

VOTING AND HISTORICAL GAMES

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF COMPUTER SCIENCE
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Michael Andrew Munie

December 2010

Preface

For a group of agents to make a good decision, we must be able to choose a good decision making method. In this thesis, we first study the problem of choosing decision making methods by taking social choice functions and providing quantitative measures on how well they might perform against a population distribution. Then we study a common, possibly the most common, decision making method online: voting to rate products. In this domain, we provide a new class of models and prove that certain convergence results hold. Finally we prove a connection between this new model and traditional voting methods.

In the first half of the thesis we develop a means of comparing various social choice functions with regard to a desired axiom by quantifying how often the axiom is violated. To this end, we offer a new framework for measuring the quality of social choice functions that builds from and provides a unifying framework for previous research. This framework takes the form of what we call a “violation graph.” Graph properties have natural interpretations as metrics for comparing social choice functions. Using the violation graph we present new metrics, such as the minimal domain restriction, for assessing social choice functions and provide exact and probabilistic results for voting rules including plurality, Borda, and Copeland. Motivated by the empirical results, we also prove asymptotic results for scoring voting rules.

These results suggest that voting rules based on pairwise comparison (ex: Copeland) are better than scoring rules (ex: Borda count). They also suggest that although we can never fulfill our desired set of axioms, the frequency of violation is so small that with even a modest number of voters we can expect to never violate our axioms.

In the second half of the thesis, we define a new class of games called *historical influence games (HIGs)*. HIGs are infinite games in which agents take turns round-robin style

choosing a value for a single-dimensional variable. The payoff at each stage to each agent is a monotonically decreasing function of the distance between two quantities: a weighted average of past values chosen by all agents, and some fixed ideal value personal to that agent. The overall payoff to the agent is the limit average of the stage payoffs. We show that myopic strategies form a subgame perfect Nash equilibrium in HIGs. We also introduce certain smoothness constraints on how the impact of a given action changes over time, constraints which define the class of *valid* HIGs. We show that for valid HIGs, under myopic play the limit average value converges to what we call the *central value*, which is the median of the agents' ideal values jointly with certain societal focal points. As a side effect we show a polarization theorem: After a finite period, almost all agents settle on one of the extreme values. Finally, we show a tight connection between valid HIGs and the class of Moulin strategy-proof voting rules in single-peaked domains.

Acknowledgements

I would like to thank my advisor, Yoav Shoham, for his advice, and and more than anything else for asking me the hard questions that I preferred to ignore. Thank you also to my committee: Ilya Segal, Jure Leskovec, Jeremy Bulow, and Balaji Prabhakar. Thank you also to my undergraduate advisors at UIUC, Steve Lavelle and Bruce Reznick, who taught me how to do research to Computer Science and Math.

I also would like to thank my fellow multiagent group members: Thuc, Chris, Nicolas, Ashton, Sam, and Rob, among others. They were a true guiding force in my PhD. They made coming into the office worthwhile instead of feeling like a chore. A special thanks to Thuc, who I've know even since meeting him at many PhD program visit weekends in 2004. I have a feeling I'll be seeing more of him after graduation.

Finally, I would like to thank my friends and family. My friends have kept me sane, distracted me when necessary, and helped me stay on track. Having them made the last six years much more enjoyable than they would have been otherwise. My family, though back in Illinois, has been with me much longer than just my graduate career, and I would like to thank them for making me who I am today.

Jean: thank you for everything. From deciding on where to go for graduate school, to finishing my thesis, you've been there every step of the way. Thank you for keeping me motivated, focused, and honest.

Contents

Preface	iv
Acknowledgements	vi
1 Introduction	1
1.1 Motivation	1
1.2 Comparing Social Choice Functions	1
1.3 Infinite Online Aggregation	2
2 Background	4
2.1 Change in Wikis	4
2.2 Manipulation Frequency in Social Choice	5
2.3 Strategy-proof Voting Domains	7
3 Evaluation of Voting Systems	9
3.1 Motivation	9
3.2 Background	10
3.2.1 Social Choice Functions: Maskin Monotonicity versus Strategy Proofness	11
3.2.2 Voting Rules	12
3.3 Related Work	13
3.4 Violation Graph	15
3.5 Empirical Results	16
3.5.1 Edge Ratio Results	17

3.5.2	Node Ratio Results	23
3.5.3	Domain Restriction Results	23
3.6	Asymptotic Results	24
3.7	Conclusion and Future Work	27
4	Dynamic Voting in Single Peaked Domains	31
4.1	Introduction	31
4.1.1	Motivation	31
4.1.2	Overview of model and results	32
4.1.3	Related work	33
4.1.4	Structure of this chapter	34
4.2	The Model	34
4.3	Myopic Strategies in Historical Influence Games	37
4.4	Central Value	38
4.4.1	Experimental Results	38
4.4.2	The Central Value	44
4.4.3	Convergence to the Central Value under Myopic Strategies	47
4.4.4	Polarization	49
4.5	Correspondence with Voting	49
4.5.1	Moulin Single-Peaked Voting Rules	49
4.6	Conclusion	51
4.7	Appendix	51
5	Conclusion and Future Work	57
	Bibliography	61

List of Tables

List of Figures

3.1	Violation Graph for Two Voters and Three Candidates under Plurality (and Plurality with Runoff)	17
3.2	Violation Graph for Two Voters and Three Candidates under Borda (and Copeland)	18
3.3	Violation Graph for Two Voters and Three Candidates under Maxmin . . .	18
3.4	Violation Graph for Three Voters and Three Candidates under Plurality . . .	19
3.5	Violation Graph for Three Voters and Three Candidates under Plurality with Runoff	19
3.6	Violation Graph for Three Voters and Three Candidates under Maxmin (and Copeland)	20
3.7	Violation Graph for Three Voters and Three Candidates under Borda	20
3.8	Violation Graph for Four Voters and Three Candidates under Plurality . . .	21
3.9	Violation Graphs for Four Voters and Three Candidates under Copeland . .	22
3.10	Comparative Edge Ratio Results for Three Candidates (Two views of the same data)	28
3.11	Comparative Edge Ratio Results for Three Voters	29
3.12	Comparative Node Ratio Results for Three Candidates	29
3.13	Proportion of Nodes Removed with Domain Restriction for Three Candidates	30
4.1	Convergence under $I(t, t') = \frac{1}{t'}$ with myopic play	39
4.2	Convergence under $I(t, t') = \frac{n\beta_{(t \bmod n)+1}}{t'}$ with myopic play	40
4.3	Convergence under $I(t, t') = (1 - \alpha_{t-1}) \prod_{j=t}^{t'-1} \alpha_j$ with myopic play	41
4.4	Convergence under $I(t, t') = (1 - \alpha_{t-1}) \prod_{j=t}^{t'-1} \alpha_j$ with myopic play	42
4.5	Convergence under $I(t, t') = (1 - \alpha_{t-1}) \prod_{j=t}^{t'-1} \alpha_j$ with myopic play	43

4.6	Non-convergence example	44
4.7	Possible mean values of the set of agent goals given $C(N)$, with $ N = 30$. The upper bound is in green, and the lower bound is in blue.	46
5.1	2-D historical influence game <i>in silico</i> experiment with two random agents .	59
5.2	2-D historical influence game <i>in silico</i> run with three random agents	59
5.3	2-D historical influence game <i>in silico</i> run with four random agents	60

Chapter 1

Introduction

1.1 Motivation

When designing a system within which a group of agents will make a decision, it is important to consider the practicality of the methods. Not only this, but there are many systems that already exist ‘in the wild’ and it is important to be able to analyze the results ex post and understand the preferences of the group as a whole. We examine two different ways of making decisions about voting and decision making. The first starts with theoretical models and then we proceed to make quantitative comparisons between them. The second starts with an observed decision making process, models it theoretically, and then proves convergence and other desirable theoretical properties about it.

1.2 Comparing Social Choice Functions

Social choice theory is concerned with the extent to which individual preferences can be aggregated into a social decision in a satisfactory manner. This is a problem that faces groups of people, in size ranging from small groups to large nations, every day. There are many strong impossibility results but nevertheless voting continues to be used. How should we consider this apparent contradiction? First we consider a new framework for quantitatively evaluating and comparing social choice functions.

There are three key dimensions on which our evaluation depends: the type of voting

‘errors’ we seek to quantify, the specific metric by which we measure the errors, and the population distribution of voters. We show that our three dimensions have simple and natural graph interpretations. By introducing the “violation graph,” we can vary the dimensions of the problem while applying the same algorithms to produce the computational results.

This framework allows to compare commonly used social choice functions via the aforementioned dimensions using violation graphs. For small and medium sized domains, we provide both exact and sampled results. Interestingly, these results show that (among other things) voting rules based on pairwise comparison, such as Copeland and Maxmin, perform better than score-based voting rules, such as plurality and Borda. Motivated by these empirical results, we also prove certain asymptotic results that help us better understand the satisfaction of monotonicity for scoring rules when the number of voters tends to infinity.

1.3 Infinite Online Aggregation

The second approach we take is to model voting behavior as seen occurring at a large scale online. Not only is the voting ongoing, but agents have different influence on the final result. In some cases the aggregation of votes is explicit. Consider the Amazon.com, in which users can at any time rate a product (possibly rating the same product many times), as well as see the aggregate rating for the product compiled from all ratings provided thus far.

In others it can be implicit. Consider the online collaborative encyclopedia Wikipedia as an example, and specifically an entry taking a stance on whether global warming is a real phenomenon. Imagine a value in the $[0,1]$ range that captures the degree to which the current content argues for the realism of global warming (0 being complete denial of it, 1 taking it as certainty, and values in between capturing more nuanced opinions). Users who are passionate about the topic care about the value over time since that reflects the cumulative impact the entry has on public opinion and policy.

To account for scenarios such as the ones above, we define a new class of games called *historical influence games (HIGs)*. Informally speaking, HIGs are infinite games in which agents take turns round-robin style choosing a value for a single-dimensional variable. The

payoff at each stage to each agent is a monotonically decreasing function of the distance between the agent's personal ideal value, and a weighted average of past values chosen by all agents. The overall payoff to the agent is the limit average of the stage payoffs.

Surprisingly, under smoothness constraints, limit average aggregated value converges to what we call the *central value*, which is the median of the agents' ideal values jointly with certain societal focal points. This point is tightly connected to the class of Moulin strategy-proof voting rules in single-peaked domains.

These two approaches: making quantitative decisions about traditional social choices rules and proving theoretical soundness of the sort of voting used in practice, are the focus of this thesis.

Chapter 2

Background

Social choice functions are qualitatively characterized¹ by two strong and well known results, the Gibbard-Satterthwaite [18, 50] and the Muller-Satterthwaite [36] theorems. The Gibbard-Satterthwaite theorem states that for every social choice function that is non-dictatorial and whose range has at least three alternatives, there exists at least one profile that can be manipulated by a single voter. In a similar result, the Muller-Satterthwaite theorem states that for every social choice function that is weakly Pareto efficient and non-dictatorial there is some profile where a monotonic change with respect to the winner would cause that alternative to lose.

There have been different approaches proposed in light of these impossibility results. We discuss three relevant threads of research ranging from the very practical to the very theoretical. The computer science community has researched convergence and change in wikis and we discuss this in Section 2.1. There has been work in economics using computer simulations to study the frequency of axiomatic violation for social choice functions and we discuss this in section 2.2. Finally we discuss strategy-proof voting in Section 2.3.

2.1 Change in Wikis

In the group of work on wikis, most papers concentrate on Wikipedia, the most prominent of all wikis. One important thread of this work has been on how to understand the properties

¹Social welfare functions were similarly characterized by Arrow's [2] seminal paper.

of the social community that has sprung up around Wikipedia. There has been comparisons to print media [19], estimations of the expertise of Wikipedia authors [55], and studies of the elite verses common users [26]. Others have researched how Wikipedia articles accrue edits proportional to their previous edits [59] and even comparisons across different language versions [48].

