

**AGENT MODELING OF
HUMAN INTERACTION:
Stability, Dynamics and Cooperation**

Thesis submitted for the degree of
“Doctor of Philosophy”

By

Omer Lev

Submitted to the Senate of the
Hebrew University of Jerusalem

August 2015

This work was carried out under the
supervision of:

Jeffrey S. Rosenschein

Abstract

This work models human behavior and interactions, particularly in domains where there are results indicating that game theoretic solutions diverge from real world dynamics. We do so by extending existing models to incorporate better heuristics and approximations of how people behave and interact.

We begin with voting and social-choice where, as the Gibbard-Satterthwaite theorem states, no reasonable voting rule is strategyproof; i.e., there will always be scenarios in which voters are better off mis-reporting their preferences and voting strategically. This suggests that in order to understand election results, we should look at Nash equilibria – states where no voter wishes to change their vote. However, not only are there many equilibria, but a significant number of them would never be the outcome of real world votes. We propose several enhancements to the standard voting model to handle these problems: assumptions on voters’ *nature* (a slight preference of voters to vote for what they truly believe); assumptions on voters’ *dynamics* (they participate in an iterative process, in which each voter changes their preference if they believe it will change the outcome); and a more complex model where voters have a view of “*plausible*” *election outcomes* based on some notion they have of how others will vote (e.g., a poll). In addition to theoretical proofs and characterizations, we also provide simulation results indicating that these models give us more realistic outcomes.

We continue with a similar problem in which Nash Equilibrium, as a solution concept, does not correspond to real world outcomes. All-pay auctions model various human endeavors such as the development of medical drugs

or crowdsourcing, in which many participants exert effort, but only one wins (e.g., a drug patent) while the rest lose their investment. A naive analysis indicates that no one should participate in such auctions, as the expected profit is zero. We show that assuming players collude (i.e., cooperate without other players' knowledge) increases the colluders' expected profit, and in some cases increases the expected profit of the non-colluders as well, due to their ignorance of the collusion. We further address the same problem via a different path, and show that assuming auction participants are not fixed, but rather have a certain probability of participating (as in, for example, crowdsourcing), the expected profit is positive as well.

We conclude with two different approaches to using graphs and networks for modeling interactions. Using *cooperative game theory* we analyze weakest link games, in which a path between a source and a target is valued as the weight of the lightest edge in the path. We show algorithms for finding stable solutions, and the NP-completeness of finding optimal coalition structure (as well as an appropriate approximation algorithm). In particular, we focus on cost of stability calculations, which determine how much more needs to be given to the game scenario to ensure its stability. We then turn to utilizing the graph in a different manner, as representing a social graph, to parts of which we wish to give group recommendations (e.g., what restaurant shall a group of friends go to, or what game should they play together). We pursue an *axiomatic approach*, in which we specify desirable features, showing that some are incompatible, while some result in a unique recommendation algorithm.

Improved models enable a better understanding of humans, capturing in each setting the most salient features that allow us to better represent interactions. Second, they enable us to design agents that react and interact with people, understanding, albeit to a limited degree, how people behave in certain settings. Finally, this enables better *mechanism design*, giving system designers better models and tools to construct their solutions.

Contents

1	Introduction	1
1.1	The Context of Modeling Human Behavior	1
1.2	Decision Making and Elections	4
1.3	All-Pay Auctions	6
1.4	Networks – Human and Otherwise	6
1.5	A Note on Papers	9
I	Decision-Making and Elections	11
2	Decision-Making Overview & Preliminaries	13
2.1	Preliminaries and Definitions	15
2.1.1	Voting Rules	17
2.1.2	Voters: Truth-Bias and Lazy-Bias	19
2.1.3	Iterative Voting	20
3	The Truth(-Bias) is Out There: Truth-Biased Voters	23
3.1	Introduction	23
3.1.1	Related Work	24
3.2	Simulation Methodology	26
3.3	Simulation Results	28
3.4	Analytical Results	32
3.4.1	Veto Voting Rule	33

CONTENTS

3.4.2	A Constructive Algorithm	34
3.4.3	k -Approval	35
3.5	Summary	36
4	Don't stop 'Til You Get Enough: Iterative Voting	39
4.1	Introduction	39
4.1.1	Related Work	41
4.2	Tie-Breaking Rules	43
4.3	Voting Rules	43
4.3.1	Veto	43
4.3.2	Other Scoring Rules	44
4.4	Voter Types	44
4.4.1	Regular Voters	44
4.4.2	Truth-Biased Voters	44
4.4.3	Lazy-Biased Voters	47
4.5	Summary	48
5	Think Local, Act Global: Local Dominance Voting	53
5.1	Introduction	53
5.1.1	Related Work	55
5.2	The Local Dominance Model	56
5.2.1	Locality	56
5.2.2	Dominance	58
5.2.3	Truth and Lazy Bias	58
5.2.4	Generality	60
5.3	Iterative Convergence	60
5.4	Truth-Bias and Lazy-Bias	62
5.5	Simulation Results	62
5.6	Summary	66

CONTENTS

II	All-Pay Auctions	69
6	All-Pay Auction Overview & Preliminaries	71
6.1	Preliminaries and Definitions	73
6.1.1	Bids	74
6.1.2	Bidders	75
6.1.3	Auctioneers	76
7	Come Together: Mergers and Collusions in All-Pay Auctions	79
7.1	Introduction	79
7.1.1	Related Work	80
7.2	Mergers	82
7.3	Collusions	83
7.3.1	A Single Group of Colluders	84
7.3.2	Optimal Number of Colluders	85
7.3.3	Auctioneer's Profits	86
7.3.4	Social Welfare	86
7.3.5	Other Equilibria	87
7.4	Response To Collusion	88
7.4.1	A Player Aware Of Collusion	88
7.4.2	Several Groups Of Colluders	89
7.5	Summary	90
8	"Showing Up is 80% of Life": Participation in All-Pay Auctions	93
8.1	Introduction	93
8.1.1	Related Work	94
8.2	Equilibrium	95
8.3	Profits	97
8.4	Special Case: A Single Participation Probability	100
8.5	Summary	101

III	Networks	103
9	Networks Overview	105
10	“You are the Weakest Link, Goodbye!”: Weakest Link Games	107
10.1	Introduction	107
10.1.1	Related Work	109
10.2	Preliminary: Cooperative Game Theory	111
10.2.1	Stability Concepts	112
10.2.2	Coalition Structures	114
10.3	Weakest Link Game Definition	114
10.4	Core and Least-Core	116
10.5	Cost of Stability and Series and Parallel Compositions	119
10.6	Optimal Coalition Structures	121
10.7	Summary	121
11	Where Do We Want to Go Today?: Group Recommendations	123
11.1	Introduction	123
11.1.1	Related Work	125
11.2	The Group Recommendation Model	126
11.3	The Axioms	128
11.3.1	The Basic Axioms	128
11.3.2	Group Power Axioms	130
11.3.3	Influence Structure Axioms	132
11.3.4	Independence of Axioms from Each Other	133
11.4	An Impossibility Result	134
11.5	A Recommendation System that Works	135
11.6	Summary	137
12	When All is Said and Done...	139
	Bibliography	143

CONTENTS

IV	Appendices	175
A	Proofs of Part I	177
A.1	Chapter 3	177
A.2	Chapter 4	186
A.3	Chapter 5	203
B	Proofs of Part II	209
B.1	Chapter 7	209
B.2	Chapter 8	214
C	Proofs of Part III	217
C.1	Chapter 10	217
C.2	Chapter 11	224
D	Information on Local Dominance Simulations	233
D.1	Distributions	234
D.2	Collected Variables	235

CONTENTS

Lord, what fools these mortals be!

WILLIAM SHAKESPEARE, *A Midsummer
Night's Dream*

Chapter 1

Introduction

1.1 The Context of Modeling Human Behavior

Understanding and analyzing human behavior has been attempted since the dawn of human societies, as rulers (and their nascent bureaucracies) have struggled to understand how they can implement their policies with little resistance, and how they can motivate people towards their desired goals. Naturally, philosophers have ventured to answer this, and diverse schools of thought – from Plato and his Greek contemporaries to Confucius in China – have tried to understand human behavior and develop “mechanisms” that would change this behavior more to their liking.

These attempts – and the attempts in the millennia that followed – were mostly based on an effort to generalize human behavior from people’s observed actions in particular times and places. This was usually done through a moral prism, often influenced by the theological views of the time.

While the philosophical efforts and debates continued unabated, the advent of the modern age during the late 18th century and early 19th century, made the need for a generalized analysis of human behavior even more acute and necessary [103, 234, 221]:

- **The large state:** Countries began increasing in size, and along with the speeding of communication, this meant many decisions were taken in central bureaucracies, where detailed knowledge of local conditions was limited, to a large extent.
- **Democracy:** Slowly, mass participation democracy was beginning to spread. This, combined with the previous point, meant there were dedicated parties whose purpose was to understand people's opinions and influence them (in contrast to appeals to particular influential individuals in the past).
- **Capitalism:** The growth of large companies, selling a wide range of products and serving customers in diverse locations and situations, meant that there was a "private interest" in understanding people's behaviors and conduct in a general way, and not just understanding their practices in particular settings (as a small shopkeeper may be concerned with).

The 19th century saw a huge growth in disciplines trying to explain various aspects of human behavior – economics, sociology, anthropology, etc. Each tried to explain humanity using a certain view of what people are and what society is. In this climate *game theory* began developing in the 20th century, gathering pace mid-century, as the cold-war and the nuclear age increased the need to analyze long-term, large scale military conflicts. Game theory tries to explain people's behavior as stemming from the interaction of rational individuals. The reasoning of individuals is not part of the model, allowing a great flexibility in modeling people, and the use of mathematical tools allows a more rigorous consideration of the effects of various policies and programs, while also forcing the formalization of various assumptions on the nature of people and their interactions.¹

¹While it is outside of the scope of this work, the main criticism against game theory, challenging the assumption of rational agents [226, 123, 227], has been attempted to be formalized within various game theoretical models as well, as various human non-rational

Game theory accords a special place to the concept of stability or *equilibrium* – a situation or state all parties to a certain scenario are content with, and in which they do not seek to modify their behavior. Among the various stability notions suggested, one of the most widely used is the *Nash Equilibrium* [171]. Nash’s model assumes participants in some situations will modify their actions according to the scenario they are in and to other participants’ actions, in order to improve the outcome for them, until some stable situation is reached (if it exists). The situation in which no participant can improve their outcome by modifying their behavior is the Nash equilibrium. Side by side, the theory of *cooperative game theory* developed to handle cases where we allow participants to cooperate, and assume they can find a way to divide their proceeds among themselves (hence these are also called “transferable utility games”).

The flip side of being able to better analyze various scenarios is that it enables better planning – by understanding what is a more desirable outcome for a system planner and by understanding the effects of various attempts to bring it about. Indeed, this embodies the concepts of *mechanism design*, in which we attempt to design a system that fulfills a set of desirable properties. Mechanism design goes hand-in-hand with work intending to flesh out what are the precise desirable properties of various complex tasks using the *axiomatic approach*, in which we try to enumerate properties which system planners might find attractive, and attempt to understand which mechanisms fit each set of properties.

Naturally, the complexity of the human condition does not allow us a grand, overarching, model of human behavior in the foreseeable future. However, in an ever expanding number of cases and situations, our models are better and better, and can ably analyze human behavior. As can be expected, one of the earlier cases where game theory showed its contributions was in scenarios which contained strict monetary compensation, simplifying greatly the issue of determining the desires of the players and their possible

actions are better understood and classified.

actions.

Other human endeavors, which could be modeled less easily, require us to attempt to delve deeper into the scenarios and more thoroughly understand peoples' reasoning and motives. This attempt is at the core of this work, as we strive to better understand several facets of human behavior. When setting up the models, we wish to provide a framework with a greater descriptive power of people's conduct while utilizing the great power given to us by the fundamental game theoretic field. When pursuing mechanism design and axiomatic approaches we relied on the cooperative aspects of the theory, while bringing to the fore the network structure of people and societies to aid in understanding human actions.

1.2 Decision Making and Elections

The fundamental problem of soliciting people's opinions on an issue, and reaching a decision based on these opinions has accompanied humans for millennia. However, as large democratic countries began appearing, in the latter half of the 18th century, a far greater interest in these decision making processes took place, symbolized by such researchers as the French Jean-Charles de Borda and Nicolas de Condorcet.

The first major issue that needed to be handled was determining what is a preferable outcome. Different approaches to this question result in different voting rules used to determine the winner in the elections. However, one can broadly portray some rules as striving towards a consensus, such as the Condorcet winner, which is one for which no other candidate can gather enough support to win against. However, as Condorcet himself noticed, this type of winner does not always exist (the so-called "Condorcet paradox"), requiring some decision mechanism to determine what is preferable when there is no Condorcet winner. Other voting techniques, such as plurality, wish to rely, on the other hand, on a strong core of supporters. This division touches on the difference of the voting problem from "regular" game theoretic

problems – we do not know voters’ strength of beliefs (or “utility functions”), only their ranking of options.

The second major issue was that people try to “game” the voting mechanism: they understand that in many cases they can guarantee an outcome more to their liking if they vote for an option that is not their true belief. In other words, we might design a perfect voting technique, taking into considerations all opinions, and delivering its perfect outcome, but as long as voters are incentivized to misreport their beliefs, it cannot be implemented. While Borda is said to have remarked “my election method is only for honest men”,² in real life scenarios this is a significant issue, and various thinkers tried to find a voting method that would be strategyproof – in which participants cannot improve their situation by being non-truthful.

However, in the 1970s, mathematicians Gibbard and Satterthwaite (building on Arrow’s impossibility theorem [14]) proved [109, 206] that there can never be a reasonable voting system that guarantees that truthfulness is the best strategy for voters. Fifteen years after this result, Bartholdi, Tovey and Trick [54] began a line of research that discusses the complexity facing voters when they try to find a good non-truthful vote that will change the outcome more to their liking. While this line of work continued [53, 73, 188, 161, 236], its focus on finding beneficial voter manipulations, raised a more fundamental issue – even disregarding complexity issues, can we say something on what results of an election will look like under some voting rule?

The answer to this question can return us also back to the issue of comparing the outcomes of voting methods. Manipulations have, in a sense, made the issue of selecting a good voting technique moot – the properties of the technique have limited application when voters are not truthful. Using the models by which we try to find election outcomes, along with the simulation framework we built, election methods can be compared, using whatever criterion of “good outcome” desired.

²There are several English translations of Borda’s remark, but the French original appeared in Sylvestre François Lacroix’s *Eloge Historique de Borda* in 1800.

1.3 All-Pay Auctions

Auctions, in general, are a way to allocate indivisible goods to those which desire these goods. There are many, many types of auctions, but among them, the all-pay auction seems a bit odd. All-pay auctions are mechanisms in which participants are vying for some object. Each of them gives a bid, and the participant with the highest bid receives the item. However, unlike most auctions, all players must pay their bids, whether they obtained the object or not. Despite this mechanism's oddity, all-pay auctions model very common human contests: when medical companies are racing to develop and patent a drug, they are all putting in a large amount of effort, which they will not get back, but only one company – the one that will patent the drug – will actually reap the benefit of the effort. In the past decade, all-pay auction mechanism have proliferated as *crowdsourcing* in various sites (e.g., Topcoder.com [138]), in which multiple participants are trying to solve a problem, but only one will manage to do so and get the benefits of their success.

Prima facie, this mechanism seems to lie in the realm of easy to analyze game theoretic problems, as it involves pure monetary rewards. However, simple analysis [55, 56] shows that participants' expected profit to be zero, i.e., there is little reason for them to actually participate. As these mechanisms actually exist in the real world, this seems to indicate a flaw in the model. The attempt to bring more real-life conditions to this model is what animates the second part of this work, as we make our participants more realistic in their strategies – they may cooperate, collude, or choose to avoid participation in an auction.

1.4 Networks – Human and Otherwise

While game theory – cooperative and non-cooperative alike – is based on the interaction between participants, hence on the existence of society, dealing with the general case of all possible interactions is often intractable. Hence,

using networks to model human relations and interactions is a common tool when trying to bring some order to the various variations of human connections. Furthermore, networks do not model only human interactions, but represent other interconnected systems: from organizational structure, describing the hierarchy and working of an organization, to communication and computer network, including the whole spectrum from global sized networks (such as the internet) to small, local networks.

Dealing with network structures allows us to survey several different areas which can be expressed in this manner. While they are different from each another, they are both joined by the centrality of the network structure to the fundamental problem, as well as in their practical appliance, which we try to address, by giving results to aid a system designer encountering such settings.

Weakest Link Coalitions

In some scenarios, a group's value is the strength of its most extreme member. Whether it is a minimal quality of various parts forming a product, the maximal allowed weight when traveling along different types of road, or the safety distance from a factory, all these problems rely on examining the various components of a process, and finding the most stringent constraint. In many cases, such as a manufacturing process or travel destinations, these components can be represented by a graph.

As we seek to choose which components to make up our output to have the best “weak” part, we are less concerned with the particular division of profit between the components, but more in creating a stable situation where all parts manage to choose the best path. Hence, using cooperative game theory, in which we look for a group that performs a task without deep consideration for the particular “profit” or utility made by each agent, but rather – looking at the “big picture”: whether or not a good yet stable situation may be reached, and if it cannot, how can we induce it externally.

Recommendation Systems

The ability to recommend items to people based on their preference of other items has been a subject of research in the past few centuries, as large scale markets presented people with a much greater variance within the various items they could purchase (e.g., books), along with the spread of mass communication methods, which allowed people to receive recommendations from people they have never met. Moreover, online markets in the past two decades have brought this problem to the fore, as companies strive to maintain buyers' interest by pushing them to purchase more. Emblematic of this attempt has been Netflix's wish to present movie recommendations to their viewers based on the movies they have indicated they enjoyed, leading to the "Netflix challenge" [57, 114], in which Netflix offered a million dollars to the best performing algorithm.

However, the "holy grail" of recommendations remains that of the family and friends in the social circle of the person being recommended to [216]. This brought about in the last decade an attempt to model a person's social sphere using a graph structure and, utilizing knowledge of friends and family preferences, attempt a recommendation. In some sense, this is similar to elections using a social component. Furthermore, in the last few years, with the growth of social networks and the general trend of online social activity (e.g., networked game play), a more complex recommendation problem arose – that of group recommendations. In group recommendations we seek to give a joint recommendation to a group of people, e.g., a group of friends searching for a restaurant to go to for a dinner together.

Group recommendations make the usage of the social structure of the recommended group almost imperative: how can a group of independently minded individuals be recommended the same as a group consisting of an influential individual to which all other group members accede? However, this field of research is so young, that it is not entirely clear what are desirable properties for such systems, a lacuna we try to correct, along with suggesting an algorithm to fulfill some of these properties.

1.5 A Note on Papers

I was not able to put all my relevant papers in this work, as in the interest of creating a more flowing, coherent structure, various different approaches would have taken too much space and would have reduced the readability of this work (for example, [4, 136, 145] and several papers in submission). Even included papers had some material excluded from them in the interest of coherence and (relative) brevity (this also necessitated moving all proofs to appendices). I wish to note the papers upon which this work was based, and deeply thank my co-authors:

Chapter 3 “Empirical Aspects of Plurality Election Equilibria” [224] (AAMAS 2013) with David R.M. Thompson, Kevin Leyton-Brown and Jeffrey S. Rosenschein; and “Beyond Plurality: Truth-Bias in Binary Scoring Rules” [175] (ADT 2015) with Svetlana Obraztsova, Evangelos Markakis, Zinovi Rabinovich and Jeffrey S. Rosenschein.

Chapter 4 “Convergence of Iterative Voting” [144] (AAMAS 2012) with Jeffrey S. Rosenschein; and “Analysis of Equilibria in Iterative Voting Schemes” [189] (AAAI 2015) with Zinovi Rabinovich, Svetlana Obraztsova, Evangelos Markakis, and Jeffrey S. Rosenschein.

Chapter 5 “A Local-Dominance Theory of Voting Equilibria” [159] (EC 2014) with Reshef Meir and Jeffrey S. Rosenschein.

Chapter 7 “Mergers and Collusion in All-Pay Auctions and Crowdsourcing Contests” [143] (AAMAS 2013) with Maria Polukarov, Yoram Bachrach and Jeffrey S. Rosenschein.

Chapter 8 “Agent Failures in All-Pay Auctions” [147] (IJCAI 2013) with Yoad Lewenberg, Yoram Bachrach, and Jeffrey S. Rosenschein.

Chapter 10 “Cooperative Weakest Link Games” [36] (AAMAS 2014) with Yoram Bachrach, Shachar Lovett, Jeffrey S. Rosenschein, and Morteza Zadimoghaddam.

Chapter 11 “An Axiomatic Approach to Group Recommendations” (in submission) with Moshe Tennenholtz.

Part I

Decision-Making and Elections

It might not have been a question of right and wrong. Which is to say that wrong choices can produce right results, and vice versa. I myself have adopted the position that, in fact, we never choose anything at all. Things happen. Or not.

HARUKI MURAKAMI, *The Second Bakery Attack*

Chapter 2

Decision-Making Overview & Preliminaries

People have been reaching decisions using voting for millennia, and people have been trying to understand and predict election results for nearly that long. However, trying to analyze elections has been hampered by the inability to trust voters are truthful in their voting. Indeed, throughout history, many commentators have noticed voters are able to improve the chance of selecting a candidate they find more favorable by misreporting their true beliefs.

However, the quest to find a voting system that is impregnable to such manipulations turned out to be futile: according to the Gibbard-Satterthwaite theorem [109, 206], every “reasonable” election system that is not a dictatorship has scenarios where voters are better off voting differently than their true beliefs. This result led to the line of research commenced by Bartholdi, Tovey, Trick and Orlin [54, 53], and expanded further by many researchers [73, 188, 187, 236, 247, 246] (see overviews in [200, 65]), which tries to assess the complexity of finding a beneficial manipulation for voters in various settings.

Despite the significant strides in this line of research, it has not helped us in analyzing election results: when many voters are strategic and manipulating, we are still at loss to understand which election results are possible

and which are not, not to mention trying to understand possible truthful preferences when presented with election results.¹ Furthermore, while game theory is often a powerful tool when analyzing situations where participants manipulate, voting scenarios have some properties that make them harder to analyze: participants only supply a ranking of their preferences, rather than the utility value of each option (as is typical in monetary settings), and there is a common assumption that participants have a good perception of others' utilities. In any case, looking at equilibria in voting settings results not only in an enormous amount of equilibria (hundreds of thousands even in small games – 10 voters, 5 candidates), but much worse from an analysis viewpoint: many Nash equilibria are practically useless for any analytic purpose, and do not represent a realistic possible end-state of an election. For example, even if all voters rank the same candidate in last place, there are Nash equilibria where it wins (e.g., in plurality, all voters voting for this candidate is a Nash equilibrium). This renders Nash equilibria useless as an analytical tool, despite their fundamental relevance as showing stable states which incorporate voters' manipulations, and makes analysis of actual election outcomes very limited.

This part will introduce several paths in the effort to make election analysis more realistic by combining theoretical analysis with simulations that help examine whether the models are producing results that resemble real-world cases. In the course of this work we have released an open-source framework for election simulations which is extensible and modular to enable various research both extending our own, or radically different from it.²

Assumption on nature of voters: This includes **truth-bias**, the idea that all things being equal, voters prefer to vote truthfully. While voters are

¹This problem is exacerbated when considering independent computerized agents. They are motivated to manipulate, and may have the computation resources to calculate the particular strategy to do so optimally.

²Code is at <https://github.com/omer1/IterativeVotingSimulator> and instructions are at <http://www.preflib.org/tools/ivs.php>.

still strategic, they have an innate preference to voting as they truly believe, and they will do so, if they are not adversely effecting the election result. The model introduced by Desmedt and Elkind [80] can be similarly framed as **lazy-bias**, in which voters choose to abstain and not vote if they believe they have no influence over the outcome. We examine if the Nash equilibria produced more meaningful results than those encountered with non-biased voters.

Assumption on voting mechanism: Using the **iterative voting** model introduced by Meir et al. [160], in which voters may amend their vote, we examine the bounds of the model by expanding it to more voting rules, different tie-breaking rules, etc. This dynamic means we can discuss “reachable” states, i.e., outcomes which can be reachable using this dynamic from a truthful starting-point.

Assumption on both voters and voting mechanism: We try to unite both approaches by examining type of voters within the iterative voting model. However, our work culminates in a more robust and general model, **local dominance**, in which the previous models are just a particular instantiation.

2.1 Preliminaries and Definitions

We shall first set out to define elections: an election \mathcal{E} is composed of several elements:

- *Candidates:* A group C of m elements.
- *Voters:* A group V of size n . Each voter $i \in V$ is associated with an element $a_i \in \pi(C)$, where $\pi(C)$ is the set of linear orders of the elements of C . This represents the truthful preferences of a particular voter. To make this similar to common game-theoretic settings, we can associate this preference order with some utility function $u_i : C \rightarrow \mathbb{R}$,

in which there is a higher utility from higher ranked candidates, but this utility function is hidden from anyone but the voter itself. We mark this preference order by \succ_i .

- *Voting function*: A function $f : (\pi(C))^n \rightarrow C$ which takes in a set of preferences from each voter and outputs a winner. See further details below.

As a voter can be strategic, it may report to the voting rule a different preference than its truthful \succ_i preference order. That is, each voter i is reporting to the voting rule $b_i \in \pi(C)$ and b_i might not equal a_i . A set of all n reported votes is a *voting profile*, which we shall generally mark with \mathbf{b} .

Definition 2.1. A strategic deviation by voter i in a profile \mathbf{b} in which each voter $j \neq i$ is reporting a vote $b_j \in \pi(C)$, is when there is a preference order $\hat{b}_i \in \pi(C)$, $\hat{b}_i \neq b_i$ such that

$$f(b_1, \dots, \hat{b}_i, \dots, b_n) \succ_i f(b_1, \dots, b_i, \dots, b_n)$$

We denote the profile \mathbf{b} without the voter of voter $i \in V$, as \mathbf{b}_{-i} .

The Gibbard-Satterthwaite theorem tells us that all reasonable voting rules have a profile for which there are voters which have a strategic deviation from the truthful profile.

Definition 2.2. A profile \mathbf{b} is a Nash equilibrium when no voter has a strategic deviation from the profile. So for any voter i ,

$$f(b_1, \dots, b_i, \dots, b_n) \succeq_i f(b_1, \dots, \hat{b}_i, \dots, b_n)$$

for any $\hat{b}_i \in \pi(C)$.

We use the terminology of Nash equilibria in voting scenarios to refer only to pure Nash equilibria, as real-world voters seem averse to the randomizations involved in a mixed Nash equilibrium (and they are not).

A voting function $f : (\pi(C))^n \rightarrow C$ is actually comprised of two parts – the voting rule $\hat{f} : (\pi(C))^n \rightarrow 2^C$ itself, and a tie-breaking rule $t : 2^C \rightarrow C$. Thus, $f = t(\hat{f})$.

Definition 2.3. A tie-breaking rule is a function $t : 2^C \rightarrow C$, which takes a subset of C and returns a single candidate. A particular type of tie-breaking rules is linear tie-breaking rules which are defined by an order $a \in \pi(C)$, and for $\hat{C} \subseteq C$, $t(\hat{C})$ is $c \in \hat{C}$ which is the highest ranked member in \hat{C} according to the order a (i.e., for all $c' \in \hat{C}$, $c' \neq c$, $c \succ_a c'$).

2.1.1 Voting Rules

Many voting rules have been suggested throughout history, though we shall focus only on a few of them. In particular, we will pay much attention to a particular family of voting rules:

Definition 2.4. A scoring rule is a voting rule defined by a vector

$$(\alpha_1, \alpha_2, \dots, \alpha_{m-1}, 0)$$

in which $\alpha_1 \geq \alpha_2 \geq \alpha_{m-1} \geq 0$. Each voter's vote contributes α_1 points to the highest ranked candidate in its reported preference, α_2 points to the second ranked candidate and so on. Finally, the candidates which scored the maximal number of points are the winners.

Many commonly used voting rules are scoring rules. In particular, we shall focus on:

Plurality The most widely-used voting rule. Equivalent to the scoring rule $(1, 0, \dots, 0)$.

Veto A voting rule diametrically opposed to plurality. Equivalent to the scoring rule $(1, 1, \dots, 1, 0)$.

k -approval/ k -veto Lying between plurality and veto, in which we fix the number of candidate getting a score of 1, or the number of those getting a score of 0. Equivalent to the scoring rule $(1, \dots, 1, 0, \dots, 0)$.

Borda The scoring rule defined by $(m - 1, m - 2, \dots, 1, 0)$.

A different concept of winner has been developed in the late 18th century, by the Marquis de Condorcet:³

Definition 2.5. A Condorcet winner is a candidate c such that for all other candidates $\hat{c} \in C$

$$|\{v \in V \mid c \succ_v \hat{c}\}| > |\{v \in V \mid \hat{c} \succ_v c\}|$$

In other words, there is a majority of voters which prefer c over any other specific candidate.

As Condorcet himself realized, the existence of a Condorcet winner is not assured. Furthermore, Fishburn [100] showed no scoring rule is *Condorcet consistent*, ensuring that if there is a Condorcet winner it will be selected. However, there are several Condorcet consistent voting rules, of which we shall only mention Maximin

Definition 2.6. The Maximin voting rule assigns each candidate a score, which is the number of voters it is guaranteed to support it against any other candidate:

$$sc(c) = \min_{\hat{c} \in C, c \neq \hat{c}} |\{v \in V \mid c \succ_v \hat{c}\}|$$

The maximin winner is the candidate with the maximal score – $\max_{c \in C} sc(c)$.

From here on, when discussing voting rules, we assume they are part of voting rule families, such that they are defined for any n voters.

³Though this is, apparently, a re-development of the idea. The concept of a Condorcet winner was developed by the 13th century Majorcan writer Ramon Llull, whose writings on elections were rediscovered in 2001.

2.1.2 Voters: Truth-Bias and Lazy-Bias

Beyond purely truthful voters (who do not strategize, but always vote their truthful preference), and strategic voters, introduced above, we deal with voters with a slightly biased nature. Unlike voters so far, their utility does not just stem from the outcome of the election, but also from what they themselves vote. We wish them to be strategic, but have some tendency when they do not have any strategic move.

Definition 2.7. *A truth-biased voter is a voter i with a truthful vote a_i , such that in a profile \mathbf{b} where it has no strategic deviation, and if*

$$f(b_1, \dots, b_i, \dots, b_n) = f(b_1, \dots, a_i, \dots, b_n)$$

reverts to its truthful vote.

More formally, using the voter's hidden, underlying, utility function, we define $\epsilon_i = \frac{1}{2} \min_{c, c' \in C} |u_i(c) - u_i(c')|$. We now define a new utility function for the voter, in which it gains ϵ_i utility if it is truthful (no other change from the utility function based on the election outcome).

A further explanation of truth-bias and its uses will be given in Chapter 3.

Similarly, we can define lazy-biased voters (introduced, in a different name and notation, in Desmedt and Elkind [80]), which are similar to truth-biased voters, except that they abstain from voting, rather than reverting to their truthful preferences.

Definition 2.8. *A lazy-biased voter is a voter such that in a profile \mathbf{b} where it has no strategic deviation, and if*

$$f(b_1, \dots, b_i, \dots, b_n) = f(b_1, \dots, b_{i-1}, b_{i+1}, \dots, b_n)$$

(i.e., without its participation), does not participate.

More formally, using the voter's hidden, underlying, utility function, we define $\epsilon_i = \frac{1}{2} \min_{c, c' \in C} |u_i(c) - u_i(c')|$. We now define a new utility function for the voter, in which it gains ϵ_i utility if it does not participate (no other change from the utility function based on the election outcome).

In our work it will be useful to define a subset of candidates which are viable options to become a winner:

Definition 2.9. *A runner-up in a voting profile \mathbf{b} when using a scoring rule as a voting rule, is a candidate that is not the winner, but if it would gain an extra point would become the winner. This means it either has the same score as the winner, but is ranked lower in the tie-breaking rule, or its score is 1 below the winner's, but it is ranked above the winner in the tie-breaking rule.*

Other form of biases by voters are, of course, possible, but we focus on these two biases, which have been, to a certain extent, explored in some previous literature.

2.1.3 Iterative Voting

This is a description of the dynamic introduced in Meir et al. [160]. In that model, voter are assumed to be myopic, so they only look at the current situation and are not planning ahead. Without any fixed order (but one-by-one⁴), voters may choose to change their vote if they have a strategic deviation from the current profile. This fundamental question for this dynamic is whether it converges to a stable state (actually, a Nash equilibrium) or not. While it is straightforward to see that allowing voters any strategic deviation may lead to cycles under iterative plurality, the original paper was able to show the following theorem:

Theorem 2.1 (Meir et al. [160], Theorem 3). *An iterative setting with the following conditions:*

- *An iterative plurality game.*
- *A linear tie-breaking rule.*
- *Voters are myopic.*

⁴Parallel updates can be easily shown to not converge, as shown in [160].

- *Voters pursue a best-response strategy, so of possible strategic deviations, the voter implement the one that will make its most highly ranked candidate the winner. Moreover, the best-response must be of a particular structure: if a voter can make some candidate the winner, it votes for this candidate (this is not a real limitation for the voter, as for any best-response strategy in plurality, this is also a best-response).*

will always reach a Nash equilibrium in $\mathcal{O}(m^2n^2)$ steps.

As mentioned, the original paper was only able to show that best-response is a required property (since then this requirement has been slightly more refined in [176, 157]).

Chapter 3

The Truth(-Bias) is Out There: Truth-Biased Voters

3.1 Introduction

Attempts to try and analyze elections following the understanding that all are manipulable (Gibbard-Satterthwaite theorem [109, 206]) have mostly focused on ways of limiting manipulations. Approaches included using the complexity of manipulations in certain voting rules [54, 53, 236]; limiting the range of possible preference orders of the voters (e.g., single-peaked preferences, for which there exist strategyproof mechanisms [95]); various randomization mechanisms [110, 186], and others. Moreover, much of this research focuses on the case where we have a single aim for manipulations, such as an attempt by a group of voters to make some specific candidate the winner, while all other voters are truthful and non-strategic.

Much less effort has been given to analyzing what happens when all voters are strategic, rather than only some. The natural tool for analyzing this would be Nash equilibria, as these are states of stability, following the strategic moves of all voters. However, using the Nash equilibrium presents a significant hurdle: the number of equilibria in a voting scenarios is staggering (in a simple plurality election with 5 candidates and only 10 voters we

have over 300,000 equilibria). An even more problematic issue is that many equilibria are useless in any analysis of election results, as they would never occur, such as the situation when using plurality, when all voters rank the same candidate last, if all of them voted for it that would be an equilibrium.

Instead, we try to make a reasonable assumption on the nature of our voters, and use it to see if we get more reasonable equilibria, and in particular, equilibria that resembles more closely what we would expect to see in real-world elections. Our restriction is assuming voters are slightly truth-biased. That is, they gain a small amount of utility by voting truthfully. This amount is small enough so that if they can make a strategic deviation and improve the election outcome in their view, they would do so. But all things being equal, voters would prefer to voter truthfully.

We approach the simplest – and most common – voting rules: scoring rules composed of two values, mainly plurality and veto, but also touching on k -approval/veto. We do so by first evaluating the model using a simulation. This has not been done prior to this work due to the huge size of the game, but we leverage recent work in compact representation of games, and in particular, action-graph games (AGGs) [121, 120] and the support-enumeration method [223], enabling us to run simulations and see all of the Nash equilibria of a particular game.

We then turn to veto and analyze it analytically,¹ and go beyond just finding the complexity of finding the Nash equilibria and their winners, but manage to find a sub-family of problems in which we are able to find a constructive algorithm to find the equilibrium, if it exists.

3.1.1 Related Work

Analyzing election equilibria has been the focus of much research, with various researchers proposing different frameworks with limits and presumptions to deal with both the sheer number of equilibria, as well as with situations

¹Following the publication of these simulation results in [224], plurality was analyzed by Obraztsova et al. [177]

where there is limited information. Early work in this area, by McKelvey and Wendell [156], allowed for abstentions and defined an equilibrium as one with a Condorcet winner. As this is a very strong requirement, such an equilibrium does not always exist, but they established some criteria for this equilibrium that depends on voters' utilities.

Myerson and Weber [170] dealt with the Nash equilibria of voting games by building a model which assumes voters only know the probability of ties occurring between each pair of candidates, and that voters may abstain (for which they have a slight preference). They showed that multiple equilibria exist, and noted problems with Nash equilibrium as a solution concept in this setting. The model was further studied and expanded in subsequent research [76, 127]. Assuming a slightly different model, Messner and Polborn [167], dealt with perturbations (i.e., the possibility that the recorded vote will be different than intended), and showed that equilibria only includes two candidates ("Duverger's law").

Looking at other limitations, Feddersen et al. [97] chose (like Laffont [137]) to limit preferences to single-peaked preferences. Others, like Hinich et al. [116], for example, chose to change the single-peak limitation to a specific probabilistic model of voters over a Euclidean space of candidates, while changing other parts of the model (such as allowing for abstentions). A somewhat different approach, taken by Messner and Polborn [166], analyzed equilibria by coalitional manipulation (hence, using a stronger equilibrium than Nash – a method also utilized by Dhillon and Lockwood [81]). However, one of the main limitations of many of the papers mentioned above is that they assume some knowledge of other players' preferences.

Looking at iterative processes makes handling the complexity of considering all voters as manipulators simpler, but this literature is more thoroughly examined in Chapter 4.

Dealing more specifically with the case of abstentions, Desmedt and Elkind [80] examined both a Nash equilibrium (with complete information of others' preferences) and an iterative voting protocol, in which every voter is

aware of the behavior of previous voters (a model somewhat similar to that considered by Xia and Contizer [235]). Their model assumed that voting has a positive cost, which encourages voters to abstain; this is similar in spirit to our model's incentive for voting truthfully, although in this case voters are driven to withdraw from the mechanism rather than to participate. However, their results in the simultaneous vote are sensitive to their specific model's properties.

Rewarding truthfulness with a small utility has been used in some research, though not in our settings. Laslier and Weibull [141] encouraged truthfulness by inserting a small amount of randomness to jury-type games, resulting in a unique truthful equilibrium. Dutta and Laslier [86] attempted to inject truthfulness directly into a voting rule combined of approval voting and veto, but only found a few existence results that show truthful equilibria exist in that case. A more general result has been shown in Dutta and Sen [87], where they included a subset of participants which, as in our model, would vote truthfully if it would not change the result. They show that in such cases, many social choice functions (those that satisfy the No Veto Power) are Nash-implementable, i.e., a mechanism exists in which Nash equilibria correspond to the voting rule. However, as they acknowledge, the mechanism is highly synthetic, and, in general, implementability does not help us understand voting and elections, as we have a predetermined mechanism.

3.2 Simulation Methodology

In order to find all the Nash equilibria in a certain game, we need to represent the game in a form that the algorithm can use. As normal form games require a space that is exponential in the number of players, this representation is not useful for games with more than a few players. There are various suggestions in the literature regarding “compact” game representations that require exponentially less space to store games of particular types, such as

congestion [199], graphical [128], and action-graph games [121]. Action-graph games (AGGs) are the most useful for our purposes, because they are very compactly expressive (i.e., if the other representations can encode a game in polynomial-space then AGGs can as well), and fast tools have been implemented for working with them.

Action-graph games achieve compactness by exploiting two structural properties in a game’s payoffs:

Anonymity A player’s payoff depends only on its own action and the *number* of players who played each action.

Context-specific independence A player’s payoff depends only on a simple sufficient statistic that summarizes the joint actions of the other players.

These two properties fit scoring rules perfectly, as they are not only anonymous, but as they require a simple adding of points, are simple to calculate. Moreover, it fits truth-biased voters as well, as anonymity allows for a player’s utility function to make use of its own strategy.

Encoding our plurality voting game as action-graph games is relatively straightforward: for each set of voters with identical truthful preference, we created one action node for each possible way of voting. For each candidate, we created an “adder” node that counts how many votes the candidate receives. Directed edges encode which vote actions contribute to a candidate’s score, and that every action’s payoff can depend on the scores of all the candidates (see Figure 3.1). We used a simple Borda score as the basic utility function values (e.g., in a 5 candidate contest, if a voter’s favorite candidate wins, the voter gets a utility of 4), to which we added a small ϵ value -10^{-6} – when a voter is being truthful.

