathbf{1}_{x < y}$ is an indicator variable with $\mathbf{1}_{x < y} = 1$ when x lexicographically precedes y , and 0 otherwise:

$$\mathcal{C}_x = \sum_{y=1}^m \left(\mathbf{1}_{x > y} T(y, x)(u_x - u_y) + \mathbf{1}_{x < y} \tilde{T}(y, x)(u_x - u_y) \right)$$

5.1.3 Sequential vs Simultaneous Updates

We consider two methods for scheduling when opinion updates take place: sequential and simultaneous. In sequential updates, voters are updated one at a time in a fixed order in each round, and they observe the most up-to-date ballots of their neighbors (which may be updated earlier in the current round, or in the previous round). In contrast, in simultaneous updates, all voters respond simultaneously to observed ballots from the previous round.

5.2 Experimental Design

Our investigation will focus only on the plurality voting rule. We first investigate the effects of the two tie-breaking schemes and update methods. As with Clough’s investigation [33] (see Section 3.3.2), we initialize a population of 169 voters in the baseline graph models: ER and BA. For tractability, we limit ourselves to 3- and 4-candidate scenarios. The positions of candidates and voters are drawn independently, uniformly at random from the interval [0,100]. The parameters of the graph models are chosen so that the resulting conditions have average out-degree approximately 8, 12, 16, 20, 24, and 28.³

In our second set of experiments, we investigate the effects of graph structure and homophily on the behavior of voters and the social welfare of the selected outcome. We focus the experiment on sequential updates and lexicographic tie-breaking, but extend the conditions to include all four graph models. Once again, parameters are chosen to produce the same set of average out-degrees.

The simulation is written in the D programming language, and compiled using DMD32 D Compiler v2.067.1 on a 64-bit Windows 7 machine. We limit each election to a maximum of 25 rounds (i.e. 25 updates for each voter), though this limit is never reached. Each data point in the first set of experiments is the average of 500 replications; each data point in the second set is the average of 800 replications.

5.3 Results

We define several metrics measured across our experiments. The quality of the elected candidate is measured as a ratio of social welfare scores. This allows us to normalize out some of the noise that is introduced by randomizing voter preferences and candidate positions.

First, we define the **Price of Honesty** (PoH) as the ratio of social welfare of the truthful outcome to that of the strategic outcome.⁴ Since both utility values are negative, the larger the PoH, the more costly the truthful outcome is, relative to the strategic outcome. In other words, a large PoH indicates that a population that voted truthfully produces an outcomes that is much worse off than if they had voted strategically. A PoH of less than

³The connection probabilities used for ER graphs were {0.045, 0.07, 0.095, 0.12, 0.145, 0.17}. The attachment factors used for BA graphs were {4, 6, 8, 10, 12, 14}.

⁴There are various names given for this metric: for example, “improvement in social welfare over truthful” [94], and “dynamic price of anarchy” [27].

1 indicates the population would have been better off if they had voted truthfully instead of strategically.

Similarly, we define the **Price of Stability** (PoS) is the ratio of social welfare of the strategic outcome to that of the optimal outcome.⁵ The optimal outcome is the outcome that yields the highest social welfare, and so the PoS has a minimal value of 1 (when the strategic outcome is always the optimal outcome). The larger the PoS, the worse the strategic outcome is, by comparison.

We also measure the percentage of voters that engage in strategic play – i.e. the fraction of voters who converge to a ballot that is not truthful – as well as the average number of updates required to reach stability.

Table 5.1 summarizes these four metrics measured on ER and BA graphs ($m = 4$). Within each graph type, there is little change in the amount of strategization, PoH nor PoS across the three conditions. Despite reaching a similar amount of strategization, simultaneous updates requires a larger number of updates to reach stability. By comparison, the differences between strategization, PoH and PoS is much larger between the two graph types. The same pattern appears in each of the other conditions. We conclude that neither the update methods nor tie-breaking mechanism has a significant impact on the behavior of the voters or the result of the voting process.

Next, we move to the second series of experiments, and the central findings of the chapter. We compare the four aforementioned metrics across the four graph models. Strategization is a major concern in elections using the plurality system. However, we show in our experiments that it actually improves the overall social welfare of the elected outcome. Throughout our experiment (> 4800 total trials), we found that the average PoH for each condition is greater than 1 in each condition; that is, in expectation, the candidate selected by strategic voting achieves a higher social welfare than that selected by truthful voting.

As one might expect, the amount of strategic play increases as voters gain access to more information as connectivity increases (see Figure 5.1). However, this gain is asymptotic and the ceiling of strategic play is reached relatively quickly. Interestingly, the ceiling is lower in graphs with homophily than those without. The rate at which strategic play increases (with edge density) is dependent on the graph type, with ER graphs reaching saturation more quickly than BA graphs.

Figure 5.2 shows the Price of Honesty and the Price of Stability under the different graph models. We include only $m = 4$ plots, but the same qualitative trends occur for

⁵Since the voter response is deterministic, we may view the outcome of the strategic voting process as unique, and this definition parallels the usual definition of Price of Stability or Price of Anarchy. If viewed as an online algorithm, this measure is analogous to the competitive ratio.

Update/Tie	% Strat	Updates	Avg PoH	Avg PoS
Erdős-Renyi Random Graph				
seq / lex	0.268 (0.0230)	57.9 (5.58)	1.2312 (0.0127)	1.0795 (0.0199)
seq / prob	0.267 (0.0255)	57.2 (5.80)	1.2287 (0.0128)	1.0776 (0.00915)
sim / lex	0.277 (0.0200)	78.6 (9.44)	1.2523 (0.0317)	1.0689 (0.0152)
Barabási-Albert Random Graph				
seq / lex	0.243 (0.0509)	44.9 (9.95)	1.2128 (0.0216)	1.0866 (0.0188)
seq / prob	0.235 (0.0519)	43.3 (9.93)	1.2133 (0.0157)	1.1002 (0.0171)
sim / lex	0.252 (0.0441)	68.6 (13.5)	1.2149 (0.0272)	1.0885 (0.0252)

Table 5.1: Effects of update and tie-breaking methods (ER and BA graphs with $m = 4$). The metrics measured are the percentage of agents casting strategic ballots, the number of updates before convergence, the Price of Honesty and the Price of Stability. The standard deviation is included in parentheses.

$m = 3$. Here we see a possible explanation for the lower strategic ceiling observed in homophilic graphs: it is simply less profitable. The PoH is consistently lower than PoS in these graphs, though they begin to converge at higher edge densities. That is, in these graphs, the social welfare of the strategic outcome is closer to that of the truthful outcome than the optimal outcome.

As strategization occurs in plurality elections, voters begin to abandon less promising candidates for the likely winners, even if they are less preferable. The net result of this behavior is that a multi-party system using the plurality rule will eventually devolve into a race between the two front running candidates. This tendency of plurality favoring 2-party systems is observable in electoral systems around the world, and is known in political science as Duverger’s Law [50].⁶

The consistency of Duverger’s Law is measured by the SF Ratio: the ratio of support for the third and second place candidates [40].⁷ Complete agreement with Duverger’s Law would mean no voters will “waste” their votes on lower ranking candidates, and will only cast their ballots in favor of the two leading candidates. This would be reflected by an SF Ratio of 0. Figure 5.3 shows the distribution of SF Ratios under different graph models, at the condition with the lowest edge density conditions (average out-degree 8). Duverger’s Law would predict that the distribution of SF Ratios be concentrated as a sharp peak near 0. In both 3- and 4- candidate elections, there is little agreement to Duverger’s Law in

⁶Canada and India are notable exceptions to this rule.

⁷The term SF Ratio refers to the second and first runner-up candidates.

most graphs, with fewer than 50% of the instances exhibiting an SR Ratio of less than 0.1 (i.e. the third place candidate enjoy less than 10% of the support of the second place candidate). If hER graphs are excluded, at least 50% account for those instances with SF Ratio of at least 0.2. It is interesting to note that in both 3- and 4- candidate elections, hER graphs stand out as showing the most agreement to Duverger’s Law. Notably, the dominant feature of these graphs is homophily, suggesting it helps voters enact Duverger’s Law, even when little information is available to an individual voter.

Figure 5.4 is a histogram showing the distributions of SF Ratios for ER and hER models, for the three lowest connectivity settings. The bars in blue represents the same data as presented in Figure 5.3, which is gathered at the lowest connectivity setting (with average out-degree 8). The orange bars shows the distribution of SF Ratios in graphs with average out-degree 12. Here, it is clear that the distribution peaks at 0, and Duverger’s Law is rapidly being restored due to an increase of information available to individual voters. In approximately 65% of the hER instances, the SF Ratio is below 0.1; in the ER graphs, the percentage increases to 80%. The trend continues as we increase the connectivity, as shown in the average out-degree 16 condition (shown as gray bars).

5.4 Convergence

In our empirical simulation, all trials converge to stability, and do so quickly. It is natural to ask whether the response model is guaranteed to reach an equilibrium in either the sequential or the simultaneous settings. Figure 5.5 sketches an undirected social network with preferences such that the voter responses result in a cycling of ballots. In this network, there are three candidates, denoted A , B , and C . The vertices of the graph are divided into four groups, labelled V_1 , V_2 , A and B . A and B are cliques on n' vertices; all voters in A have candidate A as their top preference, and correspondingly with B , for candidate B . V_1 contains n' vertices; each has preference $A \succ B \succ C$, and is connected to every vertex in B and V_2 , but not to each other. Similarly, V_2 contains n' vertices; each has preference $C \succ A \succ B$, and is connected to every vertex in A and V_1 , but not to each other. n' may be some large number, such as 10.

It is easy to see that there exist positions for the candidates such that none of the vertices in A or B will change their ballots. Each sees strong support for her favorite candidate, which ensures the most likely winning ties will involve that candidate.

Let us consider the sequential update process that updates the vertices of V_1 before V_2 . Each agent votes truthfully in the first round. In the second round, each vertex in V_1

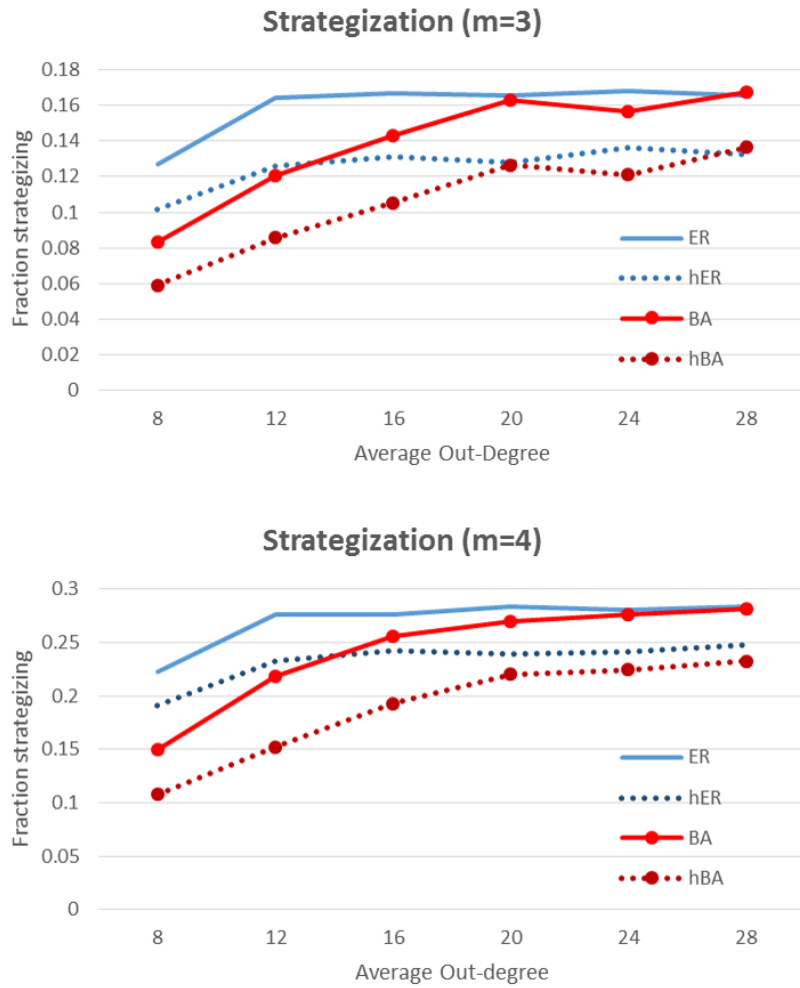


Figure 5.1: Fraction of agents strategizing (3- and 4-candidates). Note the different scales in the vertical axis.

sees n' supporters for B and C , and infers that the outcome will be a likely tie between those two candidates; each vertex switches support to their second-choice B . Each vertex in V_2 then observes a tie between A and B , and also switches to their second-choice: A . In the third round, each vertex in V_1 observes a tie between B and A , and therefore reverts to their truthful ballot in support of A . And finally, each vertex in V_2 now observes only support for their second place candidate A ; since they are unlikely to affect the election, they revert to their truthful vote of C . With all vertices returned to their initial, truthful

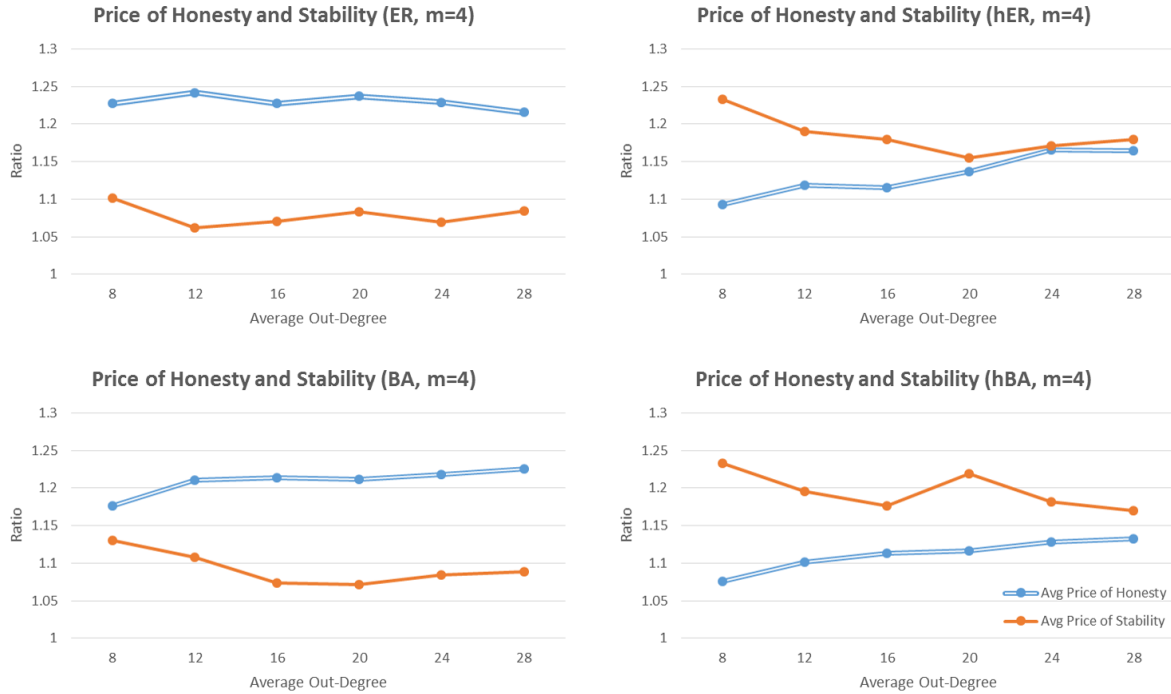


Figure 5.2: Price of Honesty and Stability under various conditions. Mann-Whitney $U < 303,000$, $n_1 = n_2 = 800$, $P < 0.01$, one-tailed, for all conditions in Figure 5.2, with two exceptions: ER (avg out-degree 8), and BA (avg out-degree 12). We obtain similar results of statistical significance on $m = 3$ conditions.

ballot, the cycle begins anew.⁸

The same counterexample works for the simultaneous update process, with V_1 and V_2 changing in alternate rounds.

Contrast this result with convergence results in the related model of Iterative Voting. By comparison, Iterative Voting occurs in the absence of a social network, where all ballots are common knowledge. Voters iteratively revise their ballots only if it alters the outcome to their benefit. Meir et al. [95] showed that Iterative Voting converges under plurality when voters respond one-at-a-time, but not when they update simultaneously. Lev and

⁸A number of positions for our candidates and voter blocs will produce this behavior. For example, consider candidates A , B , and C having positions 10, 9, 12 respectively. Let blocs B and V_1 prefer position 10 (therefore prefers candidates $A \succ B \succ C$), and A and V_2 prefer position 12 (prefers candidates $C \succ A \succ B$).

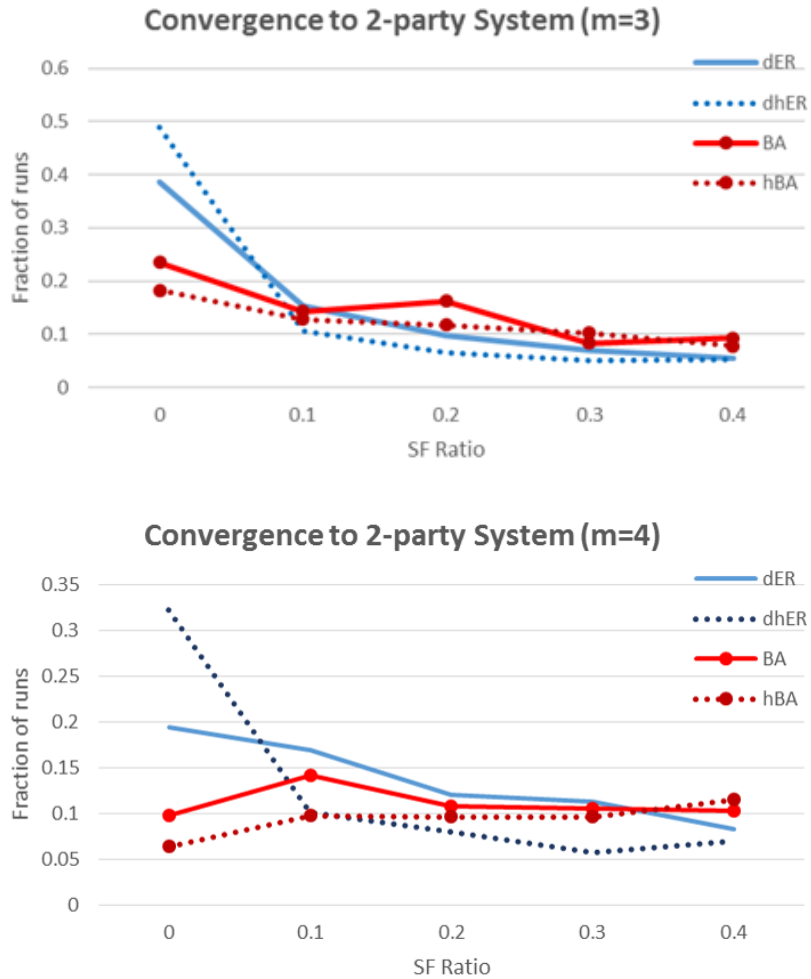


Figure 5.3: Degree of convergence to a 2-candidate system, measured as the SF ratio (3- and 4-candidates). Note the different scales in the vertical axis.

Rosenschein [84] demonstrated a similar result for veto, and showed that there is no guarantee of convergence in other scoring rules. Our result generalizes Iterative Voting to social networks, where the complete information aspect of Iterative Voting may be represented by embedding voters in a complete graph. We show that their results do not hold in the more general network setting, and that voting dynamics on a network offer a richer space of interactions between voters.

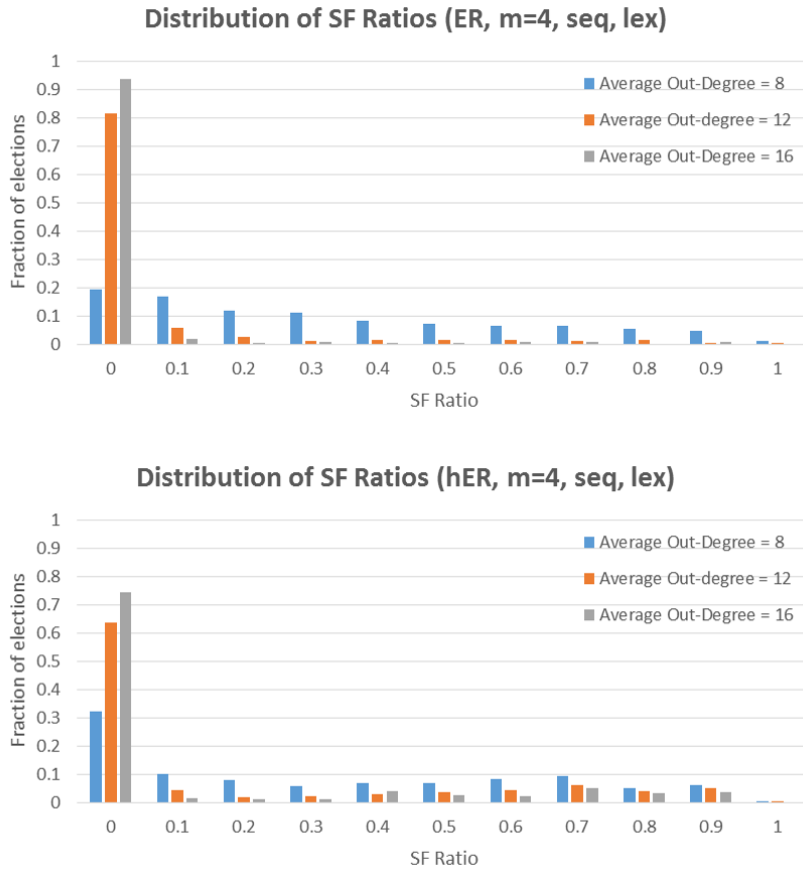


Figure 5.4: Distribution of SF Ratios show the degree with which results from each graph model conform to Duverger’s Law.

5.5 Discussion

While we have obtained empirical results for our model, the question remains as to how well it generalizes to real world scenarios. As was alluded to in the Introduction, the social network we depict with our model is a general social network. The neighbors in the network describe not merely “Facebook friends”, but include all sources of information that may be considered by a voter in deciding on her ballot. This may include close friends, trusted confidants and knowledgeable associates, but will also news feeds, political blogs, and subscriptions to any number of popular media outlets. Such institutions acts as highly-connected nodes in the social network, much like hubs in Barabási-Albert random graphs.