Another thread of research has been on how to improve the quality of Wikipedia through practical rule changes, community management, and user interface design. A large and diverse user base seems to be correlated to the quality of articles [1] while the costs of conflict and coordination [27] also go up with the population. More general studies have been done on how policy influences communities with mass participation, and looks at examples from Wikipedia to try to deduce design guidelines [28, 57]. Self-supervised information extraction combined with innovative interface design was studied in [58] and specifically designed to encourage communal content creation.

The AI community has recently begun to leverage Wikipedia as a means to solve problems in NLP and information retrieval, among other things. In [61], Wikipedia is used to evaluate ontologies. It has also been used to learn semantic relatedness and semantic interpretation in [60, 17].

2.2 Manipulation Frequency in Social Choice

While there has been research by Fedrizzi [14] and others on comparing voting systems by determining which of a set of desirable requirements for social choice functions are satisfied, the previous work most relevant here has focused on quantifying the proportion of profiles that are manipulable. Although we focus on monotonicity instead of manipulability, the quantitative nature of this work is related. The first conjecture that the proportion of manipulable profiles was small was made by Pattanaik in [43]. An asymptotic bound on unstable profiles under the plurality rule was proved by Peleg in [44], and this work was extended and sharpened by Frisrup and Kleiding [16] and also Slinko [53] where the proportion of manipulable profiles was shown to be $\Theta(\frac{1}{\sqrt{n}})$, with n being the number of voters and the constant factor depending on the number of candidates m . In a similar line of work, Nurmi [42] focuses on quantifying how often voting rules select the same winners

and how often the set of winners overlap.

Other research approached the problem using computer simulations, with early work by Nitzan [41] studying situations with up to 90 voters on plurality and Borda count by estimating the proportion of manipulable profiles. Kelly [23] empirically argues that with three candidates, as the number of voters tends towards infinity, the probability (when we assume a uniform probability distribution over profiles) of not having a Condorcet winner tends towards .088. As the number of candidates tends towards infinity, the probability of not having a Condorcet winner tends towards 1.

By looking at three metrics in addition to the proportion of manipulable profiles, Smith[54] simulated common social choice functions looking for strategy-proofness violations. A slightly different vein of work by Kelly [25, 24] examined the distribution of manipulable profile proportions for random social choice functions that were onto and non-dictatorial and, using computer simulation, found that the average proportion of manipulable profiles was almost one.

Much of the related work has assumed that each profile is equally likely (referred to as the Impartial Culture (IC) assumption) and thus interprets this proportion as a probability, as we do later in this work. However, other work has assumed that the identities of voters are anonymous as well as impartial (IAC) giving us a different distribution over profiles. Working in this distribution, Favardin [13, 12] examines both individual and coalitional manipulation (while also allowing for counter-threats) and finds that Borda count is more vulnerable to strategic manipulation than Copeland's rule. Also considering coalitional manipulation, Lepelley [29] generalizes the special cases of IAC and IC to the space of Polya-Eggenberger probability models and runs simulations to calculate violation probabilities. Slinko [47] evaluates some voting rules including approval voting, Borda, and plurality, by looking at the asymptotic average threshold coalition size and provides analytic results.

There has also been work using non-uniform distributions, *cf.*: Favardin and Slinko [12, 47].

Because the domain of all profiles violates strategy-proofness, there is a rich body of work by Kelly and others [7, 3, 39] among others, looking at various restricted domains where strategy-proofness holds. Simple majority voting has been shown by Maskin [10] to be strategy-proof in the restricted domain where the Condorcet winner exists, but the

Condorcet winner does not always exist, and the probability of its existence even tends towards 0 as the number of alternatives increases, as shown in Fishburn [15]. Recent work by Bochet [6] gives the minimal domain restriction for a non-dictatorial, Pareto optimal, and strategy-proof (or Maskin monotonic) social choice function to exist. However, the social choice functions that satisfy these requirements in this minimal domain have a very strong dictatorial flavor. A different approach is to take a known domain where strategy-proofness holds and calculate the minimal monotonic extension as in Thomson [56] and Erden [11].

We focus on completing the picture in light of previous research. While most previous work has concentrated on strategy-proofness, we focus on Maskin's monotonicity and quantifying how much it is violated. In addition to using metrics used in previous work, we also introduce the metric of domain restriction for measuring the quality of a social choice function.

2.3 Strategy-proof Voting Domains

There is a strong technical connection with social choice theory, specifically with Moulin's work on strategy-proof voting rules in single-peaked domains, where the set of strategy-proof, anonymous, and efficient voting rules is characterized exactly in [34]. As we will show in chapter 4, there is a tight connection between these voting rules and historical games. So much so, that our work can be seen as a constructive version of Moulin's characterization.

Following Moulin, there has been a string of work further pushing the bounds of single-peaked voting characterizations. [51] have dealt with location on a network lying in two dimensional space. Others, [40], [39] have shown how far single-peaked domains can be generalized and still have strategy-proofness hold.

Another related connection is to the Mean Voter Theorem [20], [5]. This theorem states that if voters have single peaked preferences in a single dimension, then the position of maximum utility for the median voter will beat any alternative in a pairwise election.

Variations of median voter theory [8, 46] have suggested models explaining why voting for extreme parties with no chance of winning can benefit the voters by allowing them

to communicate and influence the subsequent positions of the winning parties in repeated election settings. Other work has shown how voting can be a process of information aggregation, like Condorcet’s jury theorem[45, 38], and the work viewing voting as maximum likelihood estimation [9].

It is also worthy to mention a possible connection of lemma 3 to results on prediction markets. In the prediction market literature, there is no specific dynamic outlined for trading, but there has been work examining the possible equilibrium points of a market once no further trading will occur. In this setting a large (continuous) population is indexed over the unit interval, each with a subjective belief and a budget. The equilibrium price is the price at which the budget of the traders whose beliefs lie on both sides of the equilibrium is balanced. Specifically, Manski [31] showed that if the price is π_m , then the mean belief of the traders lies in the open interval $(\pi_m^2, 2\pi_m - \pi_m^2)$. Our conditions for an equilibrium are similar once the partition point are added in and counted as agents, and the bound is analogous to a continuous version of our setting for simple averaging utility. However, the major difference is that our equilibrium point, the central value, is the result of myopic dynamics, and in prediction markets it is simply a stable point. The connection here is a point for further research.

Our work on historical games also has some similarities to fictitious play [49] and the best response dynamics used in potential games [33]. For a good overview of both see [52]. Fictitious play does not apply here because we don’t have a stage game. Even if we changed the model to be a repeated normal form game, the stage game payoffs would be dependent on the history and change every round, so the Nash equilibrium would also change. Further, even if fictitious play did apply, none of the known conditions [4] for convergence apply here.

Chapter 3

Evaluation of Voting Systems

3.1 Motivation

Social choice theory is concerned with the extent to which individual preferences can be aggregated into a social decision in a satisfactory manner. The idea of preference aggregation can be implemented by designing a *social choice function*, which takes a preference profile (a collection of the orderings of each individual over alternatives that express the preference of that individual) as input and returns a socially optimal alternative. The satisfaction of our goal to aggregate individual preferences is then described by the fulfillment of a set of socially desirable conditions on the social choice function.

Unfortunately, social choice theory is sometimes called a “science of the impossible” because of the domination of various impossibility theorems. For example, the Muller-Satterthwaite theorem [36] states that every social choice function that is weakly Pareto efficient and monotonic has a dictator, whose preferred alternative is always chosen by the social choice function. One can test by a small computer program [30] that out of 3^{36} possible social choice functions with 2 voters and 3 candidates, there are only 17 monotonic functions, all of which have extremely unsatisfactory properties such as dictatorship or constant output.

One perspective is that monotonicity is too strong and should be dropped as a desirable condition. However, the strength of this condition can be viewed as a positive or a negative, depending on if the framework being used is qualitative or quantitative. Qualitatively,

we are unable to differentiate between the quality of social choice functions because they all violate our criteria. In our quantitative work, the strength is a positive because it better allows us to discriminate between social choice functions which we use to ask to what extent do different social choice functions violate this condition.

We offer a new framework for quantitatively evaluating and comparing social choice functions. There are three key dimensions on which our evaluation depends: the type of violation we seek to quantify, the specific metric by which we measure the violation, and the population distribution over which the violation exists. We choose the first dimension to be monotonicity here and will motivate this choice in the next section. As we will discuss in the paper, our three dimensions have simple and natural graph interpretations. By introducing a model called the “violation graph,” we provide a means of easily varying the dimensions of the problem while applying the same algorithms to produce the computational results.

We then demonstrate our framework using different voting rules as target social choice functions and compare them via the aforementioned dimensions using violation graphs. For small and medium sized domains, we provide both exact and sampled results. Interestingly, these results show that (among other things) voting rules based on pairwise comparison, such as Copeland and Maxmin, perform better than score-based voting rules, such as plurality and Borda. Motivated by these empirical results, we also prove certain asymptotic results that help us better understand the satisfaction of monotonicity for scoring rules when the number of voters tends to infinity.

The structure of our paper is as follows. In Section 3.2, we discuss the foundational material needed to proceed. In Sections 3.3 and 3.4 we review the related work in social choice theory and formally present the violation graph approach. Then in Section 3.5, we provide empirical results and algorithms. In Section 3.6 we prove asymptotic results that support these findings. Finally, we conclude and discuss future work in Section 3.7.

3.2 Background

To lay a foundation, we will first briefly cover the necessary material on social choice.

3.2.1 Social Choice Functions: Maskin Monotonicity versus Strategy Proofness

Formally, let $N = \{1, 2, \dots, n\}$ denote a set of voters, and O denote a finite set of alternatives. Let L be the set of strict total orders over O and denote voter i 's preference as $\succ_i$. A social choice function is a function $C : L^n \mapsto O$.

Not all social choice functions are equally desirable. There are several principles that social choice theorists have argued that social choice functions should ideally adhere to. As we mentioned, Maskin's monotonicity (sometimes called "strong monotonicity") is such a desirable property. It states that when a social choice function chooses a candidate based on some preference profile of all voters, this candidate should remain the selection under preference changes where the winner's position does not fall in each voter's profile. More formally,

Definition 1. (*Maskin Monotonicity*) A social choice function C is Maskin monotonic if for any $o \in O$ and any preference profile $\succ \in L^n$ with $C(\succ) = o$, then for any other preference profile $\succ'$ with the property that $\forall i \in N, \forall o' \in O, o \succ'_i o' \text{ if } o \succ_i o'$, it must be that $C(\succ') = o$.

There are two critical points to note about our use of Maskin's Monotonicity. First, as we mentioned in the introduction, all non-trivial social choice functions violate Maskin monotonicity at some profile. Although it may seem that the requirement of Maskin monotonicity is too strict, it is precisely the strictness that gives us great power of discrimination. If we instead used a weaker requirement like standard monotonicity [35], many social choice functions would not violate this property at all, giving us limited power to speak of the differences between voting rules. Secondly, Maskin's monotonicity also plays an important role in implementation theory [32] in that it serves as a necessary condition of Nash implementability.

A closely related condition is strategy-proofness.

Definition 2. A social choice function C is manipulable at profile $\succ$ by individual i via $\succ'$ if $C(\succ_{-i}, \succ'_i) \succ_i C(\succ)$, where $(\succ_{-i}, \succ'_i)$ is the profile resulting from replacing $\succ_i$ with $\succ'_i$ in $\succ$.

A social choice function is *strategy-proof* if it is not manipulable by any individual at any profile.

It is well known [36] that strategy-proofness and Maskin monotonicity are equivalent in unrestricted domains and that strategy-proofness is stronger in restricted domain. Strategy-proofness is a reasonable property on the basis of which to compare social choice functions, and in the next section we discuss the previous work on measuring its violation. However, Maskin Monotonicity is no less good of a candidate, and it is important to realize that it is a different diagnostic tool because the relationship of the number of profiles violating these different properties is unknown. Maskin monotonicity is especially crucial to study given its importance in social choice theory and implementation theory as we have mentioned.

Before we continue, we will describe the voting rules we will later analyze.