While various algorithms exist to find pure Nash equilibria in action-graph games [121, 75], we used the support enumeration method (SEM) [185, 223] exclusively as it allows Nash equilibria enumeration. This algorithm works by iterating over possible supports, testing each for the existence of a Nash




Figure 3.1: An action graph game encoding of a simple two-candidate plurality vote. Each round node represents an action that a voter can choose. Dashed-line boxes define which actions are open to a voter given its preferences.

equilibrium. In the worst case, this requires exponential time, but in practice SEM’s heuristics (exploiting symmetry and conditional dominance) enable it to find all the pure-strategy Nash equilibria of a game quickly.

3.3 Simulation Results

To examine how election equilibria look like, we ran 1,000 voting experiments using plurality as a voting rule, with 10 voters and 5 candidates, with preferences chosen from a uniform distribution . Using regular voters, such games have hundreds of thousands of Nash equilibria. However, adding a small truthfulness incentive ($\epsilon = 10^{-6}$) lowers these numbers significantly. Not counting permutations of voters with the same preferences, every game had 25 or fewer equilibria; counting permutations, the maximum number of equilibria was still only 146. Indeed, an overwhelming number of these games (96.2%) had fewer than 10 equilibria (27 with permutations). More

surprisingly, a few (1.1%) had no pure Nash equilibria at all,² and this number seems to vary with the parity of the voter number – when it is odd, there are more games where there is no equilibrium at all. To assess the impact of the truthfulness incentive, we also ran 50 experiments without it; every one of these games had over a hundred thousand equilibria.

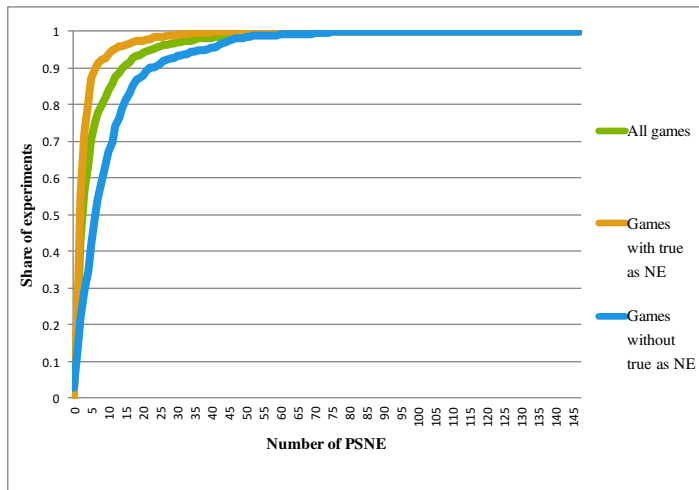


Figure 3.2: CDF showing the fraction of games having the given number of pure-strategy Nash equilibria or fewer (including permutations)

We shall examine two aspects of the results: the preponderance of equilibria with winners being the voting method’s truthful winners (which, when expanded to more voting rules, may be an interesting comparative metric between voting mechanisms), and Condorcet winners. Then, moving to the wider concept of social welfare of the equilibria, we examine the social welfare of the truthful voting rule, using Borda-like utility functions.

We first examined how likely it was for the real winners – those that would have won had all voters been truthful – to win in an equilibrium. While in 63.3% of the games truthful preferences were a Nash equilibrium, even when voters were not being truthful, the likelihood of the real winner to win in equilibrium was high: 80.4% of the games had at least one equilibrium with

²This is particularly relevant to voting procedures relying on the existence of pure Nash equilibrium, and seeking to “find” one, such as the one proposed in Meir et al. [160].




Figure 3.3: The social welfare distribution

the truthful result, and looking at the multitudes of equilibria, the average share of truthful equilibrium (i.e., result was the same as with truthful vote) was 41.56% (out of games with the truthful vote as an equilibrium, the share was 51.69%). Without the truthfulness incentive, the average share of truthful equilibrium was 21.77%.

Looking at Condorcet winners, 92.3% of games had Condorcet winners, but they were truthful winners only in 44.7% of the games (not a surprising result, as plurality is far from being Condorcet consistent). However, out of all the equilibria, the average share of equilibria with a victorious Condorcet winner was 43.49% (we ignore games without Condorcet winners). When the Condorcet winner was also the truthful winner, its average share of equilibria is 56.96%.

Looking at the wider picture (see Figure 3.4), the addition of the truthful incentive created games with very few Nash equilibria. They, very often, resulted in the truthful winner. As the number of equilibria grows, the truthful winner part becomes smaller, as the Condorcet winner part increases.

Turning to look at the social welfare of equilibria, once again, the exis-




Figure 3.4: The number of truthful and Condorcet winning equilibria, depending on total number of equilibria per experiment. Note that in the “tail”, the data is based on only a few experiments.

tence of the truthfulness incentive enables us to reach “better” equilibria. In 92.8% of the cases, the worst-case outcome was not possible at all (recall that without the truthfulness incentive, every result is possible in some Nash equilibrium), while only in 29.7% of cases, the best outcome was not possible. We note that while truthful voting led to the best possible outcome in 59% of cases, it is still, of course, dominated by best-case Nash equilibrium, i.e., there were cases where the best Nash equilibrium was not truthful (see Figure 3.3).

When looking at the distribution of welfare throughout the multitudes of equilibria, one can see that the concentration of the equilibria is around high-ranking candidates, as the average share of equilibria by candidates with an average ranking (across all voters in the election) of less than 1 was 56.38%. Even if we exclude Condorcet winners (as they are, on many occasions, highly ranked), the average ranking of less than 1 was 46.56% (excluding truthful winners resulted in 27.48% with average ranking less than

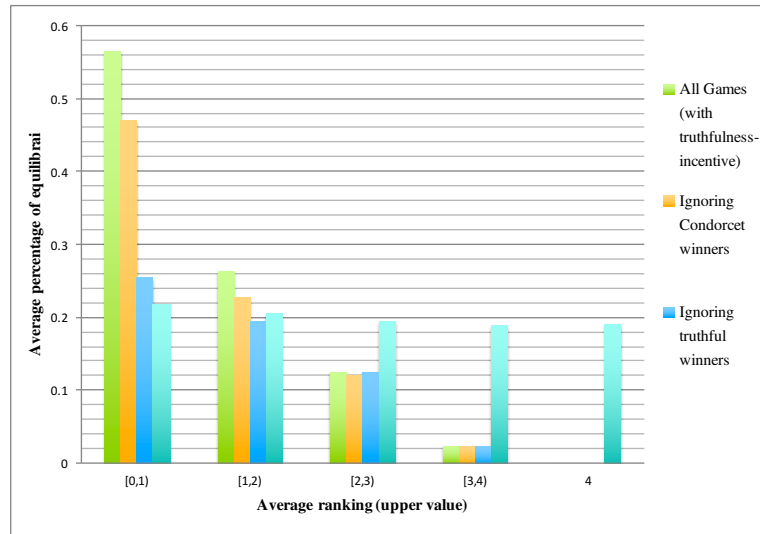


Figure 3.5: The average proportion of equilibria won by candidates with average rank of 0–1, 1–2, etc.

1). Fully 71.65%, on average, of the winners in every experiment had above (or equal) the median rank, and in more than half the experiments (52.3%) all equilibria winners had a larger score than the median. As a comparison, the numbers from experiments without the truthfulness incentive, are quite different: regardless of their average rank, candidates won, with minor fluctuations, about the same number of equilibria (57% of winners, were, on average, above or equal to the median rank).

3.4 Analytical Results

Following the publication of the above simulations, Obraztsova et al. [177] proved some analytical properties on truth-biased voters using the plurality voting rule with a fixed linear-ordered tie-breaking rule. For completeness, we mention these results before delving into our analysis for the veto voting rule.

Theorem 3.1 (Obraztsova et al. [177], Theorem 4). *Given an election with*

truth-biased voters using the plurality voting rule with a linear-ordered tie-breaking rule and a specific candidate, deciding if there is a voting profile which is a Nash equilibrium in which the candidate is victorious is NP-complete.

They also note a particular characteristic of the equilibria in these elections:

Observation 3.1 (Obraztsova et al. [177], Lemma 1). *Given an election with truth-biased voters using plurality with a linear-ordered tie-breaking rule, in any Nash equilibrium, all non-truthful voters are voting for the winner.*

The reasoning behind this observation is straightforward: if a voter is non-truthful and voting for a losing candidate, it would be better off reverting to its truthful vote, as its manipulation is not making the candidate it is voting for the winner, and reverting to the truthful vote would gain the voter at least an extra ϵ of utility.

3.4.1 Veto Voting Rule

When discussing Nash equilibria of truth-biased voters using the veto voting rule (with some tie-breaking rule), it is helpful to understand the structure of these equilibria. For that, it is handy to look at a particular runner-up candidate:

Definition 3.1. *A threshold candidate in a voting profile \mathbf{b} when using a scoring rule is a runner-up candidate that would become the winner if the current winner lost a single point.*

Notice that if a profile has runner-ups, there is a threshold candidate.

Before noting the complexity of finding a Nash equilibrium, we note a few characteristics of all Nash equilibria with truth-biased voters using the veto voting rule

Observation 3.2. *Given an election with truth-biased voters using veto and a linear-order tie-breaking rule, in any Nash equilibrium the score of the winner does not change from the score it received in the truthful vote. All non-truthful voters veto runner-ups.*

(Proof is in Appendix A)

Observation 3.3. *Given an election with truth-biased voters using veto and a linear-order tie-breaking rule, in any Nash equilibrium which is not the truthful profile, all voters except those vetoing the winners or runner-ups prefer the winner to the threshold candidate.*

(Proof is in Appendix A)

Theorem 3.2. *Given an election with truth-biased voters using the veto voting rule with linear-order tie-breaking and a specific candidate, deciding if there is a voting profile which is a Nash equilibrium in which the candidate is victorious is NP-complete.*

(Proof is in Appendix A)

3.4.2 A Constructive Algorithm

Despite the NP-completeness of the general problem, we are able to find a few conditions that make it tractable in certain cases, in which the tie-breaking rule is linear-ordered:

Condition 1 Let $t \in C$ be the candidate right below w in the tie-breaking order (i.e., the tie-breaking order is in the form $\dots \succ w \succ t \succ \dots$). Then the score in the truthful profile of t is at least as high as the truthful score of w .

Condition 2 Let $t \in C$ be the candidate right below w in the tie-breaking order (i.e., the tie-breaking order is in the form $\dots \succ w \succ t \succ \dots$). Then every voter that does not rank w last truthfully ranks w above t .

These conditions ensure that the threshold candidate would be t , and that the stability of the Nash equilibrium would not be reliant on voters trying to make the threshold candidate the winner. However, before showing the tractability of the problem with these two conditions, we show both are required.

Theorem 3.3. *Given an election with truth-biased voters using the veto voting rule and a specific candidate, assuming Condition 1 is true and Condition 2 is not, deciding if there is a voting profile which is a Nash equilibrium in which the candidate is victorious is NP-complete.*

(Proof is in Appendix A)

Showing Condition 2 without Condition 1 is NP-complete is shown in the construction in Theorem 3.2's proof.

We now turn to the proof of our constructive theorem:

Theorem 3.4. *Consider a candidate $w \in C$ and a truthful profile for which both Condition 1 and Condition 2 apply. Then finding if there is a voting profile for truth-biased voters under the veto voting rule which is a Nash equilibrium where w is the winner can be done in polynomial time (as well as finding the Nash equilibrium profile itself) of $\mathcal{O}(2n^2m)$.*

(Proof is in Appendix A)

3.4.3 k -Approval

k -approval has very similar structure to both veto (which is, in a sense, $m-1$ -approval) and plurality (1-approval). We provide 2 examples that show some changes in the characteristics of k -approval in comparison to these two voting rules.

Example 3.1. *Consider 2-approval with the lexicographical tie-breaking rule ($a \succ b \succ c \succ d \succ e$). Suppose the truthful profile is:*

- $a \succ b \succ c \succ d \succ e$.

- $e \succ d \succ a \succ c \succ b$.
- 2 voters with preference $d \succ b \succ a \succ c \succ e$.
- 2 voters with preference $a \succ d \succ b \succ c \succ e$.
- $e \succ c \succ a \succ b \succ d$.

The equilibrium profile changes the last but one voters (i.e., one out of the two identical voters with preference $a \succ d \succ b \succ c \succ e$), to $a \succ e \succ b \succ c \succ d$, and the last voter changes to $e \succ a \succ c \succ b \succ d$.

In this example, the score of the winning candidate (candidate a) in the equilibrium profile is higher than in the truthful profile. On the other hand, the score of a runner-up candidate (in this example, d , which is also the threshold candidate here) decreases in the equilibrium compared to the truthful profile score.

Example 3.2. Consider 2-approval, with the tie-breaking order $d \succ a \succ b \succ c$. The truthful preferences are as follows:

$$\begin{array}{cc} a \succ b \succ c \succ d & a \succ c \succ b \succ d \\ c \succ d \succ a \succ b & d \succ b \succ a \succ c \\ & a \succ d \succ b \succ c \end{array}$$

An equilibrium can be constructed by making only one change – the last voter changes to $a \succ b \succ d \succ c$. The score of a runner-up candidate (candidate b , which is also a threshold candidate) increases.

3.5 Summary

In this chapter we have investigated truth-bias. We started out by examining if our intuition – that voters often will revert to their truthful vote, all things being equal – has any promise. We did so by simulating many voting scenarios and looking at the resulting Nash equilibria in these settings using plurality. We saw that we get both a lower number of Nash equilibria, as well as

higher quality ones: many “silly” equilibria which occur with non truth-biased voters disappeared, while the higher ranked candidates were much more significantly represented as winners in the truth-biased equilibria. This bodes well both for using this as a requirement of synthetic voters in an environment constructed of rational agents, forcing better equilibria; and also might indicated closer relation to what people actually do, as results are more like we would expect when running elections – highly ranked candidates win, while lower ranked candidates rarely do.

However, truth-biased voters open up the possibility that there are no Nash equilibria at all for some candidate or all of them (a situation which cannot happen with regular voters³). This raises the question of the complexity of knowing this in advance. Hence, we continued with analytical work on the complexity of finding if an equilibrium exists in which a particular candidate wins in binary scoring rules (constructed only of values of 1 and 0). We also looked into various equilibria characteristics. These analytical results are summed in Table 3.1.

Conditions	Veto		Plurality	k -approval
	A	B		
Existance equilibrium with winner $w \in C$	NP-hard	P	NP-hard	NP-hard
Winner score may grow in equilibrium	No		Yes	Yes
Winner score may drop in equilibrium	No		No	No
Runner-up score may grow in equilibrium	Yes		No	Yes
Runner-up score may drop in equilibrium	Yes		No	Yes

Table 3.1: Summary of our truth-bias voter Nash equilibrium complexity and several characteristics.

A: At least one of Condition 1 or Condition 2 does not hold.

B: Both Condition 1 and Condition 2 hold.

³When all voters put the same candidate in first place, and each other candidate is put in last place by at least $\lfloor \frac{n}{m} \rfloor$ voters, no single voter can change the outcome in any scoring rule, Condorcet consistent voting rule or STV.

Chapter 4

Don't stop 'Til You Get Enough: Iterative Voting

4.1 Introduction

Understanding what are possible (and reasonable) election outcomes when voters are manipulating has been a persistent problem once it was realized (Gibbard-Satterthwaite theorem [109, 206]) that there is no point in seeking a reasonable, non-dictatorial, voting system in which voters cannot benefit from manipulation, as such a voting rule does not exist. While much research has been done on manipulations (see Introduction to Chapter 3), such as on the complexity of manipulation [54, 53, 236, 188, 247, 246], it has commonly analyzed cases where the manipulators are working to make some specific candidate the winner, rather than when all voters are strategic.

While the Nash equilibrium is a natural tool to analyze what happens when all participants manipulate, it is not a usable tool for elections, as not only are there an extraordinary amount of equilibria even in small elections (see Chapter 3), but many equilibria describe results which would never occur in real-life setting (e.g., even if all voters rank the same candidate last, using plurality it is a Nash equilibrium for all voters to vote for it).

One approach to try and describe reasonable election results is to make

certain assumptions on the nature of the voters. Such an approach, described in the previous chapter (Chapter 3) is truth-bias, which assumes voters gain a certain small amount utility when they vote their truthful preference. Another such approach is lazy-bias, presented by Desmedt and Elkind [80], which assumes that all things being equal, voters would prefer to abstain and not participate at all in the elections.

A different approach, taken here (following Meir et al. [160]), does not try to make certain assumptions on the nature of the voters, but rather on the dynamics in which manipulations are conceived and carried out. Consider, for example, a group of friends, trying to decide on a movie to watch. Once everyone's preferences are tallied and participants know the outcome for each option, a voter may realize that by altering their vote, the outcome might be more to their liking. Once this voter makes this change, a different voter realizes that under this new situation, it can change its vote as well and make the winner a different winner. As voters slowly, iteratively, adjust their votes (there is no fixed order on the voters), the question of convergence arises, i.e., does this iterative process ever end in a stable state (which is, by definition, a Nash equilibrium). Assuming that it does converge (Meir et al. [160] proved it for plurality, as well as showing it does converge if we allow for parallel updates), we can define a subset of the Nash equilibria – those that are reachable, using the iterative process, from some profile serving as a starting point (a natural choice for it, which we pursue here, is taking the truthful preferences as the starting point).

Note that whether there is an explicit iterative process or not, the iterative process may portray various elections, as voters use the information they have (e.g., polls), to understand whether they should change their vote or not.

Exploring the iterative model framework, we start by picking apart the various conditions set in Meir et al. [160]'s result, and try to examine them to the limit, to understand how applicable is the model. In this process we look at tie-breaking rules and their necessity as well as at voting rules and convergence. We find appropriate tie-breaking rules are necessary for

convergence, and more surprisingly, that it is guaranteed only for plurality (shown in [160]) and veto, while any other scoring rules can be shown to reach cycles for some profiles.

We then look into combining assumptions on the nature of voters with the assumption on their dynamic, and examine the characteristics (and existence) of Nash equilibria when the iterative process involves truth-biased voters, and well as the properties (and existence) of the Nash equilibrium when the voters are lazy-biased. In both cases we are able to present an algorithm to find all the equilibria, and also manage to portray the structure of the equilibria reached when the iterative process starts from the voters' truthful preferences.

4.1.1 Related Work

The common game-theoretic approach to iterative processes is subgame-perfect Nash equilibrium, which entails a backwards induction from potential results. This has mostly been used to analyze open serial elections (e.g., roll-calls) and sequential elections with 2 candidates [217, 155, 78]. However, these are not common, as people are notoriously bad at doing backward induction [122], and in our iterative setting, with its limited information (on other voters' preferences) and non-deterministic voting ordering it makes it even less applicable and realistic.

While we use the framework established by Meir et al. [160], the notions of an iterative approach to voting, as well as of seeking election equilibria, exist in previous research. An iterative process for reaching decisions was offered for agents in Ephrati and Rosenschein [94], but it uses a mechanism to transfer money-like value among agents, and hence is irrelevant to our voting procedures. Several researchers have considered reaching an equilibrium with an iterative (or dynamic) process, in particular when deciding on an allocation of public goods. A summary of much of that work can be found in Laffont [137], which details various approaches, including different equilibria choices (Nash, local dominant, local maximin) and methods. However,

in order to reach an equilibrium, they limit the possible preference choices to single-peaked preferences. More recently, Reijngoud and Endriss [191] explored an iterated poll setting, but as theirs had only one round where voters change their vote according to a fixed order, the type of questions discussed here does not arise.

Another branch of research deals with a process of having a player propose a change in the current state, and hold a vote on its acceptance. Such a model was used by Shepsle [214], who chose to force an equilibrium by using a combination of preference limitation and organizational limitations. De Trenqualye [77] chose to achieve an equilibrium by using a specific voting rule and Euclidean preferences. More recently, Airiau and Endriss [2] examined – theoretically and experimentally – the possibility of an equilibrium in such games, using plurality-type voting rules (the threshold can be different than 50% for a change to be accepted).

Attempting to investigate the role of knowing other players' knowledge, Chopra et al. [71] examined iterative voting with plurality, and showed the effects of limiting a player's knowledge of the other players' preferences. Another interesting model, proposed in Myerson and Weber [170], found a Nash equilibrium for scoring rules, assuming that voters have some knowledge of which candidates have a better chance of winning (based, for example, on pre-election polls), but this does not mean that every election results in an equilibrium. Further research of analyzing the Nash equilibria of voting games (without iteration) is discussed in Chapter 3, in addition to discussion of previous research on truth-biased and lazy-biased voters.

The iterative model presented in Meir et al. [160] has attracted research other than that detailed in this chapter. There has been effort to understand the need for the “best response” strategy requirement, which was analyzed in Obraztsova et al. [176], looking into strategies that might ensure convergence in an iterative voting process (and not just best-response), a direction that was also explored in Grandi et al. [112]. Brânzei et al. [66] addressed some notions of price of anarchy in the context of iterative plurality, showing

winners will have a very high truthful score. We note that following the publication of the following section on iterative veto, another proof for the same result was published by Reyhani and Wilson [196].

4.2 Tie-Breaking Rules

We start our analysis of the iterative voting model, by examining if the requirement for linear-ordered tie-breaking rules is necessary. Indeed, we show that using arbitrary voting rules is not just harmful to the convergence of plurality, but that it prevents convergence for any scoring rule.

Theorem 4.1. *An iterative scoring rule election with a deterministic tie-breaking rule, even for voters using best-response strategies and starting from the truthful state, will not converge for some preferences.*

(Proof is in Appendix A)

4.3 Voting Rules

We now seek to understand in which voting rules does the iterative model guarantee the process will converge to a stable state (i.e., a Nash equilibrium) and in which voting rules there are cases which result in a cycle, never reaching an equilibrium.

4.3.1 Veto

Definition 4.1. *A best response in the case of the Veto voting rule implies that the current (undesired) winner is vetoed.*

Theorem 4.2. *Iterative Veto elections with deterministic linear-order tie-breaking and voters which use a best-response strategy, converge even when not starting from a truthful state.*

(Proof is in Appendix A)

4.3.2 Other Scoring Rules

Apart from plurality (shown in [160]) and veto (shown above), no scoring rule family (i.e., one that can be applied for any number of candidates) converges in the iterative process.

Theorem 4.3. *Under the iterative procedure, using a best response strategy and when voters are myopic, no scoring rule apart from plurality and veto converges.*

(Proof is in Appendix A)

4.4 Voter Types

We now combine our work on iterative voting with changing our assumption on the nature of voters, and assuming some properties on these. We address the plurality voting rule, as it is the most common and widely used. We focus on analyzing equilibria that are reachable from a truthful starting point.

4.4.1 Regular Voters

While we wish to be able to point out what are the reachable Nash equilibria from the truthful starting point, as an approximation of the way people vote, the complexity of finding what these state are poses a difficulty.

Theorem 4.4. *Given a truthful profile \mathbf{a} and a profile \mathbf{b} distinct from \mathbf{a} , it is NP-complete to decide if \mathbf{b} is reachable by iterative plurality with linear-ordered tie-breaking using best-response updates, starting from \mathbf{a} .*

(Proof is in Appendix A)

4.4.2 Truth-Biased Voters

We start with a few characteristics of the Nash equilibria of truth-biased voters using iterative plurality. Recall Observation 3.1, that noted that all non-truthful voters in an equilibrium vote for the equilibrium winner.

First, we note convergence is no longer guaranteed:

Example 4.1. *Suppose tie-breaking rule is $a \succ b \succ c \succ d$, and voters' true preferences are:*

Voter 1: $a \succ b \succ c \succ d$

Voter 2: $b \succ a \succ c \succ d$

Voter 3: $c \succ d \succ a \succ b$

Voter 4: $d \succ c \succ a \succ b$

The winner is a , but voter 3 can improve their outcome by changing to $d \succ c \succ a \succ b$, and making d the winner. Then voter 1 changes its vote as well to $b \succ a \succ c \succ d$, making b the winner. However, at this point, voter 3 has no manipulation open to it, so due to truth-bias it returns to its truthful vote. At this point voter 1 has a beneficial change, and by reverting to its truthful vote, it creates a cycle, making a the winner again.

Now we continue with finding the characteristics of the Nash equilibria, when they are reached:

Lemma 4.1. *In a non-truthful Nash equilibrium with truth biased voters under iterative plurality with linear-ordered tie-breaking, starting the process with their truthful preferences, the winner will always be a runner-up candidate in the original state, with only a single voter being untruthful.*

(Proof is in Appendix A)

Note that this means that if there is an iterative path towards a state, there is a path of length 1 (the untruthful voter changes its preferences). We shall now show there is a simple, polynomial algorithm to find all the equilibria for truth-biased voters when starting from the truthful starting position.

Theorem 4.5. *Algorithm 1 finds all Nash equilibria reachable from the truthful starting point (\mathbf{a}) in an iterative plurality model with truth-biased voters and linear-ordered tie-breaking.*

Algorithm 1 Finding all truth-biased Nash equilibria

The input is the truthful profile (\mathbf{a})

```

1:  $\forall c \in \text{runnerup}(\mathbf{a}), eq[c] \leftarrow 0$     $\triangleright$  An array holding number of equilibria
   for every potential winning candidate
2:  $V, V'$  sets of all voters
3: for all  $c \in \text{runnerup}(\mathbf{a})$  do
4:   for all  $v \in V$  do
5:     if  $v \in V', \text{top}(v) \neq \text{winner}(\mathbf{a})$  and  $\text{top}(v) \neq c$  then
6:       if  $c$  is the highest ranked candidate in  $(W(\mathbf{a}) \cup H(\mathbf{a})) \setminus \{\text{top}(v)\}$ 
       then
7:          $eq[c] \leftarrow eq[c] + 1$     $\triangleright$  This voter will deviate to make  $c$  win
8:          $V' \leftarrow V' \setminus v$       $\triangleright$  If this voter deviates for  $c$ , it will not
       deviate for any other
9:       end if
10:      end if
11:      if there is  $c' \in \text{runnerup}(a)$  such that if  $c'$  has higher score in  $\mathbf{a}$ 
       than  $c$  or same as  $c$  but higher in the tie-breaking rule;  $c' \neq \text{top}(v)$ ; and
        $c' \succ_v c$  then
12:         $eq[c] \leftarrow 0$             $\triangleright$  This is a blocking voter
13:        break  $\triangleright$  No point in examining this candidate further; return
       to line 3 (next candidate)
14:      end if
15:    end for
16:  end for
       return  $eq$ 

```

(Proof is in Appendix A)

Algorithm 1 will not work for non-truthful starting position, as it works by exploring reachable positions from truthful starting states, and there are Nash equilibria for truth-biased voters that cannot be reached from a truthful

starting position.¹

4.4.3 Lazy-Biased Voters

Allowing voters to abstain and then rejoin the elections in iterative voting settings, is equivalent to letting voters choose a strategy that while may cause the outcome to change, it will not necessarily lead to a best response change, and therefore we know does not ensure convergence ([160], Theorem 3). Moreover, for lazy-biased voters, we know it cannot change the winner to a candidate more preferable to them (as otherwise they would have a voting strategy to achieve the same end). Therefore, we suggest a model in which once voters abstain they cannot rejoin the game. Hence, once a player abstains, we are left with an iterative plurality election with a potentially non-truthful starting point and one less voter. We are guaranteed that these will converge, and therefore, after a finite number of abstentions, we will reach a stable state.

As in truth-bias, we notice a structure of stable states (and Nash equilibria):

Lemma 4.2. *A stable state with lazy-biased voters under iterative plurality with linear-ordered tie-breaking, starting the process with their truthful preferences will only have a single participating voter.*

(Proof is in Appendix A)²

However, as the election Nash game is defined as a “one-off” (i.e., not repeated or turn-based game), on a fixed set of strategies, and so do not

¹For example, 4 voters with preferences $a \succ b \succ c$, and 2 voters preferring $c \succ b \succ a$. The only equilibrium reachable from the starting position is the starting position itself. However, the state in which 3 of the first 4 voters vote $b \succ a \succ c$ and the rest are truthful is a Nash equilibrium for truth-biased voters.

²This appears to be quite synthetic equilibria, stemming from the constraints of the the lazy-bias, as well as as from those of the iterative dynamic, being fundamentally myopic. Ways to address such these limitations, while conserving the lazy-bias, will be discussed in the Chapter 5.

enable us to “remove” the strategy of rejoining, notice that these stable states are not necessarily Nash equilibria:

Example 4.2. *Suppose tie-breaking rule is $a \succ b \succ c$ and voters’ true preferences are:*

Voter 1: $a \succ b \succ c$

Voter 2: $c \succ b \succ a$

Voter 3: $c \succ b \succ a$

The only Nash equilibrium here is for voter 1 to participate and the other abstain. However, in the iterative process, voter 1 cannot do anything, so it abstains, and then voter 2 or 3 abstain, leaving a single voter voting for c . This is not an equilibrium, as if voter 1 could rejoin it would change the outcome to a .

We now show a polynomial algorithm to find all the Nash equilibria for lazy-biased voters (in contrast to the regular, non-biased ones, for which there is no such algorithm):

Theorem 4.6. *Algorithm 2 finds all Nash equilibria reachable from the truthful starting point (\mathbf{a}) in an iterative plurality model with lazy-biased voters and linear-ordered tie-breaking.*

(Proof is in Appendix A)

4.5 Summary

In this chapter we examined iterative voting. We have attempted to stretch its limits and see if its various conditions are truly necessary or not. We have seen that the tie-breaking mechanism is actually critical for iterative voting’s convergence, and we have extended the convergence of iterative plurality to iterative veto, while detailing the limitations of the model, which does not converge for any scoring rules but plurality and veto.

When seeking to utilize the model in order to examine the equilibria reachable from a starting point of the iterative process, we discovered that finding these equilibria will be a difficult task, as it is NP-complete. However, when merging our assumption of voters' nature – truth or lazy bias – with the assumption made on voters' dynamics and decision process – the iterative model – we were able to find polynomial algorithms to expose the Nash equilibria to us. Moreover, we were able to characterize some of the features of the resulting Nash equilibria. These equilibria, with their somewhat synthetic structure, entice us to keep looking for a better and more robust model – in the next chapter. . .

Algorithm 2 Checking reachability of Nash equilibrium under lazy voting

Input: The initial profile \mathbf{a} , and the voter v to be checked, with preferences in the form $z \succ \dots$

- 1: **if** z cannot be a winning candidate in a Nash equilibrium – if there are voters which rank above z a candidate that is higher than z on the tie-breaking order. **then**
 - return** No
- 2: **end if**
- 3: **if** \mathbf{a} is a Nash equilibrium in the basic model (without abstentions) **then**
- 4: **if** z is a winner of \mathbf{a} **then**
 - return** Yes
- 5: **end if**
- 6: **if** $\exists \tilde{c} \neq z, v' \in V$ s.t. $\tilde{c} \succ_{v'} z \wedge z \succ_{v'} \text{winner}(\mathbf{a})$ **then**
 - return** Yes
- 7: **end if** \triangleright At this point, every voter whose top choice is not z , prefers $\text{winner}(\mathbf{a})$ to z .
- 8: $\mathcal{C}' \leftarrow \{c \in C \setminus \{z, \text{winner}(\mathbf{a})\} \text{ s. t. in } \mathbf{a}, c\text{'s score is } \geq 2 \text{ or } c\text{'s score is } 1 \text{ and } c \succ \text{winner}(\mathbf{a}) \text{ in tie-breaking.}\}$ \triangleright Potential Nash equilibrium winners.
- 9: **if** there is a voter $z \succ \dots \succ c \succ \dots \succ \text{winner}(\mathbf{a}) \succ \dots$ for some $c \in \mathcal{C}'$ **then**
 - return** Yes
- 10: **end if**
- 11: **if** there is $\tilde{c} \neq z$, and $c \in \mathcal{C}'$ such that there exists a vote in the form $\tilde{c} \succ \dots \succ c \succ \dots \succ \text{winner}(\mathbf{a}) \succ \dots$ **then**
 - return** Yes
- 12: **end if**
 - return** No
- 13: **end if**

\triangleright From now on we can assume \mathbf{a} is not a Nash equilibrium

Algorithm 2 Algorithm continued...

```

14: if  $z$  is not the winner nor a runner-up in  $\mathbf{a}$  then
    return Yes
15: end if
16: if  $z$  is a runner-up in  $\mathbf{a}$  then
17:   if For a runner up  $b \neq z$  there is a voter  $\neq v$  with preference  $\dots b \succ$ 
      $\dots \succ \text{winner}(\mathbf{a})$  then
     return Yes
18:   end if
19:   if  $|V| \geq 4$  then
     Goto line 6
20:   else return No
21:   end if
22: end if
                                      $\triangleright$  We can now assume  $z = \text{winner}(\mathbf{a})$ 
23:  $\mathbf{b} \leftarrow$  Voter profile after running iterative plurality (without abstentions)
     while preventing  $v$  from deviating  $\triangleright$  from Meir et al. [160] this is
     polynomial
24: if  $z = \text{winner}(b)$  then
     return Yes
25: end if
26: if  $z$  in not a runner-up in  $\mathbf{b}$  then
     return Yes
27: else Goto line 19 using  $\mathbf{b}$  instead of  $\mathbf{a}$ .  $\triangleright z$  a runner-up in  $\mathbf{b}$ 
28: end if

```

Chapter 5

Think Local, Act Global: Local Dominance Voting

5.1 Introduction

As the previous chapters have noted, following the realization that all “reasonable” election systems are prone to manipulation (Gibbard-Satterthwaite theorem [109, 206]), much work was devoted to manipulation techniques and complexities [54, 53, 236, 110, 186]. Less work was devoted to analyzing the outcome of elections when assuming all voters are manipulators (i.e., the Nash equilibria), due to the number of equilibria being extremely large, and the preponderance of equilibria which are not useful (e.g., it is an equilibrium in plurality when all voters have the same least liked candidate and vote for it). Thus, there has been a significant limit on the ability to analyze elections and to try and understand which candidates might emerge to be victorious based on truthful preferences (or calculate potential truthful preferences from election results).

By assuming various models on the voting environment we are able to eliminate many of these problems, with models dealing with the nature of voters (Chapter 3) or with the voting procedure, whether explicit or implicit in the voters’ thinking process (Chapter 4), and joining the two. However,

these models assume a fairly precise knowledge by voters on the election situation, such as knowing how many voters are voting for each candidate. Naturally, such knowledge is quite rare in real-world situations.

Here, we try to combine the previous models, alongside a different information model. Instead of assuming our voters have an accurate view of the election results, we assume a far more realistic option – they have a vague anchor of information, similar to an election poll, but they do not know the exact expected outcome. Furthermore, we assume they do not have an exact probabilistic knowledge on the exactness of the poll (and how likely each different variation from it is), but rather a binary understanding of the expected results – whether a situation is possible or not.

Hence, a voter faced with a poll decides how much variation from the poll does it consider realistically conceivable (the voter’s “radius”), and thus has a local “area” around the poll which it regards as the set of possible outcomes. Faced with these potential outcomes it chooses its vote (we assumed it is voting for the dominant candidate), and this changes (slightly) the poll. As all voters undergo this process, we show that in many cases this dynamic does converge, and moreover, we show that this model results in highly realistic voting patterns. Moreover, we show that there is a level of maximal variation from the poll in which most manipulations occurs which combines with the radius in which the most realistic – and, arguably, best – results happen.

Beyond this analysis, we have released our simulation code to the public, so that other researchers may use it. It is constructed in a modular way, allowing other models, voting systems, and assumptions to be made, without necessitating a massive code rewrite. We have put the links and instruction to the code on `Preplib.org` [152] at <http://www.preplib.org/tools/ivs.php>, while the code itself is at <https://github.com/omer1/IterativeVotingSimulator>.

5.1.1 Related Work

Much of the research relevant to this work has been mentioned in Chapter 4 as we use the iterative model here as well. However, we have not focused there on the information model, which is a key difference in this model.

While lack of information has been a key explanatory tool in understanding why people bother to vote when the likelihood of their vote having an impact is low [180], the models using this uncertainty have not always been applicable to the way people vote, and though they do not assume full knowledge from the voters, assume quite detailed knowledge from them. For example, Myerson and Weber [170], assumes voters use knowledge of the probability of there being a tie between any 2 candidates, and their calculation of it. Messier and Polborn [165] considered a “trembling hand” model, in which there is a probability some votes are miscounted, which, again, necessitates a significant probabilistic analysis from the voter. Furthermore, these models require the voter have a utility function with cardinal values, nullifying one of the properties that make social choice theory more applicable to the real-world, as learning the utility functions values (rather than preference orders) is quite complex, if possible (a similar requirement appears also in [215, 182, 9]). A slightly different approach to voters has been to assume voters do not wish to maximize their utility, but rather to minimize their regret. This approach has been presented in Ferejohn and Florina [99], and further discussed in Ledyard [142] and Merrill [164], and it too requires explicit utility functions. Furthermore, empirical work [61] has cast a shadow on the ability of regret minimization to explain voter behavior.

Using dominance to decide on voting manipulation has been suggested as early as Farquharson [96], though its analysis by Dhillon and Lockwood [81] showed iterative dominance to be quite limited in the settings it works in. More recently, Conitzer et al. [74] suggested a manipulation problem based on limited manipulator knowledge and dominant manipulations.

As a partial response to the empirical observation that people do not employ probabilities when deciding their strategies [226], the concept of *bounded*

rationality arose, as explaining that people act rationally only up to a certain limit (e.g., several steps ahead), and do not do so beyond that [203]. Our approach, while not, strictly speaking, a limited rationality approach, uses a similar concept in which voters are assumed to be able to foresee only a limited number of possibilities, and they choose their strategies accordingly. Recently, van Ditmarsch et al. [228] suggested a general logic of knowledge framework for discussing manipulations under limited knowledge circumstance. While still not studying equilibria, a close model to what we examine is Saari [204], which suggests – in the context of comparing 3 candidate scoring rules – considering a state feasible if it involves moving at most m voters from the truthful state, and compares scoring rules solely on those states. An approach somewhat similar to our own, of a poll defining a set of possible states, has been discussed in Reyhani et al. [197], but their dynamics assumptions are more complicated than ours, and they focus on results with 3 candidates.

5.2 The Local Dominance Model

The local dominance model contains, first and foremost, a model of locality and a strategy of dominance, both of which determine what a voter would do faced with a given state of an election. Once we have defined those, our model begins with a state (we shall assume, generally, that we start with the truthful state), and advance iteratively, using the local dominance strategy, with the hope of reaching a Nash equilibrium. We define it here only for plurality, but the extension to other scoring rules is straightforward.

5.2.1 Locality

Given voting profiles, we can define a metric on them that will allow us to measure a distance of one profile from one another. This is easily done as the profiles tell us how the score of each candidate, and we use that for our metrics. Moreover, we can consider the difference in the number of votes

for each candidate as an absolute (e.g., candidate a lost 3 votes, candidate b gained 7, etc.), or its ratio of growth (e.g, candidate c has doubled its vote, candidate d lost a third of its vote, etc.). Hence, the definition for ℓ_i are straightforward:

Definition 5.1. *Given two voting profiles with n voters and m candidates, \mathbf{x} and \mathbf{y} , for any $1 \leq i < \infty$ the distance of these profiles under the ℓ_i additive metric are:*

$$d_{\ell_i}(\mathbf{x}, \mathbf{y}) = \left(\sum_{c \in C} |\text{score}_{\mathbf{x}}(c) - \text{score}_{\mathbf{y}}(c)|^i \right)^{\frac{1}{i}}$$

For ℓ_∞ , the distance is:

$$d_{\ell_\infty}(\mathbf{x}, \mathbf{y}) = \max_{c \in C} |\text{score}_{\mathbf{x}}(c) - \text{score}_{\mathbf{y}}(c)|$$

Definition 5.2. *Given two voting profiles with n voters and m candidates, \mathbf{x} and \mathbf{y} , the distance of these profiles under the multiplicative metric is:*

$$d_{mult}(\mathbf{x}, \mathbf{y}) = \max_{c \in C} \left(\max \left(\frac{\text{score}_{\mathbf{x}}(c)}{\text{score}_{\mathbf{y}}(c)}, \frac{\text{score}_{\mathbf{y}}(c)}{\text{score}_{\mathbf{x}}(c)} \right) \right) - 1$$

Using one of these metrics, we can define voters which believe that given a data point (like a poll), the true situation is a “neighborhood” around that poll:

Definition 5.3. *For a voter with radius r using a metric d , when given a profile \mathbf{b} of $n-1$ voters, we define $S(\mathbf{b}, r)$, the set $\{\mathbf{p} \in (\pi(C))^{n-1} \mid d(\mathbf{b}, \mathbf{p}) \leq r\}$, which is the set the voter bases its strategic manipulation.*

Notice that when the radius is n , the voter bases its manipulation of all possible profiles. When the radius is 0, the voter believes the current situation is the only possible situation, and therefore, manipulates only according to it (similar to the iterative voting dynamic in Chapter 4).