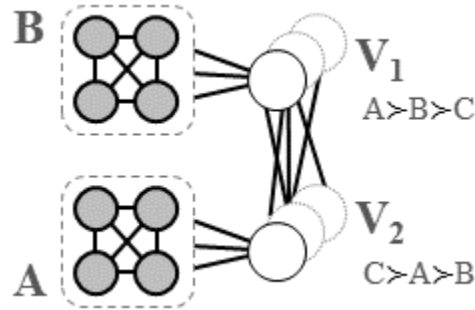


Figure 5.5: Voters need not converge to stability.

Further, as is shown in Hargittai, Gallo and Kane [67], even such social institutions are not immune to the same homophily exhibited in people.

The successive rounds of voter revision in our model represents the preliminary period preceding an election where voters may discuss and revise their opinions. In the real world, this is often accompanied by a series of preliminary polls leading up to the main election. These polls can be a major factor in strategic voting. Such polls are comprised of (tentative) ballots sampled from a random subset of the population. This is exactly the relationship captured by the (non-homophilic) Erdős-Renyi random graphs, where each voter may view the ballots of a number of other voters sampled uniformly, independently at random from the population.

With homophily being such an intrinsic property of real world networks, it is interesting to note that the graph depicted in Figure 5.5 shows a very low degree of homophily (for vertices in V_1 and V_2). This lack of homophily is necessary for the counterexample to function. Voters that are connected to likeminded voters are less likely to change their votes away from their truthful ballot. They observe many other voters declaring the same ballot, and therefore their favorite candidate is very likely to participate in winning ties. In fact, a careful analysis of the graph structure of Figure 5.5 reveals what is needed to cause a faithful voter to vote strategically: they view their own position as hopeless, and must be convinced to “pitch in” to resolve a close race between two less-favored candidates. This is in agreement with observations of political elections, such as the empirical study conducted by Cain [28].

The existence of this counterexample also gives some insights on voter sincerity, even in preliminary polls. A naive voter may believe that showing support for her favorite candidate in the polls will improve public perception of that candidate. The scenario illustrated in Figure 5.5 provides a scenario where that is not true. When voters in the V_1

and V_2 change their strategic ballots to a truthful state to support their favorite candidate, they actually cause a chain reaction that erodes support for that candidate from other voters.

Nonetheless, these results shed some light on to why there is less strategic voting in the presence of homophily, and also why the strategic outcome is (comparatively) less profitable. When voters are surrounded with those of similar opinions, it creates an Echo Chamber Effect where they view their own position as being more widely supported than it is. It causes them to be further entrenched in their current position, and they require a larger amount of conflicting evidence to change their minds. The effect causes a voter to have a harder time discerning whether their own position is in the minority, and prevents them from shifting to a more strategic choice. This, as it turns out, has a net negative effect on the social welfare of the elected outcome. Moreover, this effect may explain the relatively small number of strategic voters observed in real world elections (for example, in [20] and [61]): it is not that few voters are strategic, it may be that many voters fail to recognize the strategic opportunity due to their Echo Chamber.

5.6 Conclusion

In this chapter, we proposed a model of strategic voting on social networks, based on a natural assumption on the part of voters that their friends are representative of the population. We show that strategization leads to improved social welfare of the elected outcome in all conditions. Network structure has an effect on the social welfare of the elected outcome. However, as we observe in Figure 5.1, as edge density increases, the amount of information available to each voter also increases, and the number of strategic voters quickly saturate at a ceiling. The ceiling is independent of graph structure, but highly dependent on homophily.

It is this network homophily that causes the Echo Chamber Effect. This may offer insight on why a relatively low number of voters are strategic in real world elections. When surrounded by others with similar opinions, voters do not see an opportunity or even a need to strategize, even when their position holds little merit. This ends up hurting social welfare of the elected result.

Clough has also explored this form of strategic interaction between voters in simulated social networks [33, 34]. Coming from a political science angle, she analyzes simple 13×13 grid-based undirected graphs on 169 nodes, which is neither small-world nor scale-free. In her model, each voter responds by considering only tie-probabilities, while our model

considers all pivotal cases under different tie-breaking rules. Her work focuses solely on investigating Duverger’s Law. Her findings parallel ours: SF Ratios drop dramatically when going from 28 to 8 neighbors. Unfortunately, her model does not offer any finer levels of granularity for investigating this behavior.

Iterative voting has also been applied to social networks only very recently in Sina et al. [121], which focuses on manipulation by a chair under plurality voting. Our model differs from Sina et al.’s in that our voters individually infer the likely outcomes of the election based on their limited information, and always act upon this information (to maximize their expected utility based on tie-probabilities). By contrast, in the Iterative Voting model applied by Sina et al., agents only choose to revise their vote when they observe an exact pivotal condition in their neighborhood.

As Figure 5.5 demonstrates, our model is not guaranteed to converge to stability. However, stability is reached relatively quickly in practice. In our simulations, no instances used more than 10 rounds to reach stability. It is unclear why this is the case, and may be a direction for future work. Are such cyclic instances rare? Under what conditions can we guarantee stability? Are such conditions natural to human networks?

Another natural question to ask is, how susceptible to manipulation are voters on a social network? Will voter strategization hinder or amplify the effects of manipulations? If candidates have knowledge of the social network, what strategies may they take to improve their own odds?

Finally, it would be interesting to extend this framework to other, more interesting voting systems. Duverger’s Law applies only to plurality, so we expect to see less convergence to 2-party systems when using other voting rules. What effects will this have on strategization and social welfare? Tie-probability modeling for other systems remains an exciting open question.

Chapter 6

Heuristic Voter Models

In the previous chapter, we introduced a basic model for how voting may occur in a social network. Voters observe how their neighbors vote, and use that information to infer the likely outcomes of the final election. Of particular interest to rational voters are the pivotal outcomes; these are the only situations where a single voter's vote will alter the outcome of the election. Weighted by the likelihood of these pivotal outcomes, voters may calculate the expected utilities of casting different ballots, and thereby determine the best course of action.

While this voter model is fully rational and comprehensive, the calculation of pivot probabilities is computationally intensive – it scales poorly in larger populations and elections with many candidates; moreover, it places large cognitive burdens on the voters. In this chapter, we propose a number of heuristic models that greatly simplifies these pivot calculations. We argue that these heuristics represent natural models of boundedly rational human behavior. Simultaneously, our heuristics speed up the computation of strategic response by up to 2 orders of magnitude, allowing us to explore the strategic behavior of voters in larger populations. In particular, we examine the Micromega rule, the tendency for large political parties to favor small assemblies with large electoral districts, and vice versa. We show that population is a contributing factor to the Micromega rule in some networks.

6.1 Framework

To establish a framework for voter behavior for elections with large populations, we consider several desirable criteria for these models. We base our framework on the desiderata

presented by Meir, Lev, and Rosenschein [94], and adapt them to our domain. While their framework describes desirable criteria when modeling voters in general populations, our framework focuses on modeling voters in large populations embedded in social networks. The social network naturally restricts the availability of information to voters, and does so in an asymmetric way; one voter has different information about the election than another. Moreover, we emphasize that the voters in our framework are boundedly rational, and therefore computationally limited. This both reflects the human nature of real world voters, and also ensures our heuristics can be computed in a timely manner.

Knowledge: A voter’s knowledge of the actions of others is limited. In particular, voters are limited to what can be observed from their neighbors in the social network. Voters must infer the current state of the world based on this limited information.

Rationality: Subject to their observations, preferences, and beliefs, voters act to maximize their expected utility of the electoral outcome. In particular, while the chances of casting a pivotal vote in an election is very small, it is the only event of importance to rational voters. Their observations allow them to compare the likelihood of pivot conditions between different candidates and act accordingly.

Anonymity: Beyond readily available network properties, voters treat observations from their social contacts anonymously. Beyond the utility gained through the election of each candidate, candidates are treated anonymously as well.

Equilibrium: The model converges to an equilibrium outcome (according to some established solution concept), or readily shows it cannot exist.

Tractability: The computation of voter responses is computationally tractable for the voter. Real world voters are boundedly rational agents and frequently employ heuristics to simplify their cognitive load. While the computation or approximation of probabilities may be unavoidable, this computation should be fast, particularly for “easy” cases.

Optimistic: Voters act in the belief that their action may have an impact, even when that is not guaranteed. This is in sharp contrast to the complete information setting of Iterative Voting, where voters act strategically *only* when they know they are pivotal.

Motivated by these desiderata, we make the following assumptions in our heuristic models. We do not consider these desirable criteria; rather, we consider these to be natural ways of implementing the above desiderata in our voter heuristics.

Markovian Strategy: While voters have access to histories of past actions from their social contacts, we assume voter response is Markovian and computed as a function of current observations. While making use of past history may allow the voter to detect patterns and trends, doing so is computationally intensive and further compounds concerns of

tractability. We argue that this simplifying assumption reasonably models human behavior because humans are bounded rational agents and human memory is a limited resource.

Myopic Response: Being boundedly rational, voters are not concerned about second order effects in the network. That is, they do not consider that an adjustment in their ballot may also cause others to adjust their ballots, and this knock-on effect may be detrimental to the original voter. Instead, we assume voter responses are myopic improvements to the current situation. This is a reasonable assumption to adopt because predicting these knock-on effects will be computationally intensive. However, adopting this assumption may actually make equilibria more difficult to achieve. For instance, non-myopic agents may be able to predict actions to lead to cycling behaviors within the population, which prevent the convergence to an equilibrium outcome.

6.1.1 Voting Model

We consider the spatial model of voters and candidates proposed by Tsang and Larson [129] and described in Chapter 5. Let the population V of n voters be situated in a social network represented by a simple, directed graph $G = (V, E)$. A directed edge $(i, j) \in E$ means voter i observes voter j , and therefore, j 's actions may influence i . Let $C = \{1, 2, \dots, m\}$ denote the set of candidates. Each voter i has a preference p_i represented as an integer from the interval $[0, 100]$; each candidate supports a position drawn from the same space. If the winning candidate supports position \hat{p} , voter i derives utility $u_i(p_i, \hat{p}) = -|p_i - \hat{p}|^2$ from this outcome.

Each voter casts a ballot from the set of admissible ballots \mathcal{B} . A social choice function \mathcal{F} maps the set of submitted ballots to a unique winner from C . We focus on the plurality voting rule, where $\mathcal{B} = C$; i.e. voters mark their unique favorite candidate on their ballots. If n_i is the number of ballots supporting candidate i , then the plurality voting rule $\mathcal{F} = \arg \max_x n_x$ maps the set of ballots to the candidate receiving the most votes, breaking ties randomly.

Voting proceeds in rounds. In the first round, voters' ballots reflect their sincere top choice. In subsequent rounds, one by one, voters are allowed to revise their ballots based on observing the ballots of their out-neighbors. For a particular voter i , let the vector $\mathbf{s} = (s_1, s_2, \dots, s_m)$ denote the fraction of i 's out-neighbors supporting each candidate, with Laplace smoothing applied by adding one to the tally of support for each candidate. Each voter i computes her revised ballot $b' \in \mathcal{B}$ according to her Voter Response Function $R_i(\mathbf{s}, u_i) = b'$. We assume voter behavior is symmetric and so omit the subscript for simplicity. We also omit the parameter u_i when it is clear from context. Voting terminates

when no voter's ballot changes in a round; we say that our population has converged to an equilibrium.

6.1.2 Fully Rational Voter

The fully rational voter computes the exact pivot probability for each pair of candidates by assuming future ballots will be distributed according to a multinomial distribution with support \mathbf{s} . The probability of observing a final tally of $\mathbf{b} = (b_1, b_2, \dots, b_m)$ is

$$Pr(\mathbf{b}; n - 1; \mathbf{s}) = \frac{(n - 1)!}{b_1! b_2! \dots b_m!} \prod_{i=1}^m s_i^{b_i} \quad (6.1)$$

The voter then computes the probability $T(y, x)$ that the any two given candidates x and y are in a pivot condition. That is, by adding one ballot supporting x , the winner changes from y to x . This is calculated by enumerating all possible such pivot outcomes, and summing the probability of each outcome. Formally, we define $T(y, x)$ as follows:

$$T(y, x) = \sum_{\substack{\mathbf{b} \in \mathcal{B}^{n-1} \\ b_x = b_y > \max\{b_i : i \neq x, y\}}} Pr(\mathbf{b}, n - 1, \mathbf{s}) \quad (6.2)$$

For example, under lexicographic tie breaking, and ignoring multi-way ties, the pivot outcomes for candidates 1 and 2, when $n = 10$, $m = 4$, are $(5, 5, 0, 0)$, $(4, 4, 2, 0)$, $(4, 4, 1, 1)$, $(4, 4, 0, 2)$, and $(3, 3, 2, 2)$; each would be associated with a probability of occurrence based on the multinomial distribution. Then, she calculates the Prospect Rating for that candidate:

$$C_x = \sum_{y=1}^m T(x, y)(u_x - u_y) \quad (6.3)$$

The Voter Response Function is $R_{\text{FULL}}(\mathbf{s}) = \arg \max_x C_x$. We will refer to this model as the **Full Voter** model.

Implementation

To compute the Prospect Rating according to Equation 6.3, we must iterate over all pivot conditions indicated in the summation in Equation 6.2. In this section, we outline the

algorithm which efficiently enumerates all such pivot conditions. We begin by introducing a simple counting problem which we call the BUS PACKING PROBLEM with positive integer parameters (n, m, k) :

Given $n > 0$ passengers to be seated on $m > 0$ buses, each of capacity no more than $k > 0$. How many different ways are there to assign seats to the buses? We may assume that the seats are interchangeable, but the buses are not.

For example, a valid packing of $n = 10$ seats to $m = 3$ buses each with capacity $k = 5$, is $(5, 5, 0)$; a distinct packing would be $(5, 0, 5)$. Each counts as a single bus packing, and we do not concern ourselves with the $10!$ permutations of passengers to the 10 seats onboard the buses. Note that the Bus Packing problem may return with no valid configuration, as evidence by the parameters $n = 10, m = 2, k = 4$ (there is not enough capacity in the bus fleet to reserve that many seats). Let the Bus Packing Number $BP(n, m, k)$ be the number of such bus packings.

We implement a recursive algorithm that efficiently enumerates all valid bus packings $\{s_1, s_2, \dots, s_m\}$ for given parameters $n, m, k > 0$. At each iteration, the algorithm considers packing $s_1 = 0, 1, \dots, \min(n, k)$ seats in the first bus, and then recursively solves the problem of packing $n - s_1$ seats in the remaining $m - 1$ buses.

The case for $m = 1$ is solved as the base case (there is only one way to pack the bus), and the algorithm terminates early with an empty list if the fleet can no longer fit the requested number of seats ($m \cdot k < n$).

To implement the Full Voter Model, we must enumerate over all outcomes of the election where exactly two candidates are tied as the winner, as depicted in Equation 6.2. We make use of the BUS PACKING algorithm outlined above to achieve this. For simplicity, we assume candidates c_0 and c_1 are the winners. First, we iterate over all possible values that the co-winners may achieve $s_0 = \lceil \frac{n}{m} + 1 \rceil \dots \lfloor \frac{n}{2} \rfloor$. s_0 is upper-bounded by the scenario where all votes are split equally between c_0 and c_1 . The lower-bound occurs when votes are nearly evenly split between all candidates, with c_0 and c_1 having a slight lead; the ceiling function ensures there are few enough votes between the remaining candidates such that there are solutions where each of them receive strictly fewer than s_0 votes (to prevent a three-way tie).

For each value of $s_0 = s_1$, we must find all possible partitions of the remaining $n - 2s_0$ votes between the remaining $m - 2$ candidates such that they receive $s_0 - 1$ or fewer votes each. This exactly corresponds with the BUS PACKING PROBLEM with parameters

$(n-2s_0, m-2, s_0-1)$. Once we obtain this solution, we may simply permute the co-winners through the other candidates in $\binom{m}{2}$ combinations.

To calculate the pivot probability, we iterate over all pivot outcomes and take the sum of the probability of each outcome, according to the multinomial distribution.

6.1.3 Voter Heuristics

As we illustrate above, the Full voter model must enumerate all possible pairwise pivot conditions, calculating R_{FULL} is computationally intensive, and scales poorly as n or m increases. We can upper-bound the computation complexity of this process by examining a related counting problem, the classic STARS AND BARS PROBLEM by William Feller [60]:

Given positive integers n and m , the number of distinct m -tuples of *non-negative* integers that sum to n is given by the multiset function $\binom{n+1}{m-1}$.

Let us refer to this number as the Stars and Bars Number $SB(n, m)$. We arrive at this solution for $SB(n, m)$ by imagining n stars laid out along a line and we are invited to partition the stars into m parts by placing $m - 1$ bars as separators between the parts. There are $n + 1$ positions for the bars, and the positions may be selected multiple times – two adjacent bars indicating a part with zero stars. This gives us a solution of $\binom{n+1}{m-1}$ using the multiset function, which may also be expressed as $\binom{n+m-1}{n}$.

Notice that Feller’s STARS AND BARS PROBLEM is very similar to our BUS PACKING PROBLEM; the latter adds a restriction on the maximum number of stars that may be included in any given part. As a result, some solutions to the STARS AND BARS may be inadmissible as bus packings, but all bus packings are valid solutions to the STARS AND BARS PROBLEM. Therefore, we may use $SB(n, m)$ as an upper-bound to $BP(n, m, k)$.

We will upper-bound the runtime of a single best response calculation of the Full Voter Model by examining the number of times the probability of an outcome is queried through the multinomial distribution. Calculating the probability of an outcome based on a multinomial distribution is itself an expensive operation; in our implementation, it requires $m + 1$ invocations of the GAMMA function. For simplicity, our complexity analysis will be in terms of queries to the multinomial distribution, which allows us to encapsulate implementation details of the queries.

Let us denote this number, the number of queries to the multinomial distribution, as $\#_{\text{ties}}$. The algorithm iterates over the admissible range for the number of votes supporting

a tied winner $s_0 = \lceil \frac{n}{m} + 1 \rceil \dots \lfloor \frac{n}{2} \rfloor$. At each iteration, it will query the multinomial distribution a number of times equal to the Bus Packing Number $BP(n - 2s_0, m - 2, s_0 - 1)$, therefore the exact number of queries is given by the following expression:

$$\#_{ties} = \sum_{s_0=\lceil \frac{n}{m} + 1 \rceil}^{\lfloor \frac{n}{2} \rfloor} BP(n - 2s_0, m - 2, s_0 - 1)$$

Since $BP(n - 2s_0, m - 2, s_0 - 1) \leq SB(n - 2s_0, m - 2) = \frac{(n+m-2s_0-3)!}{(n-2s_0)!(m-3)!}$, we may rewrite

$$\#_{ties} \leq \sum_{s_0=\lceil \frac{n}{m} + 1 \rceil}^{\lfloor \frac{n}{2} \rfloor} \frac{(n + m - 2s_0 - 3)!}{(n - 2s_0)!(m - 3)!}$$

Moreover, the fraction is maximized when s_0 takes the minimal value of $\lceil \frac{n}{m} + 1 \rceil$, which allows us to upper-bound the terms of the summation independently of the index of the summation. This gives us the following upper-bound (since this is an approximation, we remove the floor and ceiling notation for clarity):

$$\begin{aligned} \#_{ties} &\leq \left(\frac{n}{2} - \frac{n}{m} - 1 \right) \frac{(n + m - \frac{2n}{m} - 5)!}{(n - \frac{2n}{m} - 2)!(m - 3)!} \\ &= \left(\frac{n}{2} - \frac{n}{m} - 1 \right) \frac{(n + m - \frac{2n}{m} - 5) \dots (n - \frac{2n}{m} - 1)}{(m - 3)!} \\ &< \left(\frac{n}{2} - \frac{n}{m} - 1 \right) \frac{(n + m - \frac{2n}{m} - 5)^{m-3}}{(m - 3)!} \end{aligned}$$

We may also use Stirling's approximation that $n! \approx \sqrt{2\pi n} \left(\frac{n}{e}\right)^n$ to simplify the denominator:

$$\#_{ties} < \left(\frac{n}{2} - \frac{n}{m} - 1 \right) \frac{(n + m - \frac{2n}{m} - 5)^{m-3}}{\sqrt{2\pi(m-3)} \left(\frac{m-3}{e}\right)^{m-3}}$$

Thus, we establish a rough upper-bound for the runtime of a single pivot probability computation between a pair of candidates in the Full Voter Model as $O(n(\frac{n}{m})^{m-3})$. Since for our purposes, we may assume $n \gg m$, we may simplify the runtime's upper-bound

as roughly $O(n^{m-2})$. There are also $\binom{m}{2}$ such pairs, so the runtime for the best response calculation requires roughly $O(m^2 n^{m-2})$ queries.

Empirical experimentation reinforces this bound. The Full Voter Model scales very poorly as either n or m increases. Since one of our desiderata is Tractability, both for the purpose of scalability, and to more accurately model human bounded rationality, we propose a number of voter heuristics that reduces the computational and cognitive load on the voters. In each of these models, the pivot probabilities are simplified to the chance that x and y are exactly tied as winners (i.e. discounting the cases where x has one fewer ballot than y).

Top-k Voter

An intuitive way of easing the voters' cognitive burden is for them to ignore unpromising candidates. This is observed in political science literature. For example, Meffert and Gschwend [91] conducted studies in the laboratory on strategic voting behavior in coalitional governments. They used fictional parties with monetary incentives based on the elected outcome, and found that participants used a number heuristics when considering their ballot. One of these heuristics was to avoid parties that did not enjoy enough popular support to contribute meaningfully to the result. We model this type of behavior directly with the Top- k voter. Here, voters consider only the $k \leq m$ candidates with the most popular support according to \mathbf{s} , breaking ties in favor of utility of victory. The voter treats the election as if only these top k candidates were participating, and computes $R_{\text{Top-}k}$ based on $\binom{k}{2}$ pivot probabilities, rather than $\binom{m}{2}$. The resulting algorithm requires only $O(k^2 n^{k-2})$ queries to the multinomial distribution, though determining the top- k supporters and permuting the entries adds a small overhead to the computation that scales with m .

Max-M Voter

The Full Voter considers the expected utility gained by supporting a candidate over all other candidates. A boundedly rational voter may employ a different measure, supporting the candidate offering the maximum marginal gain over a rival candidate. That is, the voter focuses on pairwise comparisons between candidates, picking the candidate who offers the most compelling position, and has the best chance of beating the most serious threat. Formally, consider the following utility computation in place of prospect rating \mathcal{C}_x (Equation 6.3):

$$\mathcal{D}_x = \max_{y \neq x} T(y, x)(u_x - u_y)$$

$R_{\text{Max-M}}$ = $\arg \max_x \mathcal{D}_x$ selects the candidate maximizing this marginal utility over some other candidate. The motivation behind using this alternative utility function is that it approximates the behavior of the **Full Voter**, while offering mathematical optimizations that greatly reduce computation load. Such pairwise comparisons may also be more natural for human voters to process. Rather than considering the marginal gains over every other candidate, this calculation emphasize the candidate's merits against the most salient of opposition. Indeed, political campaigns often focus on demonizing particular opponents to bolster the merits of the favored candidate.