3.2.2 Voting Rules

The specific voting rules we consider in this chapter are Plurality, Borda, Copeland, Maxmin, and Plurality with Runoff.

- **Positional scoring rules:** Let s_j be the point value of an outcome ranked in position j and refer to $s = (s_1, s_2, \dots, s_m)$ as a scoring vector. Let $d(\succ_i, o) = s_{p(\succ_i, o)}$, where $p(\succ_i, o)$ is the position of an outcome o in an individual preference $\succ_i$, and for preference profile $\succ = (\succ_1, \dots, \succ_n)$, let $d(\succ, o) = \sum_{i \in N} d(\succ_i, o)$. The winner is the outcome o^* that maximizes $d(\succ, o)$.
 - **Plurality** is a positional scoring rule with scoring vector $(1, 0, \dots, 0)$.
 - **Borda** is a positional scoring rule with scoring vector $(m-1, m-2, \dots, 0)$.
- **Plurality with Runoff** is a variation of plurality where the plurality scoring vector is used to determine the two outcomes with the highest scores. Then, after all other outcomes but the top two have been removed from all the preference orderings, the plurality scoring rule is used again to determine the winner.
- **Copeland** is a pairwise comparison based voting rule where for each pair of outcomes the plurality rule is used to determine the winner of this pairwise comparison

after removing all outcomes but the two in question. Each outcome gets one point for every pairwise comparison it wins and loses one point for each comparison it loses. The winner is the outcome with highest score.

- **Maxmin** is a voting rule based on pairwise comparison. Let $N(o, o')$ be the number of voters that prefer o to o' and the score of outcome o , equals $\min_{o' \in O \setminus o} N(o, o')$. The winner is then the outcome with the highest score.

We assume alternatives can be ordered lexicographically and that ties are broken lexicographically for all voting rules.

3.3 Related Work

Social choice functions are qualitatively characterized¹ by two strong and well known results, the Gibbard-Satterthwaite [18, 50] and the Muller-Satterthwaite [36] theorems. The Gibbard-Satterthwaite theorem states that for every social choice function that is non-dictatorial and whose range has at least three alternatives, there exists at least one profile that can be manipulated by a single voter. In a similar result, the Muller-Satterthwaite theorem states that for every social choice function that is weakly Pareto efficient and non-dictatorial there is some profile where a monotonic change with respect to the winner would cause that alternative to lose.

While there has been research by Fedrizzi [14] and others on comparing voting systems by determining which of a set of desirable requirements for social choice functions are satisfied, the previous work most relevant to this thesis has focused on quantifying the proportion of profiles that are manipulable. Although we focus on monotonicity instead of manipulability, the quantitative nature of this work is related. The first conjecture that the proportion of manipulable profiles was small was made by Pattanaik in [43]. An asymptotic bound on unstable profiles under the plurality rule was proved by Peleg in [44], and this work was extended and sharpened by Fristrup and Kleiding [16] and also Slinko [53] where the proportion of manipulable profiles was shown to be $\Theta(\frac{1}{\sqrt{n}})$, with n being the number of voters and the constant factor depending on the number of candidates m . In a similar line

¹Social welfare functions were similarly characterized by Arrow's [2] seminal paper.

of work, Nurmi [42] focuses on quantifying how often voting rules select the same winners and how often the set of winners overlap.

Other research approached the problem using computer simulations, with early work by Nitzan [41] studying situations with up to 90 voters on plurality and Borda count by estimating the proportion of manipulable profiles. By looking at three metrics in addition to the proportion of manipulable profiles, Smith[54] simulated common social choice functions looking for strategy-proofness violations. A slightly different vein of work by Kelly [25, 24] examined the distribution of manipulable profile proportions for random social choice functions that were onto and non-dictatorial and, using computer simulation, found that the average proportion of manipulable profiles was almost one.

Much of the related work has assumed that each profile is equally likely (referred to as the Impartial Culture (IC) assumption) and thus interprets this proportion as a probability, as we do in later chapters. However, other work has assumed that the identities of voters are anonymous as well as impartial (IAC) giving us a different distribution over profiles. Working in this distribution, Favardin [13, 12] examines both individual and coalitional manipulation (while also allowing for counter-threats) and finds that Borda count is more vulnerable to strategic manipulation than Copeland's rule. Also considering coalitional manipulation, Lepelley [29] generalizes the special cases of IAC and IC to the space of Polya-Eggenberger probability models and runs simulations to calculate violation probabilities. Slinko [47] evaluates some voting rules including approval voting, Borda, and plurality, by looking at the asymptotic average threshold coalition size and provides analytic results.

Because the domain of all profiles violates strategy-proofness, there is a rich body of work by Kelly and others [7, 3, 39] among others, looking at various restricted domains where strategy-proofness holds. Simple majority voting has been shown by Maskin [10] to be strategy-proof in the restricted domain where the Condorcet winner exists, but the Condorcet winner does not always exist, and the probability of its existence even tends towards 0 as the number of alternatives increases, as shown in Fishburn [15]. Recent work by Bochet [6] gives the minimal domain restriction for a non-dictatorial, Pareto optimal, and strategy-proof (or Maskin monotonic) social choice function to exist. However, the social choice functions that satisfy these requirements in this minimal domain have a very

strong dictatorial flavor. A different approach is to take a known domain where strategy-proofness holds and calculate the minimal monotonic extension as in Thomson [56] and Erden [11].

We focus on completing the picture in light of previous research. While most previous work has concentrated on strategy-proofness, we focus on Maskin's monotonicity and quantifying how much it is violated. In addition to using metrics used in previous work, we also introduce the metric of domain restriction for measuring the quality of a social choice function.

3.4 Violation Graph

The violation graph is a natural tool for comparing different social choice functions.

Definition 3. (*Violation Graph*) A violation graph is a tuple $M = (C, A, G)$. C is a social choice function. A , the violation type, is a function $V \times V \times C \mapsto \{True, False\}$ that encodes our desired violation property (such as Maskin's monotonicity) and is true exactly when the two nodes and social choice function C form a violation. $G = (V, E)$ is a graph. V is the set of all possible preference profiles, and E is the set of edges² such that if $x, y \in V$, and $A(x, y, C)$ is true, then the edge $(x, y) \in E$.

We use three metrics to quantify how much a social choice function violates the property A . The first is the proportion of nodes with degree at least one. If we assume a uniform distribution over profiles, this is the probability that a random profile is involved in a violation of A . The second metric we use is the edge ratio, i.e. the ratio of the edges in the graph to the number of edges in a fully connected graph with the same number of nodes. This can be viewed as the probability that two successive elections would provide a violation of A , or in other words, would publicly demonstrate the flaw in our voting rule. The third metric we use, the minimal domain restriction required to make the violation graph disconnected, to our knowledge has not been used before. This value shows how many profiles are truly causing the violations but is a difficult metric to compute because it requires solving the

²We also note that the edges in our violation graph can be extended to hyper edges for axioms that require more than two profiles to provide an instance of a violation.

NP-complete problem of independent set [21], which has no constant-factor approximation unless $P=NP$.

3.5 Empirical Results

To compute our desired metrics (edge ratio, node ratio, and domain restriction) on the violation graph, we apply two different computational strategies: exact computation and sampling. We purposely keep our metrics and model simple because of the computational challenges of generating the violation graph. With m alternatives and n voters, the graph grows as $O(m!^{2n})$. For everything beyond small examples this is not feasible. However, to work around this problem, we sample the space and also prove asymptotic results in the next section.

We compute the edge ratio, node ratio, and domain restriction heuristics exactly when there are three candidates and seven or fewer voters. For larger numbers of voters, we generate 1,000,000 pairs of random profiles, and for each pair, we check for an edge between the two profiles. Checking for edges is a quick operation and we are only limited by the number of edges we wish to sample. Sampling nodes is more difficult because after first choosing a random node, we must check all other nodes to see if an edge exists. This is bounded by $O(m!^n)$. To make the process more efficient, we stop checking for neighbors of a node once we find the first one. As we will see in the next section, since the edges are sparse, we check most of the nodes. Because of this, the node ratio is computed with 1000 samples.

In order to compute the domain restriction, we tested four different heuristics for domain restriction (minimum degree, maximum degree, random edge, and random node) because the NP-Complete nature of this problem made it impossible to compute exactly. In our results section, we include the results from the best heuristic—the minimum degree. This heuristic iteratively removes all the neighbors of the node with the minimum degree. The results of this heuristic are guaranteed to be an upper bound for the optimal minimum domain restriction.




Figure 3.1: Violation Graph for Two Voters and Three Candidates under Plurality (and Plurality with Runoff)

To give the reader a sense of what the violation graphs look like and their wildly varying structure, we include a sample of them for three candidates and two votes in Figures 3.1, 3.2, 3.3, three voters in Figures 3.4, 3.5, 3.6, 3.7, and four voters in Figures 3.8, 3.9. Note that we only depict the nodes with degree one or greater. We include these small cases since the graphs are suggestive of emerging structures. However, our experimental results that follow refer to the coarser measures of domain restriction and edge and node ratios.

3.5.1 Edge Ratio Results

In Figure 4, we graph the comparative behavior of the voting rules we tested. There are two key take-aways from the graphs. First, the evidence suggests that as the number of voters grows, the edge ratio goes to zero for all the voting rules we tested. As we will prove for scoring rules in Section 3.6, this ratio does in fact converge to zero as the number of voters grows to infinity. The second key point is that pairwise voting rules (Copeland and Maxmin) perform better than scoring rules (Plurality, Borda, and Plurality with Runoff) for all numbers of voters. Between the pairwise rules, Maxmin performs better than Copeland by a small amount. We do not show the graphs for other numbers of candidates due to space reasons, but these relationships also hold for candidates up to at least 20.

In addition, the data also shows another relationship. When fixing the number of voters




Figure 3.2: Violation Graph for Two Voters and Three Candidates under Borda (and Copeland)




Figure 3.3: Violation Graph for Two Voters and Three Candidates under Maxmin




Figure 3.4: Violation Graph for Three Voters and Three Candidates under Plurality




Figure 3.5: Violation Graph for Three Voters and Three Candidates under Plurality with Runoff




Figure 3.6: Violation Graph for Three Voters and Three Candidates under Maxmin (and Copeland)




Figure 3.7: Violation Graph for Three Voters and Three Candidates under Borda




Figure 3.8: Violation Graph for Four Voters and Three Candidates under Plurality




Figure 3.9: Violation Graphs for Four Voters and Three Candidates under Copeland

as we increase the number of candidates, the edge ratio also goes to zero. In Figure 3.11, the edge ratio converges, but it does so at a slower rate than when we fix the number of candidates and increase the number of voters.

3.5.2 Node Ratio Results

In Figure 3.12, we graph the node ratio for Borda, Plurality, Plurality with run-off, Copeland, and Maxmin. In contrast to our edge results, as the number of voters increases, the node ratio increases. A more striking comparison is that in previous work on the node ratio for strategy-proofness it was shown empirically and theoretically that the node ratio goes to zero as the number of voters increases. From this we can see that even though the existence of at least one violation of strategy-proofness is equivalent to the existence of at least one violation of monotonicity on the unrestricted domain, the actual node ratio for the two axioms is different, and may converge to opposite values. We will see in the next section a closer look into the meaning of the node ratio by means of the domain restriction heuristic.

We also note that according to the node ratio, pairwise rules are better than scoring rules. This is not surprising since pairwise rules such as Copeland and Maxmin are well-known as extensions of *Condorcet procedures*, which choose the Condorcet winner when it exists. One can verify that among any pair of profiles where Condorcet winners exist, there are no monotonic violation.

3.5.3 Domain Restriction Results

Domain restriction is the final metric we will examine. This allows us to refine our results on the node ratio. If a voting rule has a high node ratio but only very few nodes need to be removed to disconnect the graph (restrict the domain so that no violations occur), then the high node ratio is an artifact of a few central and well-connected nodes. If we can prohibit these profiles from occurring in practice, or we can rule them out due to low probability, a voting rule with a high node ratio but a low domain restriction might be considered a good voting rule.