5.2.2 Dominance

While we can choose various strategies for our voters, we chose to use dominance, as not only is it “obvious”, in the sense that its manipulations seem worthwhile, as it will never leave the voter worse off than it was, and while in an election it is commonly too broad to be useful (as few votes are dominated), when using voters with a limited locality, the dominance strategy becomes more viable.

Definition 5.4. *Given a voting rule f , a voter $i \in V$ with truthful preference order a_i , currently voting b_i , and a set of profiles S of $n - 1$ voters (voters $N \setminus \{i\}$), a vote c_i S -dominates b_i on S if:*

- For every profile $\mathbf{b} \in S$ $f(c_i, \mathbf{b}) \succeq_{a_i} f(b_i, \mathbf{b})$
- There is one profile $\mathbf{p} \in S$ for which $f(c_i, \mathbf{b}) \succ_{a_i} f(b_i, \mathbf{b})$

Hence, using plurality, a voter with a truthful preference of a_i , a metric d and a radius r , which is currently voting b_i , we define its local dominance strategy as:

1. Let S be the set $\{\mathbf{p} \in \pi(C) \mid d(\mathbf{b}, \mathbf{p}) \leq r\}$.
2. Let $D \subset C$ be the set of votes that S -dominate b_i .
3. If $|D| > 1$, let e be the highest ranked candidate according to a_i (if $|D| = 1$, let e be its only item).
4. If $|D| > 0$, vote for e . Otherwise, no change.

5.2.3 Truth and Lazy Bias

We combine truth-biased and lazy-biased voters into this model by letting them, as in their original definitions, fall back on their true preference (or abstaining from the vote), if they consider their situation helpless. In order to do that, we give them an additional radius k , to indicate the radius in which if they cannot influence anything, they revert to their truth/lazy habits. This

is intended to approximate real-world behavior, in which the lack of active manipulation does not necessarily go hand-in-hand with giving up. Hence:

Definition 5.5. *A truth-biased voter with radii (r, k) ($r < k$) under a metric d and a voting rule f , with a truthful preference of a_i and which is currently voting b_i , when given a profile \mathbf{b} , follows the following strategy:*

1. Let S be the set $\{\mathbf{p} \in \pi(C) \mid d(\mathbf{b}, \mathbf{p}) \leq r\}$.
2. Let $D \subset C$ be the set of votes that S -dominate b_i .
3. If $|D| > 1$, let e be the highest ranked candidate according to a_i (if $|D| = 1$, let e be its only item).
4. If $|D| > 0$, vote for e .
5. If $|D| = 0$, if there is no \mathbf{p} such that $d(\mathbf{b}, \mathbf{p}) \leq k$ and $f(b_i, \mathbf{p}) \succ f(a_i, \mathbf{p})$, the voter changes to vote a_i . Otherwise, no change.

Note the difference from a non truth-biased voter is that if there are no strategic moves for the voters, it checks if there is any reason for it to remain voting what it is in a larger radius (i.e., there is at least one state in which reverting to a truthful will harm the voter), and if there is not, it returns to its truthful preference.

The lazy-biased voter is quite similar, except that it checks if abstaining is better than staying at its current vote:

Definition 5.6. *A lazy-biased voter with radii (r, k) ($r < k$) under a metric d and a voting rule f , with a truthful preference of a_i and which is currently voting b_i , when given a profile \mathbf{b} , follows the following strategy:*

1. Let S be the set $\{\mathbf{p} \in \pi(C) \mid d(\mathbf{b}, \mathbf{p}) \leq r\}$.
2. Let $D \subset C$ be the set of votes that S -dominate b_i .
3. If $|D| > 1$, let e be the highest ranked candidate according to a_i (if $|D| = 1$, let e be its only item).

4. If $|D| > 0$, vote for e .
5. If $|D| = 0$, if there is no \mathbf{p} such that $d(\mathbf{b}, \mathbf{p}) \leq k$ and $f(b_i, \mathbf{p}) \succ f(\mathbf{p})$, the voter abstains. Otherwise, no change.

As we discuss local dominance voters, our truth-biased and lazy-biased voters will pursue it when there is a local dominant manipulation for the set $\{\mathbf{p} \in \pi(C) \mid d(\mathbf{b}, \mathbf{p}) \leq r\}$.

5.2.4 Generality

Our local dominance model is not limited to discussing elections and voters. For any normal form game, once some metric is defined on the different states, a local-dominance model can be used.

Furthermore, note that the locality and the dominance parts can be separated. While we used a dominance model in this work, one may choose different strategies to pursue using the local set. For example, an optimistic strategy such as pursuing the best case for the voter, a minimal regret strategy, etc.

5.3 Iterative Convergence

Armed with the local dominance model, we seek to find the equilibrium points when voters pursue such a strategy (note that when the voters' radius is 0, this is exactly the Nash equilibrium). However, as with the Nash equilibrium, we wish to discard the various useless equilibria, and focus on equilibria that may actually occur as a result of an election in real-life scenarios. Thus, we consider the iterative model detailed in Chapter 4, and as with it, we consider reachable results when beginning from the truthful voting state. However, while we know that when all voters' radius is 0 the iterative process converges, we do not know this for any r when using local dominance. As when $r = 0$ this is equivalent to the iterative model, we know that allowing for any tie-

breaking rule will result in non-convergence, and therefore we limit ourselves to linear-ordered tie-breaking rules.¹

To analyze the iterative process, we introduce several notations that will be useful. For each stage of the iterative process $i \in \mathbb{N}$, and a profile \mathbf{b} , let $H_i(\mathbf{b}) \subseteq C$ be the set of candidates that need exactly i more votes to become the winner. Thus, for a scoring rule f , $H_0(\mathbf{b}) = \{f(\mathbf{b})\}$, and for plurality, H_1 are the candidates which either have the same score as the winner and lose by tie-breaking, or win in the tie-breaking but have one vote less than the winner (which we mark as $winner(b)$). Let $\overline{H}_i(\mathbf{s}) = \bigcup_{i' \leq i} H_{i'}(\mathbf{s})$, i.e., all candidates that can become a winner with up to an additional i votes.

Theorem 5.1. *When using plurality with linear-ordered tie-breaking, and when all voters use local dominance strategy and have the same radius r , if the initial state was one when all voters were truthful, the iterative process converges to a stable state.*

(Proof is in Appendix A)

The proof not only shows that an equilibrium exists, it also describes exactly the way in which such equilibria are reached from the truthful state. There is always a set of “leaders” (\overline{H}_{r+1} in the case of the ℓ_1 norm). Strategic voters vote for their favorite candidate in this set, if their current candidate is not a possible winner. At some point candidates may “drop out” of the race as their gap from the winner increases, and the set \overline{H}_{r+1} shrinks. Finally, in the reached equilibrium all strategic voters vote for their best possible winners in \overline{H}_{r+1} . Furthermore, if the truthful state is not an equilibrium, voters will not vote for one of the \overline{H}_{r+1} only if the gap between the winner and the runner-up is exactly $r + 1$, and these voters prefer the current winner. Furthermore, note that the gap between the winner and the runner up can never grow larger than $r + 1$:

¹In Meir et al. [159] we show some conditions and constraints under which parallel updates may be permissible. But we focus here on the more general results of the one-by-one iterative mechanism.

Observation 5.1. *When using plurality with linear-ordered tie-breaking, and when all voters use local dominance strategy and have the same radius r , if the initial state was one when all voters were truthful, either this situation is stable or in every state \mathbf{b}^t we have $|\overline{H}_{r+1}(\mathbf{b}^t)| > 1$. Also, in the stable state either $|\overline{H}_r| = 1$ or all voters vote for possible winners. Any voter voting for $c \notin \overline{H}_{r+1}$ prefers the winner in the stable state over any other candidate in \overline{H}_{r+1} .*

(Proof is in Appendix A)

5.4 Truth-Bias and Lazy-Bias

For truth-biased and lazy-biased voters, we show, as with the regular case, that convergence is still guaranteed:

Theorem 5.2. *When using plurality with linear-ordered tie-breaking, and when all voters use local dominance strategy and are truth-biased with the same radii (r, k) , if the initial state was one when all voters were truthful, the iterative process converges to a stable state.*

(Proof is in Appendix A)

Theorem 5.3. *When using plurality with linear-ordered tie-breaking, and when all voters use local dominance strategy and are lazy-biased with the same radii (r, k) , if the initial state was one when all voters were truthful, the iterative process converges to a stable state.*

(Proof is in Appendix A)

5.5 Simulation Results

The diversity of the various voter distributions examined, the multiple strategies considered, the multiple variables collected and the sheer modularity of the framework, make these simulations particularly worthwhile, and set

them apart from previous attempts at better understanding elections via simulations. However, due to lack of space, the explanation of these is in Appendix D. We present the main results below.

The most striking result of our simulations, in a sense, was finding how the structure of almost all results resembled one another. When looking at the effects of changing the radius has on various properties (e.g., number of steps to convergence, share of stable states where the winner is the Condorcet winner, etc.), most properties have a similar form: the property increases as the radius grows, until it reaches “peak r ”, from which it slowly descends. This “peak r ” value is generally similar between elections with the same distribution and the same number of voters, while the amplitude (e.g., the actual number of steps to converge in peak r) has to do with the number of candidates, as can be seen in Figure 5.5. The reason for this shape is due to the small number of manipulations that occur with a very large radius, as each candidate has a chance of winning, and therefore voters rarely deviate from their truthful preference. Similarly, with a very small radius opportunities for manipulations are few, and therefore only few steps occur.

This type of structure also helps us to find related variables in some distributions. For example, in the 2-urn model, the share of the two top candidates (Figure 5.5) displays similar peak r to that of the number of steps it takes to reach a stable state (Figure 5.1).

We cannot, of course, include all of our results in this space, but we point out certain observations, and some of the simulation data backing it up.

Winner Quality

While Plackett-Luce distributions have a ground truth, so results can be compared to it, other distributions do not. Therefore, we try various different measures to see how changing the radius affects various different measures approximating a “good” winner. We examine the ratio of Condorcet winners, and try to approximate social welfare using the Borda score (Figure 5.3), and seeing how common it is for the winner with the highest social welfare (i.e.,




Figure 5.1: A 2-Urn distribution, showing the average number of steps it takes to reach a stable state depending on the value of r . Note the similar peak r values to those in Figure 5.5.

Borda winner) to be the outcome of a local dominance dynamic (Figure 5.4). The pattern of a growing radius (up to peak r) resulting in better winners appears over and over, in various distributions and properties. That is, allowing greater manipulation by voters results in better, high quality, winners.

Note in particular the single-peaked results (Figure 5.2), as in the single peaked distribution there is always a Condorcet winner (the median candidate), which is more commonly the outcome of a local dominance dynamic as r grows (up to peak r).

Duverger Law

“Duverger law”, nicknamed after Maurice Duverger [88], is a political-science observation on how plurality elections tend to strengthen the two top candidates over all others. Indeed, in our simulations, the strengthening of the two top candidates as r grows is observable in many distributions (not only in the urn models, which are naturally more susceptible to this property): in all distribution the share of the two top candidates in peak r is above 75%




Figure 5.2: The average of the share of stable situations in which the winner was the Condorcet winners in the single-peaked distribution (where there are always Condorcet winners).




Figure 5.3: The average rank of the winners of stable situations in the riffle distribution. Note that this uses a reverse Borda score, akin to social welfare, so a lower number indicates a better result.




Figure 5.4: The average of the share of stable situations in which the winner was the truthful Borda winner in the uniform distribution.

(Figure 5.5).

5.6 Summary

We have presented and discussed the model of local dominance in this chapter. This including setting the concept of locality, using a metric definition on the different possible states of a game (for us, using changed votes), and choosing a strategy for the voter to follow in the set of limited, local states it is considering – dominance. We then introduced into this model the dynamic of iterative voting (shown in Chapter 4), with the intention of examining the stable states reachable from a truthful starting position.

We were able to show that if voters are homogenous (e.g., with the same radius), convergence to a stable state is guaranteed, and this includes the case where all voters are truth-biased or lazy-biased. Therefore, we know we can always look at the set of reachable stable states from a truthful starting point.

Using this information, we examined multiple simulations of voting sce-




Figure 5.5: A 2-Urn distribution, showing the average share of the two top candidates, depending on the value of r , showing that near peak r we get results exhibiting Duverger’s law. Note that the peaks are reached in similar values for the same voters number.

narios, using various voter preference distribution models,² and saw real-world phenomena, such as Duverger’s law, in the simulation data. We were further able to characterize the way manipulation varied with different radii: in most cases, as the radius increases, manipulation grew and grew, and the quality of the winners grew as well (higher social welfare, more Condorcet winners, etc.), until reaching a peak at some value of the radius (“peak r ”), which depended on the number of voters, not candidates, and from that point all these variables point downward to the truthful state (when the radius equals the number of voters).

²The simulation code was released to the research community.

Part II

All-Pay Auctions

But much the commoner type of success in every walk of life and in every species of effort is that which comes to the man who differs from his fellows not by the kind of quality which he possesses but by the degree of development which he has given that quality. . .

It is the only kind of success that is open to most of us. Yet some of the greatest successes in history have been those of this. . . class

THEODORE ROOSEVELT, *An Autobiography*

Chapter 6

All-Pay Auction Overview & Preliminaries

Auctions, as a mechanism of allocating indivisible goods, have been used in the ancient world in various roles, such as allocating spoils of war [133]. However, they gained a much more prominent role since the 17th century, starting in Europe, as numerous items began to be sold in auctions, perhaps as a result of the growth in global trade, bringing various items of limited quantity which were not widely enough available for regular pricing to be used. The most common auction is the first-price auction, in which the winner is the bidder with the highest bid, which the winner pays to the auctioneer.

All-pay auctions, on the other hand, are not commonly used as a practical auction mechanism. As in first-price auctions, the winner is the bidder with the highest bid, but unlike other auctions, in all-pay auctions, all bidders pay their bid – regardless if they win or not. These type of auctions are still a way to allocate goods, but they model a real-world process, rather than suggest a mechanism to divide, for example, presents between siblings. All-pay auctions model the process in which multiple contestants are putting in effort to reach a goal, but only the first to reach it will gain it – a “winner takes all” situation. For example, employees putting in effort to receive a

bonus that is given only to the best-performing employee (“employee of the month”), or pharmaceutical companies trying to develop a drug put in effort (their “bid”, for example, in weekly hours allocated to the project), but, ultimately, only the first one to patent the medication will reap the benefit of its effort, while for the rest, the effort turns out to be futile.

Furthermore, the online market for services that has matured in the past decade, turns out to have even more use for all-pay auctions. As crowdsourcing flourishes, many of these markets are applying de-facto all-pay auctions. The Netflix challenge [57], for example, in which a million dollars was offered to the best movie recommendation algorithm, is, in effect, an all-pay auction, in which various research groups put in effort (which they did not get back if they lost), in the hope of developing the best algorithm. More systematically, sites like Topcoder (<http://www.topcoder.com>)[138] and CodeChef (<http://www.codechef.com>) are employing this type of mechanism in order to assign projects to appropriate workers.

The problem with all-pay auctions is that, *prima facie*, they should not exist. A game-analytic analysis of them [55, 56], as was done with other auctions, indicates the Nash equilibria of the auctions results in an expected utility of zero for the bidders. In other words, participants are increasing their variance for no expected gain at all. As people are notoriously risk averse [124] they would prefer not to participate in such an endeavor.¹

In order to approach this problem, we enriched the model of the bidders in our analysis. First, we allowed them to interact with each other and coop-

¹Note that all-pay auctions are not like lotteries. First, winning in lotteries is not determined by effort (or bid). While investing more in buying lottery tickets may increase one’s probability of winning, buying more than other people will still not guarantee victory. Moreover, buying one lottery ticket for low sum, one has a chance to win even when there are many other people buying large amount of tickets – their purchases do not even decrease one’s chances of winning at all. Second, in all-pay auctions, companies and rational actors are participating for a long time, without going broke – pharmaceuticals are mostly active only in all-pay auctions, and still do not go broke, which seems to indicate a positive expected profit.

erate. We examine both mergers (publicly joining together) and collusions (covertly joining together), and examine the Nash equilibria in this situation, showing that while mergers are pretty useless, collusions are quite profitable for the colluders. Somewhat surprisingly, we show that in some cases (when the number of colluders is large), the collusion leaves the other bidders – who were purposely left in the dark about it – better off! The auctioneers, on the other hand, lose in such cases, though when the number of bidders is small, they are better off than without them.

Following that, we look at the common online scenario in which the participation of bidders is not guaranteed. Assuming bidders have some probability of participating in a particular bid, we see that the expected value for the participating bidders becomes positive, encouraging their participation.

6.1 Preliminaries and Definitions

Our model is a fundamentally a simple one. We handle all-pay auctions with a single auctioned item that is valued by all participants with the same value. Thanks to Baye et al. [56] and Hillman and Riley [115], we know that we do not need to consider participants with different values, as if there is more than one bidder with the maximal valuation for the items, all those with a different valuation bid 0, and that otherwise, the bidder with the maximal valuation behaves as if its maximal valuation is the 2nd highest (and all those with a lower item valuation, once again, bid 0).

In most cases, we will look into the symmetric equilibrium, as it is not only unique, but as all bidders have a similar valuation, there is no need to assume otherwise.

Formally, each of the n bidders issues a bid of b_i , $1 \leq i \leq n$, and all bidders value the item at 1. The highest bidders win the item and divide it among themselves, while the rest lose their bid. Thus, bidder i 's utility from a combination of bids $\{b_1, \dots, b_n\}$ is given by:

$$\pi_i(b_1, \dots, b_n) = \begin{cases} \frac{1}{|\arg \max_j b_j|} - b_i & b_i \in \arg \max_j b_j \\ -b_i & b_i \notin \arg \max_j b_j \end{cases}$$

The symmetric equilibrium is a mixed equilibrium with full support on $[0, 1]$, so that each bidder's bid is distributed in $[0, 1]$ according to the same cumulative distribution function F , with the density function f (shown to uniquely exist in [56, 151]). As the bids are distributed in a continuous range, with a non-atomic distribution, we do not need to address cases of ties between them.

When there are no colluders, the setting is similar to one considered in [56, 104, 83], where various results on behavior of non-cooperating bidders have been provided. To enable us to evaluate the effect of mergers and collusion on the auction, we build on this previous analysis, which we briefly overview below.

6.1.1 Bids

With no colluders at all, the expected utility of any participant with a bid b is:

$$\pi(b) = (1 - b) \cdot \Pr(\textit{winning} \mid b) + (-b) \cdot \Pr(\textit{losing} \mid b)$$

where $\Pr(\textit{winning} \mid b)$ and $\Pr(\textit{losing} \mid b)$ are the probabilities of winning or losing the item when bidding b , respectively. In a symmetric equilibrium with n players, each of the bidders chooses their bid from a single bid distribution with a probability density function $f_n(x)$ and a cumulative distribution function $F_n(x)$. A player bidding b will only win if all other $n - 1$ players bid at most b , which occurs with probability $F_n^{n-1}(b)$. Thus, the expected utility of a player bidding b is given by:

$$\pi(b) = (1 - b)F_n^{n-1}(b) - b(1 - F_n^{n-1}(b)) = F_n^{n-1}(b) - b$$

In a mixed Nash equilibrium, all points in the support yield the same expected utility to a player, so we have $\pi(x) = \pi(y)$ for all x, y in the support. For an equilibrium with full support, this yields $\pi(0) = \pi(x)$ for all $x \in [0, 1]$. Since $\pi(0) = 0$, this means that for all bids, $F_n^{n-1}(b) = b$. Hence, we have $F_n^{n-1}(b) = \left(\int_0^b f_n(x) dx\right)^{n-1} = b$, implying that $F_n(x) = x^{\frac{1}{n-1}}$ and $f_n(x) = \frac{x^{\frac{2-n}{n-1}}}{n-1}$. Therefore, the expected bid is:

$$E(\text{bid}) = \int_0^1 x \cdot \frac{x^{\frac{2-n}{n-1}}}{n-1} dx = \frac{1}{n-1} \int_0^1 x^{\frac{1}{n-1}} dx = \frac{1}{n-1} \cdot \frac{n-1}{n} x^{\frac{n}{n-1}} \Big|_0^1 = \frac{1}{n}$$

The bid's variance is hence:

$$\begin{aligned} \text{Var}(\text{bid}) &= \int_0^1 x^2 \cdot \frac{x^{\frac{2-n}{n-1}}}{n-1} dx - \frac{1}{n^2} = \frac{1}{n-1} \int_0^1 x^{\frac{n}{n-1}} dx - \frac{1}{n^2} = \\ &= \frac{1}{n-1} \cdot \frac{n-1}{2n-1} x^{\frac{2n-1}{n-1}} \Big|_0^1 - \frac{1}{n^2} = \frac{1}{2n-1} - \frac{1}{n^2} \end{aligned}$$

That is, both the expected bid and the bid's variance monotonically decrease with n .

6.1.2 Bidders

Given the expected bids, we now more closely examine the profits of the bidders. A bidder's profit (BP) is characterized by the probabilistic density function (p.d.f.) g_{BP} below:

$$g_{BP}(z) = \begin{cases} f_n(1-z)F_n^{n-1}(1-z) & z > 0 \\ f_n(-z)(1-F_n^{n-1}(-z)) & z \leq 0 \end{cases}$$

This gives the expected bidder's profit of:

$$\begin{aligned}
E(BP) &= \int_{-1}^1 z g(z) dz = \int_{-1}^0 z \frac{1}{n-1} (-z)^{\frac{2-n}{n-1}} (1+z) dz + \\
&\quad + \int_0^1 z \frac{1}{n-1} (1-z)^{\frac{2-n}{n-1}} (1-z) dz = \\
&= -\frac{1}{n} + \frac{1}{2n-1} - \frac{1-n}{n(2n-1)} = 0
\end{aligned}$$

The bidders' profit also has the following variance:

$$\begin{aligned}
Var(BP) &= E(z^2) - E^2(z) = E(z^2) - 0 = \int_{-1}^1 z^2 g(z) dz = \\
&= \int_{-1}^0 z^2 \frac{1}{n-1} (-z)^{\frac{2-n}{n-1}} (1+z) dz + \\
&\quad + \int_0^1 z^2 \frac{1}{n-1} z^{\frac{2-n}{n-1}} (1-z) dz = \\
&= \frac{3n^2 - 5n + 2}{n(2n-1)(3n-2)}
\end{aligned}$$

Differentiating the above gives $\frac{-2n^2+4n-1}{(1-2n)^2n^2}$, which is negative for all $n \geq 2$, so the variance in the bidders' profit decreases as the number of bidders increases.

6.1.3 Auctioneers

We consider two types of auctioneers:

Sum-profit The auctioneer receives all the money paid into the system, i.e., all of the bids. Such an auctioneer can be seen in the Netflix challenge, where the company could incorporate ideas from projects that did not win, but perhaps included a novel approach.

Max-profit The auctioneer receives only the money paid by the winner – the winning bid. Our pharmaceutical example above is such a case – there is no increase to the public benefit by the parallel efforts to develop the same drug.

As the expected profit of all bidders is zero, the auctioneer's profit (AP) is equal to the total social welfare of the auction.

In the sum-profit model the auctioneer retains all the bids so its expected profit is simply the sum of expected bids:

$$E(AP) = \sum_{i=1}^n E(\text{bid}) = \sum_{i=1}^n \frac{1}{n} = 1$$

In this case, the variance in the auctioneer's utility equals $Var(AP) = \frac{n}{2n-1} - \frac{1}{n}$, monotonically increasing in n .

In contrast, in the max-profit model the auctioneer's utility is only the maximal bid, which has the following cumulative distribution function (c.d.f.) G_{AP} :

$$G_{AP}(z) = \begin{cases} F_n^n(z) = z^{\frac{n}{n-1}} & z > 0 \\ 0 & z \leq 0 \end{cases}$$

The expected profit is then given by:

$$E(AP) = \int_0^1 (1 - z^{\frac{n}{n-1}}) dz = \frac{n}{2n-1}$$

This expression is monotonically decreasing in n . Notice that this value always exceeds $\frac{1}{2}$, so the auctioneer expects to receive more utility from the auction than the utility obtained by all the winners together (as the total value of the item is 1). To find the variance we note that $E(AP^2) = \frac{n}{3n-2}$. Thus, the variance for the max-profit auctioneer is:

$$Var(AP) = \frac{n}{3n-2} - \left(\frac{n}{2n-1}\right)^2 = \frac{n(n-1)^2}{(3n-2)(2n-1)^2}$$

This expression increases with n , as was the case with the sum-profit auctioneer.

Variable	Naive approach
Expected bid [Variance]	$\frac{1}{n}$ $[\frac{1}{2n-1} - \frac{1}{n^2}]$
Bidder utility [Variance]	0 $[\frac{n-1}{n(2n-1)}]$
Sum-profit auctioneer utility [Variance]	1 $[\frac{n}{2n-1} - \frac{1}{n}]$
Max-profit auctioneer utility [Variance]	$\frac{n}{2n-1}$ $[\frac{n(n-1)^2}{(3n-2)(2n-1)^2}]$

Table 6.1: The values, in expectation, of some of the variables when there is no possibility of failure or collusion.

Chapter 7

Come Together: Mergers and Collusions in All-Pay Auctions

7.1 Introduction

All-pay auctions, despite modeling obvious real-world scenarios, have been somewhat marginalized in the vast research on auctions which flourished in the past 20 years. This stems, at least partially, from what seems to be its fundamental property: the expected profit for a bidder is zero, hence there is no good reason to participate in such a mechanism.

However, the real-world manifestations of the all-pay auction indicate this naive analysis might be flawed, and that such an analysis ignores at least some of the properties that make all-pay auctions a worthwhile activity to their bidders. Instead, we strive to find a richer model for the participants behavior, that is capable of explaining it.

One such path is including in the model one of the most fundamental human properties: cooperation. That is, allowing bidders to pool their resources and join together in a joint bid.¹ Such cooperation could be public

¹We do not discuss the division of utility between them, and that is more applicable to transferable utility games, in cooperative game theory, which is more widely discussed in Chapter 10.

and well known to other participants (a *merger*), or it could be concealed and hidden, known only to the collaborators (a *collusion*).

Once we allow bidders to collude, the structure of the auction outcome becomes radically different. We analyze what happens with a single group of merged/colluding bidders, and show its effects on the colluders, the other bidders, and the auctioneers (both types of them). We emphasize how the number of colluders has a significant effect on the fortunes of the other bidders and the auctioneers.

We then continue to explore settings which include multiple colluding groups, and show the expected outcomes, as well as analyzing what structure of collusion should auctioneers strive to achieve (naturally, the actual ability of the auctioneer to “engineer” such a structure depends on the particular real-world setting the auction is a part of). In both parts we reach results that indicate settings that encourage participation by bidders, as long as they are aware of the possibility of collusion (even if they are not colluders themselves).

7.1.1 Related Work

Research into all-pay auctions originated in political science, dealing with lobbying [115, 55], but much of the analysis (especially dealing with the Revenue Equivalence Theorem) is found in auction theory studies [133, 131]. In auctions where the winner is the bidder with highest bid, Maskin and Riley [151] showed that when bidders have the same value distribution for an item, there is a symmetric equilibrium. A prominent study of all-pay auctions in full information settings is Baye et al, [56], showing how most valuations – apart from the top two – are irrelevant to the analysis of the winner, and (using [115]), showing that in most cases, the possible equilibria are those with full support on the range from 0 to the highest (or second highest) valuation; when all players have the same valuation, a full support is the only symmetric equilibria possible. That work does not deal with cooperation among players, but helps validate our choice of focusing on full-support equilibria.

There has also been some work assuming bidders' valuation is not fixed and public knowledge, but rather is sampled from a distribution. While the case of 2 bidders has been examined in this method [10], the general family of Nash equilibria has yet to be fully characterized in this setting. However, there are advances in this direction: Showing the existence of equilibria in particular settings and conditions [134, 15], and showing behavior in setting with enforced caps on bids [108].

More recent work has extended the basic model, focusing on its applicability to crowdsourcing. DiPalantino and Vojnović [83] detailed the issues stemming from needing to choose one auction from several, and Chawla et al. [70] dealt with optimal mechanisms for crowdsourcing. Using both theoretical and empirical tools, Gao et al. [104] examined whether several stages were better for crowdsourcing, while Archak and Sundararajan [13] addressed issues on designing the award of the crowdsourcing contest.

Collusions in auctions (nicknamed “bidding rings”) were examined in the auction theory literature, differentiating between second-price and first-price auctions, due to their different ability to “self-police” each bid [133, 131, 150, 32, 22, 60, 50]. Focusing more on mergers, a model of mergers with full information was proposed in Huck et al. [119] for many auctions, including auctions where each bidder makes an investment and the gains are divided among the bidders. It explores various auction models in which there are no single winners, but rather profits are distributed among the players according to their investment relative to the others. It shows that mergers, in many domains (e.g., when there is a marked benefit to be the top bidder), are profitable for their participants, while when there is no significant benefit to being the maximal bidder, mergers may still be beneficial, but players do not coalesce around a single bid, but rather divide their resources among them. In this setting the only first-price auction is an all-pay one, and due to its full information assumptions, that work dismisses all-pay auctions as uninteresting, since the bidders' profit is always zero.

A recent paper models auction collusion using voting techniques [132].

There, bidders prefer some of the adversaries winning over other adversaries, and so collude with them. Work on partial information in auction collusions is somewhat limited and focuses on the auctioneer's ignorance rather than on bidders hiding information from one another (e.g. [98, 1]). It neither deals with all-pay auctions, nor with rival groups of colluders.

7.2 Mergers

Consider the case where k out of n bidders work together, and do not hide their cooperation from the remaining participants. These bidders coordinate their behavior, and can thus be thought of as a single player, whose strategy space is the cartesian product of the strategy space of the coordinating agents. We refer to this player as the “merged player” representing the coordinating players, and refer to the remaining non-merged players as the “singleton players”. In the Nash equilibrium, the joint player best responds to the strategies of the non-merged players, and the strategy of singleton players is the best response to the other singleton players and the merged player.

As noted earlier, only the largest submission wins the auction, so the merged player would be wasting its effort if the agents composing it were to make more than one bid. Therefore, since the merging players would only make a single submission (using one of the identities of the merging players), we may consider the joint player as a single bidder, and examine the equilibrium in the resulting game.

The utility of the “merged” player follows that of a single bidder. Therefore, we are essentially seeking a Nash equilibrium for $n - k + 1$ bidders, as the merging group would bid using a mixed strategy with full support. The resulting equilibrium is thus equivalent to the equilibrium of the auction with no mergers, but with fewer players. It follows immediately from the analysis of the setting with no mergers that the expected profit for each bidder in this equilibrium would remain zero, that the variance would grow (as it is monotonically decreasing in the number of bidders), and that the

expected bid would grow to $\frac{1}{n-k+1}$. In broad terms, this means that for non-merging players, if they win the auction, they would have a lower utility than previously, and if they lose, they would lose more (as the bids get higher). However, the chances of winning do increase, due to the lower number of actual participants.

The auctioneer's profit in the sum-profit model would not change, as the total sum of bids is still 1. However, in the max-profit model, the auctioneer's expected profit (and the social welfare) will increase, while its variance will drop.

7.3 Collusions

We now analyze the setting with colluding bidders, where the other players are not aware of their collaboration. We first focus on the case with a fixed number of colluders, and then show how the utility of each member depends on the size of a coalition. Finally, we examine the effect of collusion on the profit of the auctioneer and the social welfare.

In a sense, the analysis of collusions is a short term analysis rather than an analysis of player behavior in equilibrium. In a merger of players in an auction (be it an all-pay auction or any other auction), the collaboration is known to the other players. In contrast, the unfair advantage of colluding players stems from other players being unaware of this cooperation; in other words, the non-colluders are not best-responding to the colluders' bids. Hence, in a set of repeating auctions, other bidders might be able to figure out the existence of the collusion.

In this short term, the non-colluders believe that all bidders are operating independently, so they expect all bidders to behave according to the symmetric equilibrium. The colluders can capitalize on this behavior of the non-colluders and improve their utility. As the interaction repeats in further auctions, more and more non-colluders may become aware of the agreement between the colluders. Later, we consider possible reactions of the non-

colluders as they become aware of collusion. Once all the non-colluders gain knowledge of the agreement between the collaborating agents, the collusion becomes equivalent to mergers, analyzed above.

7.3.1 A Single Group of Colluders

Suppose we have k colluders out of n bidders, and that the remaining players are not aware of the collusion. Hence, we expect the other bidders to play according to the symmetric equilibrium. When the colluders submit a single bid b , they win if all the other $n - k$ bids are at most b and lose otherwise, so their total utility is:

$$\pi(b) = (1 - b)F_n^{n-k}(b) - b(1 - F_n^{n-k}(b)) = b^{\frac{n-k}{n-1}} - b$$

The variance for a fixed bid b is:

$$\begin{aligned} \text{Var}(b) &= (1 - b - (b^{\frac{n-k}{n-1}} - b))^2 F_n^{n-k}(b) + \\ &\quad + (-b - (b^{\frac{n-k}{n-1}} - b))^2 (1 - F_n^{n-k}(b)) = \\ &= b^{\frac{n-k}{n-1}} - b^{\frac{2(n-k)}{n-1}} \end{aligned}$$

To find the optimal utility we examine the bid where the derivative is 0:

$$\pi'(b) = \frac{n-k}{n-1} b^{\frac{1-k}{n-1}} - 1 = 0$$

This implies $b^{\frac{1-k}{n-1}} = \frac{n-1}{n-k}$, yielding the optimal bid:

$$b^* = \left(\frac{n-k}{n-1}\right)^{\frac{n-1}{k-1}}$$

To see how this may affect the profits of participants in the auction, consider the following example.

Example 7.1. *In the non-collusion setting, in an auction with two bidders the equilibrium bids are drawn from a distribution with the p.d.f. $f_2(x) = 1$, and the expected bid is $\frac{1}{2}$. If there are 3 participants, the bids are distributed*

according to the c.d.f. $F_3(x) = \sqrt{x}$, with the expected bid of $\frac{1}{3}$. The auctioneer's expected profit in the sum-profit model is 1, while in the max-profit model with 3 bidders it is $\frac{3}{5}$.

However, if 2 of these 3 bidders collude, our results show they should bid $\frac{1}{4}$, which gives the colluders the expected profit of $\frac{1}{4}$, while the outsider has the expected loss of $\frac{1}{6}$. The auctioneer's expected profit decreases in both models: in the sum-profit model it drops from 1 to $\frac{7}{12}$ and in the max-profit model from $\frac{3}{5}$ to $\frac{1}{4}F_3(\frac{1}{4}) + \int_{\frac{1}{4}}^1 bf_3(b) db = \frac{10}{24} < \frac{3}{5}$.

It is quite intuitive that the fewer the bidders that are left outside the coalition, the easier it is for colluders to out-bid them. We formally show this.

Lemma 7.1. *The colluders' bid monotonically decreases with k , and monotonically increases with n , up to $\frac{1}{e}$.*

(Proof is in Appendix B)

However, even if the collaborators optimize their bid accordingly, the number of outsiders still has a negative effect on their expected utility.

Lemma 7.2. *The colluders' expected profit decreases with n and increases with k .*

(Proof is in Appendix B)

7.3.2 Optimal Number of Colluders

We now show that not only does the overall utility of colluders increase with the size of their group, the individual share of each member also grows. This convexity implies that colluders have strong incentives to invite more players to take part in collusion.

Theorem 7.1. *The expected profit per colluder increases with k .*

(Proof is in Appendix B)

Hence, the colluders would seek to increase their numbers as much as possible. Next, we explore the effect of collusion on the auctioneer's profit and the social welfare. We show that this effect can be either positive or negative, depending on the number of colluders and the total number of bidders.

7.3.3 Auctioneer's Profits

We now show that if the total number of auction participants is large enough, collusion may be beneficial to the auctioneer in both the sum-profit model and the max-profit model.

Theorem 7.2. *In the setting with k colluders, the expected auctioneer utility is $\frac{n-k}{n} + \left(\frac{n-k}{n-1}\right)^{\frac{n-1}{k-1}}$ in the sum-profit model and $\frac{n-k}{2n-k-1} \left(1 + \frac{n-k}{n-1}\right)^{\frac{2(n-k)}{k-1}}$ in the max-profit model. The profit in both models decreases in the number of colluders and increases in the total number of participants. For sufficiently large n , they exceed the corresponding auctioneer's utilities in the setting without collusion.*

(Proof is in Appendix B)

7.3.4 Social Welfare

We now analyze social welfare in the setting with colluders. To this end, we need to calculate the expected profits of the non-colluders. Surprisingly, as the following theorem shows, in some cases they may even benefit from other players colluding. Overall, however, the presence of colluders does not affect the social welfare in the sum-profit model, and may have either a positive or a negative effect in the max-profit model, depending on the parameters of the setting.

Theorem 7.3. *The social welfare in the sum-profit model does not change due to collusion. In the max-profit model, the presence of colluders may have*

different effects on the social welfare, depending on the relation between the number of colluders and the total number of participants. In particular, the social welfare drops for settings with many participants.

(Proof is in Appendix B)

7.3.5 Other Equilibria

Our previous analysis examines the symmetric equilibrium to the all-pay auction described in the previous chapter. However, this setting also admits asymmetric equilibria [56], where at least two players have a non-atomic distribution on $[0, 1]$, while the rest of the players have a non-atomic distribution – each player i has a $0 \leq b_i \leq 1$, and has a distribution on $[b_i, 1]$ which is the same distribution on this interval as all other players (i.e., the same p.d.f.), with the rest of the probability mass concentrated at 0.² In the asymmetric case the expected player profit remains 0, and sum of bids remains 1.

We believe the symmetric equilibrium is by far the most likely to occur and thus more sensible to investigate, due to the symmetry between players in most practical settings. However, we would like to note that it is also possible to study the asymmetric case along the same lines of our analysis regarding the symmetric case.

In more detail, the analysis for mergers changes very little, as mergers will just result in a different equilibrium of fewer players, depending on the preferences of the merged players and the equilibrium structure (e.g., if the players with the non-atomic distribution on $[0, 1]$ merge, another player needs to “switch” to that distribution). The analysis of collusion becomes more elaborate. As the distribution remains the same (over the intervals) for all players, our analysis indicates that the profit from collusion is greater than or equal to the symmetric case for every bid, with the difference from the symmetric case being more significant for lower bids. Thus, the optimal bid of

²There are no equilibria other than the symmetric one and asymmetric equilibria taking this form.

the colluders in the asymmetric case is generally lower than in the symmetric case. This effect is more pronounced when there are many “cautious” bidders (i.e., those with a high b_i). Furthermore, the effects of increasing k and n on the bid remain the same as in the symmetric scenario (though the effects on colluders’ profits are smaller).

Considering players that have valuations below 1, which bid 0 in a setting without collusions, this situation does not change if they try to collude together. Of course, they can join an existing collusion ring, and, depending on the allocation mechanism between colluders, the value of winning may be below 1, which would slightly change, for example, the expression of the optimal bid. However, the effect depends on the participation of the players with valuation of 1, and therefore the effects of collusion remain exactly the same as without lower valuation participants.

7.4 Response To Collusion

The previous section examined the collusion’s direct impact, where non-colluders continue behaving as the symmetric equilibrium prescribes (which is sub-optimal in the presence of colluders). This approach is justified by the fact that this symmetric equilibrium is a mixed one, so collusion may be difficult to detect. However, after many interactions, non-colluders may notice their winning rate is different from what they would expect under the symmetric equilibrium, and suspect foul play. How would they respond to the colluding coalition? Would it make sense to collaborate with other participants and play jointly against the colluders? In this section we consider two scenarios: where there exists a single player who is aware of collusion, and where several colluding coalitions are possible.