More formally, the voter's strategic response is simply to choose x satisfying the **Arg-Max Condition**:

$$\arg \max_x \mathcal{D}'_x = \arg \max_x \max_{y \neq x} T(y, x)(u_x - u_y) \quad (6.4)$$

In order for condition 6.4 to be satisfied for x , it is sufficient (but not necessary) that, for some choice of x and y

$$T(y, x)(u_x - u_y) \geq T(y', x')(u_{x'} - u_{y'}) \quad \forall x', y' \in C; y \neq x, x' \neq y' \quad (6.5)$$

Since $u_x, u_y, u_{x'}$ and $u_{y'}$ are all constants, let $H_{xy}^{x'y'} = \frac{u_{x'} - u_{y'}}{u_x - u_y}$. For simplicity, we assume $u_x - u_y > 0$,

$$T(y, x) \geq T(y', x') H_{xy}^{x'y'} \quad \forall x', y' \in C; y \neq x, x' \neq y' \quad (6.6)$$

We call this inequality the **Domination Condition**. We may omit the subscripts and superscripts from $H_{xy}^{x'y'}$, where it is clear from context.

First, note that since $T(y, x)$ sums over all possible tallies of $n - 1$ ballots where y and x are tied as winners, both $T(y, x)$ and $T(y', x')$ sum over the same domains. We may match each term in the summation of the LHS with a corresponding term in the RHS that has the *same multinomial coefficients* (with the elements of \mathbf{b} permuted). Specifically, suppose $\mathbf{b} = (\dots b_x, \dots b_y, \dots b_{x'}, \dots b_{y'})$ is a possible tally when calculating $T(y, x)$ (i.e. where x and y are in a winning tie). Then, it corresponds to a tally $\mathbf{b} = (\dots b_x, \dots b_y, \dots b_{x'} \dots b_{y'})$ where ¹

¹To be clear, it is implied by the notation that b_x and b_x have the same indices in their respective vectors, and so on.

$$\begin{aligned}
b_{x'} &= b_x \\
b_{y'} &= b_y \\
b_x &= b_{x'} \\
b_y &= b_{y'}
\end{aligned} \tag{6.7}$$

Moreover, since these situations involve winning ties,

$$\begin{aligned}
b_{x'} &= b_{y'} \\
b_x &= b_y
\end{aligned} \tag{6.8}$$

Matched this way, it is easy to see that, for Condition 6.6 to be true, it is sufficient (but not necessary) for each term on the LHS to be at least as large as the corresponding term in the RHS (multiplied by H):²

$$\prod_{i=1}^m s_i^{b_i} \geq H \prod_{i=1}^m s_i^{b_i} \tag{6.9}$$

And since $b_i = b_i$ for all $i \neq x, y, x', y'$, we may rewrite the inequality as

$$s_x^{b_x} s_y^{b_y} s_{x'}^{b_{x'}} s_{y'}^{b_{y'}} \geq H s_x^{b_x} s_y^{b_y} s_{x'}^{b_{x'}} s_{y'}^{b_{y'}} \tag{6.10}$$

And applying the mapping specified by the Equations 6.7 and 6.8,

$$s_x^{b_x} s_y^{b_y} s_{x'}^{b_{x'}} s_{y'}^{b_{y'}} \geq H s_{x'}^{b_{x'}} s_y^{b_y} s_x^{b_x} s_{y'}^{b_{y'}} \tag{6.11}$$

$$s_x^{b_x} s_y^{b_y} s_{x'}^{b_{x'}} s_{y'}^{b_{y'}} \geq H s_{x'}^{b_{x'}} s_y^{b_y} s_x^{b_x} s_{y'}^{b_{y'}} \tag{6.12}$$

Finally, we rearrange

²For simplicity of notation, for the remainder of this section, we have omitted the +1 Bayesian smoothing, and write s_x in place of $s_x + 1$.

$$\left(\frac{s_x s_y}{s_{x'} s_{y'}}\right)^{b_x} \left(\frac{s_{x'}}{s_x}\right)^{b_{x'}} \left(\frac{s_{y'}}{s_y}\right)^{b_{y'}} \geq H \quad (6.13)$$

$$\left(\frac{s_x}{s_{x'}}\right)^{b_x - b_{x'}} \left(\frac{s_y}{s_{y'}}\right)^{b_x - b_{y'}} \geq H \quad (6.14)$$

We want to establish a lower bound of the LHS, to check it against H as a quick check that bypasses the calculations. We can break this into a number of cases called **Bounding Conditions**, depending on the values of the two ratios $\frac{s_x}{s_{x'}}$ and $\frac{s_y}{s_{y'}}$. The truth of the Bounding Condition implies the truth of Equation (6.14) for all admissible values of x , x' , y , and y' .

Case A1: $\frac{s_x}{s_{x'}} \leq 1$ and $\frac{s_y}{s_{y'}} \leq 1$.

Then, to minimize the LHS, we need to maximize both expressions $(b_x - b_{x'})$ and $(b_x - b_{y'})$. This is done by setting both $b_{x'}$ and $b_{y'}$ to their minimum values of 0, and b_x to its maximum value of $\lfloor \frac{n}{2} \rfloor$. Therefore, LHS has a lower bound of

$$\left(\frac{s_x}{s_{x'}}\right)^{\lfloor \frac{n}{2} \rfloor} \left(\frac{s_y}{s_{y'}}\right)^{\lfloor \frac{n}{2} \rfloor} \geq H$$

Case A2: $\frac{s_x}{s_{x'}} > 1$ and $\frac{s_y}{s_{y'}} > 1$.

To minimize the LHS, we need to minimize both expressions $(b_x - b_{x'})$ and $(b_x - b_{y'})$. Since these parameters are discrete values, and b_x must be strictly greater than both $b_{x'}$ and $b_{y'}$, both expressions are minimized to a value of 1.

$$\left(\frac{s_x}{s_{x'}}\right) \left(\frac{s_y}{s_{y'}}\right) \geq H$$

Case A3: $\frac{s_x}{s_{x'}} \leq 1$, and $\frac{s_y}{s_{y'}} > 1$.

To minimize the LHS, we need to maximize $(b_x - b_{x'})$ and minimize $(b_x - b_{y'})$. As in the previous cases, the first expression can be maximized by setting $b_{x'} = 0$. Additionally, we want to maximize $b_{y'}$ subject to our choice of b_x . This can be done by splitting the votes strictly between the two tied winners (with b_x each), and the runner up (with $b_{y'}$). This arrangement roughly corresponds to the following equation (ignoring parity concerns):

$$2b_x + b_{y'} = n$$

Substituting this and $b_{x'} = 0$ into the LHS of (6.14), we get

$$\left(\frac{S_x}{S_{x'}}\right)^{b_x} \left(\frac{S_y}{S_{y'}}\right)^{(3b_x-n)}$$

$$\left(\left(\frac{S_x}{S_{x'}}\right)\left(\frac{S_y}{S_{y'}}\right)^3\right)^{b_x} \left(\frac{S_y}{S_{y'}}\right)^{-n}$$

It is clear to minimize this expression, if $\left(\frac{S_x}{S_{x'}}\right)\left(\frac{S_y}{S_{y'}}\right)^3 \leq 1$, we must maximize b_x by setting it to $\frac{n}{2}$, and $b_{y'} = 0$; this gives us

$$\left(\frac{S_x}{S_{x'}}\right)^{\frac{n}{2}} \left(\frac{S_y}{S_{y'}}\right)^{\frac{n}{2}} \geq H$$

On the other hand, if $\left(\frac{S_x}{S_{x'}}\right)\left(\frac{S_y}{S_{y'}}\right)^3 > 1$, then we want to minimize b_x by equally splitting the votes between the three contenders with approximately $\frac{n}{3}$ votes each. In fact, the minimum number of winning votes is calculated as $b_x = \lfloor \frac{n-6}{3} \rfloor + 3$ for our tied winners,³ with a gap of $b_x - b_{y'} = 3 - (n \bmod 3)$, for $n \geq 6$.⁴ Therefore, the LHS is minimized as

$$\left(\frac{S_x}{S_{x'}}\right)^{\lfloor \frac{n-6}{3} \rfloor + 3} \left(\frac{S_y}{S_{y'}}\right)^{3-(n \bmod 3)} \geq H$$

Case A4: $\frac{S_x}{S_{x'}} > 1$ and $\frac{S_y}{S_{y'}} \leq 1$.

Symmetrically to the previous case, the LHS is minimized according to two sub-cases. If $\left(\left(\frac{S_x}{S_{x'}}\right)^3\left(\frac{S_y}{S_{y'}}\right)\right) \leq 1$, we want to maximize b_x by setting $b_x = \frac{n}{2}$, and we get

$$\left(\frac{S_x}{S_{x'}}\right)^{\frac{n}{2}} \left(\frac{S_y}{S_{y'}}\right)^{\frac{n}{2}} \geq H$$

³For example, if $n = 6$, then the split would be $(3, 3, 0)$, because our model does not consider a three-way tie of $(2, 2, 2)$ to be possible; by the same argument, when $n = 9$, the split is $(4, 4, 1)$. This can be thought of as reserving 3 votes for each winner, and dividing the remaining votes evenly between the three candidates, with the third candidate getting the 0, 1 or 2 remaining votes; the 3 reserved votes ensures the first two candidates have a winning tie. If we allow three-way ties, then the expression becomes $\lceil \frac{n}{3} \rceil$, which is lower upper bound for the expression.

⁴The latter expression is the difference between the votes for the tied winners and the non-winner, which as described in the above footnote, can be 1, 2, or 3 ballots. $n \bmod 3$ calculates the number of ballots “left over” for the non-winning candidate; therefore, $3 - n \bmod 3$ is the size of the gap between the non-winner and the tied winners. If three-way ties are allowed, the expression would be $2 - (n + 1) \bmod 3$.

Symmetrically, if $((\frac{s_x}{s_{x'}})^3(\frac{s_y}{s_{y'}})) > 1$, then

$$\left(\frac{s_x}{s_{x'}}\right)^{3-(n \bmod 3)} \left(\frac{s_y}{s_{y'}}\right)^{\lfloor \frac{n-6}{3} \rfloor + 3} \geq H$$

Relaxing Assumptions

Recall, in Condition 6.6, we assumed that $u_x - u_y > 0$. If $u_x - u_y < 0$, then the inequality sign in the condition is reversed:

$$T(y, x) \leq T(x', y') H_{xy}^{x'y'} \quad \forall y, x', y' \in C; y \neq x, x' \neq y' \quad (6.15)$$

And the analysis holds, with the flipped inequality sign. Condition 6.14 becomes

$$\left(\frac{s_x}{s_{x'}}\right)^{b_x - b_{x'}} \left(\frac{s_y}{s_{y'}}\right)^{b_x - b_{y'}} \leq H \quad (6.16)$$

Symmetrically to Conditions 6.14, for Condition 6.16 to hold, we must establish an *upper bound* for the LHS. This requires us to solve a *maximization* problem for the LHS, which yields the following analogs to the four **Bounding Conditions** (labelled Cases **A1** to **A4**) then become the following:

Case B1: $\frac{s_x}{s_{x'}} \leq 1$ and $\frac{s_y}{s_{y'}} \leq 1$.

To maximize the LHS, we need to minimize both coefficients, which corresponds to **Case A2**, giving us

$$\left(\frac{s_x}{s_{x'}}\right) \left(\frac{s_y}{s_{y'}}\right) \leq H$$

Case B2: $\frac{s_x}{s_{x'}} > 1$ and $\frac{s_y}{s_{y'}} > 1$.

By a similar analogy to the above, we get

$$\left(\frac{s_x}{s_{x'}}\right)^{\lfloor \frac{n}{2} \rfloor} \left(\frac{s_y}{s_{y'}}\right)^{\lfloor \frac{n}{2} \rfloor} \leq H$$

Case B3: $\frac{s_x}{s_{x'}} \leq 1$ and $\frac{s_y}{s_{y'}} > 1$.

If $\left(\frac{s_x}{s_{x'}}\right)^3 \frac{s_y}{s_{y'}} \leq 1$, then check

$$\left(\frac{S_x}{S_{x'}}\right)^{3-(n \bmod 3)} \left(\frac{S_y}{S_{y'}}\right)^{\lfloor \frac{n-6}{3} \rfloor + 3} \leq H$$

otherwise, we check

$$\left(\frac{S_x}{S_{x'}}\right)^{\frac{n}{2}} \left(\frac{S_y}{S_{y'}}\right)^{\frac{n}{2}} \leq H$$

Case B4: $\frac{S_x}{S_{x'}} > 1$ and $\frac{S_y}{S_{y'}} \leq 1$.

Symmetrically to the previous case, the LHS is maximized as

If $\frac{S_x}{S_{x'}} \left(\frac{S_y}{S_{y'}}\right)^3 \leq 1$, then check

$$\left(\frac{S_x}{S_{x'}}\right)^{\lfloor \frac{n-6}{3} \rfloor + 3} \left(\frac{S_y}{S_{y'}}\right)^{3-(n \bmod 3)} \leq H$$

otherwise, we check

$$\left(\frac{S_x}{S_{x'}}\right)^{\frac{n}{2}} \left(\frac{S_y}{S_{y'}}\right)^{\frac{n}{2}} \leq H$$

Implementation

To evaluate a voter's strategic response according to the MaxM model, we must find candidate x satisfying the **ArgMax Condition** (6.4). That is, we need to find a candidate pair (x, y) , such that the **Domination Condition** (6.6) holds for every other pair of candidates x' and y' . The heuristic shortcuts this process by providing the **Bounding Conditions** (A1 to A4, and B1 to B4). Satisfying the Bounding Condition means the Domination Condition holds for that particular set of x, x', y, y' .

If for some candidate pair (x, y) , the Bounding Conditions are satisfied for all other x', y' , then we have found the unique x satisfying the ArgMax Condition (6.4).

If one of the Bounding Conditions is violated, then we move on to try the next $y \neq x$, because we cannot offer any guarantees for this current pair (x, y) . If we have exhausted all values of y , then we try to next x .

Finally, if we cannot offer any guarantees for any x , then we revert to the general solution of performing all calculations explicitly.

By altering the objective function slightly, this heuristic allows us to incorporate a number of guard statements that check for boundary conditions. When they are triggered, they allow the algorithm to bypass a number of costly operations. In practice, the guard statements trigger with higher frequency in more complex instances (i.e. larger population n or more candidates m), though never more than half the time. The net effect is a roughly constant factor speed-up, approximately halving the time needed to compute the pivot probabilities to determine a voter’s best response. The cost is that the modified objective function does not fully replicate the decisions of the **Full Voter**.

Tie Sampling Heuristics

Rather than calculating the exact probability of a pivot condition between x and y , we may utilize sampling techniques to estimate this probability. Here, the voter may be thought of as sampling from the outcome space to consider specific pivot scenarios, and acting based only on these imaged, plausible outcomes. This gives us the **TieU**, **TieR** and **TieH** models below, each is based on the **Full** voter model.

TieU

We first enumerate all pivot outcomes using the **BUS PACKING** algorithm specified in the **Full Voter Model**. But rather than querying the multinomial distribution for each outcome, we sample l outcomes from this space, and calculate the probabilities of each of these outcomes according to Equation 6.1, and approximate $T(y, x)$ as the sum of these l probabilities. This reduces the number of queries to l , though the algorithm must still iterate over the entire pivot space, which is still an $O(m^2n^{m-2})$ operation.

TieH

Rather than sampling uniformly from the space of pivot outcomes, which requires enumerating that space fully, we use the following heuristic for sampling non-uniformly from this space. The number of ballots b_w for our tied winners is drawn uniformly from the interval $[\lceil \frac{n-2}{m} + 1 \rceil, \lfloor \frac{n}{2} \rfloor]$.⁵ To allocate the remaining ballots r , the algorithm iterates through the other candidates in random order. For each candidate, the space of admissible allocations

⁵The lower bound is obtained by reserving two ballots for the tied winners and then splitting the remaining ballots evenly between m candidates.

is the interval $[\max(0, r - m' \times \min(r, b_w - 1)), \min(r, b_w - 1)]$, where m' is the number of unallocated candidates. The algorithm draws uniformly from this interval to allocate the number of ballots for the current candidate, and updates r and m' before moving to the next candidate. The result is one possible pivot outcome. Its probability is calculated according to Equation 6.1, and $T(y, x)$ is approximated as the sum of l such probabilities. This requires exactly l queries, with negligible overhead.

TieR

We estimate $T(y, x)$ by using a Monte Carlo algorithm in the space of pivot outcomes using rejection sampling. We generate l outcomes of the election by sampling from the multinomial distribution, and reject (i.e. discard) any outcomes that do not result in a 2-way pivot condition. $T(y, x)$ is estimated as the proportion of those outcomes which result in a 2-way pivot between y and x . This requires exactly l queries, with negligible overhead.

Poisson

Myerson proposed an alternative model for elections, treating them as large Poisson Games [99]. In this interpretation, the number of voters is *uncertain* and follows a Poisson distribution with mean n_e . That is, the probability that there are k players is

$$\text{Poisson}(k|n_e) = e^{-n_e} n_e^k / k!$$

He shows that if a given voter has probability s_b of casting a particular ballot b , then the number of voters casting b is also a Poisson distribution with mean $s_b n_e$. Crucially, he shows that the number of voters casting one type of ballot is independent of the number of voters casting another ballot. In his paper, Myerson examines “one-dimensional events”, events that may be represented as rays in the outcome space. He focuses on the convergence behavior of the probability that a 2-candidate election results in a tie, allowing for abstention where voting may be costly. We will propose a voter heuristic that extends this model to a multi-candidate election.

While Myerson’s Poisson Game models a fundamentally different voting process, we extend his results for the 2-candidate election to construct a voter model with behavior similar to the Full Voter: we generalize their results to an m -candidate election, and

the voters compute and compare the pivot probabilities between pairs of candidates to determine the course of action that maximizes their expected utility.

Let n_i be the random variable representing the number of votes that candidate i receives, and s_i denote the probability that a given voter casts a ballot supporting candidate i . Recall that n_i follows a Poisson distribution with mean $s_i n_e$. Then, Myerson shows that, as n_e approaches ∞ , the probability of casting a pivotal vote⁶ in support of $c \in \{1, 2\}$, in a 2-candidate election, converges to

$$Pr(n_1 = n_2 | \mathbf{s}) \approx \frac{e^{n_e(2\sqrt{s_1 s_2} - s_1 - s_2)}}{4\sqrt{\pi n_e \sqrt{s_1 s_2}}} \frac{\sqrt{s_1} + \sqrt{s_2}}{\sqrt{s_c}}$$

This requires that $s_1 + s_2 \leq 1$, allowing for some voters to abstain. To extend this model to a multi-candidate election, we treat ballots supporting other candidates as abstentions. We also require that this be a winning tie: i.e. $n_c > n_i, \forall i \neq 1, 2$. Since n_i are drawn independently from Poisson distributions with mean $s_i n_e$, the probability that $n_c - n_i > 0$ follows a Skellam distribution, which is approximated as

$$Pr(n_c > n_i) \approx \frac{(1 + (b_c + b_i)^2)e^{-(\sqrt{b_c} - \sqrt{b_i})^2}}{2(b_c + b_i)^2} - \frac{e^{-(b_c + b_i)}}{4\sqrt{b_c b_i}} - \frac{e^{-(b_c + b_i)}}{8b_c b_i}$$

where $b_c = s_c n_e$ and $b_i = s_i n_e$, and $s_c > s_i$ [76]. If we make the simplifying assumption that the events $n_1 = n_2$ and $n_1 > n_i, \forall i \neq 1, 2$ are independent⁷, then the probability that candidate x and y are in a pivot condition is the intersection of the events where x and y are tied, and x has more votes than every other candidate $i \neq x, y$. So, we approximate $T(y, x)$ as

$$T(y, x) = Pr(n_x = n_y | \mathbf{s}) \prod_{i \in C \setminus \{x, y\}} Pr(n_x > n_i)$$

As a result, the Poisson Model requires $\binom{m}{2}$ probability calculations, each of which takes $O(m)$ computations, giving us a rough runtime of $O(m^3)$.

⁶This accounts for both a direct tie, and where c is one vote away from a tie.

⁷The two events $n_1 = n_2$ and $n_1 > n_i$ are not independent in general. For example, suppose that $s_1 = s_2 \gg s_3$. If we know that $n_1 < n_3$, then we know that n_1 is likely small, which makes the event $n_1 = n_2$ much less likely.

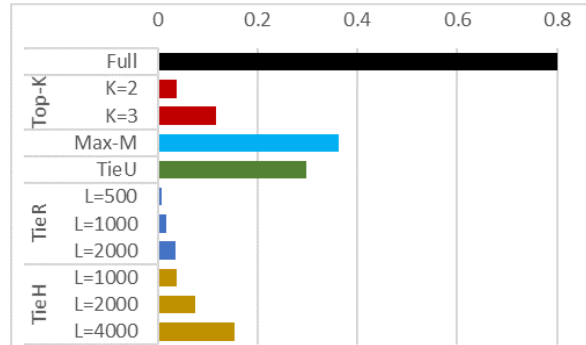


Figure 6.1: Runtime in seconds to construct one voter response ($m = 5, n = 100$).

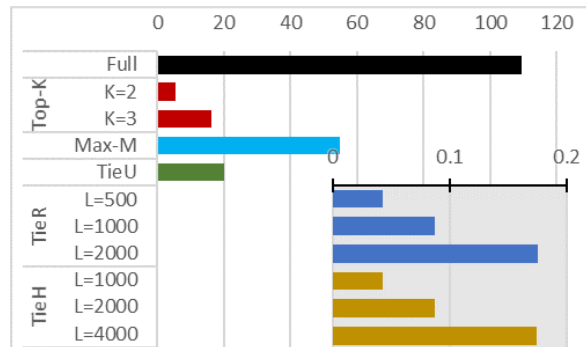


Figure 6.2: Runtime in seconds to construct one voter response ($m = 5, n = 500$). Inset show runtimes of heuristics using a different axis.