Looking at the data, we find that pairwise rules are better than scoring rules, as can be seen in Figure 3.13. Although the node ratio of Copeland is lower than Maxmin, the

structure of the graphs are such that Maxmin scores lower than Copeland using our domain restriction techniques. This can be seen in a small example by looking at the violation graphs of Maxmin in Figure 3.3, where by only deleting the two centered nodes, all the violation edges are eliminated.

Most importantly, even though the node ratio increases rapidly as the number of voters increases, most nodes do not need to be removed from the graph, and the proportion of nodes that needs to be removed tends toward zero for all the voting rules we tested.

3.6 Asymptotic Results

We can see from our experimental results that by fixing the number of candidates and increasing the number of voters, the ratio of violation edges decreases. One may start to wonder: does this ratio converge, and if so, what is the limit? We prove in this section that this ratio converges to 0 when $n \rightarrow \infty$.

Suppose the number of violation edges for some scoring rules is $VE(n, m)$ and the number of violation nodes is $VN(n, m)$, where the number of all possible edges and nodes is $E(n, m)$ and $N(n, m)$.

Our start point is a known result of Slinko[53] concerning the number of unstable³ profiles $U(n, m)$. Let $L(n, m) = \frac{U(n, m)}{N(n, m)}$.

Theorem 1. (from Slinko [53])

$$\frac{d_m}{\sqrt{n}} \leq L(n, m) \leq \frac{D_m}{\sqrt{n}},$$

where d_m and D_m are constants that depend only on m .

Theorem 2. When $n \rightarrow \infty$, we have

- $\lim_{n \rightarrow \infty} \frac{VE(n, m)}{E(n, m)} = 0,$
- $\lim_{n \rightarrow \infty} \frac{VN(n, m)}{N(n, m)} = c, \text{ where } 0 \leq c \leq 1.$

³A profile is unstable if a unilateral deviation from one voter can lead to the change of outcome.

Let $C(n, m)$ be the number of all profile $>$ in $\{1, \dots, n\}$ such that when adding a vote $>_{n+1}$ from an additional voter $n + 1$, the outcome will be changed. Then the following lemma holds:

Lemma 1. $\frac{C(n, m)}{N(n, m)} = L(n + 1, m) \leq \frac{D_m}{\sqrt{n+1}}.$

Proof. Suppose $>$ is a profile in $\{1, \dots, n\}$ such that when adding a vote $>_{n+1}$ (note that there are $m!$ possible $>_{n+1}$) from an additional voter $n + 1$, the outcome will be changed. It is easy to see that $(>, >_{n+1})$ is an unstable profile in $\{1, \dots, n, n + 1\}$ and all the unstable profiles in $\{1, \dots, n, n + 1\}$ can be generated from such $>$. We have

$$\begin{aligned} C(n, m) &= \frac{U(n + 1, m)}{m!} = \frac{L(n + 1, m)N(n + 1, m)}{m!} \\ &= L(n + 1, m)N(n, m). \end{aligned}$$

Thus, according to Theorem 1, we have

$$\frac{C(n, m)}{N(n, m)} = L(n + 1, m) \leq \frac{D_m}{\sqrt{n+1}}.$$

□

Thus when $n \rightarrow \infty$, the probability that “adding an additional vote to a random profile changes the outcome” is 0.

Proof. (Theorem 2)

- We now prove the first part of Theorem 2, the edge ratio. Suppose otherwise, there are two possibilities: $\lim_{n \rightarrow \infty} \frac{VE(n, m)}{E(n, m)} \neq 0$ or $\frac{VE(n, m)}{E(n, m)}$ does not converge. Either way, we have $VE(n, m) = cE(n, m)$ for $c \neq 0$ for some sufficiently large n . One can prove this implies $VN(n, m) = c'N(n, m)$ for $c' \neq 0$.

Now consider a monotonic violation for n voters where for voting rule f we have $f(>') = a, f(>'') = b \neq a$ and $>''$ is an improvement of $>'$ w.r.t a . When adding an additional voter $n + 1$, by the comments following Lemma 1, we know that $f(>', >'_{n+1}) = a$ and $f(>'', >''_{n+1}) = b$ (there are a small proportion of profiles that change the outcome, but this reduces the amount of violations). Further, $(>'', >''_{n+1})$

is still monotonic from $(>', >'_{n+1})$ only if the rank of a in $>''_{n+1}$ is at least the same as in $>'_{n+1}$. The number of such pairs $(>'_{n+1}, >''_{n+1})$ is strictly less than $(m!)^2$, the factor by which the number of edges increases from n voters to $n + 1$ voters.

The other case that can generate monotonic violation with $n + 1$ voters is when a previous pair of profiles without violation now becomes one by adding voter $n + 1$. One can count that the number of such pairs is at most $(\frac{2D_m}{\sqrt{n}} + \frac{D_m^2}{n})(1 - c)E(n, m)(m!)^2$. Clearly, its ratio to $E(n + 1, m)$ goes to 0 as $n \rightarrow \infty$.

To sum up, we have $\frac{VE(n+1, m)}{E(n+1, m)} = c'' \frac{VE(n, m)}{E(n, m)}$ for sufficiently large n and some $0 < c'' < 1$, which leads to $\lim_{n \rightarrow \infty} \frac{VE(n, m)}{E(n, m)} = 0$, a contradiction.

- Similarly, for the node ratio, consider a monotonic violation for n voters where for voting rule f we have $f(>) = a, f(>') = b \neq a$ and $>'$ is an improvement of $>$ w.r.t a . When adding an additional voter $n + 1$, we still have that $f(>', >'_{n+1}) = a$ and $f(>', >''_{n+1}) = b$ (with a small proportion of deviations which tends to 0 divided by the number of all profiles). We now prove that, for such a violation node $>$ for n voters, $(>', >'_{n+1})$ is a violation node for any $>'_{n+1}$ for $n + 1$ voters. In other words, we only need to prove that for any $>'_{n+1}$, we can find its monotonic improvement $>''_{n+1}$ w.r.t a . This can be easily achieved by fixing the ranking of a and permuting other candidates. similarly, we can prove that for such a violation node $>'$ for n voters, $(>', >''_{n+1})$ is still a violation node for any $>''_{n+1}$, since this amount to say that we can find some vote that can be improved to $>''_{n+1}$ w.r.t to a . Thus, we conclude that for any violation node for n voters, we can generate $m!$ violation nodes corresponding to it for $n + 1$ voters.

The same as the edge ratio, the other case that can generate monotonic violation with $n + 1$ voters is when a previous pair of profiles without violation now becomes one by adding voter $n + 1$. However we can similarly exclude this since this proportion tends to 0.

To sum up, $\frac{VN(n+1, m)}{N(n+1, m)} = \frac{m!VN(n, m)}{m!N(n, m)} = \frac{VN(n, m)}{N(n, m)}$ for sufficiently large n , which leads to $\lim_{n \rightarrow \infty} \frac{VN(n, m)}{N(n, m)} = c$ for some constant c .


□

3.7 Conclusion and Future Work


The contribution of this chapter is two-fold. First, we offer a new viewpoint on the comparative analysis of voting rules. Second, we provide empirical results on the application of our new approach to Maskin's monotonicity.

Our results show that pairwise voting rules, like Maxmin and Copeland, are well-behaved. On the other hand, the commonly used scoring rules of Plurality and Borda have a higher degree of monotonic violation by all the measures we tested. In addition, our results show that even though the majority of profiles are involved in a monotonic violation (the node ratio is increasing in voters), the violations can be eliminated by removing a small proportion of profiles. This domain restriction is not only very low, but it also tends to zero as the number of voters increase. Finally, our asymptotic results provide the first proof of the asymptotic behavior of the edge ratio, which goes to zero as the number of voters increases, and also show that the node ratio converges.

One future direction is to prove a tight bound on the node ratio when the number of voters tends to infinity and also prove a bound on the edge ratio as the number of candidates tends to infinity. We are also investigating how voting rules behave with respect to the violation of the conjunction (disjunction) of two or more axioms. Finally, the most promising direction is going beyond the metrics of node and edge ratios and leveraging the richer structure (seen in Figures 1, 2, and 3) of the violation graph itself.



(a) Up to 6 voters



(b) Up to 20 voters

Figure 3.10: Comparative Edge Ratio Results for Three Candidates (Two views of the same data)




Figure 3.11: Comparative Edge Ratio Results for Three Voters




Figure 3.12: Comparative Node Ratio Results for Three Candidates




Figure 3.13: Proportion of Nodes Removed with Domain Restriction for Three Candidates

Chapter 4

Dynamic Voting in Single Peaked Domains

4.1 Introduction

4.1.1 Motivation

HIGs model ongoing processes in which at different times agents can take action, and the payoff at a given moment is a function of the agent's type and all actions taken in the past (by all agents). Many scenarios fit this model. For example, consider a temperature-controlled room in which the individual occupants periodically go and adjust the thermostat; the temperature at a given moment is a function of all the changes performed in the past, and each agent has a privately desired temperature.

Our most direct motivation, however, is the behavior on World Wide Web-based opinion aggregation sites. Consider the Amazon website (www.amazon.com) in which users can at any time rate a DVD (possibly rating the same DVD many times), as well as see the aggregate rating for the DVD compiled from all ratings provided thus far. Users may desire that the aggregate rating have a certain value; they may want it to reflect their personal taste, or, for example, they may be members of the movie company that produced the movie. Furthermore, different users may have different influence over the rating, depending on how many times they rate it (this may vary depending on their available time and their

stake in the outcome).

In some cases the aggregation is explicit. For example, rating products on Amazon as in the above example, the Yelp business-rating website (www.yelp.com), the TripAdvisor travel-services-rating website (www.tripadvisor.com), and many others. In others it can be implicit. This was the case in the thermostat example above. Consider the online collaborative encyclopedia Wikipedia (www.wikipedia.com) as an example, and specifically an entry taking a stance on whether global warming is a real phenomenon. Imagine a value in the $[0,1]$ range that captures the degree to which the current content argues for the realism of global warming (0 being complete denial of it, 1 taking it as certainty, and values in between capturing more nuanced opinions). Users who are passionate about the topic care about the value over time since that reflects the cumulative impact the entry has on public opinion and policy.

4.1.2 Overview of model and results

To account for scenarios such as the ones above, we define a new class of games called *historical influence games (HIGs)*. Informally speaking, HIGs are infinite games in which agents take turns round-robin style choosing a value for a single-dimensional variable. The payoff at each stage to each agent is a monotonically decreasing function of the distance between the agent's personal ideal value, and a weighted average of past values chosen by all agents. The overall payoff to the agent is the limit average of the stage payoffs.

Note that HIGs do not fall trivially into existing classes of games. Certainly the historical nature of the payoff function and the turn-taking structure make them distinct from repeated games. They can be mapped into a very special class of stochastic games with infinite sets of stage games and deterministic transition functions, but this view is too general and does not provide any particular insight. It is much more natural to view HIGs as a class of their own.

Beside presenting the model of historical games, our main contributions are as follows:

- We define myopic strategies in HIGs, and show that they form a subgame perfect Nash equilibrium.

- We introduce certain smoothness constraints on how the impact of a given action changes over time, constraints which define the class of *valid* HIGs. We show that for valid HIGs, under myopic play the limit average value converges to what we call the *central value*, which is the median of the agents' ideal values jointly with certain societal focal points.
- As a side effect we show a polarization theorem: After a finite period, almost all agents settle on one of the extreme values.
- We show a tight connection between valid HIGs and the class of Moulin strategy-proof voting rules in single-peaked domains [34].

4.1.3 Related work

The strongest technical connection of which we are aware is in connection with social choice theory, specifically with Moulin's work on strategy-proof voting rules in single-peaked domains, where the set of strategy-proof, anonymous, and efficient voting rules is characterized exactly in [34]. As we will show in section 5, there is a tight connection between these voting rules and historical games. So much so, that our work can be seen as a constructive version of Moulin's characterization.

Following Moulin, there has been a string of work further pushing the bounds of single-peaked voting characterizations. [51] have dealt with location on a network lying in two dimensional space. Others, [40], [39] have shown how far single-peaked domains can be generalized and still have strategy-proofness hold.