7.4.1 A Player Aware Of Collusion

If one of the players becomes aware of k other bidders colluding, they would never submit a non-zero bid below the colluders’ bid $b^*(k)$, as they would

lose and get a negative utility. They would rather respond by either bidding 0 (thus obtaining zero utility), or placing a bid b which is higher than $b^*(k)$. In the latter case, the expected profit is:

$$\pi(b) = (1 - b)F_n^{n-k-1}(b) - b(1 - F_n^{n-k-1}(b)) = b^{\frac{n-k-1}{n-1}} - b$$

This is always positive (since $b \leq 1$ and $\frac{n-k-1}{n-1} < 1$), so it is always beneficial to bid above $b^*(k)$ rather than 0. As the optimal bid for $k + 1$ colluders is $b^*(k+1)$, which according to Lemma 7.1 is smaller than $b^*(k)$, the best bid for the responder is the smallest possible value above $b^*(k)$. Having a larger bid is less profitable, as the bidder's expected profit is monotonically decreasing in k in the interval $k > b^*(k)$.

Now, since this bid is larger than $\frac{1}{n}$, this means that in the sum-profit model, it is beneficial for the auctioneer to expose the existence of a collusion ring to some players. Similarly, in the max-profit model, the expected profit grows (especially when k is significantly smaller than n), making the leak profitable to this auctioneer type as well.

7.4.2 Several Groups Of Colluders

If there are several groups of colluders that are not aware of one another, each would bid its optimal value as prescribed by the previous analysis. By Lemma 7.1, this bid decreases with the size of a coalition, and so the smallest coalition would outbid the others and get positive (though sub-optimal) expected utility. Indeed, suppose there are m colluder groups, each with k_i colluders, and let $k_{min} = \min\{k_1, k_2, \dots, k_m\}$. The expected profit of the smallest (winning) coalition is

$$\begin{aligned} & \left(\frac{n - k_{min}}{n - 1}\right)^{\frac{n-1}{k_{min}^{n-1}} \cdot \frac{n - \sum_{i=1}^m k_i + m - 1}{n-1}} - \left(\frac{n - k_{min}}{n - 1}\right)^{\frac{n-1}{k_{min}^{n-1}}} \\ & = \left(\frac{n - k_{min}}{n - 1}\right)^{\frac{n - \sum_{i=1}^m k_i + m - 1}{k_{min}^{n-1}}} - \left(\frac{n - k_{min}}{n - 1}\right)^{\frac{n-1}{k_{min}^{n-1}}} \end{aligned}$$

As the bids are larger with a smaller number of colluders, the auctioneer prefers several small groups of colluders over a single big one. If k_{min} is large enough, it becomes worthwhile for a sum-profit auctioneer to uncover collusion rings and publicize them. Similarly, for a max-profit auctioneer, it may be worthwhile to expose the collusions (even more so than for the sum-profit one). For example, for $n > 6$, the sum-profit auctioneer would rather divide the bidders into pairs. However, for the max-profit auctioneer, it is never profitable to have all bidders be colluders. As the maximal bid of colluders is $\frac{1}{e}$ (by Lemma 7.1), and its expected profit without them is above $\frac{1}{2}$, the max-profit auctioneer “needs” non-colluders to increase its expected profit.

7.5 Summary

In this chapter we gave an potential answer to the basic question of why people and groups bother participating in all-pay auctions, as, ostensibly, their expected profit is zero. In this work we suggested that while publicly cooperating (mergers) does not help them, cooperating covertly (collusion) does not only make it worthwhile for the colluders to participate, but in cases where there is a large colluder ring, it induces a positive outcome for the non-colluder as well (hence, they would be worst-off should they learn of the collusion). Quite naturally, the effect on the auctioneers is a mirror image of the effect on the non-colluders: when the non-colluders are benefiting from the collusion, the auctioneers are worst off, and the auctioneers are benefiting from the collusion when the non-colluders are losing. The numerical results of this part are summarized in Table 7.1.

However, beyond the case of a one-off auction collusion, we are interested to see what happens when participants may suspect collusion (for example, due to their past experience). We find what the strategies of the non-colluding bidders should be. We also show that in a setting of multiple colluders, the large collusion groups are in danger from smaller ones,

Variable	No cooperation	Mergers	Collusion
Expected bid [Variance]	$\frac{1}{n}$ $[\frac{1}{2n-1} - \frac{1}{n^2}]$	$\frac{1}{n-k+1} \uparrow$ $[\frac{1}{2(n-k+1)-1} - \frac{1}{(n-k+1)^2}] \uparrow$	colluders: $(\frac{n-k}{n-1})^{\frac{n-1}{k-1}} \downarrow [0]$ non-colluders: $\frac{1}{n} [\frac{1}{2n-1} - \frac{1}{n^2}]$
Bidder utility [Variance]	0 $[\frac{3n^2-5n+2}{n(2n-1)(3n-2)}]$	0 $[\frac{3(n-k+1)^2-5(n-k+1)+2}{(n-k+1)(2(n-k+1)-1)(3(n-k+1)-2)}] \uparrow$	colluders: $(\frac{n-k}{n-1})^{\frac{n-1}{k-1}} (\frac{k-1}{n-1}) \uparrow$ $[(\frac{n-k}{n-1})^{\frac{n-k}{k-1}} - (\frac{n-k}{n-1})^{\frac{2(n-k)}{k-1}}]$ non-colluders: $\frac{k-n(\frac{n-k}{n-1})^{\frac{n-k}{k-1}}}{n(n-k)}$
Sum-profit principal utility [Variance]	1 $[\frac{n}{2n-1} - \frac{1}{n}]$	1 $[\frac{n-k+1}{2n-2k+1} - \frac{1}{n-k+1}] \downarrow$	$\frac{n-k}{n} + (\frac{n-k}{n-1})^{\frac{n-1}{k-1}}$
Max-profit principal utility [Variance]	$\frac{n}{2n-1}$ $[\frac{n(n-1)^2}{(3n-2)(2n-1)^2}]$	$\frac{n-k+1}{2n-2k+1} \uparrow$ $[\frac{(n-k+1)(n-k)^2}{(3n-3k+1)(2n-2k+1)^2}] \downarrow$	$\frac{n-k}{2n-k-1} (1 + (\frac{n-k}{n-1})^{\frac{2(n-k)}{k-1}})$

Table 7.1: The values, in expectation, of some of the variables in a non-cooperative setting, when k members merged, and when k members are colluding. Arrows indicate monotonicity of expression, as k grows.

as these will bid higher than them. For auctioneers, we find they should strive to have small collusion groups, which would improve their situation significantly (particularly for sum-profit auctioneer).

Chapter 8

“Showing Up is 80% of Life”: Participation in All-Pay Auctions

8.1 Introduction

Despite the usefulness of all-pay auctions as a model for real-life scenarios, in which much effort is expended to prepare a bid but only one option is selected (e.g., defense projects' competitions, pharmaceuticals chasing patents, etc.) they have not been a relatively widely explored mechanism. This stems, as mentioned in the previous chapters, from their basic analysis, that seems to indicate that the expected profit in such auctions is zero.

The continued existence of all-pay auctions indicates their participants do consider them to be a useful mechanism, so we need to examine our analysis. In Chapter 7 we tried to consider collusions as a method by which participants are able to have a positive utility from all-pay auctions. Here, we suggest a different approach to this same quandary.

Common all-pay auctions today exist online, in sites such as TopCoder [138] and TopChef. One of the distinctive properties of such online auctions is that participants do not know exactly whether all potential participants will ac-

tually take part in the auction. Some may drop out early, others may be too busy, etc. In any case, when the bidders make their bid, they do not know who else will participate with them.

More formally, we suggest that each participant has a certain fixed probability of the auctions it will participate in, which is known to all other participants – p_i . Bidders take this knowledge into account when they formulate their bidding strategy, resulting in different equilibria than those of the naive models. We also show how this affects the auctioneer, and imply, therefore, what all parties can do to increase their expected profits.

8.1.1 Related Work

The general research on all-pay auctions has been mentioned in Chapter 7, so here we will focus on participation.

The early major work on participation in auctions was McAfee and McMillan [154], followed by Matthews [153], which introduced bidders who are not certain of how many bidders there will actually participate at the auction. Their analysis showed that in first-price auctions, risk averse bidders prefer to know the numbers, while it is the auctioneer’s interest to hide that number. In the case of neutral bidders, their model claimed that bidders were unaffected by the numerical knowledge. Dyer et al. [89] claimed that experiments that allowed “contingent” bids (i.e., one submits several bids, depending on the number of actual participants) supported these results.

Menezes and Monteiro [163] presented a model where auction participants know the maximal number of bidders, but not how many will ultimately participate. However, the decision in their case was endogenous to the bidder, and therefore a reserve price has a significant effect in their model (though ultimately without change in expected revenue, in comparison to full-knowledge models). In contrast to that, our model,¹ which assumes a little more information is available to the bidders (they know the maxi-

¹We use a framework similar to the one in Meir et al. [162], albeit there it was used in congestion games.

mal number of bidders and the probability of failure), finds that in such a scenario, bidders are better off not having everyone show up, rather than knowing the real number of contestants appearing. Empirical work done on actual auctions [148] seems to support some of our theoretical findings (though not specifically in all-pay auction settings).

8.2 Equilibrium

We assume each bidder has its own probability for participating in the auction, with $0 \leq p_1 \leq \dots \leq p_n \leq 1$. We shall now present a symmetric equilibrium for this case (i.e., assigning the same behavior to agents with the same participation values), with a positive expected profit for the bidders.

We begin by defining a few helpful functions. First, we define

$$\lambda = \prod_{j=1}^{n-1} (1 - p_j)$$

We also define the following expressions for $k \in \{1, \dots, n-1\}$:

$$H_k(x) = \begin{cases} \frac{\lambda+x}{\prod_{j=1}^{k-1} (1-p_j)} \frac{1}{n-k} & k > 1 \\ (\lambda+x)^{\frac{1}{n-1}} & k = 1 \end{cases}$$

$$\alpha_k = \begin{cases} (1-p_k)^{n-k} \prod_{j=1}^{k-1} (1-p_j) - \lambda & k > 1 \\ (1-p_1)^{n-1} - \lambda & k = 1 \end{cases}$$

For the virtual “0” index, we use $\alpha_0 = 1 - \lambda$. Note that because the p_i s are ordered, so are the α_i s: $1 \geq \alpha_0 \geq \alpha_1 \geq \dots \geq \alpha_{n-1} = 0$.

We can now define the cumulative distribution functions (c.d.f.s) for our equilibrium. For every player $i \in \{1, \dots, n-1\}$:

$$F_i(x) = \begin{cases} 1 & x \geq \alpha_0 \\ \frac{H_1(x)+p_i-1}{p_i} & x \in [\alpha_1, \alpha_0) \\ \vdots & \vdots \\ \frac{H_k(x)+p_i-1}{p_i} & x \in [\alpha_k, \alpha_{k-1}) \\ \vdots & \vdots \\ \frac{H_i(x)+p_i-1}{p_i} & x \in [\alpha_i, \alpha_{i-1}) \\ 0 & x < \alpha_i \end{cases}$$

F_n , while very similar to F_{n-1} in its piecewise composition, has, uniquely, an atomic point in the distribution at 0 of $1 - \frac{p_{n-1}}{p_n}$, so:

$$F_n(x) = \begin{cases} 1 & x \geq \alpha_0 \\ \frac{H_1(x)+p_n-1}{p_n} & x \in [\alpha_1, \alpha_0) \\ \vdots & \vdots \\ \frac{H_k(x)+p_n-1}{p_n} & x \in [\alpha_k, \alpha_{k-1}) \\ \vdots & \vdots \\ \frac{H_{n-1}(x)+p_n-1}{p_n} & x \in (\alpha_{n-1}, \alpha_{n-2}) \\ 1 - \frac{p_{n-1}}{p_n} & x = 0 \\ 0 & x < 0 \end{cases}$$

All c.d.f.s are continuous and piecewise differentiable,² and when $p_i = p_j$, $F_i = F_j$, making this is a symmetric equilibrium. In the course of proving this is, indeed, an equilibrium, we shall calculate the expected utility of the bidders when they participate.

Behind this equilibrium lies the understanding that bidders that rarely participate will usually bid high, while those that frequently participate in auctions with less competition would more commonly bid low.

²Note that when $\prod_{j=1}^{k-1}(1-p_j) = 0$ and therefore H_k is undefined for that k , there is no range for which that H_k is used.

Theorem 8.1. *The F_i presented above are a Nash equilibrium, and each bidder's profit is λ .*

(Proof is in Appendix B)³

8.3 Profits

In the previous section we saw that when a bidder actually participates its expected profit is λ , and therefore the overall expected utility for bidder i is $p_i\lambda$ (which, naturally, decreases with n). As is to be expected, a bidder's profit rises the less reliable their fellow bidders are, or the fewer participants the auction has. However, the most reliable of the bidders does not affect the profits of the rest. If a bidder can set their own participation rate, if there is no bidder with $p_j = 1$, that is the best strategy; otherwise, its optimal probability should be $\frac{1}{2}$, as that maximizes $p_i(1 - p_i) \prod_{j=1; j \neq i}^{n-1} (1 - p_j)$.

In order to calculate the auctioneer's profit in a sum-profit model, we need to calculate the expected bid by each bidder, and for that we need to calculate the bidders' probabilistic density function (p.d.f.) For $1 \leq i \leq n-1$:

$$f_i(x) = \begin{cases} 0 & x \geq \alpha_0 \\ \frac{(\lambda+x)^{\frac{2-n}{n-1}}}{p_i(n-1)} & x \in [\alpha_1, \alpha_0) \\ \vdots & \vdots \\ \frac{(\lambda+x)^{\frac{k+1-n}{n-k}}}{p_i(n-k)(\prod_{j=1}^{k-1}(1-p_j))^{\frac{1}{n-k}}} & x \in [\alpha_k, \alpha_{k-1}) \\ \vdots & \vdots \\ \frac{(\lambda+x)^{\frac{i+1-n}{n-i}}}{p_i(n-i)(\prod_{j=1}^{i-1}(1-p_j))^{\frac{1}{n-i}}} & x \in [\alpha_i, \alpha_{i-1}) \\ 0 & x < \alpha_i \end{cases}$$

and $f_n(x) = \frac{p_{n-1}}{p_n} f_{n-1}(x)$.

Using this, we can calculate the expected bid by each bidder, for $1 \leq i \leq n-1$:

³Proof also shows that players with valuation lower than 1 will not participate.

$$\begin{aligned}
E[bid_i] &= \sum_{k=1}^i \int_{\alpha_k}^{\alpha_{k-1}} x f_i(x) dx = \\
&= \frac{1}{p_i} \left(\frac{1}{n} + \sum_{k=1}^i \frac{(1-p_k)^{n-k} \prod_{j=1}^k (1-p_j)}{(n-k)(n-k+1)} - \right. \\
&\quad \left. - \frac{(1-p_i)^{n-i} \prod_{j=1}^i (1-p_j)}{n-i} - p_i \lambda \right)
\end{aligned}$$

and $E[bid_n] = \frac{p_{n-1}}{p_n} E[bid_{n-1}]$.

This expression decreases with n , indicating, as in the regular model, that as more bidders participate, the chance of losing increases, causing bidders to lower their bids. Surprisingly, when summing over all bidders, we receive a much simpler expression, and the sum-profit auctioneer's profits are:

$$\sum_{i=1}^n p_i E[bid_i] = 1 - \lambda \left(1 + \sum_{i=1}^{n-1} p_i \right)$$

In this case, growth with n is monotonic, and hence, any addition to n is a net positive for the sum-profit auctioneer.

To calculate a max-profit auctioneer's profits, we need to first define the auctioneer's c.d.f.:

$$G(x) = \prod_{i=1}^n (p_i F_i(x) + 1 - p_i)$$

That is,

$$G(x) = \begin{cases} 1 & x \geq \alpha_0 \\ (\lambda + x)^{\frac{n}{n-1}} & x \in [\alpha_1, \alpha_0) \\ \vdots & \vdots \\ \frac{(\lambda+x)^{\frac{n-k+1}{n-k}}}{(\prod_{j=1}^{k-1} (1-p_j))^{\frac{1}{n-k}}} & x \in [\alpha_k, \alpha_{k-1}) \\ \vdots & \vdots \\ \frac{(\lambda+x)^2}{\prod_{j=1}^{n-2} (1-p_j)} & x \in [\alpha_{n-1}, \alpha_{n-2}) \\ 0 & x < 0 \end{cases}$$

This is differentiable, and hence we can find $g(x) = \frac{d}{dx}G(x)$; looking for the expected profit, we have:

$$\int_{\alpha_{n-1}}^{\alpha_0} xg(x) dx = \frac{n}{2n-1} - \lambda + \sum_{k=1}^{n-1} \left(\frac{(1-p_k)^{2n-2k-1} \prod_{j=1}^k (1-p_j)^2}{4(n-k)^2 - 1} \right)$$

This means that the max-profit auctioneer would prefer to have two reliable players ($p_n = p_{n-1} = 1$), and the other $n - 2$ bidders as unreliable as possible.

Example 8.1. Consider how four bidders interact. Our bidders have participation probability of $p_1 = \frac{1}{3}$, $p_2 = \frac{1}{2}$, $p_3 = \frac{3}{4}$ and $p_4 = 1$. Let us look at each bidder's c.d.f.s:

$$F_1(x) = \begin{cases} 1 & x \geq \frac{11}{12} \\ 3\left(\frac{1}{12} + x\right)^{\frac{1}{3}} - 2 & x \in \left[\frac{23}{108}, \frac{11}{12}\right) \\ 0 & x < \frac{23}{108} \end{cases}$$

$$F_2(x) = \begin{cases} 1 & x \geq \frac{11}{12} \\ 2\left(\frac{1}{12} + x\right)^{\frac{1}{3}} - 1 & x \in \left[\frac{23}{108}, \frac{11}{12}\right) \\ 2\left(\frac{3\left(\frac{1}{12} + x\right)}{2}\right)^{\frac{1}{2}} - 1 & x \in \left[\frac{1}{12}, \frac{23}{108}\right) \\ 0 & x < \frac{1}{12} \end{cases}$$

$$F_3(x) = \begin{cases} 1 & x \geq \frac{11}{12} \\ \frac{4}{3}\left(\frac{1}{12} + x\right)^{\frac{1}{3}} - \frac{1}{3} & x \in \left[\frac{23}{108}, \frac{11}{12}\right) \\ \frac{4}{3}\left(\frac{3\left(\frac{1}{12} + x\right)}{2}\right)^{\frac{1}{2}} - \frac{1}{3} & x \in \left[\frac{1}{12}, \frac{23}{108}\right) \\ 4\left(\frac{1}{12} + x\right) - \frac{1}{3} & x \in \left[0, \frac{1}{12}\right) \\ 0 & x < 0 \end{cases}$$

$$F_4(x) = \begin{cases} 1 & x \geq \frac{11}{12} \\ (\frac{1}{12} + x)^{\frac{1}{3}} & x \in [\frac{23}{108}, \frac{11}{12}) \\ (\frac{3(\frac{1}{12} + x)}{2})^{\frac{1}{2}} & x \in [\frac{1}{12}, \frac{23}{108}) \\ 3(\frac{1}{12} + x) & x \in (0, \frac{1}{12}) \\ \frac{1}{4} & x = 0 \\ 0 & x < 0 \end{cases}$$

The expected utility for bidder 1 is $\frac{1}{36}$, for expected bid of $\frac{14}{27}$; for bidder 2, $\frac{1}{24}$ for expected bid of 0.394; for bidder 3, $\frac{1}{16}$ for expected bid of 0.277; and for the last bidder, $\frac{1}{12}$ for expected bid of 0.207.

A sum-profit auctioneer will see an expected profit of $\frac{113}{144}$, while a max-profit one will get, in expectation, 0.49.

As a comparison, in the case where we do not allow failures, the c.d.f. of the bidders is $x^{\frac{1}{3}}$ with expected bid of $\frac{1}{4}$ and expected utility of 0. The expected profit of the sum-profit auctioneer is 1, while the expected profit of the max-profit auctioneer is $\frac{4}{7}$.

If we allow our bidders to have the same probability of failure (e.g., when failures stem from weather conditions), many of the calculations become more tractable, and we are able to further understand the scenario.

8.4 Special Case: A Single Participation Probability

Having all bidders with the same probability of participation – p – is a particular instance of the general case presented above. This can be considered, for example, if reasons of non-participation are uniform (e.g., weather). The simplification of the identical probabilities allows us to better understand the scenario by, for example, examining the variance as well.

For each bidder, the expected utility is $p(1 - p)^{n-1}$, monotonically decreasing in n . Using the c.d.f. calculated in the general case, we can also

calculate the expected utility squared, and we use it to calculate the utility variation:

$$\frac{n-1}{n(2n-1)} - \frac{(1-p)^n}{n} + \left(p + \frac{1}{2n-1}\right)(1-p)^{2n-1}$$

The variance increases with p .

The expected bid is $\frac{1}{np}(1 - (1-p)^{n-1}(1+p(n-1)))$, which is neither monotonic in n nor in p . Hence, the expected profit of the sum-profit auctioneer is $1 - (1-p)^{n-1}(1+p(n-1))$ which is monotonically increasing in p and in n . The profit variance is $\frac{np(1-(1-p)^{2n-1})}{2n-1} - \frac{(1-(1-p)^n)^2}{n}$. Note that as n grows, the auctioneer's expected revenue approaches that of the full participation case.

In the case of the max-profit auctioneer, the expected profit is $\frac{n}{2n-1} + \frac{n-1}{2n-1}(1-p)^{2n-1} - (1-p)^{n-1}$, which is monotonically increasing in p ; while not monotonic in n , for large enough n it approaches the expected revenue of the full participation case. The variance is $(1-p)^{2n-2} - \frac{2n(1-p)^{n-1}}{2n-1} + \frac{n}{3n-2} - \frac{2(n-1)^2(1-p)^{3n-2}}{(3n-2)(2n-1)} - \left(\frac{n}{2n-1} + \frac{n-1}{2n-1}(1-p)^{2n-1} - (1-p)^{n-1}\right)^2$.

8.5 Summary

In this chapter we have explored all-pay auctions in settings where participants do not always take part in every bid. Such a scenario is quite common in online crowdsourcing communities (e.g. TopCoder [138]), where there is a set of participants which pick and choose which particular auction they wish to submit a bid to. The reasons for not participating may be varied and specific to each potential bidder – from family to weather to other work projects.

We have shown that assuming non-full participation indeed makes participation of all bidders profitable, thus showing a potential reason for bidders' participation in such auctions, as without this assumption, the utility of participating is, in expectation, zero. Moreover, while in real-world setting participants may not know the actual probability of participation by their peers, the knowledge that it is profitable to join when others may not nec-

Variable	Full	Uniform participation	Individual participation probability
Expected bid	$\frac{1}{n}$	$\frac{1}{np}(1 - (1 - p)^{n-1}(1 + p(n - 1)))$	$\frac{1}{p_i}(\frac{1}{n} + \sum_{k=1}^i \frac{(1-p_k)^{n-k} \prod_{j=1}^k (1-p_j)}{(n-k)(n-k+1)} - \frac{(1-p_i)^{n-i} \prod_{j=1}^i (1-p_j)}{n-i} - p_i \lambda),$ $E[bid_n] = \frac{p_{n-1}}{p_n} E[bid_{n-1}]$
Bidder utility	0	$p(1 - p)^{n-1}$	$p_i \prod_{j=1}^{n-1} (1 - p_j)$
Sum-profit principal utility	1	$1 - (1 - p)^{n-1}(1 + p(n - 1))$	$1 - \lambda(1 + \sum_{i=1}^{n-1} p_i)$
Max-profit principal utility	$\frac{n}{2n-1}$	$\frac{n}{2n-1} + \frac{n-1}{2n-1}(1 - p)^{2n-1} - (1 - p)^{n-1}$	$\frac{n}{2n-1} - \lambda + \sum_{k=1}^{n-1} (\frac{(1-p_k)^{2n-2k-1} \prod_{j=1}^k (1-p_j)^2}{4(n-k)^2 - 1})$

Table 8.1: The values, in expectation, of some of the variables in a full participation setting; when all bidders have the same participation probability; and when each member has their own participation probability.

essarily do so may serve as enough of an incentive for them (naturally, this would change equilibrium strategies, though lack of knowledge may not allow an optimal strategy at all).

For auctioneers, this analysis presents a clear preference for diligent participants, but in the case such do not exist, their particular preferences become much more intricate, as their specific profit functions indicate (see Table 8.1).

Part III

Networks

*No man is an island, entire of itself; every
man is a piece of the continent, a part of the
main. If a clod be washed away by the sea,
Europe is the less, as well as if a
promontory were, as well as if a manor of
thy friend's or of thine own were: any man's
death diminishes me, because I am involved
in mankind, and therefore never send to
know for whom the bells tolls;
it tolls for thee.*

JOHN DONNE, *Devotions upon Emergent
Occasions*

Chapter 9

Networks Overview

In many historical overviews of graphs and graph theory it is noted that they have been used to represent objects in the real-world only since the 18th century, when Euler used them to represent the physical structure (and bridges) of the city of Königsberg. Since then, of course, their use has grown and extended to various purposes, from maps and roads to computer networks and abstract mathematical structures. However, in one sense, they have been used for much longer than 300 years: the use of graphical means to represent familial and social relationships via family trees of royal families or of religiously important figures has existed for centuries. In times in which a link to great kings and figures of the past was an important political tool, such trees were needed to strengthen a person's claim to a position of power, like the throne.

Using graphs to denote social connections in society enables us to take into consideration issues which have not been the focus of the previous parts – the way people cooperate and how they form groups.¹ This is a wide and varied field, and we have chosen to showcase here two very different approaches to utilizing this information.

One approach is using *cooperative game theory* (also commonly referred

¹Chapter 7 dealt with the possibility of cooperation in the form of player mergers and collusions, but did not analyze its formation.

to as transferable utility games), which delve into the structure of group formations, and focuses on their stability and the ability to maintain them. We chose a to implement them on a particular graph structure, which enables us both to explore using the network representation of various settings, including non-social ones. We pay particular attention to fleshing out the precise calculation of the amount of value that needs to be added to a particular setting to make it stable. That is, helping designers understand the limitation of their scenario, and aiding the analysis of how to make it stable, and the amount needed to be invested to make it so.

This desire, to let system designers, faced with the complexity of human interaction, understand their limitations – and their options – also permeates the second approach we present here. Making use of the *axiomatic approach*, we start off by enumerating various desirable properties. We then explore which of these properties cannot co-exist together, and what algorithms may implement some different combination of these properties. The setting we deal with is the world of group recommendations: it is well established that people prefer the advice of their family and friends, and therefore relying on the social network, and understanding the preference of those connected to those whom we wish to make a recommendation to, is key.

In both approaches, despite their significant differences, we maintain the importance of the network structure to the construction of solutions, as the data represented by the graph is fundamental to understanding any interaction in these settings. Furthermore, we wish to give a system designer practical tools when approaching such problems – encouraging realistic, stable results.

Chapter 10

“You are the Weakest Link, Goodbye!”: Weakest Link Games

10.1 Introduction

Consider a travel agency preparing to offer a fixed-price travel deal. The deal must include a flight to a travel destination and a hotel stay. People deciding whether to take the deal or not would examine the hotel that is being offered, and are only likely to take the package if the hotel's quality is sufficient for their taste. Similarly, if the airline's quality is not high enough, people are likely to reject the deal. A potential buyer would reject the package when either the hotel or the airline do not have the required quality. Thus the total number of buyers, and the agency's revenue, is determined by the *weakest* part of the package.

Alternatively, consider a truck driver who wants to deliver as much cargo as possible from New York to Los Angeles. Even if the truck can carry all available cargo, any path from the source to the target involves using toll roads with bridges and tunnels, each limiting the weight or height of vehicles going through them. Any road that is used places a restriction on the load

the truck can carry when passing through it. Any possible path between the source and target consists of several such roads, and is limited by its weakest link (i.e., the road with the most stringent restrictions along the path). The optimal path is the one with the best weakest link, as it allows the most amount of cargo to be transferred.

Less geographically oriented, one could consider a manufacturing process that takes various materials, and applies multiple transformations to produce a desired product. A complex manufacturing process may have several stages, and there may be several alternative methods that lead to the same final product. Each manufacturing stage has a certain environmental impact (for example, pollution with a negative impact on the environment or resulting in a perimeter of a certain distance from the factory that needs to be cleared of people), and we seek to find a manufacturing process that has the minimal negative impact, and to incentivize firms to use it over alternative, more harmful, methods.

In the above examples, the package's value depends on its weakest component. However, individual components can be composed into various packages, in ways captured by certain graph structures. If these components are controlled by self-motivated agents, how are the agents likely to share the package's total value? For example, which travel packages are likely to form? How would the toll road owners, or the hotel and airline providers, share the obtained revenues?

Many domains where self-motivated agents interact have been studied in the algorithmic game theory literature. But this setting calls for a strong collaboration between agents and negotiations between them, which are the particular strengths of cooperative game theory. In such domains, having enforceable contracts between the agents has an important impact on the equilibrium outcome that emerges.

The need for computationally tractable game-theoretic concepts is highlighted by the applicability of a “weakest-link” model in common tasks such as crowdsourcing and large projects, which are typically comprised of several

parts, while the overall quality may depend primarily on the lowest-quality part. Hence, we propose a new class of cooperative games, called cooperative *Weakest Link Games* (WLGs), which capture domains (such as the examples above) where the value a coalition can achieve is determined by its weakest member.

Our weakest link model makes use of an edge-weighted graph with designated source and target vertices, where the agents are the edges of the graph. The quality of a path from the source to the target is the minimal edge weight along the path; the value of an agent coalition is the maximal quality of all the paths contained in the coalition (i.e., all the paths that are comprised of edges that are all in the coalition).

In some cases, the need for all to cooperate may be necessary: for example, in the environmental example above, we may desire that the more harmful methods not be used. Hence, we seek a global stability. We provide a polynomial algorithm for computing the value of a coalition in a weakest link game, and show various properties and algorithms for solutions based on team stability, including calculating and quantifying the Cost of Stability [26] which measures the minimal external subsidy required to allow stable payoff allocations to exist.

In other cases, the partition of the game may result in a better outcome than a global cooperation. We explore the problem of finding the best partitioning of the agents to teams (optimal coalition structure). Though we show the problem is NP-complete, we provide a polynomial $\mathcal{O}(\log n)$ approximation for it.

10.1.1 Related Work

Cooperative game theory (see [179, 67] for a survey of cooperative games) has been used to analyze an increasingly diverse set of problems, from negotiations [172, 225, 33], voting power [210, 233, 37, 244, 208, 117] and manipulations [25, 18, 245, 140], and proof system analysis [52] to crypto-currency protocols [3, 146], shipping optimization [218], and information aggregation

[28]. More connected to our model is work having to do with resource management – be it water [82, 183], or cloud computing [173, 62] – and with cost and revenue sharing in projects [242, 135, 5, 51, 41, 34, 48, 237, 169].

The Weakest Link Game (WLG) is a class of cooperative game, and similarly to other classes such as [212, 125, 79] or cooperative game languages such as [47, 20, 92, 23, 46, 42, 19, 64, 69, 68], it is based on a graph representation, where agents control parts of the graph. However, the value function of weakest link games differs from all of these other forms. In flow games [125, 192] a coalition’s value is the maximal flow it allows between source and target, so a coalition always gains by adding another path. In contrast, in WLGs a coalition’s value is determined by a single path, so it gains nothing from adding a path unless it is better than even the best path already in it. In graph games [79] the agents are vertices, and the coalition’s value is the sum of the edges occurring between coalition members, as opposed to WLGs where we examine paths between two specific vertices.

Weakest link games are somewhat reminiscent of Connectivity Games [44], where agents are vertices and a coalition wins if it contains a path from the source to the target. WLGs are also based on paths from a source to a target, but the agents in them are the edges. Further, in WLGs the graph is weighted, and the value of a coalition depends on these weights through a max-min structure. Other forms also have different network goals from WLGs: finding an optimal project or matching [212, 16], spanning a set of vertices [20], or interdicting paths [42].

In particular, weakest link games are an instance of sub-additive games, which have been widely explored in the literature [67]. These include weighted voting games [91, 40, 90], skill games [45, 39], and MC-nets [67]. All have been explored both with respect to their cores and Cost of Stability, as well as with respect to finding optimal coalition structures for them.

The solutions we focus on are the core [111], ϵ -core and least-core [211]. The core was proposed as a characterization of payoff allocations where no agent subset is incentivized to deviate from the grand coalition and work on

its own [111]. One limitation of the core is that it can be too restrictive, as in some games no imputation fulfills its requirements. Such games can be solved by the more relaxed solutions of the ϵ -core and least-core. *Cost of Stability* (CoS), the minimal subsidy that allows stable agreements, was proposed in [26, 158] to model domains where an external party wishes to increase cooperation by offering a subsidy.¹

One key area in algorithmic game theory is team formation, and the problem of optimal coalition structure generation was widely studied [213, 205, 190, 209, 35] along with its applications, ranging from vehicle-routing tasks to sensor networks, as well as its relation to other solutions [113]. Though even restricted versions of the problem are hard [241, 205], exponential algorithms and tractable approximations have been proposed [213] and studied empirically [139].

Arguably, the state of the art method for general games [190] has a reasonable runtime on average cases, but has a worst case runtime of $\mathcal{O}(n^n)$. Many such algorithms use an oracle for computing the value of a coalition, in contrast to our approximation which relies on the restricted WLG representation. Another method [39] relies on a different representation called coalitional skill games [45], which is based on set-cover domains.

10.2 Preliminary: Cooperative Game Theory

A coalitional game is comprised of a set of n agents, $N = \{1, 2, \dots, n\}$, and a characteristic function mapping agent subsets (coalitions) to a real value

¹Various perturbations to the solutions of cooperative games have been studied, including modifying agent goals [85, 93], which is not relevant for our setting; introducing stochastic agent failures [102, 38, 49, 31] (somewhat akin to the model for agents' failure to join an all-pay auction described in Chapter 8); allowing agent manipulations and false-name attacks [238] or taxing or bribing agents [195, 243]. We have chosen to focus on the more central and widely recognized ones.

$v : 2^N \rightarrow \mathbb{R}$, indicating the total utility these agents achieve together.

There are several very common assumptions in cooperative game, which also apply in our case:

Emptiness $v(\emptyset) = 0$.

Positivity For all $C \subseteq N$, $v(C) \geq 0$.

Monotonicity For $C \subseteq D \subseteq N$, $v(C) \leq v(D)$.

A solution for cooperative game is one that both divides the parties into coalitions, as well as the value of each coalition between its participants. It consists of:

Coalition structure $CS = \{C_1, \dots, C_m\}$ for all $1 \leq i \leq m$, $C_i \subseteq N$; for all $1 \leq i \neq j \leq m$, $C_i \cap C_j = \emptyset$; and $\bigcup_{i=1}^m C_i = N$.

Payoff vector $p \in \mathbb{R}^{|N|}$ such that for all $1 \leq i \leq m$, $\sum_{j \in C_i} p_j \leq v(C_i)$.

We shall require at least a basic criterion from the value division: that it be *individual rational*, i.e., for all $i \in N$, $p(i) \geq v(\{i\})$ – no agent is better off on its own. Such divisions are called *imputations*.

Furthermore, we are particularly concerned with the grand coalition – the functioning of all agents together in a single coalition (i.e., solutions in which $CS = \{N\}$). To simplify notation, when discussing an imputation (p_1, \dots, p_n) , we denote the payoff of a coalition C as $p(C) = \sum_{i \in C} p_i$.

10.2.1 Stability Concepts

The concern for the stability of the grand coalition is a common concern in cooperative game theory, incorporating several solution concepts:

Core and Least-Core

A coalition $B \subset N$ blocks the payoff vector (p_1, \dots, p_n) if $p(B) < v(B)$, since B 's members can split from the original coalition, derive the gains of $v(B)$

in the game, and give each member $i \in B$ its previous gains p_i and still each member can get additional utility. Under a blocked payoff vector, the coalition is unstable. A solution based on this is the core [111].

Definition 10.1. *The core of a game is the set of all imputations (p_1, \dots, p_n) that are not blocked by any coalition, so that for any coalition $C \subseteq N$, we have: $p(C) \geq v(C)$.*

In some games, every imputation is blocked by some coalition, so the core can be empty. As the core is too restrictive, one possible alternative is to use relaxed stability requirements. One such model is based on the assumption that coalitions that have only a small incentive to drop-out from the grand coalition will not do so – the ϵ -core [211].

Definition 10.2. *The ϵ -core, for $\epsilon > 0$, is the set of all imputations (p_1, \dots, p_n) such that for any coalition $C \subseteq N$, $p(C) \geq v(C) - \epsilon$.*

Unlike the core, the ϵ -core always exists for a large-enough ϵ . For the value $\epsilon = \max_{C \subseteq N} p(C) - v(C)$ the ϵ -core is always non-empty. The set $\{\epsilon | \epsilon\text{-core is non-empty}\}$ is compact, and thus has a minimal element. The minimal value ϵ^* for which the ϵ -core is non-empty is called the least-core value of the game, and the ϵ^* -core is called the *least-core*.

Cost of Stability

When the core is empty, an external party interested in having the agents cooperate may offer a subsidy if the grand coalition is formed. This increases the total payoff, but does not change the core constraints, so when a large-enough subsidy is given, the perturbed game has a non-empty core. The minimal subsidy required to achieve a non-empty core can measure the degree of instability or the agents' resistance to cooperation, and is called the Cost of Stability [26].

Definition 10.3. *A game's Cost of Stability (CoS) is the minimal external subsidy that allows the game to have a non-empty core. Formally, given a*

game with characteristic function $v : 2^N \rightarrow \mathbb{R}$, the modified game v_Δ is the game with the characteristic function $v' : 2^N \rightarrow \mathbb{R}$ where $v'(N) = v(N) + \Delta$ and for every $C \subset N$ we have $v'(C) = v(C)$ (v' is a super-imputation, which is an imputation in which $v'(N) \geq v(N)$). The CoS is the minimal Δ such that v_Δ has a non-empty core.

10.2.2 Coalition Structures

In certain domains several disjoint agent coalitions may emerge, each working independently, creating a structure of coalitions [67]. When the same characteristic function $v : 2^N \rightarrow \mathbb{R}$ determines the utility obtained by each such coalition, we may seek the optimal partition of the agents maximizing the total value obtained. This problem is called the optimal coalition structure generation problem [205, 139].

Definition 10.4. *A coalition structure is a partition CS of the agents (N) into several disjoint sets (C_1, \dots, C_m) . The total value of a partition is the sum of the values of the parts, so $v(CS) = \sum_{i=1}^m v(C_i)$. The optimal coalition structure is the partition with the maximal value: $\arg \max_{CS} v(CS)$.*

10.3 Weakest Link Game Definition

Weakest Link Games (WLGs) model domains such as the examples given above, using an underlying graph structure. A Weakest Link Domain consists of a graph $G = (V, E)$ with designated source and target vertices $s, t \in V$, and an edge weight function $w : E \rightarrow \mathbb{R}_+$ mapping any edge to the “restriction” applied on it (the set W includes all different weights in the graph).

We denote the set of all paths between s and t as $R_{(s,t)}$. The value of a path $r = (e_1, \dots, e_m) \in R_{(s,t)}$, where (e_1, \dots, e_m) are the edges along the path, is the minimal edge weight along this path: $q(r) = \min_{e_j \in r} w(e_j)$. In other words, a chain of edges forming a path is only as strong as its weakest link. Given an edge subset $C \subseteq E$, we denote the set of s - t paths that consist

only of edges in C as $R_{(s,t)}^C = \{r = (e_1, \dots, e_m) \in R_{(s,t)} \mid e_j \in C \text{ for } 1 \leq j \leq m\}$.

Our game is defined over $(G = (V, E), s, t, w)$, where the agents N are the edges in the graph, so $N = E$, and we denote $|N| = |E| = n$. The characteristic function $v : 2^N \rightarrow \mathbb{R}$ maps a coalition $C \subseteq N$ to the value of the best (strongest) path that consists solely of coalition edges.

Definition 10.5. *A Weakest Link Game (WLG) is defined over a domain $(G = (V, E), s, t, w)$ where agents are edges $N = E$, and using the following characteristic function:*

$$v(C) = \max_{r \in R_{(s,t)}^C} q(r) = \max_{r \in R_{(s,t)}^C} \min_{e_j \in r} w(e_j)$$

By convention, if for a coalition $C \subseteq E$ no such path exists (i.e., $R_{(s,t)}^C = \emptyset$) we set $v(C) = 0$.

Intuitively, the value of coalition C is the highest threshold τ such that there exists a path between s and t using only edges in C with weight at least τ .