6.2 Comparison via Simulations

To benchmark these heuristics against the **Full** Voter model, we construct a framework where a voter v is queried for a strategic response based on a particular set of observations. For each trial, m candidates are generated with positions drawn uniformly at random from $[0, 100]$. The voter v and her $d = 25$ out-neighbors also draw their preferences from the same distribution. The value of d captures the amount of information available to the voter, and also the “resolution” of Figures 6.5–6.7. We chose a moderate value of $d = 25$, though we do not expect changing the value of d to significantly impact the results in this section. Each of the out-neighbors are assumed to vote truthfully, and v constructs a strategic ballot based on her Voter Response Function. The framework and algorithms are implemented using the D programming language, compiled using DMD64 D Compiler v2.071.1 on a 64-bit Ubuntu 14.04.5 server.

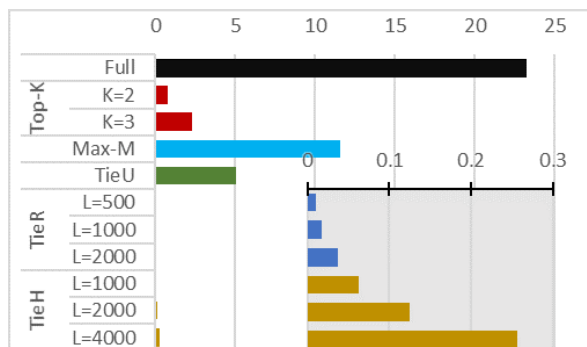


Figure 6.3: Runtime in seconds to construct one voter response ($m = 6, n = 100$). Inset show runtimes of heuristics using a different axis.

Figures 6.1 through 6.3 show the time (in seconds) required to construct one voter response under different voter models, with varying number of candidates m and size of electorate n . Each datum is averaged over 1000 trials. The top bar shows the time required for the **Full** voter. We see that the sampling heuristics, particularly **TieR** and **TieH** perform well, up to 2 orders of magnitude faster than **Full**; their runtimes are plotted in the inset using a different scale. Additionally, the runtime of **TieH** is unaffected by n . The **Poisson** Voter is not included in these benchmarks because its runtime is negligible.

We also compare the voter response of the heuristics to the response of the **Full** Voter. Figure 6.4 shows the rate of disagreement between the heuristics and the **Full** Voter. **TieH** ($L = 2000$) has a disagreement rate of 0.124, which means it computes a strategic ballot different from **Full** 12.4% of the time. Most heuristics are comparable in their accuracy, except **TieR** which performs noticeably worse. The accuracy of **TieR** may be increased by increasing the sample size l , but the gain is modest compared to the increase in runtime. It is interesting to note that even though the **Poisson** Voter is based on a fundamentally different model of voting, its accuracy is comparable to the other heuristics according to this benchmark. The results of Figure 6.4 are representative of other settings of n and m .

Poisson and **TieH** Voters appear to be our best heuristics for approximating the behavior of **Full** Voter. Next, we examine exactly *when* they disagree with **Full** Voter. For this benchmark, we consider the space of all possible observations that the voter may encounter, and determine which observations induce differences in voter response. For simplicity, we consider the case where $m = 4$ and $n = 500$. We assume the voter observes the ballots of $d = 25$ other people in her social network. We fix the voter's preference to be 0, the candidates positions to be $(10, 15, 20, 25)$; that is, v likes candidate c_1 the most, and c_4 the least. This results in a plot on 3-axes: (b_1, b_2, b_3) representing the number of

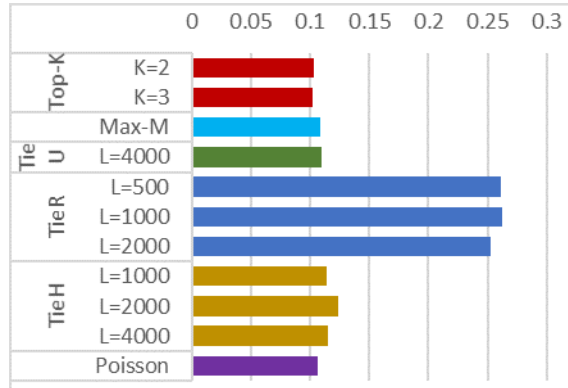


Figure 6.4: Rate of Disagreements from the Full Voter Model ($m = 5$, $n = 500$).

observed ballots supporting candidates 1, 2, and 3 (ballots supporting 4 may be inferred). We project this to the 2 dimensional heatmap shown in Figure 6.5. Along the x-axis, we plot b_1 , the number of ballots the voter observes supporting her favorite candidate; along the y-axis, b_2 , the support for her second choice. Each cell represents multiple points in the observation space. For instance, the cell $(4, 5)$ corresponds to all observations $(4, 5, b_3)$, where $b_3 \in \{0, 1, \dots, 16\}$. Cells in solid green represent conditions where Full Voter always casts a sincere ballot for candidate 1; Cells in solid white represent conditions where Full Voter never votes sincerely. Because the heatmap is a projection from a higher dimensional plot, each cell may represent more than one possible observation. When Full votes sincerely in only some of those observations, the cell is shaded in lighter green. The triangle on the bottom right, show in gray, are inadmissible conditions where the total number of ballots supporting candidate 1 and 2 exceed 25.

Naturally, as b_1 increases, Full will have a tendency to vote for 1 as well. The large triangular cutout on the left shows when candidate 2 has enough support that v will change to a strategic vote for 2. The upper left region shows situations where neither 1 nor 2 have much support, and the election is a race between 3 and 4. This plot shows that the Full Voter tends to vote sincerely when her favorite candidate has even moderate amount of support (to bolster her chance for victory), or when the race is between her top and second choices. She only votes strategically when she believes the likely winners do not include candidate 1.

Figure 6.6 uses the same axes, and highlights the conditions where TieH ($L = 1000$) disagrees with Full. Cells in red show *increased* preference for sincere voting; cells in green show *decreased* preference for sincere voting. Cells in white show the two models in general agreement (i.e. the observations represented by the cell result in the same

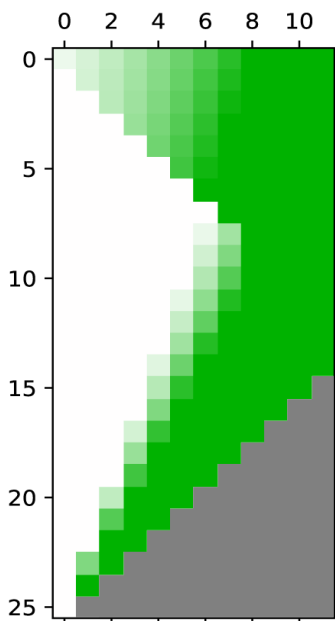


Figure 6.5: Conditions where Full votes sincerely in green.

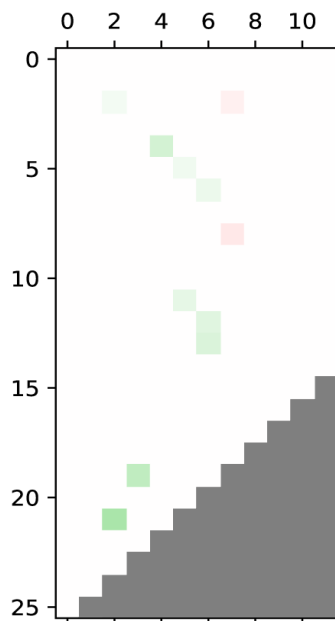


Figure 6.6: Conditions where TieH disagrees in red or green.

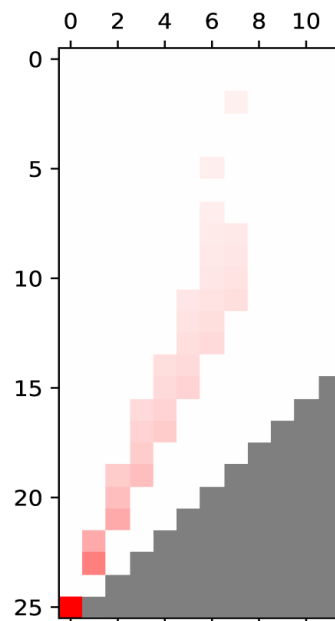


Figure 6.7: Conditions where Poisson disagrees in red or green.

number of wins for each candidate, though not necessarily for the same observations). Due to the stochastic nature of TieH, the instances of disagreement are spread out in the observation space. However, there is a trend for them to concentrate near the borders where the Full Voter transitions between different ballots. This is more clearly seen in Figure 6.7, which maps the same differences between Poisson and Full. Here, we see that Poisson systematically casts more sincere votes, and the disagreements concentrate on the transition between sincere voting and strategically voting for the second choice.

6.3 Addressing the Desiderata

We return to the desiderata proposed in Section 6.1 and consider how well our proposed heuristics implement them. Table 6.1 summarizes the degree to which each proposed heuristic fulfills the desiderata, including a column for the Full Voter Model, and a row for a heuristic’s Fidelity when compared to Full. By their design, the heuristics adhere to the Knowledge requirement, that each voter acts only upon observations gleaned from

	Full	Top-K	Max-M	TieU	TieH	TieR	Poisson
Knowledge	✓	✓	✓	✓	✓	✓	✓
Rationality	✓	Aprx	Aprx	Aprx	Aprx	Aprx	Aprx
Anonymity	✓	✓	✓	✓	✓	✓	✓
Equilibrium	*	*	*	*	*	*	*
Tractability	✗	✗	✗	*	✓	✓	✓
Optimistic	✓	✓	✓	✓	✓	✓	✓
Fidelity	–	✓	✓	✓	✓	✗	✓

Table 6.1: Adherence to the Desiderata by the Heuristics. All heuristics approximate the rational behavior of the **Full** voter. The asterisks in the **Equilibrium** row indicate each algorithm will either converge to an equilibrium or display a cyclic behavior. The symbols in the **Tractability** row provide rough indicators on the speed of the algorithms, ranging from slow (**✗**), to moderate (*), to fast (✓).

her social network. Similarly, all heuristics fulfill the **Anonymity** and **Optimistic** criteria. Our models also fulfill the **Equilibrium** desideratum in the sense that it terminates when the population converges, or exhibits cyclic behavior to show an equilibrium cannot be reached.

As explained in Section 6.1.1, the **Full Voter Model** is based on rational actions of the voter, and so fulfills the **Rationality** desideratum. The other models approximate this rational behavior to varying degrees of fidelity. With the exception of the **TieR** models, all our proposed heuristics replicate the decisions of the fully rational voter most of the time (about 88% to 90% for $m = 5$ and $n = 500$).

The main difference between our models is in the **Tractability** desideratum. As noted before, the **Full Voter Model** scales very poorly with both the number of candidates m and the number of other voters n . This is because the algorithm must explicitly enumerate every possible 2-way winning tie outcome when computing the pivot probability between two candidates, and the size of this pivot space grows very quickly with m and n . The heuristics **Top-K**, **Max-M** and **TieU** simplify this computation in different ways, but still require this enumeration. As a result, they scale poorly to larger elections. The **TieR** heuristic bypasses the need for enumerating pivot space by using Monte Carlo sampling on the entire outcome space; this is fast, but unfortunately loses accuracy quickly in larger elections due to the sparsity of pivot outcomes. On the other hand, the **TieH** heuristic generates more consistent runtimes by allow the heuristic to *not* sample uniformly from the pivot space. The cost in accuracy for making this assumption is small, while the

performance gains are enormous. Surprisingly, the `Poisson` heuristic also performs very well in our experiments. While it makes fundamentally different assumptions about the voters, it produces results that are comparable in accuracy to `TieH` in a negligible runtime. Thus, `Poisson` and `TieH` are the heuristic models that fulfill the best combination of desiderata for use in larger simulations.

6.4 Case Study: The Micromega Rule

In this section, we illustrate an application of our heuristics by studying the Micromega Rule. In particular, we focus on the effects of different population sizes, an aspect that relies on our heuristics' ability to scale to larger graphs.

In political science, the Micromega rule frequently deals with districting and systems of proportional representation. While it is intuitive that properties of the voting rules determine the qualities of successful parties, Josep Colomer posits that this influence runs both ways: that existing political parties will favor electoral rules that improve their future electoral chances. In particular, he formulated the Micromega rule, which predicts that a government comprised of few, large parties will favor smaller assemblies and larger districts, while those formed from smaller parties will favor assemblies with many seats and smaller individual districts [36]. These topics have been explored in computational social choice in [13]. They compare the outcomes from a district based election to that of a popular vote. They devise the Misrepresentation Ratio to measure this deviation, and show that misrepresentation occurs when the number of voters is small or large.

We explore a variation of the Micromega rule. We posit that large parties are more effectively able to consolidate their voter base in *large electoral districts*, while less populous districts will see more continued support for less popular candidates. It is this continued support that allows the party to remain viable in future elections. The use of voter heuristics is essential to this exploration of the Micromega rule because our hypothesis depends on population size. Without using heuristic voter responses, it would not be feasible to simulate communities of any significant size, and we would not be able to sample election results from a large enough range of community sizes of test our hypothesis.

We present our electoral districts as simulated social networks. We use directed and undirected versions of the Erdős-Rényi (ER) random graph model [56], and the Barabási-Albert (BA) preferential attachment model [9]. In the (directed) ER graph with density parameter pr , every (directed) edge exists with probability pr . To create a (directed) BA graph with attachment parameter pr , we recursively add new nodes to an existing

graph, attaching it to pr existing vertices via a (directed) edge; the existing vertices are picked with probability proportion to their (in-)degree. We consider graphs of sizes $n \in \{200, 400, 600, 800\}$. We focus on the scenario with $m = 5$ candidates. We fix the average in-degree of each node to be approximately $d = 30$.⁸ A value of $d = 30$ was selected as a reasonable number of informational sources voters may consider in such a scenario. Increasing d allows voters to sample more information from the network, and therefore give them a more precise estimate on the outcome of the election. This will likely increase the amount of strategic play in all conditions, since voters will be less optimistic a favorite candidate will recover if they are behind in the polls. However, we do not believe changes to d will significantly affect the qualitative patterns that emerge.

Based on our results from the previous section, we will use both the **Poisson** and **TieH** models for this simulation. The **Poisson** model is parameter free, though it exhibits systematic bias as compared to the **Full** voter; for the **TieH** model, we set the sample size parameter at $L = 2000$. To measure the degree of support for less popular parties, we take the **SF Ratio**, the ratio of support between the second- and first-runner up candidates. Each data point is the average of 200 replications.

6.4.1 Results

Figure 6.8 plots the average SF-Ratio for each condition, and Table 6.2 shows the actual values. Under the **TieH** model, we observe that clear downward trends in the directed ER and BA graphs — the amount of support for the third place candidate diminishes exponentially with increasing voter population, reflecting an exponentially increasing ability for voters to vote strategically. Most interesting is the difference in SF Ratios of the ER and dER graphs. Structurally, there is only one difference between these two models. In the ER graph, influence is reciprocal – if v observes another voter, then that other voter also observes v . The same is not generally true in a dER graph, which allows their voters an increased ability to communicate and propagate information through strategic play.

The most significant downward trend is observed in the dBA graph. The dBA model generates directed acyclic graphs that have a strongly hierarchical structure, where all edges are oriented toward older nodes. Information in these graphs flow from the older and higher degree nodes, toward the younger, lower degree nodes. This hierarchical structure prevents effective communication, which manifests in the high SF Ratios, particularly in smaller graphs.

⁸Importantly, the parameter pr is *doubled* when constructing directed BA graphs.

In the `Poisson` model, we observe the same strong trend in the directed BA graphs. However, there is no discernible pattern in any of the other graph types as the population size changes. It is possible that this difference is due to the systematic bias in how the `Poisson` model attempts to estimate the `Full` voter’s behavior, or that the trend observed in the directed ER graphs is relatively fragile compared to that of the directed BA graph.

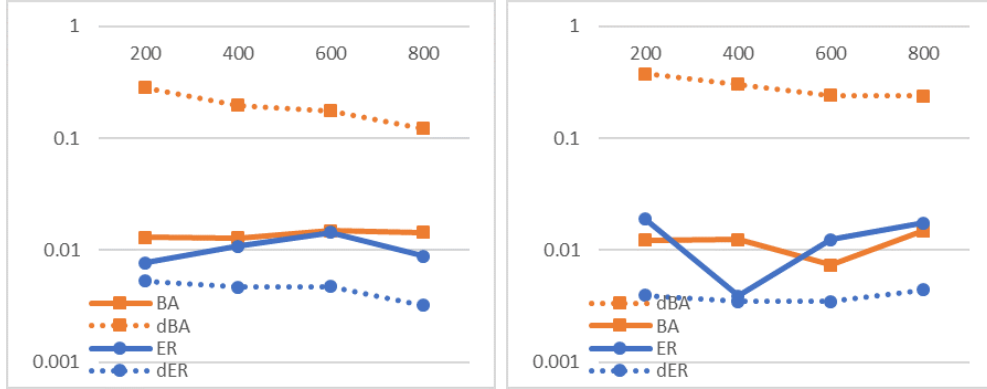


Figure 6.8: Average SF Ratios of graphs of different sizes. `TieH` voters on the left. `Poisson` voters on the right. Note y-axis is in log scale.

Model	n=200	n=400	n=600	n=800
TieH model				
ER	0.007746	0.010873	0.014477	0.008847
dER	0.005333	0.004676	0.004722	0.003231
BA	0.013058	0.012966	0.014909	0.014406
dBA	0.28263	0.19711	0.175578	0.122599
Poisson model				
ER	0.018958	0.003924	0.012530	0.017439
dER	0.003965	0.003504	0.003512	0.004412
BA	0.012314	0.012492	0.007419	0.014982
dBA	0.37516	0.30224	0.24184	0.237367

Table 6.2: Average SF-Ratios. `TieH` model on top. `Poisson` model on bottom.

It is worthy to note that the undirected BA graph shows no significant trend in either direction. Amongst the models used, the BA preferential attachment model may be the most representative of real world social networks. So while we demonstrate an alternate

cause for the Micromega rule in certain types of social networks, the result does not necessarily generalize to real world networks.

6.5 Conclusion

In this chapter, we proposed a number of heuristic voter models for strategic voting in social networks. While the Full Voter model in [129] works well in small graphs, the exact computation of expected utilities proves infeasible for larger graphs. Voters in our models are boundedly rational and our heuristics lighten their cognitive burden in ways that would be natural for a human voter. Our heuristics perform up to 2 orders of magnitude faster, and retain a high level of fidelity when compared to the Full Voter model.

We use our heuristic voter model to investigate the Micromega rule. We show that in certain networks, particularly the directed Erdős-Rényi and directed Barabási-Albert models, smaller populations offer more support for fringe candidates than larger electorates. The orientation of directed edges in the dBA graphs lends it a strict hierarchical structure, which reinforces the Micromega rule dramatically in our simulations. Other preferential attachment models such as Bollobás’s scale-free graphs offer parameters which may be tuned to allow for different degrees of hierarchy [24]. It would be interesting to explore the impact of hierarchical structure on strategic voting in future work.

Our heuristic models also pave the way to simulating strategic voting behavior in truly large scale networks. This opens up the possibility of simulating on real world datasets, where nodes number in the millions. Framing the voter’s best response decision as an optimization problem may also prove fruitful, as it allows us to leverage industry standard tools for solving optimization problems. Moreover, we may also consider broadening our model to include other scoring rules, such as Borda and k-Approval, and other social choice functions in general.

Chapter 7

Vote Timing

7.1 Introduction

Because voting is a process that takes place over time, there is an asymmetry of information that is available to earlier and later voters. The ballots cast by earlier voters inform subsequent voters. The latter may use this information to vote strategically, maximizing their chances of casting a pivotal ballot; The former may gain a first-mover advantage, establishing their favorite candidate as a lead runner by shaping what information is available to later voters. Strategic voters must decide not only which ballot to cast, but also *when* to cast their ballot.

The U.S. presidential primaries is an example of such a sequential procedure. The primaries determine each parties' presidential nominee, and are conducted as a series of elections in each state. Each state-level election determines how many delegates are sent in support of each nominee by that state. States schedule their own primary dates¹. The resulting elections are spread over several months. In 2016, primaries began in February and ended in June, in preparation for the November election [1]. Both parties and individual states recognize the importance of strategic timing. Certain time slots are highly prized by both the Republican and Democratic Parties. Both parties award bonus delegates to states holding their elections later in the primary season [18, 114].

Online polls are another domain which allows for strategic timing. These polls are used as a social choice mechanism for selecting anything from the cutest animal, to artistic direction for crowd-funded projects, to the winner of the Webby People's Voice award. A

¹Each state and party has their own set of rules.

popular implementation of online polls is the popular group scheduling platform Doodle. Doodle allows participants to approve or decline proposed time slots. Importantly, Doodle supports *open* polls, which allow voters to view the ballots cast by previous voters before committing their own, or waiting and revisiting the poll at a later time.

This type of strategic timing and voting behavior was dramatically seen in the selection of titles for the Evolution Championship Series (EVO) in 2017. EVO is an annual international fighting game tournament, the largest of its kind. To determine the title played at one of its events, the organizers ask the community to vote by contributing on an online charity platform, over a period of 15 days. Because the funds raised were updated in real time, voters effectively had access to accurate polling information throughout the process, and could strategically pick when and how they voted. The results show evidence of strategic behavior. Of the \$147,570 raised, the top two titles account for more than 90% of the funds[4], with both having established their lead by the second day[7], with large donations arriving up until the closing minutes [5].

In this chapter, we propose the Sticker Voting framework, where voters are invited to place a sticker (their ballot) on their chosen candidate (which becomes common knowledge). In addition to the potential for casting a strategic ballot, this process also invites voters to be strategic in timing their vote. We propose a model for strategic voter behavior that incorporates strategic timing, and we analyze the strategic equilibrium in a simple Sticker Plurality Voting game. Finally, we discuss how we may use our model to capture voting behavior in the real world.

7.2 Sticker Voting Model

We consider a non-sequential voting game \mathcal{G} with n voters, and m candidates \mathcal{M} . Let \mathcal{B} be the set of admissible ballots a voter may cast, and \mathcal{B}^n be the set of possible ballots cast by the population of voters. Let \mathcal{F} be a social choice function mapping \mathcal{B}^n to the set of winners, a non-empty subset of \mathcal{M} . Each voter v has a private utility function $u_v : 2^{\mathcal{M}} \rightarrow \mathbb{R}$ mapping each outcome to a utility value.