Another related connection is to the Mean Voter Theorem [20], [5]. This theorem states that if voters have single peaked preferences in a single dimension, then the position of maximum utility for the median voter will beat any alternative in a pairwise election. Other work stemming from the Condorcet jury theorem has viewed voting as a process to aggregate information [38], or as a process to perform maximum likelihood estimation [9].

One other, obviously closely related work is our own previous paper [37]. The class of games analyzed there, called *joint process games*, are a special case of historical games, and turn out to correspond to one specific Moulin voting rule. We return to this in section 5.

There has been substantial empirical work on aggregate rating sites. In connection with Wikipedia, for example, researchers have estimated the expertise of Wikipedia authors [55], the percentage of the users classified as elite [26], and even preferential-attachment-like models for updates to wiki pages [59]. However, none of this work has been analytic, let alone game theoretic.

4.1.4 Structure of this chapter

The structure of the paper is as follows.

- Section 4.2: We present the formal model of HIGs and the subclass of valid HIGs.
- Section 4.3: We define myopic strategies and show that they form a subgame perfect Nash equilibrium in HIGs.
- Section 4.4: This is the heart of the paper. We define the Central Value, show that valid HIGs converge to it under myopic strategies, and prove the polarization theorem.
- Section 4.5: We show the tight connection between valid HIGs and Moulin's class of voting rules in single-peaked domains.

4.2 The Model

We start with the most general class of games, *historical games*, then specialize them to *historical influence games (HIGs)*, and then further specialize them to *valid HIGs*. In historical games agents take turns taking action, and after each action *all* agents get a payoff according to *all* actions taken to date.

Definition 4 (Historical Game). *An historical game consists of*

- *An ordered set of agents $N = 1, \dots, n$.*
- *For each agent i , a set of actions A_i . Define*

$$H = A_1 \times A_2 \times \dots \times A_n \times A_1 \times A_2 \times \dots$$
the set of infinite round-robin action histories,

and H_t the set of H 's finite prefixes of length t , for each $t \in \mathbb{N}$. By convention, if $h \in H$ is an infinite history, $h_t \in H_t$ will denote its prefix of length t .

- An outcome space X .
- An outcome function $O : H_t \mapsto X$, for all $t \in \mathbb{N}$.
- A reward function $R_i : X \mapsto \mathbb{R}$.
- An extension of R_i to the infinite history as the limit average reward: $R_i(h) = \liminf_{t \rightarrow \infty} \frac{\sum_{j=1}^t R_i(O(h_j))}{t}$.

In the remainder of the paper we restrict our attention to a special class of historical games called historical influence games (HIGs).

Definition 5. A historical influence game (HIG) is an historical game as above, in which

- $A_i = X = [0, 1]$. (Note: Later on we will also consider the case in which $A_i = \{0, 1\}$.)
- O is a weighted average of past actions: $O(h_t) = \sum_{t'=1}^t a_{t'} \cdot I(t, t')$, where $I(t, t') : t' \leq t \in \mathbb{N} \mapsto [0, 1]$ is a time-specific weight function such that $I(t, t') \geq 0$, $\sum_{t'=1}^t I(t, t') = 1$.
- Each R_i is single peaked: There exists $g_i \in X$ (the “peak”) such that for all $x, y \neq g_i \in X$ it is the case that if either $x < y < g_i$ or $x > y > g_i$ then $R(x) < R(y) < R(g_i)$.

We deliberately presented the general class of historical games first to make clear the restrictions embodied in HIGs. Certainly one can contemplate outcomes that are not single dimensional, or reward functions that are not single peaked. We however focus on HIGs, since they already give rise to a rich theory, and match well some of the motivating applications discussed at the introduction. In fact, for much of the paper we consider even the more specific class of *valid* HIGs. The definition above assumes no relation between the influence of an action taken at a given time on outcomes at different times in the future, nor on how this influence changes in time. We next consider four restrictions on the influence functions. First is the requirement the recent past be more influential than the distant past:

Definition 6 (Diminishing Past). *An influence function I is diminishing past if for all $i < t$ it is the case that*

$$I(t, i) \leq I(t - 1, i).$$

The second requirement is that, for a sufficiently long history, any given action – even the most recent one – will have negligible influence:

Definition 7 (Diminishing Present). *An influence function I is diminishing present if,*

$$\lim_{t \rightarrow \infty} I(t, t) = 0.$$

Next we require that for sufficiently long histories, any finite initial action sequence have a negligible influence:

Definition 8 (Future Weighted). *An influence function I is future weighted if $\forall k \in \mathbb{N}$*

$$\lim_{t \rightarrow \infty} \sum_{t'=k}^t I(t, t') = 1.$$

Finally, we require that the limit of the influence assigned to each agent converge. In the following definition, $1_i(\cdot)$ is the indicator function that takes the value 1 when agent i is active and 0 otherwise:

Definition 9 (Agent Influence). *Given an influence function I , the limit influence of agent $i \in N$, denoted f_i , if it exists, is defined by*

$$f_i = \lim_{t \rightarrow \infty} \sum_{t'=1}^t 1_i(t') I(t, t')$$

If the limit exists for all agents $i \in N$ then I is said to satisfy agent limit influence.

It's not difficult to see that no one condition is implied by the others.

Proposition 3. *In any convex influence function, no three conditions among diminishing past, diminishing present, future weighted, and limit agent influence imply the fourth.*

Proof. Proofs each each of the four implications:

- Define the influence function I such that $\forall t, I(t, 1) = 1$ and for all $t' \neq 1, I(t, t') = 0$. This function satisfies diminishing past, diminishing present, and agent influence, but violates future weighted.
- Define the influence function I such that $\forall t, t' < t, I(t, t') = \frac{1}{2^{t-2}}, I(t, t) = .5$. This function satisfies diminishing past, future weighted, and agent influence, but violates diminishing present.
- Define the influence function I such that $\forall t, t' < t - 1, I(t, t') = \frac{1}{t+1}, I(t, t) = \frac{1}{t+1}, I(t, t-1) = \frac{2}{t+1}$. This function satisfies diminishing present, future weighted, and agent influence, but violates diminishing past.
- Define the influence function I as follows: Let the sequence a be defined by $a_i = 2^{j+1} - j - 2$ for $i \in \mathbb{N}, a_0 = 0$. For $a_i < t \leq a_{i+1}, I(t, t') = \frac{1}{\lfloor (t-a_i+1)/n \rfloor}$ for $t' = a_i, a_i + n, a_i + 2n, \dots$ and $I(t, t') = 0$ otherwise. This function satisfies diminishing present, diminishing past, and future weighted, but violates agent influence.

□

Definition 10 (Validity). *An outcome function is valid iff it is future weighted, diminishing present, diminishing past, and the influence of each agent exists. A HIG is valid iff its outcome function is.*

In Section 4 we present several common and natural examples of valid outcome functions.

4.3 Myopic Strategies in Historical Influence Games

Myopic strategies have been widely studied theoretically and experimentally in many settings. We study a natural definition of a myopic action selection rule in the context of historical influence games. We show that the action profiles generated by myopic play form a subgame perfect Nash equilibrium and we discuss this equilibrium behavior further in the next section.

Under myopic play agent i chooses the action that maximizes R_i given the history up to this time. We define this action selection rule as:

Definition 11 (Myopic Play). *At time t , under myopic play, the active agent i chooses the action $a = \operatorname{argmax}_{a \in A_i} R_i(h_{t-1}, a)$.*

As a check on how reasonable myopic play is we first prove that the strategy profile generated by myopic play forms a subgame perfect Nash equilibrium

Theorem 4. *The set of strategies defined by the myopic play rule form a subgame perfect Nash equilibrium.*

The proof is provided in the appendix.

4.4 Central Value

From the last section, we know that the set of strategies defined by myopic play form a subgame perfect Nash equilibrium, but what does the equilibrium and actual play look like? Somewhat unusually for a paper with a theoretical focus, we start by presenting some experiments, and then follow by showing that historical influence games always converge and characterize what they converge to.

4.4.1 Experimental Results

The *in silico* experiments were run as follows. Ten agents were generated for each experiment. The g_i (sometimes we refer to g_i as the goal) for each agent was selected uniformly at random from $[0, 1]$ and the reward function was defined as $R_i(o_t) = 1 - |g_i - O(h_t)|$. At each time step, the active agent uses myopic play to select his action. For all the following graphs, we plot the outcome over time.

First we examine a simple case; let the present outcome be the average of past actions, and thus all time steps have exactly the same influence. This corresponds to the $I(t, t') = \frac{1}{t}$. We call this influence function “simple averaging influence”. This is exactly the model used in [37] and an example of a valid influence function. The behavior is seen in Figure 4.1.

To model situations where some agents have more influence than others we can use the influence function, $I(t, t') = \frac{\beta_{(t' \bmod n)+1}}{\sum_{i=1}^t \beta_{(i \bmod n)+1}}$ where $\sum \beta_i = 1$. Using this model, we




Figure 4.1: Convergence under $I(t, t') = \frac{1}{t'}$ with myopic play




Figure 4.2: Convergence under $I(t, t') = \frac{n^{\beta_{(t \bmod n)+1}}}{t'}$ with myopic play




Figure 4.3: Convergence under $I(t, t') = (1 - \alpha_{t-1}) \prod_{j=t}^{t'-1} \alpha_j$ with myopic play

get behavior as in Figure 4.2 and call this class of influence functions “weighted averaging influence”.

To capture that in many situations current actions have more influence than past actions, we pose the following model. Let $O(h_t) = \alpha_t * O(h_{t-1}) + (1 - \alpha_t)a_t$, where $\alpha_t \rightarrow 1$, $\alpha_1 = 0$, and i is the active agent at time t . This can be captured with the following “discounted past influence” function, $I(t, t') = (1 - \alpha_{t'}) \prod_{j=t'+1}^t \alpha_j$. Different experiments using this influence function are shown in Figures 4.3, 4.4, and 4.5.

Somewhat surprisingly, the outcome appears to converge to a fixed point in all of these experiments. One can quickly verify that the simple averaging, weighted averaging, and discounted past influence functions meet the four conditions of future weighted, diminishing present, diminishing past, and existence of agent influence, and thus are valid influence functions.

We ask ourselves: what happens when a non-valid influence function is used? Testing one non-valid influence function defined by the outcome recurrence $O(h_t) = \alpha * O(h_{t-1}) +$




Figure 4.4: Convergence under $I(t, t') = (1 - \alpha_{t-1}) \prod_{j=t}^{t'-1} \alpha_j$ with myopic play




Figure 4.5: Convergence under $I(t, t') = (1 - \alpha_{t-1}) \prod_{j=t}^{t'-1} \alpha_j$ with myopic play




Figure 4.6: Non-convergence example

$(1 - \alpha)a_t$ which fails the diminishing present condition shows non-convergent behavior in Figure 4.6.

From our experiments it seems that outcome in valid historical influence games converges to a fixed point. In the rest of this section we will back up this claim theoretically, characterize the value outcome converges to, and prove a theorem on how the actions of agents polarize over time.

4.4.2 The Central Value

Theorem 5. *In a valid historical influence game the outcome converges under myopic play.*

While we defer giving the proof until later, this theorem agrees with the intuition from our *in silico* experiments. In this section we start by discussing a characterization for the

convergence point for the outcome in valid historical influence games using a simple averaging influence function. We later extend this characterization to apply in historical influence games with any valid influence function. Finally we prove a connection to a traditional aggregation method – the mean of agent goals.

We call the point the outcome converges to the *central value*. We introduce the central value by using the influence function $I(t, t') = \frac{1}{t}$.

Definition 12 (Central Value, special case). *In a historical influence game using a simple averaging influence function, let N be the set of agents and G be the set of agent goals. Let P be the set of partition points defined as $P = \{\frac{1}{n}, \frac{2}{n}, \dots, \frac{n-1}{n}\}$. The central value c is the median of the multiset union of $G \cup P$.*

For example, if we have three agents with goals $G = \{.2, .3, .4\}$ then we have the set of partition points $P = \{1/3, 2/3\}$ and the central value is $1/3$. We later show that for any fixed ordering of these agents using myopic play, the outcome will converge to $1/3$.