Example 10.1. *In Figure 10.1 the value of the grand coalition is 3, as that is the value of the weakest link in the path s - A - C - F - H - t (the edge (F, H) is the weakest link). The imputation that gives 1 to the edge (A, C) , 2 to the edge (C, F) , and 0 to all the other edges is in the core.*

Example 10.2. *In Figure 10.2 the value of the grand coalition is 2 – the path s - B - D - G - H - t (due to the edge (B, D)). However, the core is empty, as any imputation needs to have the value 2 on the path s - B - D - G - H - t and the value 1 on the path s - A - C - E - F - t , which shares no edges with the previous path, and therefore needs added value in the (super-)imputation. The CoS is 1, as the super-imputation giving 2 to the edge (G, H) and 1 to the edge (C, E) is stable.*




Figure 10.1: A WLG with a nonempty core

10.4 Core and Least-Core

We now study how agents in a weakest link game are likely to share the gains, focusing on payoff allocations that guarantee stability of the formed team, providing polynomial algorithms for computing core, ϵ -core and least-core solutions.²

By definition, a coalition's value is the weight of the lightest edge in a certain path (weakest link of maximal weight), so $v(C)$ is the weight of one of the edges in the graph, and can take at most $|W| \leq |E|$ different values.

Observation 10.1. *The value $v(C)$ of any coalition C in a WLG over the graph $G(V, E)$ is either 0 or the weight of one of the edges in the graph, so $v(C) \in W = \{w(e) | e \in E\} \cup \{0\}$.*

We use this observation to simplify many of our following calculations:

Theorem 10.1. *Computing the value $v(C)$ of a coalition C in a weakest link game can be done in polynomial time.*

²While for some WLG scenarios there is no need for stability, settings such as our environmental damage one (described in the introduction) are cases of WLGs where we specifically wish to discourage use of the other paths; hence, we require a stable imputation.




Figure 10.2: A WLG with an empty core

(Proof is in Appendix C)

Theorem 10.2. *Testing whether an imputation $p = (p_1, \dots, p_n)$ is in the core of a weakest link game can be done in polynomial time.*

(Proof is in Appendix C)

The algorithm above tests whether an imputation is in the core of a WLG. We now show that relaxed solution concepts can also be computed in polynomial time, using this algorithm as a building block.

We note that it is possible to construct a linear program (LP) with n variables [101], whose set of solutions are all the ϵ -core imputations. This LP, shown in LP 10.1, has a variable p_i for each of the agents, which represent its payoff in an imputation. The LP has 2^n constraints, one per possible coalition.

The ϵ -core is the solution to the LP, and the core is recovered when setting $\epsilon = 0$.

Similarly, the Cost of Stability is characterized by the LP given in LP 10.2, using the additional variable Δ designating the external subsidy which perturbs the value of the grand coalition.

Although it is possible to solve these LPs using the Ellipsoid method in

$$\begin{aligned} \min \epsilon \text{ s.t.:} \\ \forall C \subset N : \sum_{i \in C} p_i \geq v(C) - \epsilon; \\ \sum_{i \in I} p_i = v(I) \end{aligned}$$

LP 10.1: Linear program for the core and ϵ -core

$$\begin{aligned} \min \Delta \text{ s.t.:} \\ \forall C \subset I : \sum_{i \in C} p_i \geq v(C); \\ \sum_{i \in N} p_i = v(I) + \Delta \end{aligned}$$

LP 10.2: Linear program for the CoS

time polynomial in the size of the LP [101], we note that the size of the above LP formulations are exponential in the number of players.

Our solution to this problem uses a separation oracle, a method that takes a possible LP solution as an input and either finds a violating constraint or verifies that no such violating constraint exists. Since the Ellipsoid algorithm can run using only a separation oracle, without explicitly writing the entire LP, finding a polynomial separation oracle for an LP enables solving it in polynomial time.

Theorem 10.3. *Testing core emptiness, finding an ϵ -core imputation and finding the least core value take polynomial time for WLGs.*

(Proof is in Appendix C)

Corollary 10.1. *Calculating the Cost of Stability of a weakest link game can be done in polynomial time.³*

(Proof is in Appendix C)

³In Section 10.5 we propose a linear time algorithm for computing the CoS of WLGs for the restricted case where the underlying graph is a series-parallel graph.




Figure 10.3: Parallel composition of a graph

10.5 Cost of Stability and Series and Parallel Compositions

Having examined the general case, we provide results on how a graph's composition affects the stability of the game in our model, which rely on series-parallel graphs [84, 222].

A *two terminal graph* (TTG) is a graph with a distinguished source vertex and a distinguished target vertex. A *base graph* is a TTG that consists of a source vertex and target vertex connected directly by a single edge (i.e., the graph K_2). The parallel composition $P(G_1, G_2)$ of TTGs G_1, G_2 is the TTG generated from the disjoint union of G_1, G_2 by merging the sources of G_1, G_2 and merging their targets (Figure 10.3). The series composition $S(G_1, G_2)$ of TTGs G_1 and G_2 is the TTG generated from the disjoint union of G_1, G_2 by merging the target of G_1 with the source of G_2 (Figure 10.4).

Definition 10.6. A Series Parallel Graph (SPG) is a TTG formed by applying a sequence of parallel and series compositions starting from set of base graphs (i.e., a graph built recursively by the two composition operations over base graphs).

In WLGs, two disjoint s - t paths (i.e., parallel s - t paths) are substitutes, as either path may be used to reach the target from the source. In contrast, two disjoint edge subsets of a single simple s - t path, such as two sub-paths that are joined serially to form a full s - t path, are complements, as both




Figure 10.4: Serial composition of a graph

parts are required.

Intuitively, we expect complement agents to find it easier to cooperate, as they need each other to achieve a high value, whereas substitute agents resist cooperation as each group can achieve value on its own. We formalize and quantify this intuition using SPGs, where complementarity and substitution are easily captured by the graph's structure.

Though WLGs are defined for any graph, the restricted case of SPGs captures very natural structures: a series composition in a WLG indicates that a project has two parts and its overall success depends on the weaker component; a parallel composition indicates that either part can be used to complete the project. The travel packages example in the introduction is a direct example of an SPG domain.

We show how the resistance to cooperation, measured by the CoS, is affected by series and parallel composition. In a WLG setting, when joining graphs $\{G_i\}$, the characteristic function of the newly-formed SPG (v) can be expressed in terms of the characteristic functions of the joined graphs $\{v_i\}$: for every $C \subseteq G$, $v(C \cap G_i) = v_i(C \cap G_i)$.

Theorem 10.4. *If graph G is a parallel composition of graphs G_i , the CoS of G is $(\sum_{G_i} CoS(G_i) + v_i(G_i)) - \max_{G_i}(v_i(G_i))$.*

(Proof is in Appendix C)

Theorem 10.5. *If G is a series composition of the graphs G_i , the CoS of G is $\min_i CoS(G_i^{\min_{j \neq i}(v(G_j))})$, where $G_i^{\min_{j \neq i}(v(G_j))}$ is G_i in which all edges with weight above $\min_{j \neq i}(v(G_j))$ are lowered to that value.*

(Proof is in Appendix C)

The above theorems yield a polynomial method to compute the CoS in SPGs by recursively applying formulas on the graph's structure (CoS of a base graph is 0).⁴

10.6 Optimal Coalition Structures

The optimal coalition structure is a partition of the agents into disjoint sets that maximizes the sum of the values of the parts. Each such part has a non-zero value only if it contains some s - t path. If a single part of the partition contains more than one s - t path, it could be broken down into two sub-parts, each containing a path, which results in a higher value. Thus it seems that finding the optimal coalition structure is somewhat related to a decomposition of the agent set into sets of disjoint paths. Indeed, we first show that finding the optimal coalition structure is NP-hard using a reduction from the Disjoint Paths Problem (DPP).

Theorem 10.6. *It is NP-complete to determine whether the value of the optimal coalition structure in a weakest link game exceeds an input k .*

(Proof is in Appendix C)

We propose a polynomial approximation for this problem.

Theorem 10.7. *A polynomial time $\mathcal{O}(\log n)$ -approximation exists for the optimal coalition structure problem in weakest link games.*

(Proof is in Appendix C)

10.7 Summary

This chapter has approached the problem of weakest link games, which can be used to model various real-world scenarios using cooperative game theory.

⁴The value of a coalition in a WLG can be computed in polynomial time, and CoS of any base graph is zero.

To analyze such a setting we use cooperative game theory, which, like our previous game theoretical analysis, often strives to reach stability. However, it focuses on cooperation between agents in order to reach a particular goal. In our case, cooperation is the key element, making this the most applicable approach for weakest link games: as reaching from the source to the target is our goal, cooperation between our agents (the edges) is crucial to the setting.

Focusing on setting such as our environmental one, in which we wish to minimize the environmental effect of a process, we want to find a way to divide resources so that even unused agents are incentivized to stay unused (and therefore, non-polluting). Such a division requires a non-empty core, and we delve into the possibility of finding one, and moreover, when the core is empty, we wish to find the minimal amount the planner needs to contribute to the system in order to make it stable.

Looking at other scenarios, in which we just wish to increase the flow in the graph, we find that finding an optimal structure is NP-complete, yet are able to present an approximation algorithm for it, enabling, in many cases, a “good enough” cooperative solution.

Chapter 11

Where Do We Want to Go Today?: Group Recommendations

11.1 Introduction

The online world encourages the development of systems that allow people to effectively use it. Reputation systems, ranking systems, trust systems, recommendation systems, affiliate marketing in social networks, and more, are flowering in its midst. This recent wave of online social systems is typically associated with a large amount of data that is collected online which leads to the “big data” approach to the utilization of such information. However, the abundance of available data does not help system designers to come up with the right design for online systems in the first place. Indeed, available data is typically generated by the use of a particular system, and mining the data generated by users while interacting with one system does not provide a tool for exploring the overwhelmingly large design space. Interestingly, the main practical approach of software and hardware design, the formal specification of clear system requirements and the implementation of a system satisfying these exact requirements, has not been used often.

This classical approach, when adapted to the context of multi-agent systems, coincides with extensions of a famous tool of social choice theory and cooperative game theory, namely the *axiomatic approach*, of which one of the best known examples is the theory of social choice, in which we aim to find good aggregation of individual preference into a global ranking. A bit removed from direct extensions of social choice, one can find systems originating from personalized versions of ranking and reputation systems. Here we no longer consider the aggregation of preferences into a shared ranking, but instead seek to provide personalized rankings or recommendations to each agent.

A fundamental challenge in this context is the search for effective *trust-based recommendation systems*, in which – based on trust-relationships among the agents and the expressed opinions of a subset of them about a service or a product – a recommendation about a service or a product is provided to agents which did not evaluate it personally. The puzzling challenge of generating useful trust-based recommendation systems is amenable to an axiomatic treatment, beginning with an attempt to characterize the systems satisfying different sets of desired properties.

Here we significantly expand the body of work on the axiomatic approach for internet settings by initiating work on the axiomatic treatment of *Group Recommendation Systems*. We assume a trust-graph in which agents express who they trust, and information is provided about the opinions of some of the agents about a product/service, but we care about providing recommendations to a *group* of agents, rather than a single one (e.g., a party of friends, looking for a restaurant). Notice in this case the group may include some agents who have experienced the service/product directly and some who may have not. This can be viewed as a bridge between social choice (aggregating individual preferences) and trust-based recommendation systems. In addition to its theoretical importance, the topic of group recommendation systems is of great practical interest.

While group trust-based recommendation systems are vastly different

from (individual) trust-based recommendation systems, for reasons of exposition we will do our best to connect to the literature on the latter, adopting properties taken as axioms for individual trust-based recommendation systems in earlier work, and adding to them properties reflecting the fact we deal with group recommendations. Interestingly, putting these together leads to a powerful and illuminating impossibility result; the axioms/properties are all essential for this impossibility, as removing each one of them leads to possibility. Given this impossibility, we replace our three group related axioms by three other properties; in a second result we show an extension of a random walk system that satisfies all desired properties, and moreover, is the *only* system satisfying these properties. Together, this provides rigorous foundations to a theory of group recommendations.

11.1.1 Related Work

Most approaches to individual recommendation systems [193] (i.e., systems not attempting to recommend to groups) have proposed a model based on their observations of recommendation dynamics in real life, without setting out to achieve any particular mechanism behavior. Much work is devoted to collaborative filtering, in which an agent receives recommendations based on the views of agents with similar properties – similar to the design of the Netflix challenge [57, 114] and other projects [72]. Other work focuses on simulations and field experiments in various fields [207, 194, 17] (including movies [168], music [239, 231] and tourist recommendations [198, 24]) or improving the running time of large-scale recommender systems [240, 149, 43, 30, 29, 27]. Some models add a social graph to the recommendation system, supporting a different trust level for each agent. However, these works [230, 181, 232] use the social graph as a mechanism which propagates true, objective information to the agents, and do not consider agents’ recommendations as opinions which may depend on taste (and hence, have no fixed value of “trustworthiness”).

In the past few years, more research has been devoted to group recommen-

dations, as the scenarios where group recommendation are useful are more and more evident. Early work has simply aggregated all members' preferences [178, 105], but the common approach tries to implement a model which adds a layer of complexity beyond agents' approval of a choice, adding also the measure of the rejection of a choice by each participant [11, 130, 107]. Unlike others, Gartrell et al. [107] try to utilize the social graph for their model, but ultimately their approach is limited as it utilizes the social graph only to propagate information through it, and it is not used in the actual recommendation system.

A different approach to recommendation systems is the axiomatic one, which seeks to first describe the goals of a system, and then to find the systems that implement such goals. Such an approach has been taken in ranking systems [6, 8], including detailed analysis of specific mechanisms [7] including collaborative filtering [184]. More importantly, it has been applied to the individual recommendation system problem, first in Andersen et al. [12], and following that, in additional papers complementing it [202, 63]. One of the key strengths of this research path is in its basic model, which incorporates the influence of the social graph on agents' behavior. Alas, these papers do not deal with group recommendations, and hence with the particular needs and desired properties of this problem.

11.2 The Group Recommendation Model

The basic model (adapted from [12, 202, 63]), deals with a graph that has opinionated nodes (or voters) over some option – some are + nodes (agents which like the option), and some are - nodes (the agents which did not like the option). The rest of the nodes are nonvoters, i.e., they have no predetermined opinion. A directed edge (a, b) indicates that the agent b influences agent a 's opinion to some degree. We wish to find a mechanism that takes any group of agents (both voters and nonvoters) and gives the members of the group a single recommendation – +, -, or 0 (in case of inability to recommend).

Formally:

Definition 11.1. A voting network is a directed graph $G(N, V_+, V_-, E)$, where N are the nodes, $V_+ \subseteq N$ are the nodes which vote +, $V_- \subseteq N$ are nodes which vote -, and E are directed edges, in which parallel edges are allowed,¹ but not self loops. We say node $b \in N$ influences node $a \in N$ when there exists an edge (a, b) .

From this definition we can derive the group of voters – $V_+ \cup V_-$ and nonvoters – $N \setminus (V_+ \cup V_-)$.

Definition 11.2. A group recommendation system is a function $R_G : 2^N \rightarrow \{+, -, 0\}$, assigning a recommendation to each subset of graph nodes in the graph G .

Before proceeding to the axioms' definitions, we define our group random-walk recommendation system variant. For this we need to first define an individual random walk recommendation system, which, basically, assigns to each node the sign of the weighted average of all voters which are reachable from it.

Definition 11.3. An individual random walk recommendation system takes a voting network $G(N, V_+, V_-, E)$ and assigns each node $a \in N$ a value r_a : if $a \in V_+$ (respectively, $a \in V_-$), then $r_a = 1$ (respectively, -1). If a is a nonvoter which does not have a path to any voter, $r_a = 0$. If it does have paths to other voters, we look at the group $\text{succ}_a = \{b \mid (a, b) \in E\}$, and define $r_a = \frac{\sum_{b \in \text{succ}_a} r_b}{|\text{succ}_a|}$. Once calculated r_a (based on other r_i s), the recommendation system recommends $\text{sgn}(r_a)$ (i.e., sign of the number).

This recommendation system is, in effect, taking random walks from the recommended node, stopping at a voter. The sign of expected value of the random walks is the recommendation.

¹Equivalent to using weights, but easier to analyze in our case. This means E is, in effect, a multiset.

Group random-walk, in a sense, calculates a value for each group member based on their individual recommendations, and sums over all of the group members.

Definition 11.4. A group random walk recommendation system takes a group $C \subseteq N$ and for each $c \in C$ assigns r_c to be the value (not just the sign) of the same node under the individual recommendation system (so $-1 \leq r_c \leq 1$). It then returns $\text{sgn}(\sum_{c \in C} \text{sgn}(r_c))$.

11.3 The Axioms

The axioms and their formulation are key to this approach, and therefore we expand on their motivation and intuitive understanding, and include the formal definitions after describing a group of several axioms.

11.3.1 The Basic Axioms

These axioms were adapted from the individual recommendation case, and are quite basic, so that we believe most general-use systems which rely on the social graph for their recommendations would seek to implement them:

1. *Anonymity – No node is special.* Isomorphic graphs (including recommendation isomorphism – symmetry between + and - votes) have isomorphic recommendations.

Axiom 1. Anonymity: Let $G(N, V_+, V_-, E)$ be a voting network, and R a recommendation system. For any permutation $\pi : N \rightarrow N$ and G' , the isomorphic voting network under it, for any $C \subseteq N$ $R_G(C) = R_{G'}(\pi(C))$. Furthermore, For $G''(N, V_-, V_+, E)$ and $C \subseteq N$, $R_G(C) = -R_{G''}(C)$.

It's natural that if a group gets a certain recommendation (w.l.o.g, +), and a + voter joins the group or gets additional influence over it, the recommendation should not change.

2. *Positive Response – Adding support for a recommendation cannot reverse it.* A group which is recommended + and a + voter is added to it (or begins to influence one of its members) does not change its recommendation. If a group is recommended 0, adding a + voter to the group changes the vote to +. Furthermore, adding a + voter and a - voter (both not in the group), both influencing the same node in the group does not change the recommendation.

Axiom 2. Positive Response: Let $G(N, V_+, V_-, E)$ be a voting network, R a recommendation system, $C \subset N$, and $a \in V_+ \cap \{N \setminus C\}$ such that there is no edge $(c, a) \in E$ for any $c \in C$. Then if $R_G(C) = +$ or $R_G(C) = 0$, then both $R_G(C \cup \{a\}) = +$, and if we define G' as G with an added edge (c, a) for some $c \in C$, $R_{G'}(C) = +$.

Furthermore, let $b \in V_- \cap \{N \setminus C\}$ such that there is no edge $(v, b) \in E$ for any $v \in N$ nor edges $(v, a) \in E$ for any $v \in N$ (i.e., a and b are isolated). For any $c \in C$ define $G''(N, V_+, V_-, E \cup \{(c, a), (c, b)\})$ then $R_G(C) = R_{G''}(C)$.

As we are investigating members of a social graph, it makes sense to ignore nodes that are not in the same connectivity group as members of the group, hence:

3. *Independence of Irrelevant Stuff (IIS) – Unrelated nodes do not affect recommendation.* A group's recommendation is only dependent on the nodes that it influences. Voters are not, of course, influenced by any edges, as their opinion is already set, so removing their outgoing edges has no effect.

Axiom 3. IIS: Let $G(N, V_+, V_-, E)$ be a voting network, R a recommendation system, and $e \in E$ an edge (a, b) for a voter a (i.e., a is influenced by b). For $G'(N, V_+, V_-, E \setminus e)$, for every $C \subset N$, $R_G(C) = R_{G'}(C)$. Furthermore, if $d \in N$ is not reachable from C , then let $E_d \subseteq E$ be the set of edges outgoing or incoming from d , and define $G''(N \setminus d, V_+ \setminus d, V_- \setminus d, E \setminus E_d)$, then $R_G(C) = R_{G''}(C)$.




Figure 11.1: A star group example

11.3.2 Group Power Axioms

We now turn to axioms which try to portray the unique properties of a group recommendation system. In such systems, we want the group members to have a larger influence on the decision than external agents, though we do wish to allow external influence in some cases. In order to portray these two, somewhat conflicting, desires, we define very limited axioms, only on a particular structure of social graph – *star groups*, which (as shown in figure 11.1), are made of a certain type of voters in the group (w.l.o.g, +), which influence all of the group’s nonvoters, with - voters influencing these nonvoters from outside the group.

Definition 11.5. *A group C with n voters $\{v_1, \dots, v_n\}$ and m nonvoters $\{u_1, \dots, u_m\}$ is a star group (e.g., Figure 11.1) if:*

- *All voters $v_j \in C$ have the same label (w.l.o.g., +)*
- *For every nonvoter $u_i \in C$ and voter $v_j \in C$ there exist an edge (u_i, v_j) (so each nonvoter is connected to all voters)*
- *Every nonvoter $u_i \in C$ has an associated group, $D^i = \{t \mid \text{there is an edge } (u_i, t) \text{ for voter } t \notin C\}$. For every i , all D^i members are labeled -, and for every i, h : $D^i \cap D^h = \emptyset$*
- *Nonvoters $u_i \in C$ have no other edges.*

4. *α -centripetal* – Members of a group have more influence over the recommendation than voters outside it. If the star-group has k + voters, the recommendation will be + as long as each nonvoter is connected to less than αk - voters.

Axiom 4. *α -centripetal: A recommendation system has some $\alpha \in \mathbb{R}_+$, $\alpha \geq 1$, such that for every star group (whose members vote, w.l.o.g., +) for which for every i , $|D^i| \leq \alpha \cdot n$, the recommendation for the group is +.*

However, we do not want the group to be all powerful. When there are few + voters in the star-group, and many nonvoters and - voters, we would like to see some influence of the outside agents. Thus:

5. *(β, r) -centrifugal* – Agents outside the group may still influence it. If a star-group (whose members vote, w.l.o.g., +) has k + voters, but even more nonvoters – more than r nonvoters for each + voter; and each nonvoter is connected to many - voters – at least βk – the group recommendation would be -.

Axiom 5. *(β, r) -centrifugal: A recommendation system has some $\beta \in \mathbb{R}_+$, $\beta \geq 1$ such that for every star group for which $\frac{m}{n} \geq r$ ($r \in \mathbb{R}_+$) and for which for every i , $|D^i| \geq \beta \cdot n$, the group's recommendation is -.*

6. *Internal consistency* – If all of a group's partitions have the same recommendation, that will be recommendation of the whole. If all subgroups in a disjoint partition of a group of agents, are given the same (non-neutral) recommendation (and there are no contradicting unanimous, disjoint, non-neutral partitions), the whole group will have this recommendation as well.

Axiom 6. *Internal consistency: In a recommendation system R , for every $C \subseteq N$, for some disjoint partition $C = C_1 \cup C_2 \cup \dots \cup C_n$ for*

which $R_G(C_1) = R_G(C_2) = \dots = R_G(C_n) \neq 0$, and if all other similar partitions $C = C'_1 \cup C'_2 \cup \dots \cup C'_n$, for which $R_G(C'_1) = R_G(C'_2) = \dots = R_G(C'_n) \neq 0$ have $R_G(C_1) = R_G(C'_1)$, then $R_G(C) = R_G(C_1)$.

Ultimately, we will show the above six basic axioms are incompatible, and there exist no group recommendation system that can accommodate them. We will now consider three alternate axioms for group recommendations.

11.3.3 Influence Structure Axioms

The final three axioms, are, to a certain extent, a group-recommendation extension of axioms suggested for the non-group recommendation case [12]. They try to consider influence and the way it “moves” through the social connections, such that influence can extend beyond an immediate node (so I may influence one person, and that person may, in turn, influence another).

7. *Trust Propagation – Influence moves along the graph.* If nonvoter b has k edges to nodes influencing it, and node a is influenced by b with k edges, then the edges to b can be replaced by edges to the nodes influencing b .

Axiom 7. Trust Propagation: Consider recommendation system R , voting network $G(N, V_+, V_-, E)$, group $C \subseteq N$, and nonvoters $u, v \in N$ for which the edges leaving v (beside (u, v)) are $(v, w_1) \dots (v, w_k)$ for some $k \geq 1$. Suppose E contains k copies of (u, v) , and we construct $E' = (E \cup \{(u, w_1), \dots (u, w_k)\} \setminus \{(u, v) \cdot k\})$ and $G'(N, V_+, V_-, E')$, then $R_G(C) = R_{G'}(C)$.

8. *Scale Invariance – Influence does not care about units.* Duplicating a node’s outgoing edges (i.e., edges to the nodes influencing it) does not change recommendation.

Axiom 8. Scale Invariance: For a voting network $G(N, V_+, V_-, E)$, and a nonvoter u , the recommendations are identical for $G'(N, V_+, V_-, E \cup E')$ where E' contains k copies of each of u ’s outgoing edges.

9. *Proportional Inclusiveness* – An external influence can be described as a group-member influence. A voter outside a group, connected directly to a nonvoter inside it, has an influence over the group recommendation in proportion to its weight of influence on the nonvoter, and the nonvoter’s influence in the group. Therefore, the recommendation for a group would be the same as for a group that includes also the voters influencing a nonvoter in the original group (with a few adjustments to maintain relative power of group members, see figure 11.2).

Axiom 9. Proportional Inclusiveness: For a voting network $G(N, V_+, V_-, E)$, a group $C \subseteq N$, a nonvoter $u \in C$ and voters $v_1, \dots, v_m \in V \setminus C$ and $v_{m+1}, \dots, v_t \in C$ which are influencing it (i.e., $(u, v_i) \in E$) then the following transformation retains recommendations: Let there be k_i copies of (u, v_i) in E , and s edges $(u, *)$ in E (i.e., $s = \sum_{i=1}^t k_i$). For $1 \leq j \leq s$ we define $N^j = N \setminus \{u\}$ and define $N' = \cup_{j=1}^s N^j \cup \{v_{m+1}\}^{k_{m+1}} \cup \dots \cup \{v_t\}^{k_t}$. For each $1 \leq i \leq m$ we choose k_i nodes of type v_i (there are s copies of these in N'), and mark them $v_i^1, \dots, v_i^{k_i}$. For each N^j we define $C^j = C \setminus \{u\}$ and $E^j = E \setminus \{(*, u), (u, *)\}$ (no edges ingoing or outgoing from u), and tweak it a little: for each $c \in C$ such that there is an edge $(c, u) \in E$, we multiply s times each edge $(c, *) \in E^j$, and add k_i edges (c, v_i) for $1 \leq i \leq t$ (excluding self, of course). We define $E' = \cup_{j=1}^s E^j$, and $C' = \cup_{j=1}^s C^j \cup_{h=1}^m \cup_{r=1}^{k_h} v_h^r$.
Now, for $G'(N', V'_+, V'_-, E')$, $R_G(C) = R_{G'}(C)$.
(See axiom example in Figure 11.2, in which the nonvoter has also been eliminated using trust propagation axiom)

These axioms will prove to singularly define a group variant of the random walk recommendation algorithm.

11.3.4 Independence of Axioms from Each Other

Theorem 11.1. Axioms 1–6 and axioms 1–3, 7–9 (the sets we deal with) are all independent of one another.




Figure 11.2: Example of applying proportional inclusiveness on u : The $-$ node is now inside the group, but since its weight on u 's recommendation was half the total of influences on u , the other nodes in the group are duplicated accordingly.

(Proof is in Appendix C)

11.4 An Impossibility Result

When looking at the axioms detailed above, it is clear that axioms 1 – 3 are imperative for any basic social graph based recommendation system, whether it is an individual or a group one. However, when expanding to group recommendation systems, one appreciates the goal that the nodes inside a group will have some different – stronger – effect than those outside it, but that this effect will not be unbounded, so outside information will not completely ignored. Our axioms 4 and 5 distill this goal in a most clear-cut way (on *star groups*), being intentionally narrow and non-sweeping. To that we add axiom 6, which is an intuitive and desirable consistency requirement.

However, these requirements are not compatible:

Theorem 11.2. *No recommendation system satisfies axioms 1 – 6, i.e., is anonymous, positive responsive, IIS, is internally consistent and is α –*

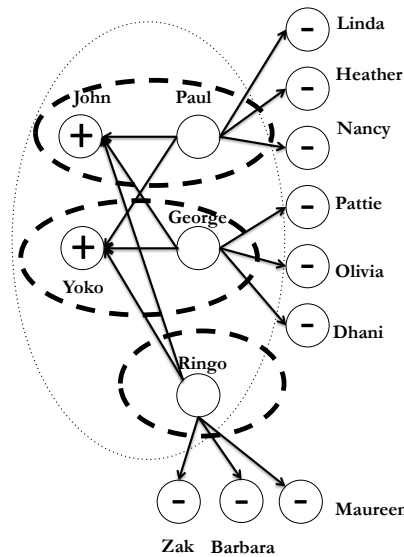


Figure 11.3: The original group (surrounded by light line) and its partition *centripetal* and (β, r) – *centrifugal* for $1 < \alpha, \beta < \infty$.

(Proof is in Appendix C)

Example 11.1. Showcasing the proof's main parts: *A group of friends are constructed as in in Figure 11.3, as John, Paul, George, Ringo and Yoko want to go to a restaurant together. John and Yoko have been there and were extremely satisfied with it. However, all the rest of Paul, George and Ringo's acquaintances have a very negative view about the place.*

Suppose $\alpha = 2$, $\beta = 2.5$ and $r = 1$. According to α -centripetality, the recommendation should be to go the restaurant. However, we can subdivide it into 3 groups (shown in Figure 11.3), which – using positive response (axiom 2) and (β, r) -centrifugality (axiom 5) – should each be advised not to go to the restaurant.

11.5 A Recommendation System that Works

Due to the impossibility result above, we must, of course, give up some of those axioms. We cannot eliminate axioms 1 – 3, as they are fundamental to

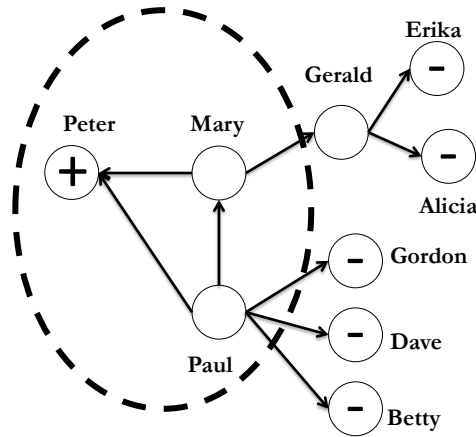


Figure 11.4: Initial state

any social graph based recommendation system. However, we replace axioms 4 – 6 with axioms 7 – 9, which, while they do not allow us the full power of the previous axioms, allow us to show that there is only a single group recommendation system which satisfies them, giving us a potential candidate for a useful, practical system that can be implemented in real-world systems.

Theorem 11.3. *The group random-walk recommendation system is the only one which satisfies axioms 1 – 3 and 7 – 9, i.e., is anonymous, positive responsive, IIS, has trust propagation, scale invariance and proportional inclusiveness.*

(Proof is in Appendix C)

Example 11.2. Showcasing the proof’s main parts: *We alter our example, now discussing Peter, Paul and Mary, which are influenced as in Figure 11.4.*

Using scale invariance and trust propagation, Gerald is no longer connected to Mary. Instead, her attachment to his other influences doubled, and two new connections have been struck (Figure 11.5).

Using the algorithm for our example means Peter contributes 1, Mary contributes 0, and Paul contributes $-\frac{2}{5}$, meaning the total outcome is still positive, and the restaurant will be recommended.




Figure 11.5: Applying scale invariance and trust propagation

11.6 Summary

This chapter has tackled the group recommendation problem: being able to recommend to a group of people based on the opinions of their peers. There is a wide and varied literature indicating that recommendations from family and friends are much better – and more accepted – than other recommendations. This work tries to solve the recommendation problem while utilizing this vital source of opinion, by taking into account the social graph and the relationships it embodies.

Attempting to characterize appropriate properties for our group recommendation systems, we had the basic idea that the group’s members should have more influence than the agents outside it. Quite surprisingly, in our view, this turned to cause an impossibility theorem. Instead, we tried to axiomatize the influence structure between agents, leading us to a unique algorithm implementing these features.

These results help shine a light on two main issues: first, the significance of the social graph, without the social graph, the recommendation system collapses to voting. What makes this a real recommendation system is precisely the use of the social network.

Second, while one may debate our choice of axioms, the need for properties that will stand at the center of the work of a system designer are imperative: without them, in the words of Yogi Berra, “If you don’t know where you’re going, you might not get there” [59].

*If you're gonna play the game, boy, you
gotta learn to play it right:
You've got to know when to hold 'em,
know when to fold 'em;
Know when to walk away,
know when to run.
You never count your money when you're
sittin' at the table;
There'll be time enough for countin' when
the dealin's done...*

KENNY ROGERS (LYRICS BY DON
SCHLITZ), *The Gambler*

Chapter 12

When All is Said and Done...

This work contains, ostensibly, a variety of results that have to do with a game theoretic approach to solving problems, helping the player know what to do in various settings. But looking closer, there are stronger connections between these various results.

Limiting abstractions While the abstraction of the human condition into easily defined problems is the basic building block of any model, due to the sheer complexity of any real-world situation, we are always at the risk of abstracting away key elements. An overly abstract model is what indicates no agent should participate in all-pay auctions; or that all voters voting for their least favorite candidate might be a reasonable stable state. We must always remain anchored in what type of phenomena we are observing in the real world (and trying to model), and beware that our models do not produce meaningless results.

Fundamental cooperation Game theory deals with behavior of agents when in an environment containing other agents. But apart from the truth-bias model (Chapter 3), all the models presented in this work focused on the cooperation and interaction between the different agents, be it through the dynamics of the interaction, the cooperation between them, or their social network, for which cooperation and interaction are intrinsic features.

Stability as a guiding principle Apart from the chapter that used the axiomatic approach (Chapter 11), all chapters contained approaches that strove to reach stability. While a desire for stability is common in many game theoretic analysis, we wish to move from stability on its own, a property describing some states, to stability as realistically describing states which are reasonably reachable by people as they would play the game. In other words, we want to strengthen the link between stable states (usually discovered by analytical work) to end states, as they are reached in real life.

Dynamic analysis One of the main tools which we can wield to make settings resemble real-world situations better, is to analyze the dynamic by which people (or agents) reach their stable state (we do this in voting in Chapters 4,5; in all-pay auctions in Chapter 7; and in designing the group recommendations axioms in Chapter 11). Finding which dynamics are helpful and which are not is part of the research process, and we have attempted here to present some of the more insightful dynamics in the fields we have explored.

Using this toolkit we approach areas for which the current game theoretical results are lackluster – in this work we have shown this on Nash equilibria in voting scenarios and in all-pay auctions, as well as its importance when we add another dimension to the social problems – the one involving networks and the precise relationship between people.

The use of this better modeling is three-fold: first, it enables a better understanding of humans, aiming to capture in each setting the most salient features that enable us to model them better. Second, it enables us to design agents that react and interact with people, understanding, albeit to a limited degree, how people behave in certain settings, and allowing them to anticipate and analyze such behavior to make them more useful and easier to engage with.

Third, better modeling allows us better mechanism design. As has been most explicitly shown in the last chapter (Chapter 11), a better understand-

ing of human behavior – and human expectations – allow us to design better systems which are able to provide desirable properties by having a deeper understanding of the expected behavior of their users.

One of the big, overarching questions that remain in such analysis is how to recognize if our models are, indeed, appropriate. In this work we have demonstrated an attempt to deal with this problem in the realm of elections, by building a large, robust simulation framework, and checking our results on it with various voter distributions (taken from the relevant literature). We have further released this framework to the public, encouraging other researchers to use it to test their (and our) results, and extend it to the various models as they see fit.

Using such simulators, and using our knowledge of nature, we need to find “good enough” abstractions, that allow us, on the one hand, to maintain a clear and coherent model of reality, and on the other, to achieve a good approximation of human behavior. Finding these abstractions, and finding realizable systems that enable us to compare models with real-life, is at the forefront of getting better results from the game theoretic approach and analysis.

תושלבע

Bibliography

- [1] Atila Abdulkadirođlu and Kim-Sau Chung. Auction design with tacit collusion. Working paper, Columbia and Northwestern, 2003.
- [2] Stéphane Airiau and Ulle Endriss. Iterated majority voting. In *Proceedings of the 1st International Conference on Algorithmic Decision Theory (ADT)*, pages 38–49, Venice, Italy, October 2009.
- [3] Robleh Ali, John Barrdear, Roger Clews, and James Southgate. Innovations in payment technologies and the emergence of digital currencies. *Bank of England Quarterly Bulletin 2014 Q3*, September 2014.
- [4] Noga Alon, Michal Feldman, Omer Lev, and Moshe Tennenholtz. How robust is the wisdom of the crowds? In *Proceedings of 24th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2055–2061, Buenos Aires, Argentina, July 2015.
- [5] John A. Aloysius. Research joint ventures: A cooperative game for competitors. *European Journal of Operational Research*, 136(3):591–602, February 2002.
- [6] Alon Altman and Moshe Tennenholtz. On the axiomatic foundations of ranking systems. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 917–922, Edinburgh, Great Britain, August 2005.

- [7] Alon Altman and Moshe Tennenholtz. Ranking systems: The pagerank axioms. In *Proceedings of the 6th ACM conference on Electronic Commerce (EC)*, pages 1–8, Vancouver, Canada, June 2005.
- [8] Alon Altman and Moshe Tennenholtz. An axiomatic approach to personalized ranking systems. *Journal of the ACM*, 57(4):1–35, 2010.
- [9] Ramon Michael Alvarez and Jonathan Nagler. A new approach for modelling strategic voting in multiparty elections. *British Journal of Political Science*, 30(1):55–75, January 2000.
- [10] Erwin Amann and Wolfgang Leininger. Asymmetric all-pay auctions with incomplete information: The two-player case. *Games and Economic Behavior*, 14(1):1–18, May 1996.
- [11] Sihem Amer-Yahia, Senjuti Basu Roy, Ashish Chawlat, Gautam Das, and Cong Yu. Group recommendation: semantics and efficiency. In *Proceedings of the VLDB Endowment*, volume 2, pages 754–765, 2009.
- [12] Reid Andersen, Christian Borgs, Jennifer Chayes, Uriel Feige, Abraham Flaxman, Adam Kalai, Vahab Mirrokni, and Moshe Tennenholtz. Trust-based recommendation systems: an axiomatic approach. In *Proceedings of the 17th international conference on World Wide Web (WWW)*, pages 199–208, Beijing, China, April 2008.
- [13] Nikolay Archak and Arun Sundararajan. Optimal design of crowdsourcing contests. In *Proceedings of the 13th International Conference on Information Systems, ICIS*, Phoenix, Arizona, 2009.
- [14] Kenneth Joseph Arrow. *Social Choice and Individual Values*. Yale University Press, 1951.
- [15] Susan Athey. Single crossing properties and the existence of pure strategy equilibria in games of incomplete information. *Econometrica*, 69(4):861–889, July 2001.

- [16] John Augustine, Ning Chen, Edith Elkind, Angelo Fanelli, Nick Gravin, and Dmitry Shiryayev. Dynamics of profit-sharing games. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 37–42, Barcelona, Spain, July 2011.
- [17] Paolo Avesani, Paolo Massa, and Roberto Tiella. A trust-enhanced recommender system application: Moleskiing. In *Proceedings of the 20th ACM symposium on Applied computing (SAC)*, pages 1589–1593, Santa Fe, New Mexico, March 2005.
- [18] Haris Aziz, Yoram Bachrach, Edith Elkind, and Mike Paterson. False-name manipulations in weighted voting games. *Journal of Artificial Intelligence Research*, 40:57–93, January 2011.
- [19] Haris Aziz, Felix Brandt, and Paul Harrenstein. Monotone cooperative games and their threshold versions. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, volume 1, pages 1107–1114, Toronto, Canada, May 2010.
- [20] Haris Aziz, Oded Lachish, Mike Paterson, and Rahul Savani. Power indices in spanning connectivity games. In *Proceedings of the 5th International Conference on Algorithmic Aspects in Information and Management (AAIM)*, pages 55–67, San Francisco, California, June 2009.
- [21] Moshe Babaioff, Nicole Immorlica, and Robert Kleinberg. Matroids, secretary problems, and online mechanisms. In *Proceedings of the 18th annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 434–443, New Orleans, Louisiana, January 2007.
- [22] Yoram Bachrach. Honor among thieves: collusion in multi-unit auctions. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, volume 1, pages 617–624, Toronto, Canada, May 2010.