We define a Sticker Voting game based on \mathcal{G} by specifying a number of voting rounds $T \geq 1$. In each round, voters may cast a ballot or choose to “Wait”; this choice is made simultaneously within each round. Once a voter casts a ballot, it is committed and irreversible. Formally, in each round, each voter plays an action from the action set $\mathcal{B} \cup \{\emptyset\}$, where \emptyset corresponds to “Wait” action. Once a voter casts a ballot $b \in \mathcal{B}$, their action space for subsequent rounds is reduced to the set $\{b\}$, representing their vote being “locked in”;

we refer to this as moving from the **controlled game** to the **uncontrolled game**. Let $H_t \in \mathcal{B}^n$ denote the set of actions played by agents in round t . The history of play prior to current round t , $\mathcal{H}_t = (H_1, H_2, \dots, H_{t-1})$ is common knowledge. The winner set is $\mathcal{F}(H_T)$, where \emptyset actions are interpreted as “Abstain”.

In round t , a voter may act according to a pure strategy function \mathcal{S} , which maps \mathcal{H}_t to an action $a^t \in \mathcal{B} \cup \{\emptyset\}$. \mathcal{S} maps to the action b if the agent entered the uncontrolled game by casting ballot $b \in \mathcal{B}$ in a prior round. We also allow voters to play mixed strategies, which map \mathcal{H}_t to a mixed strategy, i.e. a distribution over $\mathcal{B} \cup \{\emptyset\}$.

We focus on *Markovian* strategies, where the voters do not care about the history of ballots prior to the previous round $t - 1$. A Markovian strategy \mathcal{S} maps t and H_{t-1} to a mixed strategy.

7.2.1 Plurality Sticker Voting

In this chapter, we focus on the Resolute Plurality Voting Rule. Admissible ballots \mathcal{B} are the candidates \mathcal{M} . For round t , denote the standing \mathbf{s}_t as a vector whose i -th element corresponds to the number of ballots supporting candidate i in H_{t-1} , or the zero vector if $t = 1$. The social choice function \mathcal{F} maps the final votes H_T to the unique candidate i with the highest s_T^i , breaking ties uniformly at random.

We consider Markovian strategies that are also anonymous to other voters. In round t , while in the controlled game, an agent’s strategy simply maps t and \mathbf{s}_t to a mixed strategy.

7.2.2 Solution Concept

The Sticker Voting Game uses the solution concept of the Perfect Bayesian Equilibrium (PBE). PBE is a refinement of Subgame Perfect Equilibrium (SPE) for sequential games. In a SPE, players act according to strategies that form a Nash equilibrium in every subgame of the original game. PBE additionally allows players to have incomplete information, where certain nodes of the game tree are indistinguishable from each other to particular players; these are called Information Sets. Players maintain beliefs corresponding to the probability that they are in a particular node in the current Information Set; their strategies are defined according to these beliefs that are on the equilibrium path (and may depend on the history of play). In the Sticker Voting Game, Information Sets correspond to voters not knowing the types of the other voters.

In the Plurality Sticker Voting Game, the current round and tally form a tuple (t, \mathbf{s}_t) that uniquely identifies the information set for the player in the controlled game. Each information set consists of nodes representing the possible types that the remaining uncommitted voters may have. The voter has a belief over the distribution of types of the uncommitted voters.

A second set of nodes capture the uncontrolled games, with a unique node for each round t and uncontrolled tally \mathbf{s}_t .

After each round, the beliefs held by each player are updated in the Bayesian manner. That is, let $P_i(\Theta_{-i})$ be voter i 's prior belief about the the types of each other player. Then, if each player draws their type from a type space Θ , $P_i(\Theta_{-i})$ is a distribution over Θ^{n-1} . After observing a sequence of action profiles $\{\mathbf{b}^{(1)}, \mathbf{b}^{(2)}, \dots, \mathbf{b}^{(k)}\}$, corresponding to the actions of players from the first k rounds of the game, voter i 's beliefs are updated as follows:

$$P_i(\Theta_{-i}|\mathbf{b}^{(1)}, \mathbf{b}^{(2)}, \dots, \mathbf{b}^{(k)}) = \frac{P_i(\mathbf{b}^{(k)}|\mathbf{b}^{(1)}, \mathbf{b}^{(2)}, \dots, \mathbf{b}^{(k-1)}, \Theta_{-i})P_i(\mathbf{b}^{(1)}, \mathbf{b}^{(2)}, \dots, \mathbf{b}^{(k-1)}, \Theta_{-i})}{P_i(\mathbf{b}^{(1)}, \mathbf{b}^{(2)}, \dots, \mathbf{b}^{(k)})} \quad (7.1)$$

The numerator term $P_i(\mathbf{b}^{(1)}, \mathbf{b}^{(2)}, \dots, \mathbf{b}^{(k-1)}, \Theta_{-i})$ refers to the probability that the game has reached the information set of the preceding round, given assumption Θ_{-i} about the types of the players; the other term in the numerator, $P_i(\mathbf{b}^{(k)}|\mathbf{b}^{(1)}, \mathbf{b}^{(2)}, \dots, \mathbf{b}^{(k-1)}, \Theta_{-i})$ refers to the probability of observing the actions of the current round, given the preceding information set. The term in the denominator $P_i(\mathbf{b}^{(1)}, \mathbf{b}^{(2)}, \dots, \mathbf{b}^{(k)})$ is a normalizing factor that is equal to the probability of observing the given sequence of actions, regardless of voter types. Note that all three terms may be calculated from the game tree, and these calculations may be efficiently implemented via dynamic programming.

7.3 Complete Information Game

We first consider a simplified scenario with $n = 3$ voters, $\{1, 2, 3\}$, with complete information, and $m = 3$ candidates, $\{A, B, C\}$, in a $T = 2$ round game. Player 1 has preference $A \succ B \succ C$; player 2, $B \succ C \succ A$; player 3, $C \succ A \succ B$, forming a Condorcet cycle. Each player gains utility u_1 if their favorite candidate wins, u_2 utility for their second choice, and 0 for their third choice, with $u_1 > u_2 > 0$. We also require that $2u_2 > u_1$ so that conceding to one's second place alternative is better than a three-way tie. The types of all agents are public knowledge. The following table summarizes the utilities:

Voter	A	B	C
1	u_1	u_2	0
2	0	u_1	u_2
3	u_2	0	u_1

For simplicity of notation, we denote voter v 's favorite candidate as $b_{v,1}$, the second choice as $b_{v,2}$, and so on. When the v is clear from context, we omit v from the subscript. Since we focus on Plurality, where voters act by submitting a ballot that support one particular candidate, we also use $b_{v,i}$ to denote the action where v votes for this candidate (i.e. v 's i -th preference). We will actualize the utility values as $u_1 = 3$ and $u_2 = 2$.

Since agents have complete information, the agents' beliefs correspond to the actual types of the agents, and do not update throughout the game.

Analysis: Final Round

Since the types are common knowledge, we use the more general solution concept of the Subgame Perfect Equilibrium, and use backward induction to solve the game. Without lost of generality, we take the perspective of Agent 1.

We begin with the final round T . If the agent is still in control, she may find the game in a number of different states:

Case 1: 2 ballots for the same candidate.

Agent 1's vote is irrelevant, and that candidate is selected

Case 2: 2 ballots for different candidates.

Agent 1 breaks ties in favor of the better option.

Case 3: 1 ballot for A

Agent 1 also votes A and gets A as the outcome.

Case 4: 1 ballot for B

Note that this ballot must be cast by Agent 2, since Agent 3 would never vote for B . In this scenario, we can break down the utilities for the remaining players in the following table. Entries indicate the winning candidate, with the payoff for the row and column players in parentheses.

		Agent 3	
		C	A
Agent 1	A	$tie(5/3, 5/3)$	$A(3, 2)$
	B	$B(2, 0)$	$B(2, 0)$

It is clear both agents will **coordinate on action** $b_{1,1} = b_{3,2} = A$ as other actions are strictly dominated, and we may iteratively remove dominated strategies.

Case 5: Only Agent 3 has voted, for $C = b_{1,3}$

We also break down utilities here:

		Agent 2	
		B	C
Agent 1	A	$tie(5/3, 5/3)$	$C(0, 2)$
	B	$B(2, 3)$	$C(0, 2)$

By the same argument before, the two agents will coordinate on **selecting** B .

Case 6: Only Agent 2 has voted, for $C = b_{1,3}$

Since Agent 3 has not voted, this is actually **Case 1** from the perspective of Agent 3. That is, since C is Agent 3's top choice, Agent 3 will also vote C , secure it as the outcome. Agent 1's vote is irrelevant. For completeness, we include the utility breakdown as well:

		Agent 3	
		C	A
Agent 1	A	$C(0, 3)$	$A(3, 2)$
	B	$C(0, 3)$	$tie(5/3, 5/3)$

Case 7: No votes observed

Assuming each agent plays symmetric strategies, each outcome is equally likely, giving an expected utility of $5/3$.

Interestingly, **Case 4** and **Case 5** clearly show that there is no straightforward first-mover advantage in this scenario. Any agent that is the sole voter in the initial round, and votes for b_1 , will force the remaining agents to coordinate in the next round, and produce b_3 as the outcome.

Analysis: Initial Round

Now that we have the actions and expected utilities for the final round, we apply backward induction to determine the course of play in the initial round. We assume symmetric play; that is, each player v plays action $b_{v,i}$ with probability p_i , $i = \emptyset, 1, 2$, $0 \leq p_\emptyset, p_1, p_2 \leq 1$ and $p_\emptyset + p_1 + p_2 = 1$. We analyze the expected utility for Agent 1 for each action.

Case 1: Agent 1 plays A

As we have established, if Agent 1 plays A and the other agents plays \emptyset , the other agents will coordinate to select C , yielding 0 utility for Agent 1. However, Agent 1 may potentially gain an advantage if the other agents choose not to wait. The following table shows the outcomes and their payoffs for Agent 1, based on the actions of Agents 2 and 3.

		Agent 3		
		C	A	\emptyset
Agent 2	B	$tie(5/3)$	$A(3)$	$A(3)$
	C	$C(0)$	$A(3)$	$C(0)$
	\emptyset	$C(0)$	$A(3)$	$C(0)$

The expected utility for voting b_1 in the first round is

$$E(u|b_1) = \frac{5}{3}p_1^2 + 3p_2 + 3p_\emptyset p_1 \quad (7.2)$$

Case 2: Agent 1 plays B

		Agent 3		
		C	A	\emptyset
Agent 2	B	$B(2)$	$B(2)$	$B(2)$
	C	$C(0)$	$tie(5/3)$	$C(0)$
	\emptyset	$B(2)$	$B(2)$	$B(2)$

The expected utility for voting b_2 in the first round is

$$E(u|b_2) = 2(p_1^2 + p_1 p_2 + 2p_\emptyset p_1 + p_\emptyset p_2 + p_\emptyset^2) + \frac{5}{3}p_2^2$$

$$= 2p_1 + 2p_\emptyset + \frac{5}{3}p_2^2 \quad (7.3)$$

Case 3: Agent 1 plays \emptyset

		Agent 3		
		C	A	\emptyset
Agent 2	B	$B(2)$	$A(3)$	$A(3)$
	C	$C(0)$	$A(3)$	$C(0)$
	\emptyset	$B(2)$	$A(3)$	$tie(5/3)$

The expected utility for Waiting in the first round is

$$E(u|b_\emptyset) = 3p_2 + 2p_1^2 + 5p_\emptyset p_1 + \frac{5}{3}p_\emptyset^2 \quad (7.4)$$

Notice immediately that even when factoring in the possibility of multiple agents voting in the initial round, Waiting dominates voting for A . So we conclude that $p_1 = 0$.

Suppose we are at a symmetric mixed Nash Equilibrium, then Agent 1 must be ambivalent over the actions in its support (i.e. b_2 and \emptyset). So we may set equations (7.3) and (7.4) equal, and solve.

$$-p_1 p_2 - p_\emptyset p_1 - p_\emptyset p_2 + \frac{1}{3}p_\emptyset^2 - \frac{4}{3}p_2^2 = 0$$

And since $p_1 = 0$:

$$\begin{aligned} -p_\emptyset p_2 + \frac{1}{3}p_\emptyset^2 - \frac{4}{3}p_2^2 &= 0 \\ p_\emptyset^2 - 3p_\emptyset p_2 - 4p_2^2 &= 0 \\ (p_\emptyset + p_2)(p_\emptyset - 4p_2) &= 0 \\ p_\emptyset &= 4p_2 \end{aligned}$$

Surprisingly, the symmetric mixed Nash Equilibrium strategy for the initial round is for each agent to **play b_2 with probability 0.2, and Wait with probability 0.8.**

7.3.1 Rational Voter Behavior

In this simple, complete information game, rational voters will never vote for their top choice in the first round. Instead, they will vote b_2 with probability 0.2, or otherwise Wait in the first round. In the latter case, Agent 1 will vote for her favorite candidate in the second round, unless both other voters have committed their ballots and she must break a tie in her favor; or Agent 3 casts the only ballot and has voted for C, in which case Agent 1 votes for B.

7.4 Incomplete Information Game

Next, we consider an incomplete information scenario based on the simple game above. As before, we have $n = 3$ voters $\{1, 2, 3\}$ and $m = 3$ alternatives $\{A, B, C\}$. Players may be one of three types: Type A players have preference $A \succ B \succ C$; Type B, $B \succ C \succ A$; and Type C, $C \succ A \succ B$. The possible types form a Condorcet cycle, but there is no guarantee that such a cycle will exist in a particular realization of types. Nature assigns a type to each player with equal probability. Players know their own types, but do not know the types of other players. The game will be played over $T \geq 2$ rounds. We impose the same utility structure as before.

Each agent begins knowing her own type, and with the belief that the remaining agents may have types A , B or C with equal probabilities. Once an agent i commits a ballot, the other agents update their beliefs of i 's type in a Bayesian manner, according to Equation 7.1. In particular, if the vote was for c , they rule out the possibility that i is of a type that prefers c the least, because the probability that they commit such a ballot b^k , $P(b^{(k)} | \mathbf{b}^{(1)}, \mathbf{b}^{(2)}, \dots, \mathbf{b}^{(k-1)}, \Theta_{-i})$, is identically zero. We note that these belief updates have no impact on the actions of the remaining voters, as the agent who has voted can no longer affect the outcome of the election, and that no additional information is gleaned about the other (uncommitted) players.

Analysis: Final Round T

WLOG, we consider the game from the perspective of Agent 1, who is Type A. If we are in the final round of the controlled game, with tally \mathbf{s}_t , let the voters' strategy $\mathcal{S}(t, \mathbf{s}_t)$ be a mixed strategy playing b_i with probability p_i^{t, \mathbf{s}_t} , where $i \in \{1, 2, 3, \emptyset\}$. We will omit the t and/or \mathbf{s}_t from the superscript where it is clear from context. Additionally, because

voter strategies are symmetric with respect to type, we adopt the notational convenience of permuting the vector \mathbf{s}_t so that its i -th entry corresponds to the tally of the voter's i -th favorite candidate.

Playing b_3 is strictly dominated, so by the iterated removal of dominated strategies, $p_3 = 0$ in all situations. Moreover, since this is the final round, Waiting is strictly dominated by voting b_1 , so $p_\emptyset = 0$. Therefore, for any particular \mathbf{s} , $p_1^s + p_2^s = 1$. All probability values are bounded within $[0, 1]$.

Case 1: 2 ballots for the same alternative.

Agent 1's vote is irrelevant, as that alternative is selected. There are three outcomes, with utilities for Agent 1 being 3, 2 or 0.

Case 2: 2 ballots for different alternatives.

Agent 1 breaks ties in favor of the better option. There are 6 outcomes here. Agent 1 may break the tie to gain her top choice in 4 cases, and get her second choice in 2 cases.

Case 3: Agent 3 (WLOG) casts the only vote, for $A = b_{1,1}$

Agent 1 also votes A and gets A as the outcome.

Case 4: Agent 3 casts the only vote, for $B = b_{1,2}$

Agent 2 may be of one of three types. If Agent 2 is Type B , then they will also vote for B . Agent 1's vote is irrelevant, and gets a payoff of 2. The following tables break down the utility of Agent 1's actions for the other two cases:

Table 7.1: Utility breakdown if Agent 2 is Type A

		Agent 2	
		A	B
Agent 1	A	$A(3, 3)$	$B(2, 2)$
	B	$B(2, 2)$	$B(2, 2)$

Table 7.2: Utility breakdown if Agent 2 is Type C

		Agent 2	
		<i>C</i>	<i>A</i>
Agent 1	<i>A</i>	<i>tie</i> (5/3, 5/3)	<i>A</i> (3, 2)
	<i>B</i>	<i>B</i> (2, 0)	<i>B</i> (2, 0)

Since Agent 2's type is not known to Agent 1, neither action is dominant. But we can calculate the expected utility for each action.

$$\begin{aligned}
 E(u|b_1^{(0,1,0)}) &= \frac{1}{3}(3p_1^{(0,1,0)} + 2p_2^{(0,1,0)}) \\
 &+ \frac{1}{3}\left(\frac{5}{3}p_1^{(0,0,1)} + 3p_2^{(0,0,1)}\right) \\
 &+ \frac{1}{3}(2)
 \end{aligned} \tag{7.5}$$

$$\begin{aligned}
 E(u|b_2^{(0,1,0)}) &= \frac{1}{3}(2p_1^{(0,1,0)} + 2p_2^{(0,1,0)}) + \frac{1}{3}(2p_1^{(0,0,1)} + 2p_2^{(0,0,1)}) + \frac{1}{3}(0) \\
 &= \frac{1}{3}(2) + \frac{1}{3}(2) + \frac{1}{3}(2) \\
 &= 2
 \end{aligned} \tag{7.6}$$

If there is a mixed equilibrium, then Agent 1 will be ambivalent over the two choices. We set equations (7.5) = (7.6), and solve to obtain

$$p_1^{(0,1,0)} = \frac{4}{3}p_1^{(0,0,1)} - 1 \tag{7.7}$$

We set aside this equation, and carry it forward to Case 5.

Case 5: Agent 3 casts the only vote, for $C = b_{1,3}$

Agent 2 may be of one of three types. If Agent 2 is Type *C*, then they will vote for *C* and Agent 1's action is irrelevant, and they get utility 0. The following tables break down the utility of Agent 1's actions for the other two cases:

Table 7.3: Utility breakdown if Agent 2 is Type A

		Agent 2	
		A	B
Agent 1	A	A(3, 3)	<i>tie</i> (5/3, 5/3)
	B	<i>tie</i> (5/3, 5/3)	B(2, 2)

Table 7.4: Utility breakdown if Agent 2 is Type B

		Agent 2	
		B	C
Agent 1	A	<i>tie</i> (5/3, 5/3)	C(0, 2)
	B	B(2, 3)	C(0, 2)

Since Agent 2's type is not known to Agent 1, neither action is dominant. But we can calculate the expected utility for each action.

$$E(u|b_1^{(0,0,1)}) = \frac{1}{3}(3p_1^{(0,0,1)} + \frac{5}{3}p_2^{(0,0,1)}) + \frac{1}{3}(\frac{5}{3}p_1^{(0,1,0)} + 0p_2^{(0,1,0)}) + \frac{1}{3}(0) \quad (7.8)$$

$$E(u|b_2^{(0,0,1)}) = \frac{1}{3}(\frac{5}{3}p_1^{(0,0,1)} + 2p_2^{(0,0,1)}) + \frac{1}{3}(2p_1^{(0,1,0)} + 0p_2^{(0,1,0)}) + \frac{1}{3}(0) \quad (7.9)$$

As before, if we are at a mixed equilibrium, the agent must be ambivalent over the two actions. And so we set equations (7.8) = (7.9), and solve:

$$\begin{aligned} \left(\frac{4}{3}p_1^{(0,0,1)} - \frac{1}{3}p_2^{(0,0,1)}\right) + \left(-\frac{1}{3}p_1^{(0,1,0)}\right) &= 0 \\ 4p_1^{(0,0,1)} - p_2^{(0,0,1)} - p_1^{(0,1,0)} &= 0 \\ 5p_1^{(0,0,1)} - 1 - p_1^{(0,1,0)} &= 0 \end{aligned} \quad (7.10)$$

More over, we can substitute equation (7.7) into (7.10) to obtain $p_1^{(0,0,1)} = 0$. But substituting this result back into Equation 7.7, we get $p_1^{(0,1,0)} = -1$. A contradiction. So we are not at a mixed equilibrium.

We then consider the pure strategy outcomes based on the actions in Case 4 and Case 5. An agent who observes $(0, 1, 0)$ may play $p_1^{(0,1,0)} = 1$ or $p_2^{(0,1,0)} = 1$. In addition, an agent who observes $(0, 0, 1)$ has options $p_1^{(0,0,1)} = 1$ or $p_2^{(0,0,1)} = 1$. There are four possible pure strategy combinations, and we may calculate the expected payoff for each player, in each scenario. For example, consider $p_1^{(0,1,0)} = 1$ and $p_1^{(0,0,1)} = 1$, where both players will play b_1 regardless of their observation. That means, if Agent 1 observed $(0, 1, 0)$, we will reach one of three possible outcomes: we elect A , B or reach a Tie. Thus, the expected utility will be $20/9$. We repeat these calculations to formulate the outcomes in the matrix below:²

Table 7.5: Expected Utilities for Pure Strategies

		Observes $(0,0,1)$	
		$p_1^{(0,0,1)} = 1$	$p_2^{(0,0,1)} = 1$
Observes $(0,1,0)$	$p_1^{(0,1,0)} = 1$	$(\frac{20}{9}, \frac{14}{9})$	$(\frac{8}{3}, \frac{4}{3})$
	$p_2^{(0,1,0)} = 1$	$(2, 1)$	$(2, \frac{2}{3})$

Notice three of the pure strategies are dominated, leaving only the top left cell as the unique symmetric Nash Equilibrium for the final round. This corresponds to the actions of **voting for the top choice regardless of the nature of the single ballot observed**. This nets an expected utility of $\frac{20}{9}$ if Agent 1 observed a ballot for her second choice, and $\frac{14}{9}$, for her third choice.

Case 6: No agent has cast any ballots, in which case Agent 1's best response is to vote honestly and hope for the best: $p_1^{(0,0,0)} = 1$, with probability $\frac{5}{9}$ of electing A, $\frac{2}{9}$ of getting a tie, $\frac{1}{9}$ of getting B, and $\frac{1}{9}$ of getting C. This results in an expected utility of $\frac{61}{27}$.

Analysis: Preceding Round t

Now that we have an equilibrium analysis of the last round, we extend our analysis to preceding rounds via backward induction. Here, each agent has three actions, and Waiting is not a clearly dominated action: $p_1^{(0,0,1)} + p_2^{(0,0,1)} + p_\emptyset^{(0,0,1)} = 1$, and $p_1^{(0,1,0)} + p_2^{(0,1,0)} + p_\emptyset^{(0,1,0)} = 1$.