To define the general form of the central value we must use the influence of particular agents as defined earlier.

Definition 13 (Central Value, general case). *In a valid historical influence game let N be the set of agents, G be the set of agent goals, and W be the set of agent influences. Without loss of generality assume that the agents are ordered by increasing goal values, i.e. $g_1 < g_2 < \dots < g_n$. Let P be the set of partition points such that $P = \{\sum_{i=k}^n w_i | k \in \{2, 3, \dots, n\}\}$. The central value, c , is the median of $G \cup P$.*

In words, the central value is the median of the set of agent goals and the set of postfix sums of agent influences when agents are ordered by increasing goal values. For example, let $N = 1, 2, 3$, $G = \{.2, .5, .6\}$, and $W = \{.3, .4, .3\}$. Then the resulting set of partition points is $P = \{.3, .7\}$, $G \cup P = .2, .3, .5, .6, .7$, and $C(N) = .5$. Thus, as we later prove, the outcome in this historical game will converge to $.5$ under any valid influence function.

We are able to provide a further characterization of the central value – we are able to link the mean of the goals of the agents (private information) with the convergence point, the central value.

To begin, we link the arrangement of goals and partition points. We show that the number of goals $g_i > c$ is equal to the number of partition points $p_i \leq c$.




Figure 4.7: Possible mean values of the set of agent goals given $C(N)$, with $|N| = 30$. The upper bound is in green, and the lower bound is in blue.

Lemma 2.

$$|\{g_i > c : i \in N\}| = |\{p \leq c : p \in P\}|$$

The proof is provided in the appendix.

In the following lemma we clarify the connection of the central value with the mean of agent goals. The bounds between the central value and the mean of agent goals are smallest when the central value is close to 0 or 1.

Lemma 3. *Let the central value for a set of n agents be c . Let l be the greatest member of $P \cup \{0, 1\}$ such that $l \leq c$. Let r be the least member of $P \cup \{0, 1\}$ such that $r \geq c$. Then the mean of agent goals lies in $[r * c, l + (1 - l) * c]$.*

The proof is provided in the appendix.

We graph an example of the bounds for thirty agents using simple averaging influence in Figure 4.7.

4.4.3 Convergence to the Central Value under Myopic Strategies

We prove convergence of outcome to the central value through a series of lemmas employed in the final proof. The process has additional properties beside this convergence—in particular, the convergence proofs shows that as a side effect the agents' actions become polarized in time—and we return to discuss this in more detail in the next section.

The notion of the 'active window' is used in many of the following proofs and provides insight into the myopic process.

Definition 14 (Active Window). *The active window is the interval between the two values the current agent is choosing among. The active window at time t is written as,*

$$\left[\sum_{t'=1}^{t-1} I(t, t')a_{t'}, I(t, t) + \sum_{t'=1}^{t-1} I(t, t')a_{t'} \right] \quad (4.1)$$

The active window shrinks over time, and we formalize this here.

Lemma 4. *There exists a T_0 such that for all $t > T_0$, the active window contains at most one g_i .*

Proof. First, we note that the length of the active window equals $I(t, t)$ and $I(t, t)$ converges to zero by the diminishing present condition. Given the set of agent goals G , where d denotes the Euclidean distance metric, define the distance between the closest pair of goals as,

$$\rho = \min_{i,j | i,j \in G, i \neq j} d(i, j) > 0. \quad (4.2)$$

Let $\epsilon = \rho$, and by diminishing present, $\exists T$, such that $\forall t > T, I(t, t) < \epsilon$ so there can be at most one value in the active window because the length of the interval is small enough to prevent any other points from simultaneously being included. $\square$

Not only does the active window shrink over time, but the relative position of the active window and the agent goals determine the actions of the agents under myopic play.

Lemma 5. *For all times t and all agents i such that $g_i \leq \sum_{t'=1}^{t-1} I(t, t')a_{t'}$, $a_t = 0$. For all times t and all agents i such that $g_i \geq I(t, t) + \sum_{t'=1}^{t-1} I(t, t')a_{t'}$, $a_t = 1$.*

Proof. Given that R_i is a single peaked reward function, if g_i is less than the lower bound of the active window, the myopic reward will be maximized by agent i setting $a_t = 0$. Likewise, if agent i has g_i greater than the upper bound of the active window, the myopic reward will be maximized by setting $a_t = 1$. $\square$

The next lemma shows that the active window will not jump over intervals.

Lemma 6. *The active window at time $t - 1$ overlaps the active window at time t .*

Proof. The outcome after time step $t - 1$, $O(h_{t-1})$, was one of the endpoints of the previous active window by definition. If the active window at time t contains $O(h_{t-1})$ then the windows overlap. Let $\epsilon = I(t, t)$. Because $\sum_{t'=1}^t I(t, t') = 1$, then $\sum_{t'=1}^{t-1} I(t-1, t') - \sum_{t'=1}^{t-1} I(t, t') = \epsilon$. In other words, the total influence for actions 1 to $t' - 1$ has decreased exactly by ϵ . If $I(t-1, t') - I(t, t') > 0$ exactly for t' such that $a_{t'} = 1$, then the new active window is $[O(h_{t-1}) - \epsilon, O(h_{t-1})]$. Likewise, if the decrease of influence was only for actions that were 0, then the new active window is $[O(h_{t-1}), O(h_{t-1}) + \epsilon]$. If the decrease was mixed among all actions, the endpoints of the active window will be between these extremes, but the distance between endpoints will still be ϵ . Since $O(h_{t-1})$ is a member of all such intervals, the active window at $t - 1$ overlaps with the active window at t . $\square$

Lemma 7. *For all ϵ and all times t , there exists a time $t_0 \geq t$ where $|c - O(h_{t_0})| < \epsilon$.*

The proof is provided in the appendix. Now we have all the necessary lemmas to prove our main result:

Theorem 6 (Convergence). *If the central value of a historical game is c , then for all $\epsilon > 0$ there exists a time T_1 such that for all $t > T_1$, $|c - O(h_t)| < \epsilon$*

The proof is rather involved, but the basic flow of argument can be laid out simply. First, after any point in time, there exists a later point in time where the active window will be within ϵ of c . After this time, the influence of the agents to the left and right of c will not change by more than ϵ . Because the active window is within ϵ of c and thus the actions of the agents are determined by myopic play and their relative position to c , the definition of a valid influence function requires that the active window be contained within ϵ of c from this time onwards. See the appendix for the full proof.

We also note that all the previous lemmas up to this point directly apply to the case in which we have an action space of $\{0, 1\}$ rather than $[0, 1]$. Theorem 6 can also be extended to cover the discrete action space though a slight modification of the proof.

4.4.4 Polarization

When $A_i = [0, 1]$ all agents but at most one will take only the extreme actions 0 and 1 after some point in time. This provides insight on the polarized behavior which has been observed to occur on many Internet sites. For example [22] show that the reviews on online product sites frequently follow bimodal distributions.

Theorem 7 (Polarization). *In any valid historical influence game where agents have the action space $[0, 1]$ there exists a time T such that for all $t > T$, all agents but at most one will each either repeatedly play $a = 0$ or repeatedly play $a = 1$.*

The proof is given in the appendix.

4.5 Correspondence with Voting

The central value provides a surprising link between historical influence games and voting theory. We prove a natural mapping from historical influence games to single-shot voting rules and vice-versa.

4.5.1 Moulin Single-Peaked Voting Rules

[34] proved that the set of strategy-proof, anonymous, and efficient voting rules in single-peaked preference domains is exactly equal to the set of voting rules that can be expressed as a function that chooses the median of voters' preferences and a set of societal focal points α_i .

The theorem is:

Theorem 8 (Moulin, 1980). *When voters have single peaked preferences the following two statements are equivalent*

- voting rule π from $\mathbb{R}^n$ into $\mathbb{R}$ is strategy-proof, anonymous, and efficient
- there exists $(n-1)$ values $\alpha_i \in \mathbb{R} \cup \{\infty, -\infty\}$ such that $\pi(x_1, \dots, x_n) = \text{median}(x_1, \dots, x_n, \alpha_1, \dots, \alpha_{n-1})$.

The following two theorems show that for every set of weights for an influence function I , exactly one Moulin voting rule always has the same result as the historical game's convergence point. Also, for every Moulin voting rule there is exactly one set of weights such that the historical game always has the same result as the Moulin voting rule.

Theorem 9. *Given a valid historical influence game with n agents, goals $g_1 < g_2 < \dots < g_n$ and agent influences $F = \{f_1, f_2, \dots, f_n\}$, then there is exactly one Moulin voting rule defined by $(\alpha_1, \alpha_2, \dots, \alpha_{n-1}) = (f_n, f_n + f_{n-1}, \dots, \sum_{i=2}^n f_i)$ where the outcome selected by Moulin's voting rule is the same as the central value c .*

For the next theorem we define the function,

$$w(x) = \begin{cases} 1 & \text{if } x \geq 1 \\ 0 & \text{if } x \leq 0 \\ x & \text{otherwise} \end{cases}$$

Theorem 10. *Given a Moulin voting rule defined by $(\alpha_1, \alpha_2, \dots, \alpha_{n-1})$, there is exactly one set of influences $f_n = w(a_1)$, $f_i = w(a_{n-i+1}) - w(a_{n-i})$, and $f_1 = 1 - w(a_1)$, such that all historical influence games with these influences converge to the outcome selected by Moulin's voting rule.*

See the appendix for the proofs.

This means that not only will historical influence games converge, but also their convergence point is the only strategy-proof, anonymous, and efficient result when the weights of the players are exogenously given. This provides strong indication that historical influence games are well justified as models for joint group decision making. Furthermore, the information revelation requirements for Moulin's voting rules are stronger than for historical influence games. Even when it is common knowledge that all players are using myopic play, an outside observer would not always be able to infer the exact goal of agent i – some agent's action may be 1 for the entire process even if $g_i \ll 1$.

An interesting side-effect of these theorems is that there are parameterized families of influence functions that have a one to one correspondence with Moulin's voting rules. The weighted averaging influence function $I(t, t') = \frac{\beta_{(t' \bmod n)+1}}{\sum_{i=1}^t \beta_{(i \bmod n)+1}}$ is one example. Every set of β_i corresponds to a unique Moulin voting rule and vice versa. However, not all families of influence sets are parameterized in this fashion. The simple averaging influence function $I^{(t)}(t') = \frac{1}{t'}$ corresponds to a single Moulin voting rule where $a_i = \frac{i}{n}$, and the discounted past influence function has a one to many relationship with Moulin rules.

4.6 Conclusion

We have defined a new class of games, Historical Games that model infinite games where players take turns and actions have residual effects over time. We specialized this class to Historical Influence Games and even further to valid HIGs. To do this we required the influence over time to satisfy smoothness properties.

We showed that under myopic play the limit outcome converges to the central value, a mean of agent's peak preferences and societal focal points. As a side effect we showed that agent actions will polarize in the limit, which matches with observational data on online rating sites. Finally, we show a tight connection between valid HIGs and the class of Moulin strategy-proof voting rules in single-peaked domains [34].

A number of interesting open problems remain here, of which we highlight two. First, it is still unknown whether the set of valid influence functions is equivalent to the set of converging HIGs. Second, *in silico* experiments suggest convergence also holds in certain higher dimension preference spaces even though no strategy-proof voting rules exist in these domains. It would be interesting to understand the conditions for convergence in these higher dimension domains.

4.7 Appendix

Theorem 4. *The set of strategies defined by the myopic play rule form a subgame perfect Nash equilibrium.*

Proof. Let $A = (a_1, a_2, \dots, a_n)$ be the set of strategies defined by the myopic play rule that map from all possible action histories to actions. Consider any a'_i that agent i is considering. If agent i has $g_i = C(N)$ then it follows that under myopic play this agent receives the maximum possible average reward so there is no profitable deviation.