- [23] Yoram Bachrach. The least-core of threshold network flow games. In *Proceedings of the 36th international conference on Mathematical foundations of computer science (MFCS)*, pages 36–47, Warsaw, Poland, August 2011.
- [24] Yoram Bachrach, Sofia Ceppi, Ian A. Kash, Peter Key, Filip Radlinski, Ely Porat, Michael Armstrong, and Vijay Sharma. Building a personalized tourist attraction recommender system using crowdsourcing. In *Proceedings of the 13th international conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1631–1632, Paris, France, May 2014.
- [25] Yoram Bachrach, Edith Elkind, and Piotr Faliszewski. Coalitional voting manipulation: A game-theoretic perspective. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 49–54, Barcelona, Spain, July 2011.
- [26] Yoram Bachrach, Edith Elkind, Reshef Meir, Dmitrii Pasechnik, Michael Zuckerman, Jörg Rothe, and Jeffrey S. Rosenschein. The cost of stability in coalitional games. In *Proceedings of the 2nd International Symposium on Algorithmic Game Theory (SAGT)*, pages 122–134, Paphos, Cyprus, October 2009.
- [27] Yoram Bachrach, Yehuda Finkelstein, Ran Gilad-Bachrach, Liran Katzir, Noam Koenigstein, Nir Nice, and Ulrich Paquet. Speeding up the xbox recommender system using a euclidean transformation for inner-product spaces. In *Proceedings of the 8th ACM Conference on Recommender systems (RecSys)*, pages 257–264, Foster City, California, October 2014.
- [28] Yoram Bachrach, Thore Graepel, Gjergji Kasneci, Michal Kosinski, and Jurgen Van Gael. Crowd iq: aggregating opinions to boost performance. In *Proceedings of the 11th International Conference on Au-*

onomous Agents and Multiagent Systems (AAMAS), volume 1, pages 535–542, Valencia, Spain, May 2012.

- [29] Yoram Bachrach and Ralf Herbrich. Fingerprinting ratings for collaborative filtering—theoretical and empirical analysis. In *Proceedings of the 17th International Symposium on String Processing and Information Retrieval (SPIRE)*, pages 25–36, Los Cabos, Mexico, October 2010.
- [30] Yoram Bachrach, Ralf Herbrich, and Ely Porat. Sketching algorithms for approximating rank correlations in collaborative filtering systems. In *Proceedings of the 16th International Symposium on String Processing and Information Retrieval (SPIRE)*, pages 344–352, Saariselkä, Finland, August 2009.
- [31] Yoram Bachrach, Ian A. Kash, and Nisarg Shah. Agent failures in totally balanced games and convex games. In *Proceedings of the 8th Workshop on Internet and Network Economics (WINE)*, pages 15–29, Liverpool, Great Britain, December 2012.
- [32] Yoram Bachrach, Peter Key, and Morteza Zadimoghaddam. Collusion in VCG path procurement auctions. In *Proceedings of the 6th international conference on Internet and network economics (WINE)*, pages 38–49, Palo Alto, California, December 2010.
- [33] Yoram Bachrach, Pushmeet Kohli, and Thore Graepel. Rip-off: playing the cooperative negotiation game. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, volume 3, pages 1179–1180, Taipei, Taiwan, May 2011.
- [34] Yoram Bachrach, Pushmeet Kohli, and Vladimir Kolmogorov. Team coverage games. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, volume 3, pages 1197–1198, Taipei, Taiwan, May 2011.

- [35] Yoram Bachrach, Pushmeet Kohli, Vladimir Kolmogorov, and Morteza Zadimoghaddam. Optimal coalition structures in cooperative graph games. In *Proceedings of the 27th National Conference on Artificial Intelligence (AAAI)*, pages 81–87, Bellevue, Washington, July 2013.
- [36] Yoram Bachrach, Omer Lev, Shachar Lovett, Jeffrey S. Rosenschein, and Morteza Zadimoghaddam. Cooperative weakest link games. In *Proceedings of the 13th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 589–596, Paris, France, May 2014.
- [37] Yoram Bachrach, Evangelos Markakis, Ezra Resnick, Ariel D. Procaccia, Jeffrey S. Rosenschein, and Amin Saberi. Approximating power indices: theoretical and empirical analysis. *Journal of Autonomous Agents and Multi-Agent Systems*, 20(2):105–122, March 2010.
- [38] Yoram Bachrach, Reshef Meir, Michal Feldman, and Moshe Tennenholtz. Solving cooperative reliability games. In *Proceedings of the 28th Conference in Uncertainty in Artificial Intelligence (UAI)*, pages 27–34, Barcelona, Spain, July 2011.
- [39] Yoram Bachrach, Reshef Meir, Kyomin Jung, and Pushmeet Kohli. Coalitional structure generation in skill games. In *Proceedings of the 24th National Conference on Artificial Intelligence (AAAI)*, Atlanta, Georgia, 2010. 2010.
- [40] Yoram Bachrach, Reshef Meir, Michael Zuckerman, Jörg Rothe, and Jeffrey S. Rosenschein. The cost of stability in weighted voting games. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, volume 2, pages 1289–1290, Budapest, Hungary, May 2009.
- [41] Yoram Bachrach, Maria Polukarov, and Nicholas R. Jennings. The good, the bad and the cautious: safety level cooperative games. In

Proceedings of the 6th Workshop on Internet and Network Economics (WINE), pages 432–443, Palo Alto, California, June 2010.

- [42] Yoram Bachrach and Ely Porat. Path disruption games. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1123–1130, Toronto, Canada, May 2010.
- [43] Yoram Bachrach, Ely Porat, and Jeffrey S. Rosenschein. Sketching techniques for collaborative filtering. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2016–2021, Pasadena, California, July 2009.
- [44] Yoram Bachrach, Ely Porat, and Jeffrey S. Rosenschein. Sharing rewards in cooperative connectivity games. *Journal of Artificial Intelligence Research*, 47:281–311, June 2013.
- [45] Yoram Bachrach and Jeffrey S. Rosenschein. Coalitional skill games. In *Proceedings of the 7th international joint conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1023–1030, Estoril, Portugal, May 2008.
- [46] Yoram Bachrach and Jeffrey S. Rosenschein. Power in threshold network flow games. *Autonomous Agents and Multi-Agent Systems*, 18(1):106–132, 2009.
- [47] Yoram Bachrach, Jeffrey S. Rosenschein, and Ely Porat. Power and stability in connectivity games. In *Proceedings of the 7th international joint conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 999–1006, Estoril, Portugal, May 2008.
- [48] Yoram Bachrach, Rahul Savani, and Nisarg Shah. Cooperative max games and agent failures. In *Proceedings of the 13th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 29–36, Paris, France, May 2014.

- [49] Yoram Bachrach and Nisarg Shah. Reliability weighted voting games. In *Proceedings of the 6th International Symposium on Algorithmic Game Theory (SAGT)*, pages 38–49, Aachen, Germany, October 2013.
- [50] Yoram Bachrach, Morteza Zadimoghaddam, and Peter Key. A cooperative approach to collusion in auctions. *ACM SIGecom Exchanges*, 10(1):17–22, March 2011.
- [51] Yoram Bachrach, Michael Zuckerman, and Jeffrey S. Rosenschein. Effort games and the price of myopia. *Mathematical Logic Quarterly*, 55(4):377–396, 2009.
- [52] Yoram Bachrach, Michael Zuckerman, Michael Wooldridge, and Jeffrey S. Rosenschein. Proof systems and transformation games. *Annals of Mathematics and Artificial Intelligence*, 67(1):1–30, January 2013.
- [53] John J. Bartholdi III and James B. Orlin. Single transferable vote resists strategic voting. *Social Choice and Welfare*, 8(4):341–354, 1991.
- [54] John J. Bartholdi III, Craig A. Tovey, and Michael A. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3):227–241, 1989.
- [55] Michael R. Baye, Dan Kovenock, and Casper G. de Vries. Rigging the lobbying process: An application of the all-pay auction. *The American Economic Review*, 83(1):289–294, March 1993.
- [56] Michael R. Baye, Dan Kovenock, and Casper G. de Vries. The all-pay auction with complete information. *Economic Theory*, 8(2):291–305, August 1996.
- [57] James Bennett and Stan Lanning. The Netflix prize. In *Proceedings of KDD Cup and Workshop*, 2007.
- [58] Sven Berg. Paradox of voting under an urn model: The effect of homogeneity. *Public Choice*, 47(2):377–387, 1985.

- [59] Yogi Berra and Dave Kaplan. *When You Come to a Fork in the Road, Take It!: Inspiration and Wisdom From One of Baseball's Greatest Heroes*. Hachette Books, 2002.
- [60] Omer Biran and Françoise Forges. Core-stable rings in auctions with independent private values. *Games and Economic Behavior*, 73(1):52–64, September 2011.
- [61] André Blais, Robert Young, Christopher Fleury, and Miriam Lapp. Do people vote on the basis of minimax regret? *Political Research Quarterly*, 48(4):827–836, December 1995.
- [62] Gideon Blocq, Yoram Bachrach, and Peter Key. The shared assignment game and applications to pricing in cloud computing. In *Proceedings of the 13th international conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 605–612, Paris, France, May 2014.
- [63] Christian Borgs, Jennifer Chayes, Adam Tauman Kalai, Azarakhsh Malekian, and Moshe Tennenholtz. A novel approach to propagating distrust. In *Proceedings of the 6th Workshop on Internet and Network Economics (WINE)*, pages 87–105, Palo Alto, California, June 2010.
- [64] Ronen I. Brafman, Carmel Domshlak, Yagil Engel, and Moshe Tennenholtz. Transferable utility planning games. In *Proceedings of the 24th National Conference on Artificial Intelligence (AAAI)*, pages 709–714, Atlanta, Georgia, July 2010.
- [65] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, March 2016.
- [66] Simina Brânzei, Ioannis Caragiannis, Jamie Morgenstern, and Ariel D. Procaccia. How bad is selfish voting? In *Proceedings of the 27th National Conference on Artificial Intelligence (AAAI)*, pages 138–144, Bellevue, Washington, July 2013.

- [67] Georgios Chalkiadakis, Edith Elkind, and Michael Wooldridge. *Computational Aspects of Cooperative Game Theory*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan and Claypool, October 2011.
- [68] Georgios Chalkiadakis, Valentin Robu, Ramachandra Kota, Alex Rogers, and Nicholas R. Jennings. Cooperatives of distributed energy resources for efficient virtual power plants. In *Proceedings of the 10th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 787–794, Taipei, Taiwan, May 2011.
- [69] Tallichetty S. Chandrashekar and Yadati Narahari. Procurement network formation: A cooperative game approach. In *Planning Production and Inventories in the Extended Enterprise*, volume 152 of *International Series in Operations Research & Management Science*, pages 185–206. Springer, 2011.
- [70] Shuchi Chawla, Jason D. Hartline, and Balasubramanian Sivan. Optimal crowdsourcing contests. In *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 856–868, Kyoto, Japan, 2012. SIAM.
- [71] Samir Chopra, Eric Pacuit, and Rohit Parikh. Knowledge-theoretic properties of strategic voting. In *Proceedings of 9th European Conference on Logics in Artificial Intelligence (JELIA)*, pages 18–30, Lisbon, Portugal, September 2004.
- [72] William W. Cohen and Wei Fan. Web-collaborative filtering: recommending music by crawling the web. *Computer Networks*, 33(1-6):685–698, June 2000.
- [73] Vincent Conitzer, Tuomas Sandholm, and Jérôme Lang. When are elections with few candidates hard to manipulate? *Journal of the ACM*, 54(3):1–33, June 2007.

- [74] Vincent Conitzer, Toby Walsh, and Lirong Xia. Dominating manipulations in voting with partial information. In *Proceedings of the 25th National Conference on Artificial Intelligence (AAAI)*, pages 638–643, San Francisco, California, August 2011.
- [75] Constantinos Daskalakis, Grant Schoenebeck, Gregory Valiant, and Paul Valiant. On the complexity of Nash equilibria of action-graph games. In *Proceedings of the 20th annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2009.
- [76] Francesco De Sinopoli. On the generic finiteness of equilibrium outcomes in plurality games. *Games and Economic Behavior*, 34(2):270–286, February 2001.
- [77] Pierre de Trenchay. An extension of Bowen’s dynamic voting rule to many dimensions. *Social Choice and Welfare*, 15(1):141–159, 1998.
- [78] Eddie Dekel and Michele Piccione. Sequential voting procedures in symmetric binary elections. *Journal of Political Economy*, 108(1):34–55, February 2000.
- [79] Xiaotie Deng and Christos Harilaos Papadimitriou. On the complexity of cooperative solution concepts. *Mathematics of Operations Research*, 19(2):257–266, May 1994.
- [80] Yvo Desmedt and Edith Elkind. Equilibria of plurality voting with abstentions. In *Proceedings of the 11th ACM conference on Electronic Commerce (EC)*, pages 347–356, Cambridge, Massachusetts, June 2010.
- [81] Amrita Dhillon and Ben Lockwood. When are plurality rule voting games dominance-solvable? *Games and Economic Behavior*, 46(1):55–75, January 2004.
- [82] Ariel Dinar, Aharon Ratner, and Dan Yaron. Evaluating cooperative game theory in water resources. *Theory and Decision*, 32(1):1–20, 1992.

- [83] Dominic DiPalantino and Milan Vojnović. Crowdsourcing and all-pay auctions. In *Proceedings of the 10th ACM conference on Electronic commerce*, pages 119–128, Palo Alto, California, July 2009.
- [84] Richard James Duffin. Topology of series-parallel networks. *Journal of Mathematical Analysis and Applications*, 10(2):303–318, April 1965.
- [85] Paul E. Dunne, Wiebe van der Hoek, Sarit Kraus, and Michael Wooldridge. Cooperative Boolean games. In *Proceedings of the 7th international joint conference on Autonomous Agents and Multiagent Systems (AAMAS)*, volume 2, pages 1015–1022, Estoril, Portugal, May 2008.
- [86] Bhaskar Dutta and Jean-François Laslier. Costless honesty in voting. in 10th International Meeting of the Society for Social Choice and Welfare, Moscow, 2010.
- [87] Bhaskar Dutta and Arunava Sen. Nash implementation with partially honest individuals. *Games and Economic Behavior*, 74(1):154–169, January 2012.
- [88] Maurice Duverger. *Political Parties: Their Organization and Activity in the Modern State*. Methuen Publishing, 1954.
- [89] Douglas Dyer, John H. Kagel, and Dan Levin. Resolving uncertainty about the number of bidders in independent private - value auctions: An experimental analysis. *The RAND Journal of Economics*, 20(2):268–279, 1989.
- [90] Edith Elkind, Georgios Chalkiadakis, and Nicholas R. Jennings. Coalition structures in weighted voting games. In *Proceedings of the 18th European Conference on Artificial Intelligence (ECAI)*, pages 393–397, 2008.
- [91] Edith Elkind, Leslie Ann Goldberg, Paul W. Goldberg, and Michael Wooldridge. On the computational complexity of weighted voting

- games. *Annals of Mathematics and Artificial Intelligence*, 56(2):109–131, 2009.
- [92] Edith Elkind, Leslie Ann Goldberg, Paul W. Goldberg, and Michael Wooldridge. A tractable and expressive class of marginal contribution nets and its applications. *Mathematical Logic Quarterly*, 55(4):362–376, August 2009.
- [93] Ulle Endriss, Sarit Kraus, Jérôme Lang, and Michael Wooldridge. Designing incentives for Boolean games. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, volume 1, pages 79–86, Taipei, Taiwan, May 2011.
- [94] Eithan Ephrati and Jeffrey S. Rosenschein. Deriving consensus in multiagent systems. *Artificial Intelligence*, 87(1-2):21–74, November 1996.
- [95] Bruno Escoffier, Jérôme Lang, and Meltem Öztürk. Single-peaked consistency and its complexity. In *Proceedings of the 18th European Conference on Artificial Intelligence (ECAI)*, pages 366–370, Amsterdam, The Netherlands, 2008. IOS Press.
- [96] Robin Farquharson. *Theory of Voting*. Yale University Press, 1969.
- [97] Timothy J. Feddersen, Itai Sened, and Stephen G. Wright. Rational voting and candidate entry under plurality rule. *American Journal of Political Science*, 34(4):1005–1016, November 1990.
- [98] Jonathan S. Feinstein, Michael K. Block, and Frederick C. Nold. Asymmetric information and collusive behavior in auction markets. *The American Economic Review*, 75(3):441–460, June 1985.
- [99] John A. Ferejohn and Morris P. Fiorina. The paradox of not voting: A decision theoretic analysis. *American Political Science Review*, 68:525–536, June 1974.

- [100] Peter C. Fishburn. Paradoxes of voting. *American Political Science Review*, 68(2):537–546, June 1974.
- [101] Robert Fletcher. *Practical methods of optimization*. Wiley, 2000.
- [102] Katsushige Fujimoto, Ivan Kojadinovic, and Jean-Luc Marichal. Axiomatic characterizations of probabilistic and cardinal-probabilistic interaction indices. *Games and Economic Behavior*, 55(1):72–99, April 2006.
- [103] Francis Fukuyama. *The Origins of Political Order*. Farrar, Straus and Giroux, 2011.
- [104] Xi Alice Gao, Yoram Bachrach, Peter Key, and Thore Graepel. Quality expectation-variance tradeoffs in crowdsourcing contests. In *Proceedings of the 26th National Conference on Artificial Intelligence (AAAI)*, Toronto, Canada, June 2012.
- [105] Inma Garcia, Laura Sebastia, Eva Onaindia, and Cesar Guzman. A group recommender system for tourist activities. In *Proceedings of the 10th International Conference on E-Commerce and Web Technologies (EC-Web)*, pages 26–37, Linz, Austria, September 2009.
- [106] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Series of Books in the Mathematical Sciences. W. H. Freeman, 1979.
- [107] Mike Gartrell, Xinyu Xing, Qin Lv, Aaron Beach, Richard Han, Shivakant Mishra, and Karim Seada. Enhancing group recommendation by incorporating social relationship interactions. In *Proceedings of the 16th ACM international conference on Supporting group work (GROUP)*, pages 97–106, 2010.
- [108] Arie Gavious, Benny Moldovanu, and Aner Sela. Bid costs and endogenous bid caps. *The RAND Journal of Economics*, 33(4):709–722, 2002.

- [109] Allan Gibbard. Manipulation of voting schemes. *Econometrica*, 41(4):587–602, July 1973.
- [110] Allan Gibbard. Manipulation of schemes that mix voting with chance. *Econometrica*, 45(3):665–681, April 1977.
- [111] Donald Bruce Gillies. *Some Theorems on n-Person Games*. PhD thesis, Princeton University, 1953.
- [112] Umberto Grandi, Andrea Loreggia, Francesca Rossi, Kristen Brent Venable, and Toby Walsh. Restricted manipulation in iterative voting: Condorcet efficiency and borda score. In *Proceedings of 3rd International Conference of Algorithmic Decision Theory (ADT)*, pages 181–192, Brussels, Belgium, November 2013.
- [113] Gianluigi Greco, Enrico Malizia, Luigi Palopoli, and Francesco Scarcello. On the complexity of the core over coalition structures. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 216–221, Barcelona, Spain, July 2011.
- [114] Anil Gürsel and Sandip Sen. Producing timely recommendations from social networks through targeted search. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, volume 2, pages 805–812, Budapest, Hungary, May 2009.
- [115] Arye L. Hillman and John G. Riley. Politically contestable rents and transfers. *Economics & Politics*, 1(1):17–39, March 1989.
- [116] Melvin J. Hinich, John O. Ledyard, and Peter C. Ordeshook. Nonvoting and the existence of equilibrium under majority rule. *Journal of Economic Theory*, 4(2):144–153, April 1972.
- [117] Nicolas Houya and William S. Zwicker. The geometry of voting power: weighted voting and hyper-ellipsoids. *Games and Economic Behavior*, 84:7–16, March 2014.

- [118] Jonathan Huang and Carlos Guestrin. Riffled independence for ranked data. In *Proceedings of the 23rd Annual Conference on Neural Information Processing Systems (NIPS)*, Vancouver, Canada, December 2009.
- [119] Steffen Huck, Kai A. Konrad, and Wieland Müller. Merger and collusion in contests. *Journal of Institutional and Theoretical Economics*, 158(4):563–575, December 2002.
- [120] Albert Xin Jiang and Kevin Leyton-Brown. Bayesian Action-Graph Games. In *Proceedings of the 24th Conference on Neural Information Processing Systems (NIPS)*, 2010.
- [121] Albert Xin Jiang, Kevin Leyton-Brown, and Navin A.R. Bhat. Action-graph games. *Games and Economic Behavior*, 71:141–173, 2011.
- [122] Eric J. Johnson, Colin Camerer, Sankar Sen, and Talia Rymon. Detecting failures of backward induction: Monitoring information search in sequential bargaining. *Journal of Economic Theory*, 104(1):16–47, May 2002.
- [123] Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292, March 1979.
- [124] Daniel Kahneman and Amos Tversky. Choices, values, and frames. *American Psychologist*, 39(4):341–350, April 1984.
- [125] Ehud Kalai and Eitan Zemel. Generalized network problems yielding totally balanced games. *Operations Research*, 30(5):998–1008, September 1982.
- [126] Richard M. Karp. Reducibility among combinatorial problems. In Raymond E. Miller and James W. Thatcher, editors, *Proceedings of a Symposium on the Complexity of Computer Computations*, pages 85–103, Yorktown Heights, New York, March 1972.

- [127] Kei Kawai and Yasutora Watanabe. Inferring strategic voting. Technical report, Northwestern, 2010.
- [128] Michael J. Kearns, Michael L. Littman, and Satinder P. Singh. Graphical models for game theory. In *Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence (UAI)*, 2001.
- [129] Orit Kedar. *Voting for Policy, Not Parties*. Cambridge Studies in Comparative Politics. Cambridge University Press, January 2014.
- [130] Jae Kyeong Kim, Hyea Kyeong Kim, Hee Young Oh, and Young U. Ryu. A group recommendation system for online communities. *International Journal of Information Management*, 30(3):212–219, 2010.
- [131] Paul Klemperer. *Auctions: Theory and Practice*. The Toulouse Lectures in Economics. Princeton University Press, 2004.
- [132] Bettina Klose and Dan Kovenock. Moderating alliances. Fourth Congress of the Game Theory Society (GAMES), Istanbul, Turkey, July 2012.
- [133] Vijay Krishna. *Auction Theory*. Academic Press, 2002.
- [134] Vijay Krishna and John Morgan. An analysis of the war of attrition and the all-pay auction. *Journal of Economic Theory*, 72(2):343–362, February 1997.
- [135] Lech Kruś and Piotr Bronisz. Cooperative game solution concepts to a cost allocation problem. *European Journal of Operational Research*, 122(2):258–271, April 2000.
- [136] David Kurokawa, Omer Lev, Jamie Morgenstern, and Ariel D. Procaccia. Impartial peer review. In *Proceedings of 24th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 582–588, Buenos Aires, Argentina, July 2015.

- [137] Jean-Jacques Laffont. Incentives and the allocation of public goods. In Alan J. Auerbach and Martin Feldstein, editors, *Handbook of Public Economics*, volume 2 of *Handbook of Public Economics*, chapter 10, pages 537–569. Elsevier, April 1987.
- [138] Karim R. Lakhani, David A. Garvin, and Eric Lonstein. Topcoder (a): Developing software through crowdsourcing. *Harvard Business Review*, 2010.
- [139] Kate Larson and Tuomas Sandholm. Anytime coalition structure generation: average case study. *Journal of Experimental and Theoretical AI*, 12(1):23–42, 2000.
- [140] Ramoni O. Lasisi and Abibat A. Lasisi. The Shapley value in voting games: Computing single large party’s power and bounds for manipulation by merging. In *Proceedings of the 28th International Florida Artificial Intelligence Research Society (FLAIRS)*, pages 55–60, Hollywood, Florida, May 2015.
- [141] Jean-François Laslier and Jörgen W. Weibull. An incentive-compatible Condorcet jury theorem. *The Scandinavian Journal of Economics*, 115(1):84–108, January 2013.
- [142] John O. Ledyard. The paradox of voting and candidate competition: A general equilibrium analysis. In G. Horwich and J. Quirk, editors, *Essays in Contemporary Fields of Economics*, pages 54–80. Purdue University Press, 1982.
- [143] Omer Lev, Maria Polukarov, Yoram Bachrach, and Jeffrey S. Rosenschein. Mergers and collusion in all-pay auctions and crowdsourcing contests. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 675–682, St. Paul, Minnesota, May 2013.

- [144] Omer Lev and Jeffrey S. Rosenschein. Convergence of iterative voting. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, volume 2, pages 611–618, Valencia, Spain, June 2012.
- [145] Omer Lev, Moshe Tennenholtz, and Aviv Zohar. An axiomatic approach to routing. In *Proceedings of the 15th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, pages 173–180, Pittsburgh, Pennsylvania, June 2015.
- [146] Yoad Lewenberg, Yoram Bachrach, Yonatan Sompolinsky, Aviv Zohar, and Jeffrey S. Rosenschein. Bitcoin mining pools: A cooperative game theoretic analysis. In *Proceedings of the 14th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 919–927, Istanbul, Turkey, May 2015.
- [147] Yoad Lewenberg, Omer Lev, Yoram Bachrach, and Jeffrey S. Rosenschein. Agent failures in all-pay auctions. In *Proceedings of 23rd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 241–247, Beijing, China, August 2013.
- [148] Dennis Lu and Jing Yang. Auction participation and market uncertainty: Evidence from canadian treasury auctions. In *37th Annual Meeting of the Canadian Economics Association*, 2003.
- [149] Rangan Majumder. Fast collaborative filtering through sketch function based approximations. US Patent 7,624,095, November 2009.
- [150] Robert C. Marshall and Leslie M. Marx. Bidder collusion. *Journal of Economic Theory*, 133(1):374–402, March 2007.
- [151] Eric Maskin and John Riley. Uniqueness of equilibrium in sealed high-bid auctions. *Games and Economic Behavior*, 45(2):395–409, November 2003.

- [152] Nicholas Mattei and Toby Walsh. Preflib: A library of preference data. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT)*, pages 259–270, Brussels, Belgium, November 2013.
- [153] Steven Matthews. Comparing auctions for risk averse buyers: A buyer’s point of view. *Econometrica*, 55(3):633–646, May 1987.
- [154] Randolph Preston McAfee and John McMillan. Auctions with a stochastic number of bidders. *Journal of Economic Theory*, 43(1):1–19, October 1987.
- [155] Richard D. McKelvey and Richard G. Niemi. A multistage game representation of sophisticated voting for binary procedures. *Journal of Economic Theory*, 18(1):1–22, June 1978.
- [156] Richard D. McKelvey and Richard E. Wendell. Voting equilibria in multidimensional choice spaces. *Mathematics of Operations Research*, 1(2):144–158, May 1976.
- [157] Reshef Meir. Plurality voting under uncertainty. In *Proceedings of the 29th Conference on Artificial Intelligence (AAAI)*, pages 2103–2109, Austin, Texas, January 2015.
- [158] Reshef Meir, Yoram Bachrach, and Jeffrey S. Rosenschein. Minimal subsidies in expense sharing games. In *Proceedings of the 3rd International Symposium on Algorithmic Game Theory (SAGT)*, pages 347–358, Athens, Greece, October 2010.
- [159] Reshef Meir, Omer Lev, and Jeffrey S. Rosenschein. A local-dominance theory of voting equilibria. In *Proceedings of the 15th ACM conference on Economics and Computation (EC)*, pages 313–330, Palo Alto, California, June 2014.

- [160] Reshef Meir, Maria Polukarov, Jeffrey S. Rosenschein, and Nicholas R. Jennings. Convergence to equilibria of plurality voting. In *Proceedings of the 24th National Conference on Artificial Intelligence (AAAI)*, pages 823–828, Atlanta, July 2010.
- [161] Reshef Meir, Ariel D. Procaccia, and Jeffrey S. Rosenschein. A broader picture of the complexity of strategic behavior in multi-winner elections. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 991–998, Estoril, Portugal, May 2008.
- [162] Reshef Meir, Moshe Tennenholtz, Yoram Bachrach, and Peter Key. Congestion games with agent failures. In *Proceedings of the 26th National Conference on Artificial Intelligence (AAAI)*, pages 1401–1407, Toronto, Canada, July 2012.
- [163] Flavio M. Menezes and Paulo K. Monteiro. Auctions with endogenous participation. *Review of Economic Design*, 5(1):71–89, March 2000.
- [164] Samuel Merrill. Strategic voting in multicandidate elections under uncertainty and under risk. In Manfred J. Holler, editor, *Power, Voting, and Voting Power*, pages 179–187. Physica-Verlag Heidelberg, 1982.
- [165] Matthias Messner and Mattias K. Polborn. Robust political equilibria under plurality and runoff rule. Working Paper 288, IGIER (Innocenzo Gasparini Institute for Economic Research), April 2005.
- [166] Matthias Messner and Mattias K. Polborn. Strong and coalition-proof political equilibria under plurality and runoff rule. *International Journal of Game Theory*, 35(2):287–314, 2007.
- [167] Matthias Messner and Mattias K. Polborn. Miscounts, Duverger’s law and Duverger’s hypothesis. Technical Report 380, IGIER (Innocenzo Gasparini Institute for Economic Research), Bocconi University, 2011.

- [168] Bradley N. Miller, Istvan Albert, Shyong K. Lam, Joseph A. Konstan, and John Riedl. Movielens unplugged: experiences with an occasionally connected recommender system. In *Proceedings of the 8th international conference on Intelligent user interfaces (IUI)*, pages 263–266, Miami, Florida, January 2003.
- [169] Lázló Monostori, Paul Valckenaers, Alexandre Dolgui, Hervé Panetto, Mietek Brdys, and Balázs Csanád Csáji. Cooperative control in production and logistics. *Annual Reviews in Control*, 39:12–29, 2015.
- [170] Roger B. Myerson and Robert J. Weber. A theory of voting equilibria. *The American Political Science Review*, 87(1):102–114, March 1993.
- [171] John Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286–295, September 1951.
- [172] Scott A. Neslin and Leonard Greenhalgh. Nash’s theory of cooperative games as a predictor of the outcomes of buyer-seller negotiations: An experiment in media purchasing. *Journal of Marketing Research*, 20(4):368–379, November 1983.
- [173] Dusit Niyato, Athanasios V. Vasilakos, and Zhu Kun. Resource and revenue sharing with coalition formation of cloud providers: Game theoretic approach. In *Proceedings of the 11th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid)*, pages 215–224, Newport Beach, California, May 2011.
- [174] Hannu Nurmi. An assessment of voting system simulations. *Public Choice*, 73(4):459–487, June 1992.
- [175] Svetlana Obraztsova, Omer Lev, Vangelis Markakis, Zinovi Rabinovich, and Jeffrey S. Rosenschein. Beyond plurality: Truth-bias in binary scoring rules. In *Proceedings of the 4th International Conference on Algorithmic Decision Theory (ADT)*, pages 451–468, Lexington, Kentucky, September 2015.

- [176] Svetlana Obraztsova, Evangelos Markakis, Maria Polukarov, Zinovi Rabinovich, and Nicholas R. Jennings. On the convergence of iterative voting: How restrictive should restricted dynamics be? In *Proceedings of the 29th Conference on Artificial Intelligence (AAAI)*, pages 993–999, Austin, Texas, January 2015.
- [177] Svetlana Obraztsova, Evangelos Markakis, and David Robert Martin Thompson. Plurality voting with truth-biased agents. In *Proceedings of the 6th International Symposium on Algorithmic Game Theory (SAGT)*, pages 26–37, Aachen, Germany, October 2013.
- [178] Mark O’Connor, Dan Cosley, Joseph A. Konstan, and John Riedl. Polylens: a recommender system for groups of users. In *Proceedings of the 7th conference on European Conference on Computer Supported Cooperative Work (ECSCW)*, pages 199–218, Bonn, Germany, September 2001.
- [179] Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory*. MIT Press, 1994.
- [180] Guillermo Owen and Bernard Grofman. To vote or not to vote: The paradox of nonvoting. *Public Choice*, 42(3):311–325, March 1984.
- [181] Jordi Palau, Miquel Montaner, Beatriz López, and Josep Lluís de la Rosa. Collaboration analysis in recommender systems using social networks. In *Proceedings of 8th Cooperative Information Agents International Workshop (CIA)*, volume 3191 of *Lecture Notes in Computer Science*, pages 137–151, Erfurt, Germany, September 2004.
- [182] Thomas R. Palfrey and Howard Rosenthal. A strategic calculus of voting. *Public Choice*, 41(1):7–53, 1983.
- [183] Irene Parrachino, Ariel Dinar, and Fioravante Patrone. Cooperative game theory and its application to natural, environmental, and water

- resource issues: 3. application to water resources. Policy Research Working Papers 4074, World Bank, December 2006.
- [184] David M. Pennock, Eric Horvitz, and C. Lee Giles. Social choice theory and recommender systems: Analysis of the axiomatic foundations of collaborative filtering. In *Proceedings of the 17th National Conference on Artificial Intelligence (AAAI)*, pages 729–734, Austin, Texas, August 2000.
- [185] Ryan W. Porter, Eugene Nudelman, and Yoav Shoham. Simple search methods for finding a Nash equilibrium. *Games and Economic Behavior*, 63:642–662, 2009.
- [186] Ariel D. Procaccia. Can approximation circumvent Gibbard-Satterthwaite? In *Proceedings of the 24th National Conference on Artificial Intelligence (AAAI)*, pages 836–841, Atlanta, Georgia, July 2010.
- [187] Ariel D. Procaccia and Jeffrey S. Rosenschein. Junta distributions and the average-case complexity of manipulating elections. *Journal of Artificial Intelligence Research*, 28:157–181, 2007.
- [188] Ariel D. Procaccia, Jeffrey S. Rosenschein, and Aviv Zohar. Multi-winner elections: Complexity of manipulation, control and winner-determination. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1476–1481, Hyderabad, India, January 2007.
- [189] Zinovi Rabinovich, Svetlana Obraztsova, Omer Lev, Evangelos Markakis, and Jeffrey S. Rosenschein. Analysis of equilibria in iterative voting schemes. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, pages 1007–1013, Austin, Texas, January 2015.

- [190] Talal Rahwan and Nicholas R. Jennings. An improved dynamic programming algorithm for coalition structure generation. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, volume 3, pages 1417–1420, Estoril, Portugal, May 2008.
- [191] Annemieke Reijngoud and Ulle Endriss. Voter response to iterated poll information. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, volume 2, pages 635–644, Valencia, Spain, June 2012.
- [192] Ezra Resnick, Yoram Bachrach, Reshef Meir, and Jeffrey S. Rosenschein. The cost of stability in network flow games. In *Proceedings of the 34th International Symposium on Mathematical Foundations of Computer Science (MFCS)*, pages 636–650, Novy Smokovec, High Tatras, Slovakia, August 2009.
- [193] Paul Resnick and Hal R. Varian. Recommender systems. *Communications of the ACM*, 40(3):56–58, March 1997.
- [194] Paul Resnick and Richard Zeckhauser. Trust among strangers in internet transactions: Empirical analysis of eBay’s reputation system. In Michael R. Baye, editor, *The Economics of the Internet and E-commerce (Advances in Applied Microeconomics)*, volume 11, pages 127–157. Emerald Group Publishing Limited, 2002.
- [195] Anja Rey and Jörg Rothe. Bribery in path-disruption games. In *Proceedings of the 2nd International Conference on Algorithmic Decision Theory (ADT)*, pages 247–261, Piscataway, New Jersey, October 2011.
- [196] Reyhaneh Reyhani and Mark C. Wilson. Best-reply dynamics for scoring rules. In *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI)*, volume 2, pages 672–677, Montpellier, France, June 2012.

- [197] Reyhaneh Reyhani, Mark C. Wilson, and Javad Khazaei. Coordination via polling in plurality voting games under inertia. In *Proceedings of COMSOC*, pages 359–370, Kraków, Poland, September 2012.
- [198] Francesco Ricci. Travel recommender systems. *IEEE Intelligent Systems*, 17(6):55–57, 2002.
- [199] Robert W. Rosenthal. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.
- [200] Jörg Rothe, editor. *Economics and Computation*. Springer, 2015.
- [201] Jörg Rothe and Lena Schend. Control complexity in bucklin, fallback, and plurality voting: an experimental approach. In *Proceedings of the 11th international conference on Experimental Algorithms (SEA)*, pages 356–368, Bordeaux, France, July 2012.
- [202] Ola Rozenfeld and Moshe Tennenholtz. Consistent continuous trust-based recommendation systems. In *Proceedings of the 5th Workshop on Internet and Network Economics (WINE)*, pages 113–124, Rome, Italy, December 2009.
- [203] Ariel Rubinstein. *Modelling Bounded Rationality*. MIT Press, 1998.
- [204] Donald G. Saari. Susceptibility to manipulation. *Public Choice*, 64:21–41, 1990.
- [205] Tuomas Sandholm, Kate Larson, Martin Andersson, Onn Shehory, and Fernando Tohmé. Coalition structure generation with worst case guarantees. *Artificial Intelligence*, 111(1-2):209–238, July 1999.
- [206] Mark Allen Satterthwaite. Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217, April 1975.

- [207] J. Ben Schafer, Joseph Konstan, and John Riedi. Recommender systems in e-commerce. In *Proceedings of the 1st ACM conference on Electronic Commerce (EC)*, pages 158–166, Denver, Colorado, November 1999.
- [208] Abigail See, Yoram Bachrach, and Pushmeet Kohli. The cost of principles: analyzing power in compatibility weighted voting games. In *Proceedings of the 13th international conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 37–44, Paris, France, May 2014.
- [209] Travis C. Service and Julie A. Adams. Approximate coalition structure generation. In *Proceedings of the 24th National Conference on Artificial Intelligence (AAAI)*, pages 854–859, Atlanta, Georgia, July 2010.
- [210] Lloyd S. Shapley and Martin Shubik. A method for evaluating the distribution of power in a committee system. *The American Political Science Review*, 48(3):787–792, September 1954.
- [211] Lloyd S. Shapley and Martin Shubik. Quasi-cores in a monetary economy with nonconvex preferences. *Econometrica*, 34(4):805–827, October 1966.
- [212] Lloyd S. Shapley and Martin Shubik. The assignment game I: The core. *International Journal of Game Theory*, 1(1):111–130, 1971.
- [213] Onn Shehory and Sarit Kraus. Methods for task allocation via agent coalition formation. *Artificial Intelligence*, 101(1-2):165–200, May 1998.
- [214] Kenneth A. Shepsle. Institutional arrangements and equilibrium in multidimensional voting models. *American Journal of Political Science*, 23(1):27–59, February 1979.
- [215] Jonathan Silberman and Garey Durden. The rational behavior theory of voter participation. *Public Choice*, 23(1):101–108, 1975.

- [216] Rashmi R. Sinha and Kirsten Swearingen. Comparing recommendations made by online systems and friends. In *Proceedings of the 2nd DELOS Network of Excellence Workshop on Personalisation and Recommender Systems in Digital Libraries*, Dublin, Ireland, June 2001.
- [217] Birgitte Sloth. The theory of voting and equilibria in noncooperative games. *Games and Economic Behavior*, 5(1):152–169, January 1993.
- [218] Dong-Wook Song and Photis M. Panayides. A conceptual application of cooperative game theory to liner shipping strategic alliances. *Maritime Policy & Management*, 29(3):285–301, 2002.
- [219] Hossein Azari Soufiani, David C. Parkes, and Lirong Xia. Preference elicitation for general random utility models. In *Proceedings of the 26th Conference on Neural Information Processing Systems (NIPS)*, pages 126–134, Lake Tahoe, Nevada, December 2012.
- [220] Yves Sprumont. The division problem with single-peaked preferences: A characterization of the uniform allocation rule. *Econometrica*, 59(2):509–519, March 1991.
- [221] Jonathan Steinberg. *Bismarck: A Life*. Oxford University Press, 2012.
- [222] K. Takamizawa, Takao Nishizeki, and Nobuji Saito. Linear-time computability of combinatorial problems on series-parallel graphs. *Journal of the ACM*, 29(3):623–641, July 1982.
- [223] David Robert Martin Thompson, Samantha Leung, and Kevin Leyton-Brown. Computing nash equilibria of action-graph games via support enumeration. In *Proceedings of the 7th International Conference on Internet and Network Economics (WINE)*, 2011.
- [224] David Robert Martin Thompson, Omer Lev, Kevin Leyton-Brown, and Jeffrey S. Rosenschein. Empirical aspects of plurality election equilibria. In *Proceedings of the 12th International Conference on Autonomous*

- Agents and Multiagent Systems (AAMAS)*, pages 391–398, St. Paul, Minnesota, May 2013.
- [225] William Thomson. Cooperative models of bargaining. In Robert J. Aumann and Sergiu Hart, editors, *Handbook of Game Theory with Economic Applications*, volume 2. Elsevier, 1994.
- [226] Amos Tversky and Daniel Kahneman. Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–1131, September 1974.
- [227] Amos Tversky and Daniel Kahneman. The framing of decisions and the psychology of choice. *Science*, 211(4481):453–458, January 1981.
- [228] Hans van Ditmarsch, Jérôme Lang, and Abdallah Saffidine. Strategic voting and the logic of knowledge. In *Proceedings of the 14th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, pages 196–205, Chennai, India, January 2013.
- [229] Toby Walsh. An empirical study of the manipulability of single transferable voting. In *Proceedings of the 19th European Conference on Artificial Intelligence (ECAI)*, pages 257–262, Lisbon, Portugal, August 2010.
- [230] Frank Edward Walter, Stefano Battiston, and Frank Schweitzer. A model of a trust-based recommendation system on a social network. *Journal of Autonomous Agents and Multi-Agent Systems (JAAMAS)*, 16(1):57–74, February 2008.
- [231] Jun Wang and Marcel J. T. Reinders. Music recommender system for wi-fi walkman. ICT Group Technical Report Series ICT-2003-01, Delft University of technology, 2003.
- [232] Chung wei Hang and Munindar P. Singh. Trust-based recommendation based on graph similarity. In *Proceedings of the 13th International Workshop on Trust in Agent Societies*, pages 71–81, Toronto, Canada, May 2010.