²While this matrix resembles a normal form game, it is only analogous to one. The rows and columns represent information states that the players find themselves in, and the actions they may take. The cell represents the payoff to the player for a particular pure strategy the agents symmetrically pursue.

If Agent 1 takes the Wait action \emptyset , she proceeds into the information state $(t + 1, \mathbf{s}^+)$ of the controlled game, where \mathbf{s}^+ is obtained from \mathbf{s}_t by adding a number of ballots up to an including the number of uncommitted voters, representing new ballots cast this turn by the other voters. If Agent 1 casts a ballot b , then she enters into the uncontrolled game $(t + 1, \mathbf{s}^+)$ (see Section 7.4.2).

Case 1: 2 ballots for the same alternative.

Agent 1's vote is irrelevant.

Case 2: 2 ballots for different alternatives.

Agent 1 breaks ties in favor of the better option.

Case 3: Agent 3 (WLOG) casts the only vote, for $A = b_{1,1}$

Agent 1 also votes A and gets A as the outcome.

Case 4: Agent 3 casts the only vote, for $B = b_{1,2}$

As before, we may lay out the possible actions of each agent, based on the possible types of Agent 2 (recall if Agent 2 is type B, the outcome is decided regardless of the actions Agent 1):

Table 7.6: Utility breakdown if Agent 2 is Type A

		Agent 2		
		A	B	\emptyset
Agent 1	A	$A(3, 3)$	$B(2, 2)$	$A(3, 3)$
	B	$B(2, 2)$	$B(2, 2)$	$B(2, 2)$
	\emptyset	$A(3, 3)$	$B(2, 2)$	$*(H, H)$

Table 7.7: Utility breakdown if Agent 2 is Type C

		Agent 2		
		C	A	\emptyset
Agent 1	A	$tie(5/3, 5/3)$	$A(3, 2)$	$A(3, 2)$
	B	$B(2, 0)$	$B(2, 0)$	$B(2, 0)$
	\emptyset	$B(2, 0)$	$A(3, 2)$	$*(H, L)$

Importantly, the outcome designated as $*$ represents the outcome computed in the inductive step for the next round, where the expected utility for a player who observes a ballot for her second choice is H , or is L if a ballot for her last choice is observed ($H > L$, and $H > 2$). If the current round is $t = T - 1$, then $H = \frac{20}{9}$ and $L = \frac{14}{9}$.

As before, we can write equations for expected utilities and solve to show that $b_2^{(0,1,0)}$ is dominated by $b_\emptyset^{(0,1,0)}$, if $H \geq 2$. We solve the remaining equalities in conjunction with Case 5 below.

$$E(u|b_1^{(0,1,0)}) = \frac{1}{3}(3p_1^{(0,1,0)} + 2p_2^{(0,1,0)} + 3p_\emptyset^{(0,1,0)}) + \frac{1}{3}\left(\frac{5}{3}p_1^{(0,0,1)} + 3p_2^{(0,0,1)} + 3p_\emptyset^{(0,0,1)}\right) + \frac{1}{3}(2) \quad (7.11)$$

$$\begin{aligned} E(u|b_2^{(0,1,0)}) &= \frac{1}{3}(2) + \frac{1}{3}(2) + \frac{1}{3}(2) \\ &= 2 \end{aligned} \quad (7.12)$$

$$E(u|b_\emptyset^{(0,1,0)}) = \frac{1}{3}(3p_1^{(0,1,0)} + 2p_2^{(0,1,0)} + Hp_\emptyset^{(0,1,0)*}) + \frac{1}{3}(2p_1^{(0,0,1)} + 3p_2^{(0,0,1)} + Hp_\emptyset^{(0,0,1)}) + \frac{1}{3}(2) \quad (7.13)$$

Since we know that $p_1^{(0,1,0)} + p_2^{(0,1,0)} + p_\emptyset^{(0,1,0)} = 1$ and $p_1^{(0,0,1)} + p_2^{(0,0,1)} + p_\emptyset^{(0,0,1)} = 1$, we can simplify Equations 7.11 and 7.13 as follows:

$$E(u|b_1^{(0,1,0)}) = \frac{1}{3}(2 + p_1^{(0,1,0)} + p_\emptyset^{(0,1,0)}) + \frac{1}{3}\left(\frac{5}{3} + \frac{4}{3}p_2^{(0,0,1)} + \frac{4}{3}p_\emptyset^{(0,0,1)}\right) + \frac{1}{3}(2) \quad (7.14)$$

$$E(u|b_\emptyset^{(0,1,0)}) = \frac{1}{3}(2 + p_1^{(0,1,0)} + (H - 2)p_\emptyset^{(0,1,0)}) + \frac{1}{3}(3 - p_1^{(0,0,1)} + (H - 3)p_\emptyset^{(0,0,1)}) + \frac{1}{3}(2) \quad (7.15)$$

Let us examine the expected utilities $E(u|b_2^{(0,1,0)})$ and $E(u|b_\emptyset^{(0,1,0)})$. Suppose the former is greater than the latter. Then $p_2^{(0,1,0)} > 0$, but by Equation 7.15, $E(u|b_\emptyset^{(0,1,0)}) > 2$, and we reach a contradiction. Suppose that the two are equal, then $p_2^{(0,0,1)} = 0$, which is also a contradiction. Therefore, we may assume that playing $b_2^{(0,1,0)}$ is a dominated strategy, and $p_2^{(0,1,0)} = 0$.

Case 5: Agent 3 casts the only vote, for $C = b_{1,3}$

Agent 2 may be of one of three types. If Agent 2 is Type C , then they will vote for C and Agent 1's action is irrelevant, and they get utility 0. The following tables break down the utility of Agent 1's actions for the other two cases:

Table 7.8: Utility breakdown if Agent 2 is Type A

		Agent 2		
		A	B	\emptyset
Agent 1	A	$A(3, 3)$	$tie(5/3, 5/3)$	$A(3, 3)$
	B	$tie(5/3, 5/3)$	$B(2, 2)$	$B(2, 2)$
	\emptyset	$A(3, 3)$	$B(2, 2)$	$*(L, L)$

Table 7.9: Utility breakdown if Agent 2 is Type B

		Agent 2		
		B	C	\emptyset
Agent 1	A	$tie(5/3, 5/3)$	$C(0, 2)$	$C(0, 2)$
	B	$B(2, 3)$	$C(0, 2)$	$B(2, 3)$
	\emptyset	$B(2, 3)$	$C(0, 2)$	$*(L, H)$

We formulate expected utilities as before. We utilize Gambit [88] to solve this subgame for the $t = T - 1$ case, and find that $p_2^{(0,0,1)} = 0$. Using this information (see Section 7.4.1 for details), we may solve the system of equations exactly to obtain

$$p_\emptyset^{(0,1,0)} = -\frac{3H - 3L - 1}{24H - 3L - 71} \quad (7.16)$$

$$p_\emptyset^{(0,0,1)} = \frac{3H - 3L - 8}{24H - 3L - 71} \quad (7.17)$$

We can substitute this solution into the Expected Utility calculations in Equations (7.15):

$$E(u|b_\emptyset^{(0,1,0)}) = \frac{(H-3)}{3}p_\emptyset^{(0,1,0)} + \frac{(H-2)}{3}p_\emptyset^{(0,0,1)} + \frac{7}{3} \quad (7.18)$$

$$= -\frac{(H-3)}{3} \frac{3H-3L-1}{24H-3L-71} + \frac{(H-2)}{3} \frac{3H-3L-8}{24H-3L-71} + \frac{7}{3} \quad (7.19)$$

$$= \frac{4(41H-6L-121)}{3(24H-3L-71)} \quad (7.20)$$

This gives the expected utility of 2.34 for the optimal action (the mixed strategy described above) when observing the vote vector $(0, 1, 0)$ in the second last round. We may also derive the expected utility for Waiting when observing vote vector $(0, 0, 1)$ from the above tables.

$$E(u|b_\emptyset^{(0,0,1)}) = \frac{(L-2)}{3}p_\emptyset^{(0,1,0)} + \frac{(L-3)}{3}p_\emptyset^{(0,0,1)} + \frac{5}{3} \quad (7.21)$$

$$= -\frac{(L-2)}{3} \frac{3H-3L-1}{24H-3L-71} + \frac{(L-3)}{3} \frac{3H-3L-8}{24H-3L-71} + \frac{5}{3} \quad (7.22)$$

$$= \frac{117H-19L-333}{3(24H-3L-71)} \quad (7.23)$$

In particular, for $t = T - 1$ of the controlled game, when observing $(0, 1, 0)$, **Agent 1 should vote b_1 with probability $p_1^{(0,1,0)} = 64/67$ (and Wait otherwise) for an expected utility of 2.34. When observing $(0, 0, 1)$, she should vote b_1 with probability $p_1^{(0,0,1)} = 49/67$ for an expected utility of 1.53.**

Case 6: No ballots observed.

If no ballots are observed, all agents are in the same information set, and we may assume they act symmetrically. We denote the probability that they play their top choice, second choice and Wait as p_1 , p_2 , and p_\emptyset , respectively.

If Agent 1 Waits, then with probability p_\emptyset^2 , we enter the next round with the tally $(0, 0, 0)$, which gives an expected utility of N ($N = \frac{61}{27}$ in round $T - 1$). With probability $2p_\emptyset(1-p_\emptyset)$, we enter the next round with one other ballot cast (uniformly randomly selected between the candidates); each of these outcomes gives an expected utility of 3, H , and L . Finally, with probability $(1-p_\emptyset)^2$, both other agents cast their ballots. There are 9 possible outcomes (all equally likely); Agent 1 gains her top choice in 5 cases, her second choice in

3 cases, and her last choice in 1 case. This gives an expected utility of $\frac{7}{3}$. Therefore, the expected utility of waiting is

$$E(u|b_\emptyset^{(0,0,0)}) = Np_\emptyset^2 + 2p_\emptyset(1 - p_\emptyset)\frac{3 + H + L}{3} + (1 - p_\emptyset)^2\frac{7}{3} \quad (7.24)$$

If Agent 1 votes for b_1 , then with probability p_\emptyset^2 , we enter the uncontrolled game $(t + 1, (1, 0, 0))$, with expected utility U_1 (see Section 7.4.2). With probability $2p_\emptyset(1 - p_\emptyset)$, one other agent has blindly voted, resulting in the vote vector $(2, 0, 0)$ (utility = 3), $(1, 1, 0)$ (utility = $\frac{8}{3}$)³, or $(1, 0, 1)$ (utility = 1). Finally, with probability $(1 - p_\emptyset)^2$, both other agents have blindly voted, giving a utility of $\frac{61}{27}$.

Thus, the expected utility for this action is

$$\begin{aligned} E(u|b_1^{(0,0,0)}) &= U_1p_\emptyset^2 + 2p_\emptyset(1 - p_\emptyset)\frac{1}{3}\left(3 + \frac{8}{3} + 1\right) + (1 - p_\emptyset)^2\frac{61}{27} \\ &= U_1p_\emptyset^2 + \frac{40}{9}p_\emptyset(1 - p_\emptyset) + (1 - p_\emptyset)^2\frac{61}{27} \end{aligned} \quad (7.25)$$

By a similar set of calculations, we get the expected utility for casting a b_2 ballot is

$$\begin{aligned} E(u|b_2^{(0,0,0)}) &= U_2p_\emptyset^2 + 2p_\emptyset(1 - p_\emptyset)\frac{1}{3}\left(2 + \frac{8}{3} + \frac{4}{3}\right) + (1 - p_\emptyset)^2\frac{61}{27} \\ &= U_2p_\emptyset^2 + 4p_\emptyset(1 - p_\emptyset) + (1 - p_\emptyset)^2\frac{49}{27} \end{aligned} \quad (7.26)$$

where U_2 is the expected utility from the uncontrolled game $(t + 1, (0, 1, 0))$, and $U_2 < U_1$. Notice $E(u|b_2^{(0,0,0)})$ is smaller than $E(u|b_1^{(0,0,0)})$ for all values of p_\emptyset . Therefore, we may assume $p_2^{(0,0,0)} = 0$, and $p_1^{(0,0,0)} + p_\emptyset^{(0,0,0)} = 1$.

Let us consider the difference of expected utility from the remaining two options:

$$E(u|b_1^{(0,0,0)}) - E(u|b_\emptyset^{(0,0,0)})$$

³Note that the Condorcet cycle is important here: if the remaining voter is Type *C*, she would strategically vote for *A*.

$$= (U_1 - N + \frac{2}{3}(H + L) - \frac{68}{27})p_0^2 + (\frac{70}{27} - \frac{2}{3}(H + L))p_0 - \frac{2}{27} \quad (7.27)$$

Clearly, if $p_0 = 0$, this would result in a negative value and $E(u|b_1^{(0,0,0)}) < E(u|b_0^{(0,0,0)})$, which is a contradiction. So we know that regardless of the values of H and L , there is a non-zero probability that an agent Waits.

If $t = T - 1$, then $N = U_1 = \frac{61}{27}$ and $H + L = \frac{34}{9}$, which zeroes out the p_0^2 term, and (7.27) becomes $\frac{2}{27}(p_0 - 1)$. Therefore, $p_0 = 1$ and Agent 1 waits.

We carry forward the induction to $t = T - 2$. $N = 61/27$ $U_1 = 2.1739$ and $H + L = 3.8723$. Equation 7.27 becomes $1/27(-2 + 0.2986p_0 - 0.6033p_0^2)$, which is negative for all values of p_0 . Thus, $E(u|b_1^{(0,0,0)}) < E(u|b_0^{(0,0,0)})$, and so Agent 1 waits as well. Trend continues in further rounds of induction.

Therefore, regardless of the number of rounds in the election, the rational voter always **Waits until the last round in the process before casting a sincere ballot for their top choice**. For this arrangement of candidates and voter preferences, Sticker Voting is equivalent to a simultaneous vote.

7.4.1 Utilities for Round t

The expected utilities for playing b_1 , b_2 or b_0 in round t , upon observing a single ballot for C can be calculated as follows:

$$\begin{aligned} E(u|b_1^{(0,0,1)*}) &= \frac{1}{3}(\frac{5}{3}p_1^{(0,1,0)*} + 0p_2^{(0,1,0)*} + 0p_0^{(0,1,0)*}) \\ &\quad + \frac{1}{3}(3p_1^{(0,0,1)*} + \frac{5}{3}p_2^{(0,0,1)*} + 3p_0^{(0,0,1)*}) \end{aligned} \quad (7.28)$$

$$\begin{aligned} E(u|b_2^{(0,0,1)*}) &= \frac{1}{3}(2p_1^{(0,1,0)*} + 0p_2^{(0,1,0)*} + 2p_0^{(0,1,0)*}) \\ &\quad + \frac{1}{3}(\frac{5}{3}p_1^{(0,0,1)*} + 2p_2^{(0,0,1)*} + 2p_0^{(0,0,1)*}) \end{aligned} \quad (7.29)$$

$$\begin{aligned} E(u|b_0^{(0,0,1)*}) &= \frac{1}{3}(2p_1^{(0,1,0)*} + 0p_2^{(0,1,0)*} + \frac{14}{9}p_0^{(0,1,0)*}) \\ &\quad + \frac{1}{3}(3p_1^{(0,0,1)*} + 2p_2^{(0,0,1)*} + \frac{14}{9}p_0^{(0,0,1)*}) \end{aligned} \quad (7.30)$$

At this point, we may use Gambit to solve the game for the $T - 1$ round numerically. We get the following mixed Nash equilibrium: $p_1^{(0,1,0)} = 0.96$, $p_2^{(0,1,0)} = 0$, $p_0^{(0,1,0)} = 0.045$,

and $p_1^{(0,0,1)} = 0.73$, $p_2^{(0,0,1)} = 0$, $p_\emptyset^{(0,0,1)} = 0.27$. This leads to an expected utility of 2.31 for a player who observes a ballot for her second choice, or of 1.53 for a player who observes a ballot for her last choice.

In other words, in the second-to-last round, an agent plays a mixed strategy between playing her top choice and waiting. The probability of waiting is higher if she observes a ballot supporting her last choice.

More importantly, this informs us that playing b_2 is always dominated by another strategy, when observing both $(0, 1, 0)$ and $(0, 0, 1)$. This allows us to calculate the exact solution. If we assume that $p_2^{(0,0,1)*} = 0$, we may substitute

$$p_1^{(0,0,1)} + p_\emptyset^{(0,0,1)} = 1p_1^{(0,1,0)} + p_\emptyset^{(0,1,0)} = 1 \quad (7.31)$$

into the previous expected utilities:

$$\begin{aligned} E(u|b_1^{(0,1,0)*}) &= \frac{1}{3}(3) + \frac{1}{3}\left(\frac{5}{3}p_1^{(0,0,1)*} + 3p_\emptyset^{(0,0,1)*}\right) + \frac{1}{3}(2) \\ &= \frac{1}{3}(3) + \frac{1}{3}\left(\frac{5}{3} + \frac{4}{3}p_\emptyset^{(0,0,1)*}\right) + \frac{1}{3}(2) \\ &= \frac{4}{9}p_\emptyset^{(0,0,1)*} + \frac{20}{9} \\ E(u|b_\emptyset^{(0,1,0)*}) &= \frac{1}{3}(3p_1^{(0,1,0)*} + \frac{20}{9}p_\emptyset^{(0,1,0)*}) \\ &\quad + \frac{1}{3}(2p_1^{(0,0,1)*} + \frac{20}{9}p_\emptyset^{(0,0,1)*}) + \frac{1}{3}(2) \\ &= \frac{1}{3}\left(3 - \frac{7}{9}p_\emptyset^{(0,1,0)*}\right) + \frac{1}{3}\left(2 + \frac{2}{9}p_\emptyset^{(0,0,1)*}\right) + \frac{1}{3}(2) \\ &= -\frac{7}{27}p_\emptyset^{(0,1,0)*} + \frac{2}{27}p_\emptyset^{(0,0,1)*} + \frac{7}{3} \\ E(u|b_1^{(0,0,1)*}) &= \frac{1}{3}\left(\frac{5}{3}p_1^{(0,1,0)*} + 0p_\emptyset^{(0,1,0)*}\right) + \frac{1}{3}(3p_1^{(0,0,1)*} + 3p_\emptyset^{(0,0,1)*}) \\ &= \frac{5}{9}p_1^{(0,1,0)*} + 1 \\ E(u|b_\emptyset^{(0,0,1)*}) &= \frac{1}{3}(2p_1^{(0,1,0)*} + \frac{14}{9}p_\emptyset^{(0,1,0)*}) \\ &\quad + \frac{1}{3}(3p_1^{(0,0,1)*} + \frac{14}{9}p_\emptyset^{(0,0,1)*}) \end{aligned}$$

$$\begin{aligned}
&= \frac{1}{3}\left(2 - \frac{4}{9}p_\emptyset^{(0,1,0)*}\right) + \frac{1}{3}\left(3 - \frac{13}{9}p_\emptyset^{(0,0,1)*}\right) \\
&= -\frac{4}{27}p_\emptyset^{(0,1,0)*} - \frac{13}{27}p_\emptyset^{(0,0,1)*} + \frac{5}{3}
\end{aligned}$$

If we assume the equilibrium strategy is a mixed strategy comprised of the remaining actions, then we may also set $E(u|b_1^{(0,1,0)*}) = E(u|b_\emptyset^{(0,1,0)*})$, and $E(u|b_1^{(0,0,1)*}) = E(u|b_\emptyset^{(0,0,1)*})$, and solving gives us the system of equations:

$$\begin{aligned}
7p_\emptyset^{(0,1,0)*} + 10p_\emptyset^{(0,0,1)*} &= 3 \\
-11p_\emptyset^{(0,1,0)*} + 13p_\emptyset^{(0,0,1)*} &= 3
\end{aligned}$$

This solves to give us the exact solution that verifies with the empirical solution provided by Gambit, $p_\emptyset^{(0,1,0)*} = 3/67$ and $p_\emptyset^{(0,0,1)*} = 18/67$.

Using this same method allows us to compute the exact solution for any values for expected utility obtained for taking the Wait action for any given round. Let H (L) be the expected utility gained by waiting when observing $(0, 1, 0)$ ($(0, 0, 1)$), respectively. The only changes are to the utility calculations for $E(u|b_\emptyset^{(0,1,0)*})$ and $E(u|b_\emptyset^{(0,0,1)*})$ (Equation (7.30)), as follows:

$$\begin{aligned}
E(u|b_\emptyset^{(0,1,0)*}) &= \frac{1}{3}(3p_1^{(0,1,0)*} + 2p_2^{(0,1,0)*} + Hp_\emptyset^{(0,1,0)*}) \\
&+ \frac{1}{3}(2p_1^{(0,0,1)*} + 3p_2^{(0,0,1)*} + Hp_\emptyset^{(0,0,1)*}) + \frac{1}{3}(2) \\
&= \frac{1}{3}(3p_1^{(0,1,0)*} + Hp_\emptyset^{(0,1,0)*}) \\
&+ \frac{1}{3}(2p_1^{(0,0,1)*} + Hp_\emptyset^{(0,0,1)*}) + \frac{1}{3}(2) \\
&= \frac{1}{3}(3 + (H - 3)p_\emptyset^{(0,1,0)*}) \\
&+ \frac{1}{3}(2 + (H - 2)p_\emptyset^{(0,0,1)*}) + \frac{1}{3}(2) \\
&= \frac{(H - 3)}{3}p_\emptyset^{(0,1,0)*} \\
&+ \frac{(H - 2)}{3}p_\emptyset^{(0,0,1)*} + \frac{7}{3}
\end{aligned}$$

$$\begin{aligned}
E(u|b_\emptyset^{(0,0,1)*}) &= \frac{1}{3}(2p_1^{(0,1,0)*} + 0p_2^{(0,1,0)*} + Lp_\emptyset^{(0,1,0)*}) \\
&+ \frac{1}{3}(3p_1^{(0,0,1)*} + 2p_2^{(0,0,1)*} + Lp_\emptyset^{(0,0,1)*}) \\
&= \frac{1}{3}(2p_1^{(0,1,0)*} + Lp_\emptyset^{(0,1,0)*}) \\
&+ \frac{1}{3}(3p_1^{(0,0,1)*} + Lp_\emptyset^{(0,0,1)*}) \\
&= \frac{1}{3}(2 + (L - 2)p_\emptyset^{(0,1,0)*}) \\
&+ \frac{1}{3}(3 + (L - 3)p_\emptyset^{(0,0,1)*}) \\
&= \frac{(L - 2)}{3}p_\emptyset^{(0,1,0)*} + \frac{(L - 3)}{3}p_\emptyset^{(0,0,1)*} + \frac{5}{3}
\end{aligned}$$

We set $E(u|b_1^{(0,1,0)*}) = E(u|b_\emptyset^{(0,1,0)*})$, and $E(u|b_1^{(0,0,1)*}) = E(u|b_\emptyset^{(0,0,1)*})$, and solve:

$$\begin{aligned}
(H - 3)p_\emptyset^{(0,1,0)*} + (H - \frac{10}{3})p_\emptyset^{(0,0,1)*} &= -\frac{1}{3} \\
(L - \frac{1}{3})p_\emptyset^{(0,1,0)*} + (L - 3)p_\emptyset^{(0,0,1)*} &= -\frac{1}{3}
\end{aligned}$$

which gives the solution

$$\begin{aligned}
p_\emptyset^{(0,1,0)} &= -\frac{3H - 3L - 1}{24H - 3L - 71} \\
p_\emptyset^{(0,0,1)} &= \frac{3H - 3L - 8}{24H - 3L - 71}
\end{aligned}$$

7.4.2 The Uncontrolled Subgames

We say Agent 1 enters the uncontrolled game node $(t + 1, \mathbf{s})$ when she has chosen to cast a ballot in round t , resulting in the tally \mathbf{s} (which includes her ballot and other ballots submitted simultaneously in round t).