By diminishing past the influence of any finite history of actions will go to zero. Without loss of generality, assume that agent i is such that $g_i < C(N)$. Thus, by Lemma 5, and Theorem 6 after any finite prefix of actions, there exists a T_0 such that for all $t > T_0$, agent i 's action under the myopic play rule is 0. Let h be the infinite history generated by strategy profile A and h' be the infinite history generated by strategy profile A' . Thus the average reward for player i when the strategy profile is A is $R_i(h)$. If for all times t except when agent i is acting, strategy profile $A' = (a_1, a_2, \dots, a'_i, \dots, a_n)$ causes players to take the same actions as under A , it then follows that,

$$R_i(A) = \lim_{k \rightarrow \infty} \frac{\sum_{j=1}^k R_i(h'_j)}{k}, \quad (4.3)$$

$$= \lim_{k \rightarrow \infty} \frac{\sum_{j=T_0}^k R_i(h'_j)}{k}, \quad (4.4)$$

$$R_i(A) - R_i(A') = \lim_{k \rightarrow \infty} \frac{\sum_{j=T_0}^k R_i(O(h'_i)) - R_i(O(h_i))}{k}, \quad (4.5)$$

and finally, because after T_0 agent i 's action is always 0 under strategy a_i , then any under $a'_i \neq a_i$ will result in a history of actions that can have zero or more ones. Thus, because the influence function is the same, $O(h_i) < O(h'_i)$, giving us that $R_i(A) \geq R_i(A')$ because of single-peakedness.

Now suppose that a' causes an agent other than agent i to play a different action in some round. Given that the length of the active window approaches zero and that the agents use the myopic play rule, all agents with g_i less than the active window will play zero and all agents with g_i above the active window will play one. For agent i to get a higher reward under action a' the active window would have to be less than c . All agents currently less than c can not cause the active window to move closer to zero since they all have converged to playing all zeros. No agent $\geq c$ will ever cause the active window to decrease below

c , as it would lower their reward in the round they choose zero, violating myopic play. Thus, agent i can never cause the outcome to decrease and can never increase his utility by deviating. Because there is no profitable deviation from all possible histories, A forms a Subgame Perfect Nash equilibrium. $\square$

Lemma 2.

$$|\{g_i > c : i \in N\}| = |\{p \leq c : p \in P\}|$$

Proof. If the central value is a goal value, let y be the number of goals greater than c and let z be the number of partition points less than or equal to c . By the definition of the median and definition 13 we have,

$$\begin{aligned} z + (n - 1 - y) &= y + (n - 1 - z) \\ z &= y, \end{aligned}$$

If c is a partition point we follow similar reasoning and the conclusion still holds. $\square$

Lemma 3. *Let the central value for a set of n agents be c . Let l be the greatest member of $P \cup \{0, 1\}$ such that $l \leq c$. Let r be the least member of $P \cup \{0, 1\}$ such that $r \geq c$. Then the mean of agent goals lies in $[r * c, l + (1 - l) * c]$.*

Proof. Let y be the number of goals greater than c and let z be the number of partition points less than or equal to c . By lemma 2 we know $y = z$. Thus by the definition of the central value, l is the y th partition point and equal to the sum of influence for agents with goals greater than c . Likewise, r is the sum of influence for agents with goals greater than or equal to c .

To prove the right bound, consider the case where all agents with $g_i > c$ have $g_i = 1$ and all other agents with $g_i \leq c$ have $g_i = c$, this gives us a mean of agents goals equal to $l \times 1 + (1 - l) \times c$.

To prove the left bound, consider the case where all agents with $g_i \geq c$ have $g_i = c$ and all other agents have $g_i = 0$. This gives us a mean of agents goals equal to $r \times c$. $\square$

Lemma 7. *For all ϵ and all times t , there exists a time $t_0 \geq t$ where $|c - O(h_{t_0})| < \epsilon$.*

Proof. Let t_1 be the greatest of t and the time from lemma 4 such that the active window only contains one g_i . The outcome at time $t_0 > t_1$, can be written $O(h_{t_0}) = A + B$, where

$$A = \sum_{t'=1}^{t_1} a_{t'} \times I(t_0, t')$$

$$B = \sum_{t'=t_1+1}^{t_0} a_{t'} \times I(t_0, t').$$

By the future weighted property, we know that $\sum_{t'=t+1}^{t_1} I(t_1, t') \rightarrow 1$. By the diminishing present condition we know that $I(t, t) \rightarrow 0$. Thus for $\epsilon' = \epsilon/2$, there exists a time t_0 such that $A \leq \epsilon'$ and the size of the active window is less than ϵ' .

By lemma 5, for all $t > t_1$, the actions of agents outside the active window will be 0 or 1. Without loss of generality, assume the active window is less than c . Thus B is within ϵ' of the sum of influence for those agents with goals greater than or equal to c .

Let $X = \min\{p \in P \cup \{0, 1\} | p \geq c\}$. By lemma 2 and the definition of the central value, we know that X is the sum of influence of agents with $g_i \geq c$ who will all choose action 1 after time t_1 .

If $X = c$, then at time t_0 , $|c - O(h_{t_0})| < \epsilon$.

If $X > c$, then at some point in time between t_1 and t_0 the active window will overlap c . The outcome at that time will one of the endpoints of the active window and because the length of the active window is less than ϵ' , we have $|c - O(h_{t_0})| < \epsilon$. Because of lemma 6 the active window can not jump over c . $\square$

Theorem 6 (Convergence). *If the central value of a historical game is c , then for all $\epsilon > 0$ there exists a time T_1 such that for all $t > T_1$, $|c - O(h_t)| < \epsilon$*

Proof. Let $\epsilon' = \min(\frac{\rho}{2n}, \frac{\epsilon}{2})$, where ρ is as defined in 4. Define $F_{right} = \sum_{i|g_i > c} f_i$. By the convergence of each agent's influence to f_i , we know for ϵ' there exists a T_1 such that for all $t > T_1$, $|\sum_{i|g_i > c} \sum_{t'=1}^t 1_i(t') I(t, t') - F_{right}| < \epsilon'$. Define F_{left} in a similar fashion for those agents to the left of central value. Let $t_0 > T_1$ be the first time after T_1 where the active window is within ϵ of c as from lemma 7.

By lemma 5, for every time between t_0 and $t_0 + n$, the agents to the left of c will have $a = 0$, and those to the right will have $a = 1$. We want to show that the change in the influence on the right and left is bounded and the distance between the closest pair of point is $\rho > n\epsilon'$, the active window will not cross any other points between t_0 and $t_0 + n$.

For contradiction, assume that compared to this point in time, $O(h_{t_0+n})$ decreased by more than ϵ' . By myopic play, all players to the right of $O(h_{t_0})$ would still always play $a = 1$, and all the players to the left would play 0 or 1 depending on how far $O(h_{t_0+n})$ decreased.

By the diminishing past condition, the influence associated with any action can never increase over time, only decrease. Thus, because the actions under myopic play rule out a decrease due to actions after time t_0 , a decrease in $O(h_{t_0+n})$ can only be caused by a reduction of the influence of $a = 1$ actions that occurred before time t_0 .

We have already shown that F_{right} , the influence of the players to the right, can not deviate by more than ϵ' . If the total influence of the right players before t_0 decreases by ϵ' , then since influence must sum to 1, the influence of all the actions for the left players must be ϵ' . Under myopic play, the actions after t_0 must all be $a \geq 0$, and the total influence for actions before time t_0 can not increase or decrease. Note under myopic play, the agent with at c , if one exists, will have actions balancing to c . Furthermore, because of the diminishing past condition, the influence for any of these individual time steps will not decrease, thus, since the sum is constant, it is also the case that none will increase. Thus, it is not possible for $O(h_{t_0+n}) - O(h_{t_0}) > \epsilon'$, and an analogous argument holds for proving that $O(h_{t_0}) - O(h_{t_0+n}) < \epsilon'$.

To summarize, we've shown that there exists a time such that once the active window is within ϵ of c , it will never change by more than ϵ . $\square$

Theorem 7 (Polarization). *In any valid historical influence game where agents have the action space $[0, 1]$ there exists a time T such that for all $t > T$, all agents but at most one will each either repeatedly play $a = 0$ or repeatedly play $a = 1$.*

Proof. The proof of this result is a simple application of results previously proved. By theorem 6 there exists a time such that for all later times, the outcome will be within ϵ of the central value. This result combined with lemmas 4 and 5 imply that all agents (except

the agent, if there is one, with $g_i = c$) will take the same actions from this time forward. $\square$

Theorem 9. *Given a valid historical influence game with n agents, goals $g_1 < g_2 < \dots < g_n$ and agent influences $F = \{f_1, f_2, \dots, f_n\}$, then there is exactly one Moulin voting rule defined by $(\alpha_1, \alpha_2, \dots, \alpha_{n-1}) = (f_n, f_n + f_{n-1}, \dots, \sum_{i=2}^n f_i)$ where the outcome selected by Moulin's voting rule is the same as the central value c .*

Proof. First, the Moulin voting rule defined selects the median of the α values and the peak rewards of the agents which is exactly the central value.

For the other direction, assume without loss of generality that there is a second Moulin voting rule defined by α' that differs only in that $\alpha_i < \alpha'_i < \alpha_{i+1}$ and selects the same outcome for all inputs. Choosing $g_1, g_2, \dots, g_{n-i} < \alpha_i$, and $\alpha_i < g_{n-i+1}, \dots, g_n < \alpha'_i$ we see that the Moulin voting rule defined under α chooses α_i while the voting rule defined under α' chooses α_n . Thus there is only one such voting rule $\square$

Theorem 10. *Given a Moulin voting rule defined by $(\alpha_1, \alpha_2, \dots, \alpha_{n-1})$, there is exactly one set of influences $f_n = w(a_1)$, $f_i = w(a_{n-i+1}) - w(a_{n-i})$, and $f_1 = 1 - w(a_1)$, such that all historical influence games with these influences converge to the outcome selected by Moulin's voting rule.*

Proof. First, the historical influence game defined selects the median of the α values (projection into the space $[0, 1]$) and the peak rewards of the agents which is exactly the central value.

For the other direction, assume without loss of generality that there is a second set of influences defined by f' that differs only in that $f_i < f'_i < f_{i+1}$ and selects the same outcome for all inputs. Choosing $g_1, g_2, \dots, g_{n-i} < f_i$, and $f_i < g_{n-i+1}, \dots, g_n < f'_i$ we see that the historical influence game defined under f chooses f_i while the historical influence game defined under f' chooses f_n . Thus there is only one such set of influences. $\square$

Chapter 5

Conclusion and Future Work

In this thesis we examined how societies should make group decisions given that ‘fair’ voting is impossible. We took two different approaches. The first was to take traditional axioms of fairness for voting and quantify how frequently they were violated under a few different metrics. The second approach was to study voting-like processes that are commonly found in real life and analysis their properties in a game theoretic style.

Our main tool in quantifying the frequency of violations was the object we introduced, the violation graph. By viewing the space of all preference profiles as a set of points, we can apply any relevant property (like Maskin’s monotonicity) and voting rule to this set of points to generate a violation graph. This graph has edges exactly between the pairs of preference profiles that provide a violation instance. By transforming different voting rules into violation graphs, we can compare how much they violate our chosen property. Looking at the graph features gives us a means of comparing different voting rules.

By using the violation graph we were able to provide empirical results on Maskin’s monotonicity. Our results show that pairwise voting rules, like Maxmin and Copeland, are well-behaved. On the other hand, the commonly used scoring rules of Plurality and Borda have a higher degree of monotonic violation by all the measures we tested. In addition, our results show that even though the majority of profiles are involved in a monotonic violation (the node ratio is increasing in voters), the violations can be eliminated by removing a small proportion of profiles. This domain restriction is not only very low, but it also tends to zero as the number of voters increase. Finally, our asymptotic results provide the first proof

of the asymptotic behavior of the edge ratio, which goes to zero as the number of voters increases, and also show that the node ratio converges.