- [233] Mika Widgrén. Voting power in the EC decision making and the consequences of two different enlargements. *European Economic Review*, 38(5):1153–1170, May 1994.
- [234] Sean Wilentz. *The Rise of American Democracy*. Norton, 2005.
- [235] Lirong Xia and Vincent Conitzer. Stackelberg voting games: Computational aspects and paradoxes. In *Proceedings of the 24th National Conference on Artificial Intelligence (AAAI)*, pages 805–810, Atlanta, Georgia, 2010.
- [236] Lirong Xia, Michael Zuckerman, Ariel D. Procaccia, Vincent Conitzer, and Jeffrey S. Rosenschein. Complexity of unweighted coalitional manipulation under some common voting rules. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 348–353, Pasadena, California, July 2009.
- [237] Jian Xue, Laijun Zhao, Longzhen Fan, and Ying Qian. An inter-provincial cooperative game model for air pollution control in china. *Journal of the Air & Waste Management Association*, 65(7):818–827, July 2015.
- [238] Makoto Yokoo, Vincent Conitzer, Tuomas Sandholm, Naoki Ohta, and Atsushi Iwasaki. Coalitional games in open anonymous environments. In *Proceedings of the 20th National Conference on Artificial Intelligence (AAAI)*, volume 2, pages 509–514, Pittsburgh, Pennsylvania, July 2005.
- [239] Kazuyoshi Yoshii, Masataka Goto, Kazunori Komatani, Tetsuya Ogata, and Hiroshi G. Okuno. An efficient hybrid music recommender system using an incrementally trainable probabilistic generative model. *IEEE Transactions on Audio, Speech and Language Processing*, 16(2):435–447, February 2008.

- [240] Kai Yu, Shenghuo Zhu, John Lafferty, and Yihong Gong. Fast non-parametric matrix factorization for large-scale collaborative filtering. In *Proceedings of the 32nd international ACM conference on Research and development in information retrieval (SIGIR)*, pages 211–218, Boston, Massachusetts, July 2009.
- [241] D. Yun Yeh. A dynamic programming approach to the complete set partitioning problem. *BIT Numerical Mathematics*, 26(4):467–474, 1986.
- [242] Yingxue Zhao, Shouyang Wang, T.C. Edwin Cheng, Xiaoqi Yang, and Zhimin Huang. Coordination of supply chains by option contracts: A cooperative game theory approach. *European Journal of Operational Research*, 207(2):668–675, December 2010.
- [243] Yair Zick, Maria Polukarov, and Nicholas R. Jennings. Taxation and stability in cooperative games. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 523–530, St. Paul, Minnesota, May 2013.
- [244] Yair Zick, Alexander Skopalik, and Edith Elkind. The Shapley value as a function of the quota in weighted voting games. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, volume 1, pages 490–95, Barcelona, Spain, July 2011.
- [245] Michael Zuckerman, Piotr Faliszewski, Yoram Bachrach, and Edith Elkind. Manipulating the quota in weighted voting games. *Artificial Intelligence*, 180-181:1–19, April 2012.
- [246] Michael Zuckerman, Omer Lev, and Jeffrey S. Rosenschein. An algorithm for the coalitional manipulation problem under maximin. In *Proceedings of the 10th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 845–852, Taipei, Taiwan, May 2011.

- [247] Michael Zuckerman, Ariel D. Procaccia, and Jeffrey S. Rosenschein. Algorithms for the coalitional manipulation problem. In *Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 277–286, San Francisco, California, 2008.

Part IV

Appendices

Appendix A

Proofs of Part I

A.1 Chapter 3

Observation 3.2. *Given an election with truth-biased voters using veto, in any Nash equilibrium the score of the winner does not change from the score it received in the truthful vote. All non-truthful voters veto runner-ups.*

Proof. Suppose in a Nash equilibrium profile \mathbf{b} the winner is candidate w . If a voter truthfully vetoes w , it has no reason to deviate – no other vote can change this outcome, and it gains an ϵ of utility by being truthful. If a voter is not truthful, if it is vetoing w , yet w still wins, w will still win if the voter reverts to being truthful, but the voter will gain an ϵ of utility. Hence, only voters which vetoed w in the truthful profile veto it in \mathbf{b} .

If a voter is not truthful in a Nash equilibrium, that means vetoing its truthful choice would change the outcome. As it is not vetoing w in its truthful vote, the change in the election's result stems from the point that would be gained by the currently vetoed candidate. If it becomes the winner, that can only be if the candidate was a runner-up. \square

Observation 3.3. *Given an election with truth-biased voters using veto, in any Nash equilibrium which is not the truthful profile, all voters except those vetoing the winners or runner-ups prefer the winner to the threshold candidate.*

Proof. Suppose in a Nash equilibrium profile \mathbf{b} the winner is candidate w . As this is a non-truthful profile, there is at least one voter that can change the outcome if it returns to being truthful. Hence, there are runner-up candidates and thus, a threshold candidate.

If a voter is not vetoing w or runner-ups, it can change the result by vetoing w and making the threshold voter the winner (as it does not veto a runner-up, the candidate that stands to gain a point will not become the winner). \square

Theorem 3.2. *Given an election with truth-biased voters using the veto voting rule and a specific candidate, deciding if there is a voting profile which is a Nash equilibrium in which the candidate is victorious is NP-complete.*

Proof. As the proof is quite While membership in NP is trivial, completeness requires several steps. We will construct a reduction from exact-cover by 3-sets (X3C).

Definition A.1. *Exact cover by 3 sets (X3C) is a problem with the input of $3m$ elements $U = \{u_1, \dots, u_{3m}\}$ and a set of sets $S = \{S_1, \dots, S_n\}$ such that for $1 \leq i \leq n$: $S_i \subset U$, $S_i = \{u_{i_1}, u_{i_2}, u_{i_3}\}$. We wish to know if there is a set $T \subseteq S$ such that $|T| = m$ and $\cup_{S \in T} S = U$.¹*

Taking a X3C instance, we construct an instance of our problem. In order to aid us we denote by \mathcal{S} the members of S ordered as usual – $S_1 \succ S_2 \succ \dots \succ S_n$; similarly we use \mathcal{U} for the ordering of U . $\bar{\mathcal{S}}$ marks the opposite direction – $S_n \succ S_{n-1} \succ \dots \succ S_1$, and similarly for $\bar{\mathcal{U}}$.

The set of candidates C will be the members of S and U , to which we add two new candidates w and t . Our tie-breaking rule is $w \succ t \succ \mathcal{S} \succ \mathcal{U}$

The set of voters V consists of two blocks of voters described in Table A.1, along with 3 additional blocks described below:

Block 3 For every $u_i \in U$, we have:

- m votes of the form: $\mathcal{U} \setminus \{u_i\} \succ \mathcal{S} \succ w \succ t \succ u_i$;

¹NP-completeness was shown in Garey and Johnson [106].

Block 1					
$\mathcal{S} \setminus \{S_1\},$	$\bar{\mathcal{U}}$...	$\mathcal{S} \setminus \{S_{k-1}\},$	$\bar{\mathcal{U}}$...
$\mathcal{U},$	$\bar{\mathcal{S}} \setminus \{S_2\}$...	$\mathcal{U},$	$\bar{\mathcal{S}} \setminus \{S_k\}$...
$w,$	w	...	$w,$	w	...
$S_1,$	S_2	...	$S_{k-1},$	S_k	...
$t,$	t	...	$t,$	t	...

Block 2				
...	$\mathcal{U} \setminus \{u_{i_1}\}$	$\mathcal{U} \setminus \{u_{i_2}\}$	$\mathcal{U} \setminus \{u_{i_3}\}$...
...	$\mathcal{S} \setminus \{S_i\}$	$\mathcal{S} \setminus \{S_i\}$	$\mathcal{S} \setminus \{S_i\}$...
...	w	w	w	...
...	t	t	t	...
...	u_{i_1}	u_{i_2}	u_{i_3}	...
...	S_i	S_i	S_i	...

Table A.1: Two of the blocks of voters for the NP-completeness proof

- $n - 2m - 1$ votes of the form: $\bar{\mathcal{S}} \succ \bar{\mathcal{U}} \setminus \{u_i\} \succ w \succ t \succ u_i$.

Block 4 For every $S_i \in \mathcal{S}$, we have:

- m votes of the form: $\mathcal{S} \setminus \{S_i\} \succ \mathcal{U} \succ w \succ t \succ S_i$;
- $n - 2m - 1$ votes of the form: $\bar{\mathcal{U}} \succ \bar{\mathcal{S}} \setminus \{S_i\} \succ w \succ t \succ S_i$.

Block 5 $n - m$ votes of the form: $t \succ \mathcal{S} \succ \mathcal{U} \succ w$.

In the truthful profile, w is not the winner, but rather u_1 . We wish to prove there is an equilibrium in which w is the winner if and only if there is a solution to the X3C problem.

Given the constructed truthful profile, if a Nash equilibrium profile \mathbf{b} exists with w as a winner, t is the threshold candidate in \mathbf{b} . Since w will retain the same score as in the truthful profile and beats all other candidates in the tie-breaking rule, and since t must be the threshold candidate, we know t must have $n - m$ points.

If there is a $T = \{S'_1, \dots, S'_m\} \subseteq S$ which is a solution to the X3C problem, we have an equilibrium in which w is the winner: the voters from Block 1 whose penultimate candidate is $S'_i \in T$ will veto S'_i . The voters in Block 2 who veto $S'_i \in T$ instead veto their penultimate candidates $u_{i_{1/2/3}}$. In such a situation all candidates are vetoed by $n - m$ voters (apart from those in $S \setminus T$, which are vetoed by $n - m + 2$ voters), and therefore w is the winner. All non-truthful voters are vetoing runner-ups which they prefer less than w or t . Hence, changing their vote will make the candidate they currently veto the winner, and as they would rather have w win, they do not change their vote. Furthermore, all voters from Blocks 1 through 4 that do not veto a runner-up candidate can only deviate so that t becomes a winner. Since they prefer w to t , none of them will have an incentive to do so. Finally, none of the voters in Block 5 can change the election outcome and will therefore remain truthful.

Now assume that there is no solution to the X3C problem. Since t needs to gain some points to be a threshold candidate, at least m voters from Block 1 will veto the S_i 's (the only candidates less-preferred than w). However, in order for them not to revert to their truthful vote, those S_i 's need to be runner-up candidates, so all votes in Block 2 who would truthfully veto those S_i 's, need to veto their respective u_i 's instead. In addition, those u_i 's need to be runner-ups as well (or those Block 2 votes will revert to the truthful vote), and as they are ranked below S in the tie-breaking rule, they need to have $m - n$ vetoes in order to be runner-ups. This means that each u_i is vetoed only once in Block 2. So we have m (or more) S_i 's containing exactly one copy of each u_i ; i.e., we found an exact cover of U , contradicting the assumption that X3C has no solution. \square

Theorem 3.3. *Given an election with truth-biased voters using the veto voting rule and a specific candidate, assuming Condition 1 is true and Condition 2 is not, deciding if there is a voting profile which is a Nash equilibrium in which the candidate is victorious is NP-complete.*

Proof. As with Theorem 3.2, we will construct a reduction from X3C. We

will use the same notation, where S is the set of sets and U is the set of elements in an instance of X3C, and convert the members of these sets into distinct candidates. In this case, in addition to the candidates from S and U , we will introduce four special candidates w, t, p_1 and p_2 . The tie-breaking preference order shall be $w \succ t \succ p_1 \succ p_2 \succ S \succ U$, where candidates from S and U appear in their lexicographic order.

We now construct the set of voters, grouped into blocks according to their truthful preferences. In each block we only explicitly describe the order of a few least-preferred candidates. All candidates that are not explicitly mentioned in a profile, appear in an arbitrary order, and are marked by \dots .

Block 1 A set of n voters, one for each candidate in S , with a preference order of the form $\dots \succ t \succ w \succ S_i \succ p_1$.

Block 2 A set of $n - m$ voters with a preference order of the form $\dots \succ t \succ w \succ p_2$ and one additional voter with profile of the form $\dots \succ w \succ t \succ p_2$.

Block 3 For each $\{u_{i_1}, u_{i_2}, u_{i_3}\} = S_i \in S$ a set of $n - m + 2$ voters. Three with profiles of the form $\dots \succ w \succ t \succ u_{i_j} \succ S_i$, where $j \in \{1, 2, 3\}$, and all others of the form $\dots \succ w \succ t \succ S_i$.

Block 4 For each $u_i \in U$ a set of $n - m - 1$ voters with profiles of the form $\dots \succ w \succ t \succ u_i$.

Block 5 A set of $n - m - 1$ voters with profiles of the form $\dots \succ w \succ t$.

Block 6 A set of $n - m$ voters with profiles of the form $\dots \succ t \succ w$.

If there is a profile \mathbf{b} that is a Nash equilibrium in which w wins, w 's score does not change from its truthful score, yet by our construction t 's score is above w 's by 1 (fulfilling Condition 1). Hence, for w to become a winner in \mathbf{b} , t has to receive only one additional veto and will also be a threshold candidate.

Consider the voters of Block 2. $n - m$ of them prefer t to w , and would not veto the former when w is the winner. Yet, if they are non-truthful in \mathbf{b} , they have to veto a runner up. As a result, none of them can deviate from their truthful profile in equilibrium. On the other hand, the last voter of the block can deviate and veto t .

Similarly, thanks to the low score of p_1 in the truthful profile, up to m voters from Block 1 can deviate in the equilibrium and stop vetoing p_1 . Moreover, if less than m do so, that indicates, p_1 is not a runner-up, hence one of the voters vetoing p_1 , can devote to veto w making t the winner, hence exactly m are deviating. These newly vetoed candidates have to be less preferred than w by the deviating voters, i.e., from the set S . As a result, there are m candidates $S_i \in S$ that are being vetoed by the voters from Block 1 in the equilibrium profile \mathbf{b} .

These *chosen* S_i 's, however, need to be runner-up candidates. To achieve that, exactly 3 candidates that veto S_i 's in Block 3 must deviate in the equilibrium profile \mathbf{b} . These can only be the voters with preference profiles of the form $\dots \succ w \succ t \succ u_{i_j} \succ S_i$, where $j \in \{1, 2, 3\}$.

Since no voter in Block 4 can deviate, those voters from Block 3 that deviate to veto u_{i_j} s must be vetoing runner-ups, hence the total number of times that u_{i_j} is being vetoed is equal to $n - m$. This can happen only if each $u_{i_j} \in U$ is vetoed exactly once by voters from Block 3.

As a result, the sub-set S_i 's that are vetoed by voters in Block 1 constitutes a solution to the given X3C instance. The opposite direction, that is, constructing a Nash equilibrium profile given a solution to the X3C instance, is trivial. \square

Theorem 3.4. *Consider a candidate $w \in C$ and a truthful profile for which both Condition 1 and Condition 2 apply. Then finding if there is a voting profile for truth-biased voters under the veto voting rule which is a Nash equilibrium where w is the winner can be done in polynomial time.*

Proof. The proof is based on a polynomial reduction to the max-flow problem in a graph. We will construct a graph (and later correct the flow) in a way




Figure A.1: Palm sub-structure for potential deviators

that the set of flow-saturated edges will indicate the feasibility of obtaining a Nash equilibrium. Furthermore, positive flow at certain nodes in the graph will indicate a switch in the voters' equilibrium ballots from their truthful profile.

Given the truthful voting profile, we will construct the graph as follows. Vertices will be associated with each candidate and each voter; we also add a *source* and a *sink* node. The set of graph vertices will therefore be $\{source, sink\} \cup C \cup V$.

The set of edges, E , in the graph will consist of three subsets.

Potential deviators Edges that link voters and potentially vetoed candidates: for voter v_i whose least preferable candidate is some $r \in C \setminus w$, i.e., the preference order is in the form $\dots \succ w \succ c_1 \succ \dots \succ c_l \succ r$, the graph shall contain directed edges with *unit* flow capacity (r, v_i) (indicating flow from r to v_i) and the directed edges $(v_i, c_1), \dots, (v_i, c_l)$. The resulting palm-leaf sub-structure is depicted in Figure A.1. It essentially captures the ability of the voter to change its veto in a manner that will benefit w without hurting the voter's utility (of course, there are plenty of such sub-graphs in the graph, and vertices may be part of several such structures).

Sustainable deviations Edges from the *source* node: these edges and capacities reflect the number of additional points a candidate may absorb until it becomes a runner-up candidate competing with w . As we know

what w 's score in the Nash equilibrium will be, we know how many points a candidate may gain without having more points than w . For candidates $c \in C$ which can gain points without overshadowing w , the graph includes the edge $(source, c)$ with the capacity of that number of points (i.e., if $c \succ w$ in the tie-breaking rule, the capacity of the edge is w 's score in the truthful score - c 's score in the truthful profile - 1; if $w \succ c$ in the tie-breaking rule, the capacity of the edge is w 's score in the truthful score - c 's score in the truthful profile).

Necessary deviations Edges to the *sink* node: these edges and capacities reflect the number of additional veto votes a candidate needs to receive in order to reduce its score enough so that it does not score higher than w . For candidates $c \in C$ which must lose points so they do not overshadow w , the graph includes the edge $(c, sink)$ with the capacity of that number of points (i.e., if $c \succ w$ in the tie-breaking rule, the capacity of the edge is c 's score in the truthful profile - w 's score in the truthful score - 1; if $w \succ c$ in the tie-breaking rule, the capacity of the edge is c 's score in the truthful profile - w 's score in the truthful score)

If there is a Nash equilibrium in which w wins, all candidates connected to the sink must lose a number of points equal at least to that capacity. Hence, a flow in which the capacity of all candidates gaining points (by no longer longer being vetoed) flows to the voters which give them that points, and passing through those voters to the candidates these voters deviated to vetoing produces a flow of at least the capacity of the incoming edges to the *sink*. Therefore, if the max-flow of the graph is less than that flow we know there is no such Nash equilibrium. Now, we need to show that if all edges to the *sink* are saturated in a maximum flow, we will show that a Nash equilibrium profile \mathbf{b} , whose winner is w , can be recovered from the flow.

Let $x : E \rightarrow \mathbb{N}$ be a maximal acyclic, integer flow through the constructed graph. Such a flow can be obtained in polynomial time (polynomial in number of voters and candidates). Furthermore, all edges from a candidate node

to a voter node that have positive flow on them will be saturated (as their capacity is 1).

We will now modify the flow, while maintaining its total capacity, to maximize the flow through the *source* outgoing edges, and minimize the flow through voter nodes. Since we will later associate a flow through a voter node with the voter deviating from the truthful vote, minimizing the flow through voter nodes will reflect and ensure that the voting profile recovered from it will be truth biased (i.e., no unnecessary manipulations, as in such cases some voters will be better off reverting to their truthful preference).

Let $D = \{c \in C \mid \exists e = (\text{source}, c) \in E\}$ be the set of all nodes to which the source is directly connected. Notice that D is a subset of candidate nodes. Let $q \in D$ be a node for which there is a voter v such that $(v, q) \in E$ is with a non-zero flow (i.e., $x((v, q)) > 0$), and the edge (source, q) is not saturated. In other words, we have a path $(\text{source} = n_0, n_1, \dots, n_l = q)$ such that $e_k = (n_{k-1}, n_k) \in E$ for all $1 \leq k \leq l$ and $x(e_k) > 0$. Notice that since x is an integer flow and all edges between candidate nodes and voter nodes have unit capacity, all the edges of the path have a unit flow apart from the initial edge from the source to n_1 .

We will construct an augmented flow \hat{x} by canceling the flow through this path, replacing it with an additional unit flow from the *source* to q until no such q exist. More formally, as long as such q exist, we set initially $\hat{x} = x$ and then modify it, by setting $\hat{x}(e_k) = 0$ for all $1 \leq k \leq l$ and $\hat{x}((\text{source}, n_1)) = x((\text{source}, n_1)) - 1$ and $\hat{x}((\text{source}, q)) = x((\text{source}, q)) + 1$. We then repeat the modification procedure, if necessary, for \hat{x} .

Notice that following such a flow modification procedure the total flow from the *source* to the *sink* does not change. Following it, we are assured that for node $q \in D$ either the edge (source, q) is saturated, or q has no positive incoming flow from the voter nodes.

We now construct the Nash Equilibrium voting profile from the flow. For every voter $v_i \in V$, if there is an edge (v_i, c_j) for some $c_j \in C$ such that $x(v_i, c_j) > 0$, we let v_i change its vote to veto c_j . Otherwise, v_i votes

truthfully. This is an equilibrium, as a voter has a flow to a candidate if the edge to that candidate from the *source* is saturated (meaning that candidate is a runner-up), or the candidate has an edge directed at the *sink*, i.e., it must be vetoed to lower the score of this candidate.

Complexity of max-flow is $\mathcal{O}(2n^2m)$ (as for a graph $G(V, E)$ with a weight function f , it is $\mathcal{O}(E \max |f|)$, and in our case, there are $2nm$ edges, and the max flow is n). The process of finding the actual equilibrium is done by going for up to n edges connected to the source, and for every one of them going, at most, over all edges, which is exactly the same complexity. \square

A.2 Chapter 4

Theorem 4.1. *An iterative scoring rule election with a deterministic tie-breaking rule, even for voters using best-response strategies and starting from the truthful state, will not converge for some preferences.*

Proof. Our examples will be somewhat complex, as we deal with a large family of voting rules. In some cases the best response strategy is obvious, as there is only one choice that results in making a specific candidate the winner. In other cases there may be multiple options to reach the same outcome, hence we used a “natural” definition for scoring rules, in which players taking off points from the current winner will give it zero points and award the new winner the maximal score possible. Also, all things are equal, we assume voters will prefer to be as close as possible to their truthful preferences. However, even if one does not use such a definition for best response, cycles are still created – only longer, as they may go through several more steps than detailed here.

First, we deal with scoring rules in which at least three candidates do not receive maximal scores (i.e., $\alpha_{m-2} < \alpha_1$). We have at least four candidates, a , b , c and d . Our tie-breaking rule is structured as follows:

- c tied with others except b wins.

- b tied with others except d wins.
- d tied with others except a or c , wins.

We have two voters:

Voter 1: $a \succ b \succ c \succ d$

Voter 2: $c \succ d \succ b \succ a$

We can add several dummy candidates so that the score given by voter 1 to b is less than is given to a , and the score voter 2 gives to d is less than given to c (and dummy voters, making these dummy candidates irrelevant as potential winners). The winner in this truthful state is c (either he is the sole winner, or through a tie with a). The only option for improving the result for voter 1 is to make b victorious, changing its preference to $b \succ a \succ d \succ c$. Voter 2 can improve the result by changing its preference to $d \succ c \succ a \succ b$, making d the winner (possibly through winning the tie between b and d). The best option available to voter 1 is to return to its original preference order, making a , its favorite, the winner. However, now voter 2 will return to its original preference as well, as it ensures the victory of c , its own most preferred candidate.

If there are only two candidates that receive less than the maximal score, then we use a different setting, one with six candidates. Our tie-breaking rule follows:

- b wins when a, b, c, d are tied.
- a wins when a, c, d are tied.
- c wins when a and c are tied and when a, c, d, e are tied
- d wins in other ties that include it.
- f wins in other ties that do not include d .
- e win in other ties that do not include d or f .

	a	b	c	d	e
a	b	c	d	e	d
b	c	d	a	a	a
c	d	a	b	a	a
d	e	a	a	a	a
e	d	a	a	a	a

Table A.2: Non-linear tie-breaking rule for veto

Let us look at two voters:

$$\text{Voter 1: } a \succ b \succ c \succ d \succ e \succ f$$

$$\text{Voter 2: } b \succ c \succ a \succ d \succ e \succ f$$

The winner here is candidate b (since a, b, c, d are tied). However, when voter 1 changes its stated preference to $a \succ c \succ d \succ e \succ f \succ b$, then a , its favorite, becomes the winner (since a, c, d are tied). Voter 2 can only improve this situation by changing its stated preference to $a \succ c \succ d \succ e \succ b \succ f$, making c victorious. Voter 1 can now improve the situation by returning to its original preference, making a the winner. In this case, voter 2 will also return to its original preference, as that will make its favorite candidate, b , win.

If there is only one candidate that receives the less-than-maximal score, this is the Veto voting rule, for which there is a similar, but simpler, example. We shall use two voters, and we can describe the voting rule and tie-breaking rule using Table A.2, marking the victor according to whom the voters chose to veto.

In our case, the voters' real preferences are:

$$\text{Voter 1: } c \succ b \succ d \succ e \succ a$$

$$\text{Voter 2: } b \succ d \succ c \succ e \succ a$$

The truthful starting point would result in b being the winner. As voter 1 would rather that c win, it will move to veto b . Following that, voter 2

would move to veto b as well, as that would result in d winning. Voter 1, which would rather that c win, will return to vetoing a , and as voter 2 would rather that b be victorious, it would return to vetoing a as well, returning to our original starting point. \square

Theorem 4.2. *Iterative Veto elections with deterministic linear-order tie-breaking and voters which use a best-response strategy, converge even when not starting from a truthful state.*

Proof. Suppose there is an iterative election G that includes a cycle. We shall mark an arbitrary state in the cycle as G_0 , and enumerate the rest of the cycle accordingly. Note that G_0 is not necessarily the opening state of the election.

Definition A.2. $score_i(x)$ is defined as the score of candidate x in game state G_i . $max(G_i)$ is defined as the score of the winning candidate in G_i .

Lemma A.1. *If there is a cycle, then for $j < i$, $max(G_i) \leq max(G_j) + 1$, and if $max(G_i) = max(G_j) + 1$, there is only one candidate with that score.*

Proof. Proving by induction, the base case is trivial. Assuming it is true after $h - 1$ steps, proving it for step h : examining G_{h-1} , if there was a $j < h - 1$ for which $max(G_{h-1}) = max(G_j) + 1$, there is a single winner in G_{h-1} , which loses a point, and therefore the winner in G_h will have, at most, $max(G_{h-1})$ points. Thus, $max(G_h) \leq max(G_{h-1})$, and the claim stems from its truth for G_{h-1} .

If for every $j < h - 1$ $max(G_{h-1}) \leq max(G_j)$, the maximal score in G_h will rise by at most one point, i.e., $max(G_h) \leq max(G_j) + 1$ for all $j < h$. Furthermore, if it indeed grows, there is only a single candidate with that number of points (as only the candidate that got an extra point has this score). If the maximal score in G_h did not grow, $max(G_h) \leq max(G_{h-1})$, and claim is true from its correctness for G_{h-1} . \square

Notice that since we can choose G_0 arbitrarily from the cycle, due to the last lemma, $max(G_0) + 1 \geq max(G_i) \geq max(G_0) - 1$, otherwise, there will

be no possibility for the cycle to return to its starting point.

Lemma A.2. *There can be at most $n \cdot (m - 2)$ consecutive steps in which the voter changed their veto from candidate a to candidate b , and candidate a became the winner.*

Proof. Every time a voter changes its veto, it indicates that the vetoed candidate is preferable to the current winner; that is, the winner is a candidate less and less liked as the game progresses. Since there are n voters and, at most, $m - 1$ candidates that are worse than the current one, and as the voter will not chose to make the very worst candidate the winner, there are $n \cdot (m - 2)$ steps. \square

We shall deal, first of all, with the easiest case, solved by the lemma above, when there is always only one candidate with the winning score (the tie-breaking rule is never used). In this case, at every step, the old winner loses a point, and the new winner gains a point. This is the case dealt with in Lemma A.2, and as the number of steps is limited, there can be no cycle.

Having dealt with that case, let us take a closer look at G_0 , which we can define as one of the states in which there is more than one candidate with a maximal score. Note that there must be more than one of these states, since if there was a single winner in G_i and more than that in G_{i+1} , a candidate received a point and did not become a unique winner, i.e., its score in G_i was, at most, $\max(G_i) - 2$. Since this is a cycle, there must be a step in which it returns to that score (if it is G_{i+1} , then for a cycle to happen, the same candidate will need to rise again so the voter that increased its score in G_i will veto it again).

Lemma A.3. *For every state G_i in which there is more than one candidate scoring $\max(G_i)$, $\max(G_i) = \max(G_0)$, $|\{x \mid \text{score}_i(x) = \max(G_i)\}| = |\{x \mid \text{score}_0(x) = \max(G_0)\}|$ and $|\{x \mid \text{score}_i(x) = \max(G_i) - 1\}| = |\{x \mid \text{score}_0(x) = \max(G_0) - 1\}|$. This means the number of candidates with the maximal score remains fixed, as does the number of candidates with maximal score -1 . Furthermore, these are the same candidates, switching between the*




Figure A.2: When only one candidate has maximal score




Figure A.3: When multiple candidates have the maximal score

Diagrams showing why there is a limit on the increase and decrease of maximal score. When there is a state with only one candidate with maximal score, the maximal score will either remain the same with a single winner (move type B) or decrease (move type A). If it is a state where there are several candidates with the maximal score, the maximal score will either increase (move type D) while creating a situation with a single winner or the maximal score will remain the same (move type C). This also illustrates why the score cannot go down much if there is a cycle – it can only increase by one in the whole cycle; at no point can we reach a maximal score 2 points higher than another.

two scores: $\{x \mid \text{score}_i(x) \in \{\max(G_i), \max(G_i) - 1\}\} = \{x \mid \text{score}_0(x) \in \{\max(G_0), \max(G_0) - 1\}\}$.

Proof. According to Lemma A.1, if $\max(G_i) = \max(G_0) + 1$, there is a single candidate with the winning score, and this lemma does not handle this case. Suppose $\max(G_i) = \max(G_0) - 1$; according to the same lemma, this means there is only one candidate with the winning score in G_0 , which we defined as a state having at least two.

At any step in the game, one candidate loses a point and another gains it. Hence, if the number of those with the maximal score and maximal -1 score is not the same as in G_0 , some candidate lost (or gained) a point, which has a score lower than maximal -1 in G_0 . However, as the maximal score will never be $\max(G_0) - 1$ (otherwise, according to Lemma A.1, there would only be one candidate with winning score in G_0), there is no way in the cycle for the candidate to be vetoed when it has a score of $\max(G_0) - 1$, and get a lower score. As no candidate that has a score of $\max(G_0)$ or $\max(G_0) - 1$ can get a smaller score, the group of candidates with these scores stays fixed throughout the cycle. \square

Let B be the group of candidates who changed places in states in which $\max(G_i) = \max(G_0)$ (i.e., $B = \{x \mid \exists i \text{ such that } \text{score}_i(x) = \max(G_0) \text{ and } \exists j \text{ such that } \text{score}_j(x) = \max(G_0) - 1\}$). Let $z \in B$ be the lowest ranked candidate according to the linear tie-breaking rule, in B . Since z changes its score, there is a state G_i where z has the score $\max(G_0)$ and is vetoed, i.e., z is the winner if G_i . This means there is no other candidate from B with the score $\max(G_0)$. As the number of candidates with $\max(G_0)$ does not change (according to Lemma A.3), this means that at every state G_j in which $\max(G_j) = \max(G_0)$, there is only a single candidate from B with $\max(G_0)$ points, and it always wins (due to the tie-breaking rule). This means the candidate getting the point at every stage is the one that becomes the winner – which, as noted in Lemma A.2, is a finite process, contradicting the endless cycle. \square

Theorem 4.3. *Under the iterative procedure, using a best response strategy and when voters are myopic, no scoring rule apart from plurality and veto converges.*

Proof. Different rules require different proofs, and we consequently divide our proof:

Part 1: Scoring Rules with 2 Values (k -approval)

Scoring rules with only 2 values which are not plurality or veto are equivalent to k -approval for $k > 1$ and $k < m - 1$. We shall show it for $k = 2$ and $m = 4$. Using dummy candidates, this can be extended to any size of k .

Consider the tie-breaking rule $a \succ b \succ c \succ d$ and the voters:

Voter 1: $a \succ b \succ c \succ d$

Voter 2: $d \succ b \succ c \succ a$

The winner in this case is b , with a score of 2. But voter 1 can change its vote to $a \succ c \succ d \succ b$, making a the winner (thanks to the tie-breaking). But now voter 2 can change to $c \succ d \succ b \succ a$, making c the winner. Now, voter 1, by reverting to its truthful preference makes a the winner again (thanks to tie-breaking rules), and, finally, by reverting to its true preference, voter 2 returns us to the original state, making b the winner and creating the cycle.

Part 2: Scoring Rules with 3 Values

Oddly, this is the most complicated case. Our scoring rule has the values α_1 , α_2 and 0. Let us first examine the case when the scoring rule for 4 candidates is $(\alpha_1, \alpha_2, 0, 0)$ and $\alpha_1 < 2\alpha_2$. Our tie-breaking rule is $a \succ b \succ c \succ d$, and the voters are:

Voter 1: $b \succ d \succ c \succ a$

Voter 2: $a \succ d \succ c \succ b$

Hence the winner is d , with a score of $2\alpha_2$. Voter 2 can now change its vote to $a \succ c \succ b \succ d$, making a the winner with α_1 points (as d changes to

α_2). Now voter 1 can change to $c \succ b \succ d \succ a$, making c the winner, with $\alpha_1 + \alpha_2$ points. By reverting to its truthful vote, voter 2 can make its favorite candidate, a , the winner again, to which voter 1 retaliates by returning to its truthful vote as well, creating a cycle, with d as the winner.

We now create a cycle for the case $\alpha_1 \geq 2\alpha_2$. Our tie-breaking rule is $a \succ b \succ c \succ d$, and the voters are:

Voter 1: $a \succ c \succ d \succ b$

Voter 2: $d \succ c \succ b \succ a$

This means a is the winner with α_1 points, thanks to the tie-breaking rule. Voter 2 now changes to $c \succ d \succ b \succ a$, making c the winner with $\alpha_1 + \alpha_2$ points. Voter 1 retaliates by changing to $a \succ d \succ b \succ c$, making a the winner again. By returning to its truthful preference, voter 2 makes the winner d , its favorite candidate, with $\alpha_1 + \alpha_2$ points. However, by completing the cycle and returning to its truthful preference, voter 1 makes its favorite, a , the winner again.

We now look at the case where the scoring rule for 4 candidates is $(\alpha_1, \alpha_2, \alpha_2, 0)$. The tie-breaking rule is $a \succ b \succ c \succ d$, and the voters are:

Voter 1: $c \succ d \succ b \succ a$

Voter 2: $a \succ b \succ d \succ c$

The winner is either a or b (depending on if $\alpha_1 < 2\alpha_2$ or not). Voter 1 reacts by voting $d \succ c \succ b \succ a$, making d the winner with $\alpha_1 + \alpha_2$ points, to which voter 2 reacts by voting $a \succ b \succ c \succ d$, making a or b the winner. However, by reverting to its truthful preference, voter 1 makes c the winner with a score of $\alpha_1 + \alpha_2$. Voter 2 completes the cycle by reverting to its truthful preference.

The last case is when the scoring rule for 4 candidates is $(\alpha_1, \alpha_1, \alpha_2, 0)$. First, suppose $\alpha_1 \geq 2\alpha_2$. The tie-breaking rule is $a \succ b \succ c \succ d$, and the voters are:

Voter 1: $a \succ c \succ b \succ d$

Voter 2: $c \succ d \succ b \succ a$

Voter 3: $a \succ b \succ d \succ c$

The winner is a with $2\alpha_1$ points, which is changed to b (with $2\alpha_1 + \alpha_2$ points) when voter 2 changes to $b \succ c \succ d \succ a$. Now voter 3 makes a the winner with a score of $2\alpha_1$ when it votes $a \succ d \succ c \succ b$. Voter 2 now returns to being truthful, making c the winner, with $2\alpha_1 + \alpha_2$ points, causing voter 3 now reverts to being truthful, making a the winner again.

Now, suppose $\alpha_1 < 2\alpha_2$. The tie-breaking rule is $a \succ b \succ c \succ d$, and the voters are:

Voter 1: $d \succ b \succ a \succ c$

Voter 2: $d \succ c \succ a \succ b$

Voter 3: $b \succ a \succ c \succ d$

Voter 4: $d \succ b \succ c \succ a$

The winner is b with a score of $3\alpha_1$, enticing voter 4 to vote $d \succ c \succ b \succ a$, making d the winner with the same score. Voter 3 now changes its vote to $c \succ b \succ a \succ d$, making c the winner with $3\alpha_1$ points. By reverting to being truthful, voter 4 makes d the winner again, which voter 3 changes back to b when it returns to being truthful as well.

Part 3: Scoring Rules with More Values

We shall show this on voting rules with 4 values and 4 candidates. This is easily extendable to more candidates and values by adding dummy candidates (and voters who do not get to manipulate) as necessary.

The scoring rule is $(\alpha_1, \alpha_2, \alpha_3, 0)$ and is strictly monotonic. The voters are:

Voter 1: $a \succ b \succ c \succ d$

Voter 2: $c \succ d \succ b \succ a$

c is the winner here with $\alpha_1 + \alpha_3$ points (a has only α_1 points; b has $\alpha_2 + \alpha_3$; and d has α_2). Voter 1 can change the outcome by changing its vote to $b \succ a \succ d \succ c$ reducing c to only α_1 points, and raising b to $\alpha_1 + \alpha_3$. Voter 2 retaliates by changing to $d \succ c \succ a \succ b$, making d the winner with $\alpha_1 + \alpha_3$ points, and reducing b to α_1 . Voter 1 improves its situation by reverting to its truthful preferences, which make a the winner (with $\alpha_1 + \alpha_3$ points). Voter 2 completes the cycle by also reverting to its truthful preferences, making its favorite candidate, c , the winner. \square

Theorem 4.4. *Given a truthful profile \mathbf{a} and a profile \mathbf{b} distinct from \mathbf{a} , it is NP-complete to decide if \mathbf{b} is reachable by iterative plurality using best-response updates, starting from \mathbf{a} .*

Block 1	Block 2				Block 3				Block 4	Block 5
$g_1 \dots g_1 \dots g_n \dots g_n$	$d_1 \dots d_1 \dots d_m \dots d_m$	$g_1 \dots g_1 \dots g_n \dots g_n$	$d_1 \dots d_1 \dots d_m \dots d_m$	$g_1 \dots g_1 \dots g_n \dots g_n$	$d_1 \dots d_1 \dots d_m \dots d_m$	$g_1 \dots g_1 \dots g_n \dots g_n$	$d_1 \dots d_1 \dots d_m \dots d_m$	$t \dots t$	stand-ins	
$u_1 \dots u_k \dots u_1 \dots u_k$	$g_{1_1} \dots g_{1_1} \dots g_{m_1} \dots g_{m_1}$	$u_1 \dots u_k \dots u_k \dots u_1$	$t \dots t$	$u_1 \dots u_k \dots u_1 \dots u_k$	$G \setminus g_1 \dots G \setminus g_n$	$u_1 \dots u_k \dots u_1 \dots u_k$	$G \setminus g_1 \dots G \setminus g_n$	arbitrary		
$\vdots \dots \vdots \dots \vdots \dots \vdots$	$w_1 \dots w_1 \dots w_m \dots w_m$	$\vdots \dots \vdots \dots \vdots \dots \vdots$	$t \dots t$	$\vdots \dots \vdots \dots \vdots \dots \vdots$	$G \setminus g_1 \dots G \setminus g_n$	$w_1 \dots w_1 \dots w_m \dots w_m$	$G \setminus g_1 \dots G \setminus g_n$	order		
$u_k \dots u_1 \dots u_k \dots u_1$	$G \setminus g_1 \dots G \setminus g_n$	$u_k \dots u_1 \dots u_1 \dots u_k$	$t \dots t$	$G \setminus g_1 \dots G \setminus g_n$	$w_1 \dots w_1 \dots w_m \dots w_m$	$G \setminus g_1 \dots G \setminus g_n$	$G \setminus g_1 \dots G \setminus g_n$	over		
$G \setminus g_1 \dots G \setminus g_n$	$t \dots t$	$G \setminus g_1 \dots G \setminus g_n$	$t \dots t$	$G \setminus g_1 \dots G \setminus g_n$	$w_1 \dots w_1 \dots w_m \dots w_m$	$G \setminus g_1 \dots G \setminus g_n$	$G \setminus g_1 \dots G \setminus g_n$	all		
$D \cup W \dots D \cup W$	arbitrary order over all others				arbitrary order over all others					others

Table A.3: NP-Completeness proof profiles: truthful profile. Recall, $|w_i| = l_i$.