Round $t + 1$	T	T-1	T-2
Utility	2.26	2.17	2.18

In particular, we are interested in the uncontrolled game $(t + 1, (1, 0, 0))$. If $t + 1 = T$, then we know (due to symmetry) both remaining agents will vote for their top preferences. This gives an expected utility $\frac{61}{27}$ as may be expected.

However, in prior rounds $t + 1 < T$, the remaining agents may be able to coordinate if they happen to vote sequentially. This only matters if the remaining agents have types B and C (a 2 in 9 chance), and depends on the probability of them waiting upon observing the controlled information state $t + 1, \mathbf{s}$. As a result, the expected utility of entering this uncontrolled game is

$$E(u|(t + 1, (1, 0, 0))) \tag{7.32}$$

$$= \frac{1}{27}(2p_{\emptyset}^{t+1,(0,1,0)}p_{\emptyset}^{t+1,(0,0,1)} + 8p_{\emptyset}^{t+1,(0,1,0)} - 10p_{\emptyset}^{t+1,(0,0,1)} + 61) \tag{7.33}$$

where $p_{\emptyset}^{t+1,(0,1,0)}$ and $p_{\emptyset}^{t+1,(0,0,1)}$ are inductively calculated for round $t + 1$ by Equations (7.16) and (7.17). The following table shows the expected utility of entering the uncontrolled game $(t + 1, (1, 0, 0))$, i.e. by casting a sincere ballot in round t after observing no ballots. Notice all are strictly less than $\frac{61}{27}$.

7.5 Discussion & Conclusion

In our two simple instances of Sticker Voting, we observe that rational voter behavior differs dramatically. In the complete information game, voters will play a mixed strategy in the first round, playing either their *second* choice or Waiting; if they chose to Wait, they will break any ties in their favor in the final round, or otherwise vote sincerely. In the incomplete information game, voters will always exercise the Wait option until they reach the final round, during which they vote sincerely.

It is interesting to contrast the two behaviors. The voters in the complete information game know that the other players are rivals, and therefore understand that there is a first-mover disadvantage if they are greedy. Yet there is also an incentive to concede early to secure acceptable compromise. In the incomplete information game, the voter is unsure

as to the nature of the other players. However, more likely than not, one of the other players has the same type as her, so there is an opportunity to signal cooperation. But any incentive to do this is outweighed by the shrewdness of Waiting until the final round, where any other players with the same type as her will naturally coordinate their votes out of self interest. Additionally, in sharp contrast with the complete information game, voting second choice is never exercised as an option.

This type of strategic timing and strategic voting is consistent with observations in the community vote organized by the Evolution Championship Series (EVO) in 2017. EVO is an annual, open format, international fighting game tournament. The tournament features a selection of video games currently popular in the gaming community. Because of the rapidly evolving nature of the medium, the organizers change the featured titles regularly based on community interest.

For the 2017 tournament, the organizers used a voting mechanism to choose one of the titles, in a manner similar to our Sticker Voting model. Voters were directed to cast their ballots by making charitable donations to one of several specially setup accounts on the online charity platform generosity.com, with 100% of the proceeds going to Make-a-Wish International⁴. The winner, naturally, would be the title with the most donations. Donations are capped at \$10,000 per person, and corporate sponsors were not allowed [3]. Since the current funds for each donation drive is clearly displayed on the website, voters have highly accurate information on the interim standings of each candidate [4]. Moreover, many gaming news sites also covered the changing standings prominently [7].

The winner, Ultimate Marvel vs. Capcom 3, received \$71,690 of donations, followed closely by Pokken Tournament at \$66,906 [5]. The other titles garnered substantially fewer donations. The third place title earned about 10% of the second place title, with the top two titles accounting for over 90% of donations. This would be consistent with human models of bounded rationality in all-in auctions. The final tallies are listed in Table 7.10 [4].⁵

Early on in the process, Killer Instinct amassed a sizable lead in the early hours [5]. However, by day 2, UMvC3 and Pokken had already emerged as the main contenders [7]. In a dramatic display of strategic timing, one donor put in \$2,000 in the closing minutes of the election, securing the win for UMvC3 and causing some supporters for Pokken to withdraw their donations[5]. These observations may be interpreted through the lens of

⁴From their website, “the mission of Make-A-Wish International is to grant the wishes of children with life-threatening medical conditions to enrich the human experience with hope, strength and joy.”

⁵There is a discrepancy in donations due to a number of donors canceled their funds when their favorite title failed to win.

Title	Money Raised	Donors
Ultimate Marvel vs. Capcom 3	\$71,640	1,486
Pokken Tournament	\$62,431	742
Killer Instinct	\$6,116	159
Windjammers	\$4,000	50
ARMS	\$1,337	74
Skullgirls	\$920	39
Super SFII Turbo	\$831	22
Nidhogg	\$158	11
Mortal Kombat XL	\$137	8

Table 7.10: EVO 2017 “People’s Choice” Voting Results, as Measured in Charitable Donations (USD). Retrieved March 4, 2017.

our model. Enthusiasm in the Killer Instinct community sought to establish an early lead and solidify its position as a major contender, but were overtaken by UMvC3 and Pokken by the second day. From day 2 onward, as predicted by strategic behavior, the majority of donations went toward either the first or second place candidates. The generous last minute donation could be seen as a dramatic example of strategic timing. By placing such a large donation so close to the end of the election, the donor all but ensured that the action would be pivotal.

Returning to our results, the result of our incomplete information game is in line with the results of Dekel and Piccione [45]. In their model, voters must commit to voting in one of two rounds. More importantly, this decision is made prior to the election, and prior to realizing their own preferences. This is an important distinction from our model, which allows voters to make their decisions dynamically as the election unfolds. For instance, we allow our voters to undertake different voting strategies if they observe other voters committing their ballots in a certain way (for instance, if another voter supports your favorite option); this is a more realistic generalization of the Dekel-Piccione model. Nonetheless, we find that our results mirror theirs — in their model, they find that rational voters will always vote in the second (i.e. final) round.

Battaglini, Morton and Palfrey [16, 17] explore a similar model with 2 candidates, in which voters cast ballots in a fixed sequence. Moreover, voting is costly, and voters may earn a small amount of utility by choosing to abstain. Voters must choose between passing up on this bonus to help the group select the better alternative, or abstaining and trusting in the decisions of others. In their analysis and laboratory experiments, they also remark that later voters benefit from informational effects revealed by earlier voters; while their

model is fundamentally different from ours, their observation parallels our own conclusion.

Finally, in Sandholm and Vulkan’s bargaining game with deadlines [117], rational agents will wait until the final moment before their deadline before acting. Yet, these results appear to be at odds with the incentives offered by the Republican and Democratic Parties in the U.S., who award bonus delegates to states voting later in the primary season.

Our solution is based on the 3-player voting game. We chose this number because this is the smallest number of voters where interesting strategic behavior could occur. A larger number of players enable richer strategies to emerge. For instance, in a 5-player game, it becomes possible for a player’s favorite candidate to become a necessary loser of the election, and therefore, *must* vote strategically.⁶ Even for our 3-player game, a full solution requires complicated calculations, and so performing a similar full analysis of a larger game seems intractable. Thus, it is important to consider to what degree our conclusions may generalize to a larger game. We may still expect Waiting to confer some strategic advantage; or rather, that voting early is likely to allow other players to subsequently gain the upper hand. It is believable that for most preference configurations, that Waiting until the last round remains an equilibrium strategy. However, this solution for the rational voter seem unintuitive when applied to human voters. In real world Sticker Voting venues and in online polls, we do not expect to see all (or even, a majority of) voters deliberating until the last minute to cast their ballots. We know that humans are impatient and place diminishing value on future payoffs. Are these important qualities to model in Sticker Voting? Human voters also place importance on the expressiveness of voting – they gain satisfaction from having expressed their opinion through voting sincerely. It would be interesting to conduct experiments similar to Battaglini, Morton and Palfrey [17] to elicit data on human voting behavior when using the Sticker Voting mechanism.

Additionally, we have made several assumptions about the preference structure and voter behavior for tractability of analysis. What happens when we relax these assumptions? The Condorcet cycle in the preference structure is an important element in at least one of the calculations in the model (see Footnote 3). Do the results hold if such cycles are rare in practice?

One possible model of bounded rationality that may be applied to Sticker Voting is the Quantal Response Equilibrium (QRE) model [89], where players have a nonzero probability of playing each action, defined as a function of the expected payoff of that action. For instance, in the logit equilibrium (LQRE), the probability of playing an action a with expected utility $u(a|\mathbf{a}_{-i})$ where other players are using strategies \mathbf{a}_{-i} is defined as

⁶This is also possible in 3-player games with a lexicographic voting rule. However, exogenous asymmetry between the candidates makes this less attractive to study.

$$Pr(a|\mathbf{a}_{-i}) = \frac{e^{\lambda u(a|\mathbf{a}_{-i})}}{Z}$$

with sharpness parameter λ and normalization constant Z . QRE has also been extended to extensive form games, where the agents' future actions are treated as mixed strategies defined inductively [89]. Other models include incorporating a future discount factor for utilities, or varying the weights of ballots cast during different phases of the election.

Alternatively, it may be interesting to consider a setting where some proportion of voters are impulsive, and will commit to a ballot early in the voting process. How will the presence of such voters affect the behavior of the strategic voters? Will their actions cause a collapse in the “Waiting” equilibrium?

Finally, it would be interesting future work to investigate other models of deliberative agents in Sticker Voting setting. For instance, agents may also make use of history to infer the types of other agents, allowing them to update their beliefs of the distribution of types in population of uncommitted voters, and therefore strategize accordingly.

Chapter 8

Conclusion

As social networks play ever increasing roles in our daily lives, there is an increasing need to understand their effects on human behavior. Our social network filters and shapes the information we receive about the world, and thereby affect the way we act in various circumstances. We focused on a particular social choice scenario — voting. We considered intelligent voters who vote strategically according to the information presented to them, and examined how social network structure may shape their behavior.

This thesis began with the intention to *advance our understanding of the interactions between social network structure and strategic voter behavior, by examining the question of how different network structures alter the flow of information through a network, and thereby alter the aggregate outcomes of independent and strategic voters over time.* We examined how network structure may affect the propagation of information in a social network of skeptical agents, leading to either the convergence or divergence of opinions (Chapter 4). We proposed a framework in the form of several desiderata for modeling voters embedded within such a social network, who must act strategically, but only according to information available to them (Chapter 6). We constructed voter models based on this framework (Chapters 5 and 6), and showed that they qualitatively replicated real world phenomena such as Duverger’s Law (Chapter 5), and the Micromega Rule (Chapter 6). We also showed network homophily leads to the Echo Chamber Effect (Chapter 5), which reduced the rate of strategization in the population, and thereby produces less suitable candidates. Moreover, we showed that by using heuristic models, we could scale our simulation to large populations (Chapter 6). Finally, we also examined how agents may exploit the informational effects of the voting process by strategically timing their ballots (Chapter 7).

8.1 Summary of Main Results

In the following, we summarize the main results of each chapter:

Chapter 4: Opinion Dynamics. We examined a model for opinion dynamics for skeptical agents. In opinion dynamics, agents each possess an initial personal opinion, and influence each others' opinions through repeated interactions with their social network neighbors. We introduced an opinion dynamics model for skeptical agents, where the influence an agent has on another agent is a function of the difference between their opinions: The greater the difference, the greater the skepticism between the agents, and therefore, the less effect they will have on each others' opinions. We studied the degree to which an opinion will spread through a social network by examining the opinions of agents at convergence — i.e. once they have reached a stable state. We showed that agents will quickly converge to an early but interim consensus before taking coordinated action to migrate to a final collective opinion. We found that, when faced with polarizing “extremists” in the networks, high connectivity helped maintain moderate opinions in large parts of the community. Moreover, in homophilic networks where agents were *less* skeptical, opinions may stratify at several distinct, moderate levels at equilibrium, representing a diversity of opinions that was stabilized by more extreme opinions from both ends of the spectrum.

Chapter 5: Voting in Social Networks. We presented a model for strategic voters embedded in a social network. This model extended Iterative Voting to an environment of incomplete information. Voters must utilize only information from their social networks to predict the likely outcomes of the election, and select a ballot that maximizes their expected utility. That is, they are willing to abandon their favorite candidate for a more promising alternative in situations where the latter is much more likely to win. We showed that in such a network, strategization is a good thing: Strategization increased as voters have access to more information, and selects candidates were more likely to produce a higher social welfare for the population. On the other hand, networks that exhibited homophily reduced the level of strategization within the community, and thereby produced less suitable candidates. In popular media, this effect is called “The Echo Chamber Effect”, which is a natural tendency for people to connect with people who are similar to themselves. When voters surround themselves with similar people, it hides opportunities for strategic voting, and so they end up casting ineffective, “wasted” ballots. This may be a contributing factor to the low rate of strategic voting observed in real world elections. Additionally, strategization may lead to the elimination of less popular candidates, as voters revise their votes to less preferred but more promising candidates. This phenomenon is known as Duverger’s Law in political science, and we showed that it does not hold in sparse network structures.

Chapter 6: Heuristic Voter Models. We proposed a general framework for mo-

deling strategic voters who operate with incomplete information, based on the structure of a social network. We defined several boundedly-rational heuristics based on this framework. We analyzed the computational complexity of each heuristic, and gauged their performance according to their fidelity and speed. According to those results, we put forth two heuristics that are both fast and accurate — **TieH**, based on a biased sampling technique; and **Poisson**, which extends Myerson’s Large Poisson Games [99] to a multiple candidate election. To illustrate the effectiveness of our techniques, we applied our heuristics to explore the Micromega rule — an observation in political science that large political parties favor small assemblies. We found that the size of electoral districts is a contributing factor to the Micromega rule in some networks. Fringe candidates retained more support in smaller districts, while larger parties dominated in larger districts.

Chapter 7: Vote Timing. We proposed a new voting mechanism called Sticker Voting, where ballots are cast by placing stickers on favored candidates. Unlike Iterative Voting, the selection of a ballot is a permanent and irrevocable commitment. Moreover, it differs from many other voting methods because the act of voting reveals information to other players, which induces an asymmetry of information available to subsequent voters. Voters may strategize through both the choice of the submitted ballot and the timing of its submission. We introduced and analyzed a model for strategic voter behavior in Sticker Voting for small games with 3 players. We found its equilibrium behaviors in these settings, speculated on how it may generalize in larger settings, and discussed how it reflects human voting behavior.

8.2 Future Directions

While we have explored a number of different aspects of the interplay between social choice and social networks, there are yet many exciting research directions that begin at the intersection of these two fields.

Strategic Models for Other Voting Mechanisms. This thesis focused on plurality voting. Therefore, a natural direction of exploration is to study strategization in other voting systems such as Borda, Bucklin or Single Transferable Vote. While Armstrong and Larson generalize our strategic model to k-Approval [11], it is not clear how to represent strategic voter responses in elections using general position scoring rules. Rational voters cast a ballot based on expected tie probabilities, inferred through their observations. Throughout our thesis, we have presented several reasonable mechanisms by which a sampling of plurality ballots may be transformed into a distribution over the possible

outcomes of the election. For this computation to be remotely feasible, the space of admissible ballots must be small. This no longer holds true for other, richer, voting systems, and therefore, more sophisticated heuristics and computational techniques must be developed to handle these scenarios. What does strategic play look like in other voting systems with incomplete information? How will different network structures affect voter behavior in other voting systems? Do certain voting systems elect more capable candidates in different communities?

Strategic Network Formation. Real world networks are dynamic and changing entities: collaborations and friendships grow and extend the network, while outdated connections fade away. In this thesis, we examined how agents in a network exchange information, influence each other, and exercise strategic behavior. But the structure of the network itself also evolves over time. While economic networks have been studied quite extensively by researchers (see [74] for example), allowing agents to alter the network adds another dimension of strategic play in social choice games. How can agents manipulate the network structure to their advantage? How does competition between agents affect the network formation process? Will this form of strategic behavior benefit or harm the social welfare of mechanisms operating in those networks?

Strategic Timing Models. In Chapter 7, we proposed a model for Sticker Voting where voters were free to strategically choose when to commit their ballots. However, the analysis for even a simple 3-player scenario proved to be a complex endeavor. In order to analyze larger games, we must adopt behavioral models for voters where their actions are computed via heuristics of bounded-rationality. These behavioral models may also allow more human-like behaviors to emerge. One possible model of human behavior we may use is the Quantal Response Equilibrium (QRE) model [89], where any action has a nonzero probability of being played, even if it is not optimal. Therefore, even if waiting until the last round is the optimal strategy, there is a chance that a number of voters will vote early anyways, possibly starting a strategic cascade in the remaining voters. This “impatience” factor may be explicitly accounted for as well, by adding a deliberation cost for remaining active in the game. Finally, we may also change how time is considered in our model: Rather than having discrete time steps, we may consider each round to be defined by the moment that the next agent chooses to act. Redefining the problem this way may prove fruitful for further mathematical analyses.

Real World Data Sets. Hand in hand with the arrival of “big data” comes veritable treasure troves of data about how individuals connect with each other in society, and how individual behaviors and opinions evolve over time. The proprietary nature of this data means this resource has remained largely untapped by researchers of computational social choice. While we must be mindful of the many privacy and ethical concerns when

dealing with such personal pieces of information, responsible access to this data may yield new insight into how individuals influence each other over time. For instance, we may validate the actions of our strategic voter models from Chapters 5 and 6 against human behavior extracted from popular social media platforms alongside information about the social network structure. It may be possible to infer some of this information from publicly available data. For example, sentiment analysis techniques may be applied to Tweets to evaluate the (publicly disclosed, even if implicitly) political allegiance of Twitter users, and the underlying social network may be approximated by examining the list of Twitter Followers. Monitoring the evolution of these opinions over time may reveal how social influence occurs in situ. Moreover, data on *when* humans commit social choice decisions will inform behavioral models of strategic timing being developed in Chapter 7. Indeed, data from public social choice platforms like Doodle show evidence of strategic timing [135, 103]. Similar data may be available from other online platforms that have been harnessed for social choice purposes, such as generosity.com discussed in Chapter 7. Finally, simple laboratory experiments may be devised to tease out ground truth on this aspect of human behavior as well.

Social Influence Evaluation. Many researchers have already examined the problem of detecting influential individuals in the network, or the more specific application to maximize influence in networks (see [77] and [78], for instance). This problem is of obvious benefit to marketers and lobbyists. However, many of these metrics are based on analyzing network structure alone, prior to and independent of exposure to the social stimuli. In this thesis, we examined several models of how this influence is effected in a network and how it materialized as distinct action. This paves the way for an action-based model of social influence, where its impact may be measured through actions taken by the agents rather than inferred through network structure alone. When coupled with community detection techniques, this may be used to enhance the identification of leaders and influential figures in communities. Moreover, this action-based model of influence allows us to detect social pressure exerted within social network whose structure are shaped by more innocuous forces, with applications in policing and detecting instances of cyberbullying.

Social Predictions Equipped with a model of social influence and strategic action, we may then apply our knowledge toward the ultimate goal of predicting social choice outcomes on our network. Our model may be used to take social network structure and social influence into account, and refine the accuracy of forecasting models. We would be able to tease out important and lasting social changes, from more ephemeral fashions, fads and viral phenomena. When coupled with tools such as sentiment analysis, our system may be able to formulate automated predictions that react dynamically to new information and chaotic events. Alternatively, if we consider voting as an objective process, where agents

are cooperating to uncover some hidden ground truth based on noisy signals, then we may also leverage social network structure to help us refine the results of the social choice mechanisms; indeed, we have already made interesting headway in this line of inquiry [127, 73], though this work is outside the scope of this thesis.

In closing, social networks encode the nature of human relationships within a community and contain insights into how information flows through a population. By understanding the nature of these interactions, and how they inform strategic behavior, this thesis sheds light on the underlying interplay between network structure and strategic voting behavior, and paves the way to further work at the intersection of social networks and social choice.

References

- [1] 2016 presidential primary election schedule. <https://www.washingtonpost.com/graphics/politics/2016-election/primaries/schedule/>. Accessed: 2016-02-01.
- [2] Evo 2016. https://en.wikipedia.org/wiki/Evo_2016. Accessed: 2017-03-04.
- [3] Evolution 2017 games revealed! 8 games locked in, vote for the 9th! <http://shoryuken.com/2017/01/24/evolution-2017-games-revealed-8-games-locked-in-vote-for-the-9th/>. Accessed: 2017-03-04.
- [4] Explore | generosity. https://www.generosity.com/explore?filter_title=evo#/search. Accessed: 2017-03-04.
- [5] Minor backlash erupts after last-minute donation secured umvc3 for evo 2017. <http://segmentnext.com/2017/02/09/backlash-donation-umvc3-evo-2017/>. Accessed: 2017-03-04.
- [6] Pólya urn model - wikipedia. https://en.wikipedia.org/wiki/P%C3%B3lya_urn_model. Accessed: 2017-04-16.
- [7] Ultimate marvel vs. capcom 3 leads the race for the ninth game at evo, one title has failed to break \$100... see the current stats here. <https://www.eventhubs.com/news/2017/jan/25/ultimate-marvel-vs-capcom-leads-race-ninth-game-evo-one-title-has-failed-break-100k/>. Accessed: 2017-03-04.
- [8] Daron Acemoglu and Asuman Ozdaglar. Opinion dynamics and learning in social networks. *Dynamic Games and Applications*, 1(1):3–49, 2011.