It still remains to be seen how to best leverage the structure present in the violation graph. In this thesis we concentrated on node and edge ratios but by leveraging the richer structure of the violation graph itself, we hypothesize that more granular metrics can be constructed.

In the second half of this thesis we defined Historical Games which model situations where players repeatedly take actions that have residual effects on an outcome shared by all players. However, the reward of each player is determined by an individual single-peaked reward function. We specialized this class to Historical Influence Games and even further to valid HIGs by requiring the influence to satisfy certain smoothness properties.

We showed that under myopic play the limit outcome converges to the central value, a mean of agent's peak preferences and societal focal points. As a side effect we showed that agent actions will polarize in the limit, which matches with observational data on online rating sites. Finally, we show a tight connection between valid HIGs and the class of Moulin strategy-proof voting rules in single-peaked domains [34].

A number of interesting open problems remain here, of which we highlight two. First, it is still unknown whether the set of valid influence functions is equivalent to the set of converging HIGs. Second, *in silico* experiments suggest convergence also holds in certain higher dimension preference spaces even though no strategy-proof voting rules exist in these domains. In figures 5.1, 5.2, and 5.3 we show example behavior for small numbers of agents. In fact, a key to this problem is the construction of single peaked hierarchy of reward functions. In situations where the rewards of agents are separable across dimensions, the outcome trivially converges to the central value in each dimension. When the reward function is symmetric in the sense that lines of indifference follow ellipses, the outcome appears to converge, but to what point is still unknown. We can further expand this single peaked hierarchy to convex and then non-convex reward functions, and the behavior there is also unknown.

Given that social choice functions in two dimensional spaces with symmetric single-peaked reward functions have no strategy proof solution, it is fascinating that the analog historical games appears to converge to a stable point.




Figure 5.1: 2-D historical influence game *in silico* experiment with two random agents




Figure 5.2: 2-D historical influence game *in silico* run with three random agents




Figure 5.3: 2-D historical influence game *in silico* run with four random agents

Bibliography

- [1] Ofer Arazy, Wayne Morgan, and Raymond Patterson. Wisdom of the Crowds: Decentralized Knowledge Construction in Wikipedia. *SSRN eLibrary*, 2006.
- [2] K.J. Arrow. A difficulty in the concept of social welfare. *Journal of Political Economy*, 58(4):328–346, 1950.
- [3] Martin Barbie, Clemens Puppe, and Atilla Tasnadi. Non-manipulable domains for the Borda count. *Economic Theory*, 27:411–430, 2006.
- [4] Ulrich Berger. Fictitious play in $2 \times n$ games. *Journal of Economic Theory*, 120(2):139 – 154, 2005.
- [5] Duncan Black. On the rationale of group decision-making. *The Journal of Political Economy*, 56(1):23–34, 1948.
- [6] Olivier Bochet and Ton Storcken. Maximal domains for strategy-proof or Maskin monotonic choice rules. Research Memoranda 002, Maastricht, 2008.
- [7] Donald E. Campbell and Jerry S. Kelly. Information and preference aggregation. *Social Choice and Welfare*, 17:3–24, 2000.
- [8] Micael Castanheira. Why vote for losers? *Journal of the European Economic Association*, 1(5):1207–1238, 09 2003.
- [9] Vincent Conitzer, Matthew Rognlie, and Lirong Xia. Preference functions that score rankings and maximum likelihood estimation. In *Proceedings of the Twenty-First International Joint Conference on Artificial Intelligence (IJCAI-09)*, 2000.

- [10] Partha Dasgupta and Eric Maskin. On the robustness of majority rule. *Journal of the European Economic Association*, 6(5):949–973, 09 2008.
- [11] Orhan Erdem and M. Sanver. Minimal monotonic extensions of scoring rules. *Social Choice and Welfare*, 25(1):31–42, October 2005.
- [12] Pierre Favardin and Dominique Lepelley. Some further results on the manipulability of social choice results. *Social Choice and Welfare*, 26:485–509, 2006.
- [13] Pierre Favardin, Dominique Lepelley, and Jerome Serais. Borda rule, Copeland method and strategic manipulation. *Review of Economic Design*, 7:213–228, 2002.
- [14] Mario Fedrizzi, Janusz Kacprzyk, and Hannu Nurmi. How different are social choice functions : a rough sets approach. *Quality and Quantity*, 30(1):87–99, 1996.
- [15] P.C. Fishburn. *The Theory of Social Choice*. Princeton University Press, 1973.
- [16] Peter Fristrup and Hans Kleiding. A note on asymptotic strategy-proofness. *Economics Letters*, 31:307–312, 1989.
- [17] Evgeniy Gabrilovich and Shaul Markovitch. Wikipedia-based semantic interpretation for natural language processing. *J. Artif. Int. Res.*, 34(1):443–498, 2009.
- [18] Allan Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41(4):587–601, July 1973.
- [19] Alexander Halavais and Derek Lackaff. An analysis of topical coverage of wikipedia. *Journal of Computer-Mediated Communication*, 13(2):429–440, 2008.
- [20] Melvin J. Hinich and Michael C. Munger. *Analytical Politics*. Cambridge University Press, Cambridge, UK, 1997.
- [21] D. Hochbaum. *Approximation Algorithms for NP-Hard Problems*. PWS Publishing Co., 1995.
- [22] Nan Hu, Paul A. Pavlou, and Jennifer Zhang. Can online reviews reveal a product’s true quality?: empirical findings and analytical modeling of online word-of-mouth

- communication. In *Proceedings of the 7th ACM conference on Electronic commerce*, EC '06, pages 324–330, New York, NY, USA, 2006. ACM.
- [23] Jerry S. Kelly. Condorcet winner proportions. *Social Choice and Welfare*, 3(1):311–314, 1986.
- [24] Jerry S. Kelly. Almost all social choice rules are highly manipulable, but a few aren't. *Social Choice and Welfare*, 10(2):161–175, April 1993.
- [25] J.S. Kelly. Conjectures and unsolved problems: 8. interjacency. *Social Choice and Welfare*, 6:331–335, 1989.
- [26] A. Kittur, E. H. Chi, B. A. Pendleton, B. Suh, and T. Mytkowicz. Power of the few vs. wisdom of the crowd: Wikipedia and the rise of the bourgeoisie. 2007.
- [27] Aniket Kittur, Bongwon Suh, Bryan A. Pendleton, and Ed H. Chi. He says, she says: conflict and coordination in wikipedia. In *CHI '07: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 453–462, New York, NY, USA, 2007. ACM.
- [28] Travis Kriplean, Ivan Beschastnikh, David W. McDonald, and Scott A. Golder. Community, consensus, coercion, control: cs*w or how policy mediates mass participation. In *GROUP '07: Proceedings of the 2007 international ACM conference on Supporting group work*, pages 167–176, New York, NY, USA, 2007. ACM.
- [29] Dominique Lepelley and Fabrice Valognes. Voting rules, manipulability and social homogeneity. *Public Choice*, 116:165–184, 2003.
- [30] Fangzhen Lin and Pingzhong Tang. Computer-aided proofs of Arrow's and other impossibility theorems. In *AAAI*, pages 114–119, 2008.
- [31] Charles F. Manski. Interpreting the predictions of prediction markets. *Economics Letters*, 91(3):425 – 429, 2006.
- [32] Eric Maskin and Tomas Sjoström. *Implementation theory*, volume 1 of *Handbook of Social Choice and Welfare*. Elsevier, March 2002.

- [33] Dov Monderer and Lloyd S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124 – 143, 1996.
- [34] H. Moulin. On strategy-proofness and single peakedness. *Public Choice*, 35(4):437–455, January 1980.
- [35] H. Moulin. *Axioms of Cooperative Decision Making*. Cambridge University Press, 1988.
- [36] Eitan Muller and Mark A. Satterthwaite. The equivalence of strong positive association and strategy-proofness. *Journal of Economic Theory*, 14(2):412–418, April 1977.
- [37] Michael Munie and Yoav Shoham. Joint process games: From ratings to wikis. In *AAMAS '10: Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 2010.
- [38] Roger B. Myerson. Extended poisson games and the condorcet jury theorem. *Games and Economic Behavior*, 25(1):111 – 131, 1998.
- [39] Klaus Nehring and Clemens Puppe. Efficient and strategy-proof voting rules: A characterization. *Games and Economic Behavior*, 59(1):132 – 153, 2007.
- [40] Klaus Nehring and Clemens Puppe. The structure of strategy-proof social choice – part i: General characterization and possibility results on median spaces. *Journal of Economic Theory*, 135(1):269 – 305, 2007.
- [41] Shmuel Nitzan. The vulnerability of point-voting schemes to preference variation and strategic manipulation. *Public Choice*, 47(2):349–370, January 1985.
- [42] Hannu Nurmi. Discrepancies in the outcomes resulting from different voting schemes. *Theory and Decision*, 25(2):193–208, 1998.
- [43] Prasanta K. Pattanaik. Strategic voting without collusion under binary and democratic group decision rules. *The Review of Economic Studies*, 42, 1975.

- [44] Bezalel Peleg. A note on manipulability of large voting schemes. *Theory and Decision*, 11:401–412, 1979.
- [45] Thomas Piketty. The information-aggregation approach to political institutions. *European Economic Review*, 43(4-6):791–800, April 1999.
- [46] Thomas Piketty. Voting as communicating. *Review of Economic Studies*, 67(1):169–91, January 2000.
- [47] Geoffrey Pritchard and Arkadii Slinko. On the average minimum size of a manipulating coalition. *Social Choice and Welfare*, 27:263–277, 2006.
- [48] Morten Rask. The reach and richness of wikipedia: Is wikinomics only for rich countries? *First Monday*, 13(6).
- [49] Julia Robinson. An iterative method of solving a game. *The Annals of Mathematics*, 54(2):296–301, 1951.
- [50] M. 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, April 1975.
- [51] James Schummer and Rakesh V. Vohra. Strategy-proof location on a network. *Journal of Economic Theory*, 104(2):405 – 428, 2002.
- [52] Yoav Shoham and Kevin Leyton-Brown. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, December 2008.
- [53] Arkadii Slinko. On asymptotic strategy-proofness of the plurality and run-off rules. *Social Choice and Welfare*, 19:313–324, 2002.
- [54] David A. Smith. Manipulability measure of common social choice functions. *Social Choice and Welfare*, 16:639–661, 1999.
- [55] Sander Spek. Wikipedia: organisation from a bottom-up approach, Nov 2006.

- [56] William Thomson. Monotonic extensions on economic domains. *Review of Economic Design*, 4(1):13–33, 1999.
- [57] Fernanda B. Viegas. The visual side of wikipedia. In *HICSS '07: Proceedings of the 40th Annual Hawaii International Conference on System Sciences*, page 85, Washington, DC, USA, 2007. IEEE Computer Society.
- [58] Daniel S. Weld, Fei Wu, Eytan Adar, Saleema Amershi, James Fogarty, Raphael Hoffmann, Kayur Patel, and Michael Skinner. Intelligence in wikipedia. In *AAAI'08: Proceedings of the 23rd national conference on Artificial intelligence*, pages 1609–1614. AAAI Press, 2008.
- [59] Dennis M. Wilkinson and Bernardo A. Huberman. Cooperation and quality in wikipedia. In *WikiSym '07: Proceedings of the 2007 international symposium on Wikis*, pages 157–164, New York, NY, USA, 2007. ACM.
- [60] Eric Yeh, Daniel Ramage, Christopher Manning, Eneko Agirre, and Aitor Soroa. Wikiwalk: Random walks on wikipedia for semantic relatedness. In *ACL-IJCNLP TextGraphs-4 Workshop 2009*, 2009.
- [61] J. Yu, J. A. Thom, and A. Tam. Ontology evaluation using wikipedia categories for browsing. In *CIKM'07 Proceedings of the sixteenth ACM conference on Conference on information and knowledge management*, pages 223–232, November 2007.

Michael Andrew Munie

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Yoav Shoham) Principal Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Jure Leskovec)

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Ilya Segal)

Approved for the University Committee on Graduate Studies