- [9] Réka Albert and Albert-László Barabási. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74(1):47, 2002.
- [10] Noga Alon, Moshe Babaioff, Ron Karidi, Ron Lavi, and Moshe Tennenholtz. Sequential voting with externalities: herding in social networks. In *Proc. of the 13th ACM conference on Economics and Computation*, page 36, 2012.
- [11] Ben Armstrong and Kate Larson. Approval in the echo chamber. *EXPLORE Workshop, International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2017.
- [12] Hossein Azari, David Parks, and Lirong Xia. Random utility theory for social choice. In *Advances in Neural Information Processing Systems*, pages 126–134, 2012.
- [13] Yoram Bachrach, Omer Lev, Yoad Lewenberg, and Yair Zick. Misrepresentation in district voting. In *Proc. of 25th the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 81–87, 2016.
- [14] Marc Barthélemy. Spatial networks. *Physics Reports*, 499(1):1–101, 2011.
- [15] John J Bartholdi, Craig A Tovey, and Michael A Trick. How hard is it to control an election? *Mathematical and Computer Modelling*, 16(8):27–40, 1992.
- [16] Marco Battaglini. Sequential voting with abstention. *Games and Economic Behavior*, 51(2):445–463, 2005.
- [17] Marco Battaglini, Rebecca Morton, and Thomas Palfrey. Efficiency, equity, and timing of voting mechanisms. *American Political Science Review*, 101(03):409–424, 2007.
- [18] Richard E. Berg-Andersson. The green papers: Democratic detailed delegate allocation - 2012. <http://www.thegreenpapers.com/P12/D-Alloc.phtml>. Accessed: 2016-02-05.
- [19] Duncan Black. On the rationale of group decision-making. *The Journal of Political Economy*, pages 23–34, 1948.
- [20] André Blais. Why is there so little strategic voting in Canadian plurality rule elections? *Political Studies*, 50(3):445–454, 2002.
- [21] Herbert Bless, Klaus Fiedler, and Fritz Strack. *Social cognition: How individuals construct social reality*. Hove and New York: Psychology Press, 2004.

- [22] Marián Boguñá, Romualdo Pastor-Satorras, Albert Díaz-Guilera, and Alex Arenas. Models of social networks based on social distance attachment. *Physical Review E*, 70(5):056122, 2004.
- [23] Paolo Boldi, Francesco Bonchi, Carlos Castillo, and Sebastiano Vigna. Voting in social networks. In *Proc. of the 18th ACM Conference on Information and Knowledge Management*, pages 777–786. ACM, 2009.
- [24] Béla Bollobás, Christian Borgs, Jennifer Chayes, and Oliver Riordan. Directed scale-free graphs. In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 132–139. Society for Industrial and Applied Mathematics, 2003.
- [25] Robert M Bond, Christopher J Fariss, Jason J Jones, Adam DI Kramer, Cameron Marlow, Jaime E Settle, and James H Fowler. A 61-million-person experiment in social influence and political mobilization. *Nature*, 489(7415):295–298, 2012.
- [26] Christian Borgs, Jennifer Chayes, Adam Tauman Kalai, Azarakhsh Malekian, and Moshe Tennenholtz. A novel approach to propagating distrust. In *Internet and Network Economics*, pages 87–105. Springer, 2010.
- [27] Simina Brânzei, Ioannis Caragiannis, Jamie Morgenstern, and Ariel Procaccia. How bad is selfish voting? In *Proc. of 27th Conference on Artificial Intelligence (AAAI)*, 2013.
- [28] Bruce E Cain. Strategic voting in Britain. *American Journal of Political Science*, pages 639–655, 1978.
- [29] Steven Callander. Bandwagons and momentum in sequential voting. *The Review of Economic Studies*, 74(3):653–684, 2007.
- [30] Arthur Carvalho and Kate Larson. A consensual linear opinion pool. *Proc. of 23th the International Joint Conference on Artificial Intelligence (IJCAI)*, 2013.
- [31] Yoon-Sik Cho, Greg Ver Steeg, and Aram Galstyan. Co-evolution of selection and influence in social networks. *Proc. of 25th Conference on Artificial Intelligence (AAAI)*, 2011.
- [32] Samir Chopra, Eric Pacuit, and Rohit Parikh. Knowledge-theoretic properties of strategic voting. In *Logics in Artificial Intelligence*, pages 18–30. Springer, 2004.
- [33] Emily Clough. Strategic voting under conditions of uncertainty: A re-evaluation of duverger’s law. *British Journal of Political Science*, 37(02):313–332, 2007.

- [34] Emily Clough. Talking locally and voting globally duverger’s law and homogeneous discussion networks. *Political Research Quarterly*, 60(3):531–540, 2007.
- [35] James Coleman, Herbert Menzel, and Elihu Katz. *Medical Innovations: A Diffusion Study*. Bobbs Merrill, 1966.
- [36] Josep M Colomer. It’s parties that choose electoral systems (or, duverger’s laws upside down). *Political studies*, 53(1):1–21, 2005.
- [37] M de Condorcet. Essay on the application of analysis to the probability of majority decisions. *Paris: Imprimerie Royale*, 1785.
- [38] Vincent Conitzer. The maximum likelihood approach to voting on social networks. In *51st Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, volume 13, pages 1482–1487. IEEE, 2013.
- [39] Rhodes Cook. Cq roundtable: Gop wants a revamp of primary process. *Congressional Quarterly*, 1997.
- [40] Gary W Cox. *Making votes count: Strategic coordination in the world’s electoral systems*, volume 7. Cambridge Univ Press, 1997.
- [41] Guillaume Deffuant. Comparing extremism propagation patterns in continuous opinion models. *Journal of Artificial Societies and Social Simulation*, 9(3), 2006.
- [42] Guillaume Deffuant, Frédéric Amblard, Gérard Weisbuch, and Thierry Faure. How can extremism prevail? a study based on the relative agreement interaction model. *Journal of Artificial Societies and Social Simulation*, 5, 2002.
- [43] Guillaume Deffuant, Frederic Amblard, and Gerard Weisbuch. Modelling group opinion shift to extreme: the smooth bounded confidence model. *2nd ESSA Conference (Valladolid, Spain)*, 2004.
- [44] Eddie Dekel and Michele Piccione. Sequential voting procedures in symmetric binary elections. *Journal of Political Economy*, 108(1):34–55, 2000.
- [45] Eddie Dekel and Michele Piccione. The strategic dis/advantage of voting early. *American Economic Journal: Microeconomics*, 6(4):162–179, 2014.
- [46] Yvo Desmedt and Edith Elkind. Equilibria of plurality voting with abstentions. In *Proc. of the 11th ACM Conference on Electronic Commerce*, pages 347–356. ACM, 2010.

- [47] Swapnil Dhamal and Y Narahari. Scalable preference aggregation in social networks. In *First AAAI Conference on Human Computation and Crowdsourcing*, 2013.
- [48] Jean-Paul Doignon, Aleksandar Pekeč, and Michel Regenwetter. The repeated insertion model for rankings: Missing link between two subset choice models. *Psychometrika*, 69(1):33–54, 2004.
- [49] Emile Durkheim. *De la division du travail social: étude sur l'organisation des sociétés supérieures*. F. Alcan, 1893.
- [50] Maurice Duverger. Factors in a two-party and multiparty system. *Party Politics and Pressure Groups*, pages 23–32, 1972.
- [51] David Easley and Jon Kleinberg. *Networks, Crowds, and Markets: Reasoning about a Highly Connected World*. Cambridge University Press, 2010.
- [52] Catherine Eckel and Charles A Holt. Strategic voting in agenda-controlled committee experiments. *The American Economic Review*, pages 763–773, 1989.
- [53] Bruce Edmonds. Assessing the safety of (numerical) representation in social simulation. *3rd European Social Simulation Association conference (ESSA)*, 2005.
- [54] Nahed Eltantawy and Julie B Wiest. The arab spring| social media in the egyptian revolution: reconsidering resource mobilization theory. *International Journal of Communication*, 5:18, 2011.
- [55] Ulle Endriss, Svetlana Obraztsova, Maria Polukarov, and Jeffrey S Rosenschein. Strategic voting with incomplete information. *Proc. of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 236–242, 2016.
- [56] Paul Erdős and Alfréd Rényi. On the strength of connectedness of a random graph. *Acta Mathematica Hungarica*, 12(1):261–267, 1961.
- [57] Piotr Faliszewski, Edith Hemaspaandra, and Lane A Hemaspaandra. How hard is bribery in elections? *Journal of Artificial Intelligence Research*, pages 485–532, 2009.
- [58] Hui Fang, Jie Zhang, and Nadia M. Thalmann. A trust model stemmed from diffusion theory for opinion evaluation. *Proc. of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2013.

- [59] Michal Feldman, Nicole Immorlica, Brendan Lucier, and S Matthew Weinberg. Reaching consensus via non-bayesian asynchronous learning in social networks. *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, 28:192–208, 2014.
- [60] William Feller. *An Introduction to Probability Theory and Its Applications*, volume 1. John Wiley & Sons New York, 1968.
- [61] Dan S Felsenthal and Avraham Brichta. Sincere and strategic voters: An Israeli study. *Political Behavior*, 7(4):311–324, 1985.
- [62] Daniel W. Franks, Jason Noble, Peter Kaufmann, and Sigrid Stagl. Extremism propagation in social networks with hubs. *Adaptive Behavior - Animals, Animats, Software Agents, Robots, Adaptive Systems*, 16(4), 2008.
- [63] Serge Gaspers, Victor Naroditskiy, Nina Narodytska, and Toby Walsh. Possible and necessary winner problem in social polls. In *Proc. of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 613–620. International Foundation for Autonomous Agents and Multiagent Systems, 2014.
- [64] John G Geer, Amy Carter, James McHenry, Ryan Teten, and Jennifer Hoef. Experimenting with the balancing hypothesis. *Political Psychology*, pages 49–63, 2004.
- [65] Allan Gibbard. Manipulation of voting schemes: a general result. *Econometrica: Journal of the Econometric Society*, pages 587–601, 1973.
- [66] Mark S Granovetter. The strength of weak ties. *American Journal of Sociology*, pages 1360–1380, 1973.
- [67] Eszter Hargittai, Jason Gallo, and Matthew Kane. Cross-ideological discussions among conservative and liberal bloggers. *Public Choice*, 134(1-2):67–86, 2008.
- [68] Noam Hazon, Raz Lin, and Sarit Kraus. How to change a group’s collective decision. *Proc. of 23th the International Joint Conference on Artificial Intelligence (IJCAI)*, 2013.
- [69] Rainer Hegselmann and Ulrich Krause. Opinion dynamics and bounded confidence models, analysis and simulation. *Journal of Artificial Societies and Social Simulation*, 5(3), 2002.

- [70] Edith Hemaspaandra, Lane A Hemaspaandra, and Jörg Rothe. The complexity of online manipulation of sequential elections. *Journal of Computer and System Sciences*, 80(4):697–710, 2014.
- [71] Melvin J Hinich and Michael C Munger. *Analytical politics*. Cambridge University Press, 1997.
- [72] Peter D Hoff, Adrian E Raftery, and Mark S Handcock. Latent space approaches to social network analysis. *Journal of the American Statistical Association*, 97(460):1090–1098, 2002.
- [73] A. Tsang R. Cohen J. A. Doucette, H. Hosseini and K. Larson. Voting with social networks: Truth springs from argument amongst friends. *EXPLORE Workshop, International Conference on Autonomous Agents and Multi-Agent Systems*, 2015.
- [74] Matthew O Jackson and Alison Watts. The evolution of social and economic networks. *Journal of Economic Theory*, 106(2):265–295, 2002.
- [75] E. T. Jaynes. *Probability Theory: the Logic of Science*. Cambridge University Press, 2003.
- [76] Govinda M Kamath, Eren Şaşoğlu, and David Tse. Optimal haplotype assembly from high-throughput mate-pair reads. In *Information Theory (ISIT), 2015 IEEE International Symposium on*, pages 914–918. IEEE, 2015.
- [77] David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *Proc. of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 137–146. ACM, 2003.
- [78] David Kempe, Jon M Kleinberg, and Éva Tardos. Influential nodes in a diffusion model for social networks. In *Proc. of 23rd International Colloquium on Automata, Languages and Programming (ICALP)*, volume 5, pages 1127–1138. Springer, 2005.
- [79] Adam DI Kramer. The spread of emotion via facebook. In *Proc. of the SIGCHI Conference on Human Factors in Computing Systems*, pages 767–770. ACM, 2012.
- [80] Ulrich Krause. Soziale dynamiken mit vielen interakteuren. eine problemskizze. *Modellierung und Simulation von Dynamiken mit vielen interagierenden Akteuren*, 1997.
- [81] Ulrich Krause. A discrete nonlinear and non-autonomous model of consensus formation. *Communications in Difference Equations*, 2000.

- [82] Ziva Kunda. The case for motivated reasoning. *Psychological Bulletin*, 108(3):480–498, 1990.
- [83] Pierre Simon Laplace. *Essai philosophique sur les probabilités*. Paris Bachelier, 1840.
- [84] Omer Lev and Jeffrey S Rosenschein. Convergence of iterative voting. In *Proc. of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, volume 2, pages 611–618. International Foundation for Autonomous Agents and Multiagent Systems, 2012.
- [85] R Duncan Luce. *Individual choice behavior: A theoretical analysis*. Wiley, 1959.
- [86] Colin L Mallows. Non-null ranking models. i. *Biometrika*, pages 114–130, 1957.
- [87] André C. R. Martins. Continuous opinions and discrete actions in opinion dynamics problems. *International Journal of Modern Physics*, 19, 2008.
- [88] Richard D. McKelvey, Andrew M. McLennan, and Theodore L. Turocy. Gambit: Software tools for game theory. Technical report, Version 0.2006.01.20, 2006. <http://econweb.tamu.edu/gambit/>.
- [89] Richard D McKelvey and Thomas R Palfrey. Quantal response equilibria for extensive form games. *Experimental Economics*, 1(1):9–41, 1998.
- [90] Miller McPherson, Lynn Smith-Lovin, and James M Cook. Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, pages 415–444, 2001.
- [91] Michael F Meffert and Thomas Gschwend. Strategic voting under proportional representation and coalition governments: a simulation and laboratory experiment. *Rationalitätskonzepte, Entscheidungsverhalten und ökonomische Modellierung*, 7, 2009.
- [92] Michael F Meffert and Thomas Gschwend. Polls, coalition signals and strategic voting: An experimental investigation of perceptions and effects. *European Journal of Political Research*, 50(5):636–667, 2011.
- [93] Michael F Meffert, Thomas Gschwend, et al. Strategic voting under proportional representation and coalition governments: a simulation and laboratory experiment. *Unpublished manuscript.(Sonderforschungsbereich 504 No. 07-55)*, 2007.
- [94] Reshef Meir, Omer Lev, and Jeffrey S Rosenschein. A local-dominance theory of voting equilibria. In *Proc. of the 15th ACM conference on Economics and Computation*, pages 313–330. ACM, 2014.

- [95] Reshef Meir, Maria Polukarov, Jeffrey S Rosenschein, and Nicholas R Jennings. Convergence to equilibria in plurality voting. In *Proc. of 24th Conference on Artificial Intelligence (AAAI)*, pages 823–828, 2010.
- [96] Jacob L Moreno. *Who shall survive*, volume 58. JSTOR, 1934.
- [97] Rebecca B Morton and Kenneth C Williams. Information asymmetries and simultaneous versus sequential voting. *American Political Science Review*, 93(01):51–67, 1999.
- [98] Roger B Myerson. Population uncertainty and poisson games. *International Journal of Game Theory*, 27(3):375–392, 1998.
- [99] Roger B Myerson. Large poisson games. *Journal of Economic Theory*, 94(1):7–45, 2000.
- [100] Roger B Myerson and Robert J Weber. A theory of voting equilibria. *American Political Science Review*, 87(01):102–114, 1993.
- [101] John F Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36(1):48–49, 1950.
- [102] Svetlana Obraztsova, Omer Lev, Maria Polukarov, Zinovi Rabinovich, and Jeffrey S Rosenschein. Farsighted voting dynamics. In *AGT@ IJCAI workshop, Buenos Aires, Argentina*, 2015.
- [103] Svetlana Obraztsova, Maria Polukarov, Zinovi Rabinovich, and Edith Elkind. Doodle poll games. pages 876–884, 2017.
- [104] Eric Pacuit and Rohit Parikh. The logic of communication graphs. In *Declarative Agent Languages and Technologies II*, volume 3476, pages 256–269. Springer, 2004.
- [105] H. V. Parunak, T. C. Belding, R. Hilscher, and S. Brueckner. Modeling an managing collective cognitive convergence. *Proc. of the 7th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2008.
- [106] Robin L Plackett. The analysis of permutations. *Applied Statistics*, pages 193–202, 1975.
- [107] Charles R Plott and Michael E Levine. A model of agenda influence on committee decisions. *The American Economic Review*, 68(1):146–160, 1978.

- [108] G. Pólya and F. Eggenberger. Über die statistik verketteter vorgänge. *Zeitschrift Angew. Math. Mech.*, 3(279–289), 1923.
- [109] Julia Poncela-Casasnovas, Mario Gutiérrez-Roig, Carlos Gracia-Lázaro, Julian Vicens, Jesús Gómez-Gardeñes, Josep Perelló, Yamir Moreno, Jordi Duch, and Angel Sánchez. Humans display a reduced set of consistent behavioral phenotypes in dyadic games. *Science Advances*, 2(8):e1600451, 2016.
- [110] Oleksandr Pryymak, Alex Rogers, and Nicholas R. Jennings. Efficient opinion sharing in large decentralised teams. *Proc. of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2012.
- [111] Zinovi Rabinovich, Svetlana Obraztsova, Omer Lev, Evangelos Markakis, and Jeffrey S Rosenschein. Analysis of equilibria in iterative voting schemes. In *Proc. of 29th Conference on Artificial Intelligence (AAAI)*, volume 15, pages 1007–1013. Citeseer, 2015.
- [112] Annemieke Reijngoud and Ulle Endriss. Voter response to iterated poll information. In *Proc. of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, volume 2, pages 635–644. International Foundation for Autonomous Agents and Multiagent Systems, 2012.
- [113] Katharina Reinecke, Minh Khoa Nguyen, Abraham Bernstein, Michael Näf, and Krzysztof Z Gajos. Doodle around the world: online scheduling behavior reflects cultural differences in time perception and group decision-making. In *Proceedings of the 2013 Conference on Computer Supported Cooperative Work*, pages 45–54. ACM, 2013.
- [114] Republican National Committee. 2/27/96 Nicholson to Chair National GOP Task Force on Presidential Primary Schedule. Press release, 1996.
- [115] Amirali Salehi-Abari and Craig Boutilier. Empathetic social choice on social networks. In *Proc. of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 693–700. International Foundation for Autonomous Agents and Multiagent Systems, 2014.
- [116] Laurent Salzarulo. A continuous opinion dynamics model based on the principle of meta-contrast. *Journal of Artificial Societies and Social Simulation*, 9, 2006.
- [117] Tuomas Sandholm, Nir Vulkan, et al. Bargaining with deadlines. In *AAAI/IAAI*, pages 44–51, 1999.

- [118] Mark Allen Satterthwaite. Strategy-proofness and arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217, 1975.
- [119] Carl A Sheingold. Social networks and voting: the resurrection of a research agenda. *American Sociological Review*, pages 712–720, 1973.
- [120] Herbert A. Simon. Bandwagon and underdog effects and the possibility of equilibrium predictions. *Public Opinion Quarterly*, pages 245–253, 1954.
- [121] Sigal Sina, Noam Hazon, Avinatan Hassidim, and Sarit Kraus. Adapting the social network to affect elections. In *Proc. of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 705–713. International Foundation for Autonomous Agents and Multiagent Systems, 2015.
- [122] David Strang and Sarah Soule. Diffusion in organizations and social movements: From hybrid corn to poison pills. *Annual Review of Sociology*, 24:265–290, 1998.
- [123] Smarth Swarup, Andrea Apolloni, and Zsuzsanna Fagyal. A model of norm emergence and innovation in language change. *Proc. of the 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2011.
- [124] Maor Tal, Reshef Meir, and Ya’akov Kobi Gal. A study of human behavior in online voting. In *Proc. of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 665–673. International Foundation for Autonomous Agents and Multiagent Systems, 2015.
- [125] Ferdinand Tönnies. *Gemeinschaft und Gesellschaft: Abhandlung des Communismus und des Socialismus als empirischer Culturformen*. Fues, 1887.
- [126] Silvano Colombo Tosatto and Marc Van Zee. Bridging social network analysis and judgment aggregation. In *International Conference on Social Informatics*, pages 17–33. Springer, 2014.
- [127] Alan Tsang, John Doucette, and Hadi Hosseini. Voting with social influence: Using arguments to uncover ground truth. In *Proc. of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2015.
- [128] Alan Tsang and Kate Larson. Opinion dynamics of skeptical agents. In *Proc. of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 277–284. International Foundation for Autonomous Agents and Multiagent Systems, 2014.

- [129] Alan Tsang and Kate Larson. The echo chamber: Strategic voting and homophily in social networks. In *Proc. of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 368–375. International Foundation for Autonomous Agents and Multiagent Systems, 2016.
- [130] Alan Tsang and Kate Larson. If you like it, then you shoulda put a sticker on it. *EXPLORE Workshop, International Conference on Autonomous Agents and Multi-Agent Systems*, 2017.
- [131] Hans Van Ditmarsch, Jérôme Lang, and Abdallah Saffidine. Strategic voting and the logic of knowledge. In *Proc. of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, volume 3, pages 1247–1248. International Foundation for Autonomous Agents and Multiagent Systems, 2012.
- [132] Duncan J Watts and Steven H Strogatz. Collective dynamics of small-world networks. *Nature*, 393(6684):440–442, 1998.
- [133] Ling Heng Wong, Phillipa Pattison, and Garry Robins. A spatial model for social networks. *Physica A: Statistical Mechanics and its Applications*, 360, 2006.
- [134] Lirong Xia and Vincent Conitzer. Stackelberg voting games: Computational aspects and paradoxes. In *Proc. of 24th Conference on Artificial Intelligence (AAAI)*, 2010.
- [135] James Zou, Reshef Meir, and David Parkes. Strategic voting behavior in doodle polls. In *Proc. of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, pages 464–472. ACM, 2015.