Block 1	Block 2				Block 3				Block 4				Block 5
$t \dots t \dots t \dots t$	$w_1 \dots w_1 \dots w_m \dots w_m$	$w_1 \dots w_1 \dots w_m \dots w_m$	$w_1 \dots w_1 \dots w_m \dots w_m$	$w_1 \dots w_1 \dots w_m \dots w_m$	$w_1 \dots w_1 \dots w_m \dots w_m$	$w_1 \dots w_1 \dots w_m \dots w_m$	$w_1 \dots w_1 \dots w_m \dots w_m$	$t \dots t$	$d_{ w_1 } \dots d_m$	$t \dots t$	$d_{ w_1 -1} \dots t$	$t \dots t$	as was
$g_1 \dots g_1 \dots g_n \dots g_n$	$d_1 \dots d_1 \dots d_m \dots d_m$	$g_1 \dots g_1 \dots g_n \dots g_n$	$d_1 \dots d_1 \dots d_m \dots d_m$	$g_1 \dots g_1 \dots g_n \dots g_n$	$d_1 \dots d_1 \dots d_m \dots d_m$	$g_1 \dots g_1 \dots g_n \dots g_n$	$d_1 \dots d_1 \dots d_m \dots d_m$	$d_1 \dots d_1 \dots d_m \dots d_m$	$d_{ w_1 -1} \dots t$	$t \dots t$	$t \dots t$	$t \dots t$	
$u_1 \dots u_k \dots u_1 \dots u_k$	$g_{1_1} \dots g_{1_1} \dots g_{m_1} \dots g_{m_1}$	$u_1 \dots u_k \dots u_k \dots u_1$	$t \dots t$	$u_1 \dots u_k \dots u_1 \dots u_k$	$G \setminus g_1 \dots G \setminus g_n$	$u_1 \dots u_k \dots u_1 \dots u_k$	$G \setminus g_1 \dots G \setminus g_n$	$u_1 \dots u_k \dots u_1 \dots u_k$	$G \setminus g_1 \dots G \setminus g_n$	$G \setminus g_1 \dots G \setminus g_n$	$G \setminus g_1 \dots G \setminus g_n$	$G \setminus g_1 \dots G \setminus g_n$	
$\vdots \dots \vdots \dots \vdots \dots \vdots$	$w_1 \dots w_1 \dots w_m \dots w_m$	$\vdots \dots \vdots \dots \vdots \dots \vdots$	$t \dots t$	$\vdots \dots \vdots \dots \vdots \dots \vdots$	$G \setminus g_1 \dots G \setminus g_n$	$w_1 \dots w_1 \dots w_m \dots w_m$	$G \setminus g_1 \dots G \setminus g_n$	$w_1 \dots w_1 \dots w_m \dots w_m$	$G \setminus g_1 \dots G \setminus g_n$	$G \setminus g_1 \dots G \setminus g_n$	$G \setminus g_1 \dots G \setminus g_n$	$G \setminus g_1 \dots G \setminus g_n$	
$u_k \dots u_1 \dots u_k \dots u_1$	$G \setminus g_1 \dots G \setminus g_n$	$u_k \dots u_1 \dots u_1 \dots u_k$	$t \dots t$	$G \setminus g_1 \dots G \setminus g_n$	$w_1 \dots w_1 \dots w_m \dots w_m$	$G \setminus g_1 \dots G \setminus g_n$	$G \setminus g_1 \dots G \setminus g_n$	$G \setminus g_1 \dots G \setminus g_n$	$G \setminus g_1 \dots G \setminus g_n$	$G \setminus g_1 \dots G \setminus g_n$	$G \setminus g_1 \dots G \setminus g_n$	$G \setminus g_1 \dots G \setminus g_n$	
$D \cup W \dots D \cup W$	as was				as was				as was				

Table A.4: NP-Completeness proof profiles: target profile. Recall, $|w_i| = l_i$.

Proof. To show that the problem is in NP, it is enough to provide as a certificate the sequence of best-response updates that leads from profile \mathbf{a} to

profile **b**. One could then check that this is a valid sequence. We know the sequence is of polynomial length from Meir et al. [160], Theorem 3.

To prove NP-hardness, we provide a reduction from the Hitting Set (HS) problem, a well known NP-complete problem,² which is the following: we are given a set of ground elements $G = \{g_1, \dots, g_n\}$, a family of subsets of G , $W = \{w_1, \dots, w_m\}$, $w_i \subseteq G$, $|w_i| = l_i$ and a number $k \leq n$. We need to determine if there is a hitting set $U \subset G$, so that $|U| \leq k$, and $\forall i \in [m]$, $U \cap w_i \neq \emptyset$. This is a well-known NP-complete problem.

We assume that we are given an instance that satisfies: for $1 \leq i \leq m$ $|w_1| \geq |w_i|$; $|w_1| \geq 3$; $m \geq n$ (we can always pad an instance by replicating a set to satisfy this). These three assumptions do not impact the complexity of the HS problem.

Given such an instance of the HS problem, we proceed by constructing an instance of our problem, i.e., a truthful profile **a**, and a matching (non-truthful) profile **b**, so that a sequence of iterative best-response updates going from **a** to **b** exists if and only if the HS instance has a solution.

Given an HS instance as above, we associate one candidate with each element of G and one candidate with each element of W . In addition, we introduce k candidates u_1, \dots, u_k , corresponding to the (up to) k elements of U . Finally, we also add a set D of m dummy candidates, $D = \{d_1, \dots, d_m\}$, and a special target candidate t . Overall, there are $n + 2m + k + 1$ candidates in our election with the following tie-breaking order: $d_1 \prec \dots \prec d_m \prec u_1 \prec \dots \prec u_k \prec w_1 \prec \dots \prec w_m \prec g_1 \prec \dots \prec g_n \prec t$. We slightly abuse the notation so that each w_j refers both to the set from the HS instance and the corresponding candidate in our instance – and similarly for the element candidates g_i .

We will now introduce five blocks of voters with preferences as depicted in Table A.3. Notice that Block 1 and Block 3 contain nk voters each, Block 2 has $\sum |w_i| \leq mn$ voters, and Block 4 has exactly m voters. These cardinalities will be used in the later stages of the proof. Block 5 (the *stand-*

²NP-completeness shown in Karp [126].

ins in Table A.3) is incidental and its voters are only necessary to create an initial balance among the candidates in the truthful profile. These voters all have the order of preference $U \succ G \succ t \succ W \succ D$, broken for each voter only by shifting one particular candidate to be the top choice. We have as many votes in Block 5 as required to ensure that after counting all truthful votes in all the Blocks, each candidate receives exactly $2k + n$ in the truthful profile **a**. This also means that given the tie-breaking rule, t is the winner in **a**.

We will show that ascertaining reachability of the reported profile depicted in Table A.4 is equivalent to solving the original HS problem.

To see this, assume first that there is a solution U , with $|U| \leq k$, to the HS instance. Then, we can associate (possibly with replications in case that $|U| < k$) to each candidate u_i an element in G , say $g(u_i)$, so that for $\widehat{U} = \bigcup_{i \in [k]} \{g(u_i)\}$, it holds that $\forall j \in [m], \widehat{U} \cap w_j \neq \emptyset$.

Let each voter in Block 1 that has u_i $1 \leq i \leq k$, as their second choice and $g(u_i)$ as their first choice (there is exactly one such voter for each i), change their vote in sequence starting in the order from 1 to k . As a result of the best-response updates, each candidate u_i will in turn receive an additional vote, while $g(u_i)$ will lose at least one vote. In this new profile, the winner will be u_k , due to tie-breaking, with all candidates from U having $2k + n + 1$ votes, while those in the set $\widehat{U} \subseteq G$ will have at most $2k + n - 1$ votes. Denote this new profile by **c**.

Now, in the profile **c**, consider those voters in Block 2 that have $g(u_i)$ as their second choice, for each element $g(u_i) \in \widehat{U}$. Note that because U is a solution to the HS instance, this implies that for each sub-block of Block 2, having w_j as a third choice, with $1 \leq j \leq m$, there is a value of i and a voter from this sub-block where $g(u_i)$ is a second choice for this voter. For all these voters, it is not a best response to vote for their second choice $g(u_i)$, since in the previous round of updates all elements from \widehat{U} lost a vote and due to tie-breaking they cannot become a winner with a single step. Instead, the best response for these voters is to vote for w_j , thus “unlocking” the

candidates of W . Let us choose one such voter for each w_j and let them change their vote in sequence. This changes profile \mathbf{c} to a profile \mathbf{d} , where all candidates in D can no longer become a winner with a single deviation, all candidates in U and in W have $2k + n + 1$ votes, and candidates in G have at most $2k + n$ votes each (some of them have $2k + n - 1$).

In the next round of updates, we will prevent all candidates in U and G from ever again becoming a possible winner by essentially “running” a competition between the candidates in W and the target candidate t . That is, we choose a voter sequence that will grant $W \cup \{t\}$, an ever increasing number of votes, eliminating any other candidate from becoming a best-response. The effect of this voting sequence will eventually be the emergence of the profile \mathbf{b} .

Let us first allow one voter from Block 4 to change their vote. The voter will naturally shift t to be the top choice. This will give the target candidate t $2k + n + 1$ votes as well, completely preventing all candidates from G from ever becoming a winner, since all of them have less than $2k + n + 1$ votes and lose in tie-breaking to t . We can now cycle, repeatedly through all candidates from W selecting for each w_j a voter from Block 2 with w_j as the third (truthful) choice, which now is the best-response for that voter. At the end of each cycle iteration we will grant one more voter from Block 4 the possibility to change their vote. There will be at most $|w_1|$ such cycles. Notice that some cycles will be shorter in the sense that there will be some sub-blocks where all voters will have already voted for their corresponding candidate from W (since w_1 may have a strictly higher cardinality from the rest of the sets). Finishing this process we will have the voters from Block 2 and Block 4 vote as they are intended in the target profile \mathbf{b} . The voters from Block 1 and Block 3 will still be voting either for a candidate from U or a candidate from G , and the voters from Block 5 remain as they are. In this intermediate voting profile, \mathbf{e} , candidates from W will have at most $2k + n + \sum |w_j|$ votes, and so will the target candidate t .

We can now reach profile \mathbf{b} from \mathbf{e} , if we alternate between allowing a

voter from Block 1 and a voter from Block 3 (that still do not vote for w_1 or t) to change their votes. Notice that the best-response top-choice for these voters is indeed either w_1 or t . This will result in Block 1 and Block 3 transforming their votes into those prescribed by \mathbf{b} , completing the vote modification sequence from \mathbf{a} . Notice, additionally, that in \mathbf{b} the winning candidate is the target candidate t , with $2k + n + nk + \sum |w_j|$ votes and no voter can change the outcome, i.e., \mathbf{b} is an equilibrium.

Finally, for the other direction, assume that there is no solution to the underlying HS instance. It is then easy to check that there is no possible sequence of votes that can “unlock” at least one voter in Block 2 for each w_j . Hence, this makes the targeted profile \mathbf{b} unreachable. \square

Lemma 4.1. *In a non-truthful Nash equilibrium with truth biased voters under iterative plurality starting the process with their truthful preferences, the winner will always be a runner-up candidate in the original state, with only a single voter being untruthful.*

Proof. In the iterative process, voters only vote for candidates which will win with an extra point. As the score of the winner does not drop in iterative plurality ([160], proof of Theorem 3), this means the winner needs to be a candidate to be voted for, hence it should be a winner or runner-up at every stage of the iterative process.

As all untruthful voters vote for the winner, and it only need a single point over its true score to become the winner, any additional untruthful voters have an incentive to return to their truthful vote, as they cannot change the result. \square

Theorem 4.5. *Algorithm 1 finds all Nash equilibria reachable from the truthful starting point (\mathbf{a}) in an iterative plurality model with truth-biased voters.*

Proof. We shall first show why the algorithm only outputs equilibrium profiles, and then that there are no equilibria that it misses.

Exploiting Lemma 4.1, every equilibrium the algorithm finds is made of a voter that can change its vote to c , making it the winner, and, if all other

voters remain the same, has no incentive to deviate to a different candidate. Furthermore, no other voters will deviate in retaliation – if there are, it means they can deviate to a candidate c' which can win over c and which they prefer over c , and these deviations are found by line 11.

Now, suppose there is an equilibrium resulting in candidate c winning. According to the previously proven lemmas, c is a runner-up in \mathbf{a} , and there is only one voter that will deviate. Therefore, that voter must be found with the algorithm's line 6. Since that equilibrium will only be eliminated if a voter is found in line 11, and such a voter will indeed destroy an equilibrium (as it will have an incentive to deviate), the algorithm will find all equilibria.

Since there are two nested loops, each counting through a subset of candidates and voters, complexity is $\mathcal{O}(mn)$. \square

Lemma 4.2. *A stable state with lazy-biased voters under iterative plurality starting the process with their truthful preferences will only have a single participating voter.*

Proof. Suppose a stable state has been reached. Any voter not voting for the winner will benefit from abstaining, and left with voters which are all voting for the same candidate, all will benefit from abstaining except the last one. \square

Theorem 4.6. *Algorithm 2 finds all Nash equilibria reachable from the truthful starting point (\mathbf{a}) in an iterative plurality model with lazy-biased voters.*

Proof. As the algorithm covers all possible cases, we simply explain every "yes" and "no" response: for "yes" we detail the sequence of best-response moves that achieves them, for "no" we explain why. We denote the winner in some profile \mathbf{b} as $winner(\mathbf{b})$.

We begin with the case where the starting state (\mathbf{a}) is a Nash equilibrium in the basic model, i.e., without abstentions. Note that in this case, as in all Nash equilibria of the basic case, all voters' best-response (potentially, excepting a voter for the winner) is to abstain.

Line 4 indicates the sequence where all voters that do not support the winner abstain (as this is a Nash equilibrium, they do not deviate to a different candidate) and then all voters supporting the winner except v abstain as well, until v is the only one left.

Line 6 indicates the sequence where all voters not voting for $winner(a)$ except v and the noted \tilde{c} voter abstain, and then all but 2 voters for $winner(a)$ abstain as well. Then the \tilde{c} supporter deviates to support z , and all other voters except v will now abstain.

Line 9 indicates the sequence where all voters not supporting $winner(a)$ or c except v and the voter supporting c over $winner(a)$ abstain, then $winner(a)$ supporters abstain until they are only as many as their score in \mathbf{a} (or their score +1, depending on tie-breaking rule). Then the z voter preferring c deviates to make c the winner, and then all voters abstain except v and 2 c -supporting voters (which include a z supporting one which deviated). This voter now reverts to z , making it the winner, and the voters except v abstain.

Line 11 indicates the sequence where all voters not supporting $winner(a)$ or c except v and the noted \tilde{c} voter abstain, then $winner(a)$ supporters abstain until they are only as many as their score in \mathbf{a} (or their score +1, depending on tie-breaking rule). Then the \tilde{c} voter deviates to c , making it the winner, and then our voter deviates to return $winner(a)$ to its victorious position (at this point, we know our voter prefers $winner(a)$ to others). Then all voters except ours abstain, and then our voter deviates back to z .

If the algorithm returns no, this means all voters find that anyone they support over $winner(a)$ (that is not their first preference) has only one point and loses to $winner(a)$ according to tie-breaking. This means no voter has any move except abstention, and hence, z can never become the winner.

We now turn to the case where the starting position is not a Nash equilibrium in the basic sense.

Line 14 indicates the sequence where we can simply have regular run of iterative plurality (without abstentions), when not allowing v to participate. We call the resulting state \mathbf{c} . z cannot become a winner (as it was not even

a runner-up). If v prefers $winner(c)$ over any runner-ups, then we look at the run of iterative plurality, and replace the last deviation of a voter to $winner(c)$ with a deviation of v . Now, all voters except v can abstain. If v prefers a runner-up b , we let all voters for candidates that are not $b, winner(c)$ or voter v to abstain, and then let v deviate. Since v has deviated to the winner, all except v can abstain.

Line 17 indicates the sequence where if there is such a voter which voted for z , it deviates to b , and we can now follow the same sequence as for line 14 in the paragraph above, as z is no longer a runner-up. If there is not (all of them prefer $winner(a)$), the noted voter deviates, making b the winner, and now one of the z voters deviates to $winner(a)$, making z no longer a runner-up, and again, we can now follow the same sequence as for line 14 in the paragraph above.

In line 19 we revert to the Nash equilibrium part, as the only voter wishing to deviate is v – without v this is a Nash equilibrium. However, in the case of 3 voters we can answer directly (line 20) – since v is the runner-up, the 2 other voters have voted for the winner and will not deviate.

Line 24 indicates the sequence where after reaching \mathbf{b} , all voters except v abstain.

Line 26 indicates the sequence where we pursue, after reaching \mathbf{b} a similar strategy to line 14. If v prefers $winner(b)$ over any runner-ups, then we look at the run of iterative plurality, and replace the last deviation of a voter to $winner(b)$ with a deviation of v . Now, all voters except v can abstain. If v prefers a runner-up b , we let all voters for candidates that are not $b, winner(b)$ or voter v to abstain, and then let v deviate. Since v has deviated to the winner, all except v can abstain.

Finally, in line 27 we are at a situation where, at most, only v wants to deviate, as in Line 19.

Complexity stems from line 11, which loops twice over candidates and then checks all the voters, reaching $\mathcal{O}(m^2n)$. \square

A.3 Chapter 5

Theorem 5.1. *When using plurality and when all voters use local dominance strategy and have the same radius r , if the initial state was one when all voters were truthful, the iterative process converges to a stable state.*

Proof. Our proof is for ℓ_1 , but can be easily modified for ℓ_i for $1 < i$ and ℓ_∞ .

As a useful guide for the proof, note that, in essence, for the voters, the set of potential winners they consider in profile \mathbf{b} is in $\overline{H}_{r+1}(\mathbf{b})$, as the winners in $S(\mathbf{b}_{-v}, r)$ are in $\overline{H}_r(\mathbf{b})$ (up to r votes are changing), and when taking into account the voter's own actions, a candidate that was up to $r+1$ points from the winner can become the winner by a voter's move.

If the truthful state \mathbf{a} is stable, then we are done. Thus assume it is not.

Let \mathbf{b}^t be the voting profile after t steps from the initial truthful vote $\mathbf{b}^0 = \mathbf{a}$. Let $b_i \rightarrow b'_i$ be a move of voter i at state $\mathbf{s} = \mathbf{b}^t$ to state $\mathbf{s}' = \mathbf{b}^{t+1}$. We claim that the following hold throughout the game:

1. $b_i \notin \overline{H}_{r+1}(\mathbf{s}')$, i.e., once a candidate is deserted, it is no longer a possible winner.
2. $b'_i \prec_i b_i$, i.e., voters always “compromise” by voting for a less-preferred candidate.
3. $\max_{c \in C} \text{score}_{\mathbf{s}'}(c) \geq \max_{c \in C} \text{score}_{\mathbf{s}}(a)$, i.e., the score of the winner never decreases.
4. $\overline{H}_{r+1}(\mathbf{s}') \subseteq \overline{H}_{r+1}(\mathbf{s})$, i.e., the set of possible winners can only shrink.

We prove this by a complete induction.

1. If this is the first move of i then $b_i = a_i$: as the voter moved to b'_i , this means it $S(\mathbf{s}_{-i}, r)$ -dominates b_i . Hence, there is a profile $\mathbf{p} \in S(\mathbf{s}_{-i}, r)$ for which $\text{winner}(b'_i, \mathbf{p}) \succ_i \text{winner}(b_i, \mathbf{p})$ and there was no profile $\mathbf{p}' \in S(\mathbf{s}_{-i}, r)$ where $\text{winner}(b_i, \mathbf{p}') \succ_i \text{winner}(b'_i, \mathbf{p}')$. Hence, $a_i \notin \overline{H}_{r+1}(\mathbf{s}')$ (as then there would have been such \mathbf{p}').

If this is not the first move of voter i , if $b_i \in \overline{H}_{r+1}(\mathbf{s}')$ this means that $b_i, b'_i \in \overline{H}_r(\mathbf{s})$, hence there is at least one scenario in $S(\mathbf{s}, r)$ in which candidate $\hat{c} \in C$, the winner is \mathbf{s} , is tied with b_i , and by moving to b'_i ensures \hat{c} will win. Similarly, there is a scenario where voter i 's move makes b'_i win instead of \hat{c} , so $b'_i \succ_i b_i$. However, let $\mathbf{b}^{t'}$ for $t' < t$ be the last time voter i moved. According to induction assumption 4, $\overline{H}_{r+1}(\mathbf{s}) \subseteq \overline{H}_{r+1}(\mathbf{b}^{t'})$, hence b_i and b'_i were in $\overline{H}_{r+1}(\mathbf{b}^{t'})$. As b_i dominated voter i 's previous vote, so did b'_i , but according to the local dominance strategy, the most preferred candidate is taken, so $b_i \succ_i b'_i$, reaching a contradiction.

2. If this is the first move of i then this is immediate. Otherwise, let $\mathbf{b}^{t'}$ for $t' < t$ be the last time voter i moved. According to induction assumption 4, $\overline{H}_{r+1}(\mathbf{s}) \subseteq \overline{H}_{r+1}(\mathbf{b}^{t'})$, hence b_i and b'_i were in $\overline{H}_{r+1}(\mathbf{b}^{t'})$. If $b'_i \succ_i b_i$, then if b_i dominated voter i 's previous vote, so did b'_i , but according to the local dominance strategy, the most preferred candidate is taken, so $b_i \succ_i b'_i$, reaching a contradiction.
3. As we proved in 1, $b_i \notin \overline{H}_{r+1}(\mathbf{s}')$, hence b_i was not the winner in \mathbf{s} , so voter i 's move did not change $winner(\mathbf{s})$'s score, and the score of the winner in \mathbf{s}' must be the same or more.
4. Since by 3 the score of the winner never decreases, the only way to expand \overline{H}_{r+1} is to add a vote to a candidate not in \overline{H}_{r+1} . Suppose such a move $d_j \rightarrow d'_j$ occurred in $\mathbf{b}^{t'}$. According to 1, d_j is not in $\overline{H}_r(\mathbf{b}^{t'})$, so no candidate not in $\overline{H}_{r+1}(\mathbf{b}^{t'})$ can dominate it (both make no difference to the result of the elections in $S(\mathbf{b}^{t'}, r)$).

Finally, by property 2, each voter moves at most $m - 1$ times before the game converges □

Observation 5.1. *When using plurality and when all voters use local dominance strategy and have the same radius r , if the initial state was one when all voters were truthful, either this situation is stable or in every state \mathbf{b}^t we*

have $|\overline{H}_{r+1}(\mathbf{b}^t)| > 1$. Also, in the stable state either $|\overline{H}_r| = 1$ or all voters vote for possible winners. Any voter voting for $c \notin \overline{H}_{r+1}$ prefers the winner in the stable state over any other candidate in \overline{H}_{r+1} .

Proof. If $|\overline{H}_{r+1}(\mathbf{b}^t)| = 1$, there are no longer any strategic moves (as no vote can change the outcome), so we only look at the final vote, which we shall call state \mathbf{s} , at which the vote $b_i \rightarrow b'_i$ happened. Since in the following state $|\overline{H}_{r+1}| = 1$, we know $|\overline{H}_r(\mathbf{s})| = 1$, but in this case, a vote to b'_i will dominate b_i only if $b'_i \in \overline{H}_r(\mathbf{s})$ (otherwise, it is increasing the score of a weak candidate, which does not “empty” H_{r+1}), but that is only the winner of \mathbf{s} , which does not dominate b_i (as voting for b_i will never harm this candidate), contradicting the existence of the move.

If in a stable state \mathbf{b} $|\overline{H}_r(\mathbf{b})| > 1$, there is at least one case in $S(\mathbf{b}, r)$ where two candidates are tied. Hence a voter not voting for a possible winner has a dominating strategy to vote for the candidate it prefers in the tie. A voter not voting for a candidate in $H_{r+1}(\mathbf{b})$ is satisfied with the winner, otherwise it would have a dominating strategy to vote for the candidate it prefers there. \square

Theorem 5.2. *When using plurality and when all voters use local dominance strategy and are truth-biased with the same radii (r, k) , if the initial state was one when all voters were truthful, the iterative process converges to a stable state.*

Proof. Our proof is for ℓ_1 , but can be easily modified for ℓ_i for $1 < i$ and ℓ_∞ . We shall follow the proof for Theorem 5.1 quite closely, and use its 1-4 properties which, together, implied convergence.

If the truthful state \mathbf{a} is stable, then we are done. Thus assume it is not. Therefore, we know that in the opening state, $|\overline{H}_{r+1}| > 1$. We wish to show that properties 3 and 4 of the proof of Theorem 5.1 (winner’s score never decreases, and set of potential winner can only narrow) remain true, and as they ensure convergence, that is enough.

Suppose a truth biased voter i is making a change to their truthful vote a_i for the first time in the iterative process when the profile is \mathbf{s} when

the profile is \mathbf{s} and it is voting for $b_i \neq a_i$. Hence, from Observation 5.1, $|\overline{H}_{r+1}(\mathbf{s})| > 1$. If $b_i \in \overline{H}_{r+1}(\mathbf{s})$ we know it is the favorite candidate of voter i in that group, and hence it is tied with at least one other candidate in a radius of $r + 1 \leq k$, and hence, it would not move to its truthful voter. So $b_i \notin \overline{H}_{r+1}(\mathbf{s})$, hence moving from b_i to a truthful voter will not decrease the score of the winner. Moreover, it cannot bring a_i into \overline{H}_{r+1} , since any voter which ranked a_i first and left it (including voter i) did this when a_i 's score was $r + 1$ or more points behind the winner. Since voters leave candidates only when they're $r + 1$ points from the winner score, and when voter i moved to vote for b_i it was at most r points from the winner, and is no longer so, the score of the winner has increased. Hence, even if all voters which rank a_i first will return to vote for a_i , it will not be within a radius $r + 1$ of the winner, so \overline{H}_{r+1} will not be affected.

Notice this means that even after the truthful move $|\overline{H}_{r+1}(\mathbf{s})| > 1$, and hence we can repeat these claim not only for the first truthful move, but for all of them. \square

Theorem 5.3. *When using plurality and when all voters use local dominance strategy and are lazy-biased with the same radii (r, k) , if the initial state was one when all voters were truthful, the iterative process converges to a stable state.*

Proof. Our proof is for ℓ_1 , but can be easily modified for ℓ_i for $1 < i$ and ℓ_∞ . We shall follow the proof for Theorem 5.1 quite closely, and use its 1-4 properties which, together, implied convergence.

Suppose in the truthful state \mathbf{a} , $|\overline{H}_{r+1}| = 1$. Voters for the winner will choose to abstain only if the difference in score between the winner and the runner up is more than k . However, once the difference is $k + 1$, no other supporter of the winner will abstain, yet no voter will be able to make a strategic move, since this means that still $|\overline{H}_{r+1}| = 1$. Hence that would be a stable state.

Now assume in the truthful opening state, $|\overline{H}_{r+1}| > 1$. We wish to show that properties 3 and 4 of the proof of Theorem 5.1 (winner's score never

decreases, and set of potential winner can only narrow) remain true, and as they ensure convergence, that is enough.

Suppose a lazy biased voter i is abstaining for the first time in the iterative process when the profile is \mathbf{s} and it is voting for b_i . Hence, from Observation 5.1, $|\overline{H}_{r+1}(\mathbf{s})| > 1$. If $b_i \in \overline{H}_{r+1}(\mathbf{s})$ we know it is the favorite candidate of voter i in that group, and hence it is tied with at least one other candidate in a radius of $r+1 \leq k$, and hence, the voter would not abstain. So $b_i \notin \overline{H}_{r+1}(\mathbf{s})$, hence moving from b_i will not decrease the score of the winner and it cannot bring any candidate into \overline{H}_{r+1} .

Notice this means that even after the truthful move $|\overline{H}_{r+1}(\mathbf{s})| > 1$, and hence we can repeat these claim not only for the first truthful move, but for all of them. \square

Appendix B

Proofs of Part II

B.1 Chapter 7

Lemma 7.1. *The colluders' bid monotonically decreases with k , and monotonically increases with n , up to $\frac{1}{e}$.*

Proof. We have to show that the derivatives of $b^*(\cdot)$ with respect to k and n are negative and positive, respectively. We have:

$$(b^*(k))' = -\frac{\left(\frac{n-k}{n-1}\right)^{\frac{n-1}{k-1}}}{(k-1)^2} \left((n-k) \ln\left(\frac{n-k}{n-1}\right) + k - 1 \right)$$

For any $1 \leq k < n$, the first multiplicative term is positive, we only need to examine the sign of the second term, and so it suffices to show that $(n-k) \ln\left(\frac{n-k}{n-1}\right) + k - 1 > 0$. Using the standard logarithm inequality $\ln(1+z) \geq \frac{z}{1+z}$, we obtain the required result:

$$(n-k) \ln\left(\frac{n-k}{n-1}\right) + k - 1 \geq (n-k) \frac{1-k}{n-k} + k - 1 = 0$$

Now, differentiating w.r.t. n , we get

$$(b^*(n))' = \frac{\left(\frac{n-k}{n-1}\right)^{\frac{n-1}{k-1}} (n-k) \ln\left(\frac{n-k}{n-1}\right) + k - 1}{(n-k)(k-1)}$$

By the same inequality as above,

$$(b^*(n))' \geq \frac{\left(\frac{n-k}{n-1}\right)^{\frac{n-1}{k-1}}(k-1) + k - 1}{(n-k)(k-1)} = \frac{1 - \left(\frac{n-k}{n-1}\right)^{\frac{n-1}{k-1}}}{n-k} > 0$$

for any $1 \leq k < n$, as required.

Finally, we rewrite b^* as $\left(1 - \frac{k-1}{n-1}\right)^{\frac{1}{k-1}}$ and note that for a fixed k and $n \rightarrow \infty$, we have that $b^* \rightarrow (e^{-(k-1)})^{\frac{1}{k-1}} = e^{-1}$, completing the proof. \square

Lemma 7.2. *The colluders' expected profit decreases with n and increases with k .*

Proof. The overall expected profit for colluders when bidding optimally is:

$$\pi(b^*) = (b^*)^{\frac{n-k}{n-1}} - b^* = \left(\frac{n-k}{n-1}\right)^{\frac{n-k}{k-1}} \cdot \frac{k-1}{n-1}$$

Differentiating this w.r.t. n gives

$$\frac{\left(\frac{n-k}{n-1}\right)^{\frac{n-k}{k-1}} \ln\left(\frac{n-k}{n-1}\right)}{n-1}$$

which is negative as the first multiplicative term in the numerator is positive, and the logarithm of $\frac{n-k}{n-1} < 1$ is negative. Thus, the profit is monotonically decreasing in n . Now, taking the derivative w.r.t. k results in

$$\frac{\left(\frac{n-k}{n-1}\right)^{\frac{n-k}{k-1}} \left(-\ln\left(\frac{n-k}{n-1}\right)\right)}{k-1}$$

This expression is positive using the same argument as before, and so the total expected profit of colluders increases with their number, k . \square

Theorem 7.1. *The expected profit per colluder increases with k .*

Proof. The individual expected profit for each member of the coalition is:

$$h(k) = \frac{\pi(b^*)}{k} = \left(\frac{n-k}{n-1}\right)^{\frac{n-k}{k-1}} \cdot \frac{k-1}{k(n-1)}$$

The derivative w.r.t. k is given by:

$$h'(k) = -\frac{(n-1)\frac{n-k}{n-1}\frac{n-k}{k-1}(k(n-1)\ln(\frac{n-k}{n-1}) + (k-1)^2)}{(k-1)k^2}$$

It suffices to show that the last multiplicative term in the numerator is negative, or, equivalently, that:

$$\ln\left(\frac{n-k}{n-1}\right) < -\frac{(k-1)^2}{k(n-1)}$$

To this end, we use the standard logarithm inequality $\ln(1+x) \leq x$. As required we have:

$$\ln\left(\frac{n-k}{n-1}\right) = \ln\left(1 + \frac{1-k}{n-1}\right) \leq \frac{1-k}{n-1} < -\frac{(k-1)^2}{k(n-1)}$$

□

Theorem 7.2. *In the setting with k colluders, the expected auctioneer utility is $\frac{n-k}{n} + \left(\frac{n-k}{n-1}\right)^{\frac{n-1}{k-1}}$ in the sum-profit model and $\frac{n-k}{2n-k-1} \left(1 + \frac{n-k}{n-1}\right)^{\frac{2(n-k)}{k-1}}$ in the max-profit model. The profit in both models decreases in the number of colluders and increases in the total number of participants. For sufficiently large n , they exceed the corresponding auctioneer's utilities in the setting without collusion.*

Proof. The expected profit of a sum-profit auctioneer is given by replacing the bids of $\frac{1}{n}$ for each of the k colluders with a joint single bid of $\left(\frac{n-k}{n-1}\right)^{\frac{n-1}{k-1}}$. This results in a total bid sum of $\frac{n-k}{n} + \left(\frac{n-k}{n-1}\right)^{\frac{n-1}{k-1}}$. Thus, collusion is obviously profitable for the auctioneer whenever the colluders' bid is larger than $\frac{k}{n}$. This, broadly speaking, is common for smaller k and larger n (as, by Lemma 7.1, the bid increases with n and decreases with k).

For the max-profit model, we examine the maximal bid's distribution, defined by the c.d.f. G_{AP} as follows:

$$G_{AP}(z) = \begin{cases} 0 & z < \frac{n-k}{n-1} \frac{n-1}{k-1} \\ \left(\frac{n-k}{n-1}\right)^{\frac{n-k}{k-1}} & z = \frac{n-k}{n-1} \frac{n-1}{k-1} \\ z^{\frac{n-k}{n-1}} & z > \frac{n-k}{n-1} \frac{n-1}{k-1} \end{cases}$$

Where $G_{AP}(z)$ is not constant, its derivative is $\frac{n-k}{n-1} z^{\frac{1-k}{n-1}}$, so the expected auctioneer's profit is:

$$\begin{aligned} E(AP) &= \left(\frac{n-k}{n-1}\right)^{\frac{2n-k-1}{k-1}} + \int_{\left(\frac{n-k}{n-1}\right)^{\frac{n-1}{k-1}}}^1 \frac{n-k}{n-1} z^{\frac{n-k}{n-1}} dz = \\ &= \frac{n-k}{2n-k-1} \left(1 + \left(\frac{n-k}{n-1}\right)^{\frac{2(n-k)}{k-1}}\right) \end{aligned}$$

We compare this with the expected auctioneer's profit in the case of no collusion. To do so, we rewrite it as follows:

$$\frac{n-k}{2n-k-1} + \frac{(n-1)^2}{(2n-k-1)(n-k)} \left(\left(1 - \frac{k-1}{n-1}\right)^{n-1}\right)^{\frac{2}{k-1}}$$

The value of the above expression wobbles for low n and k , but for a fixed k and increasing n (i.e., $n \rightarrow \infty$) it approaches $\frac{n-k}{2n-k-1} + \frac{(n-1)^2}{(2n-k-1)(n-k)} e^{-2}$. That is, while the expected profit without collusion is edging close to $\frac{1}{2}$, with colluders, the profit is closing in on $(\frac{1}{2} + \frac{1}{2e^2})$, and there exists a number of participants n for which the profit with collusion is strictly greater than without collusion. For a fixed n , the profit is monotonically decreasing in k , which fits with earlier results indicating a lower bid from colluders, as their cohort grows. \square

Theorem 7.3. *The social welfare in the sum-profit model does not change due to collusion. In the max-profit model, the presence of colluders may have different effects on the social welfare, depending on the relation between the number of colluders and the total number of participants. In particular, the social welfare drops for settings with many participants.*

Proof. We calculate the expected profit for a non-colluding player, which is defined by a p.d.f. g below, depending on the colluders' bid b^* :

$$g(z) = \begin{cases} f(-z)(1 - F_n^{n-k}(-z)) & -1 \geq z < -b^*(k) \\ f(-z) & -b^*(k) \leq z \leq 0 \\ f(1-z)F_n^{n-k}(1-z) & 0 < z < 1 - b^*(k) \end{cases}$$

The non-colluder's expected profit is:

$$\begin{aligned} E(z) &= \int_{-1}^{-b^*} \frac{z}{n-1} (-z)^{\frac{2-n}{n-1}} (1 - (-z)^{\frac{n-k-1}{n-1}}) dz + \\ &+ \int_{-b^*}^0 \frac{z}{n-1} (-z)^{\frac{2-n}{n-1}} dz + \\ &+ \int_0^{1-b^*} \frac{z}{n-1} (1-z)^{\frac{2-n}{n-1}} (1-z)^{\frac{n-k-1}{n-1}} dz = \frac{k - n(b^*)^{\frac{n-k}{n-1}}}{n(n-k)} \end{aligned}$$

This expression may be positive or negative, depending on k and n . As the colluders' bid does not exceed $\frac{1}{e}$, for small k the expression takes a negative value. However, when k is rather large with regard to n (e.g., when k is roughly $\frac{n}{2}$), it is positive. That is, the non-colluders may benefit from collusion, not despite not being aware of it, but due to their lack of awareness.

Summing up the expected profits of all the parties in the sum-profit model results in the same social welfare as in the case of mergers or of no bidder cooperation:

$$\frac{k}{n} - \frac{n-k}{n-1} \frac{n-k}{k-1} + \frac{n-k}{n-1} \frac{n-k}{k-1} - \frac{n-k}{n-1} \frac{n-1}{k-1} + \frac{n-k}{n} + \frac{n-k}{n-1} \frac{n-1}{k-1} = 1$$

However, in the max-profit model the results are more ambiguous. For very large n , looking coarsely at the non-colluders' expected profit, we see that when we have $n - k$ such players, the sum of their expected losses

approaches $-\frac{1}{e}$. We already know that in this scenario the expected profit of colluders is zero, so we need to examine this in relation to the changes in the profits of the auctioneer. In this case, the auctioneer's profit approaches $\frac{1}{2} + \frac{1}{2e^2}$, so the social welfare drops below $\frac{1}{2}$, which is lower than what would happen without colluders. \square

B.2 Chapter 8

Theorem 8.1. *The F_i presented above are a Nash equilibrium, and each bidder's profit is λ .*

Proof. When bidder i bids according to this distribution, i.e., $b \in [\alpha_k, \alpha_{k-1})$ for $1 \leq k \leq i$:

$$\begin{aligned}
\pi_i(b) &= (1-b) \prod_{j=1; j \neq i}^{n-1} (p_j F_j(b) + 1 - p_j) - \\
&\quad - b(1 - \prod_{j=1; j \neq i}^{n-1} (p_j F_j(b) + 1 - p_j)) = \\
&= \prod_{j=1; j \neq i}^{n-1} (p_j F_j(b) + 1 - p_j) - b = \\
&= \prod_{j=1}^{k-1} (1 - p_j) \prod_{j=k; j \neq i}^n H_k(x) - b = \\
&= H_k^{n-k}(b) \prod_{j=1}^{k-1} (1 - p_j) - b = \\
&= \frac{\lambda + b}{\prod_{j=1}^{k-1} (1 - p_j)} \prod_{j=1}^{k-1} (1 - p_j) - b = \\
&= \lambda
\end{aligned}$$

If bidder i bids outside their support, i.e., $b \in [\alpha_k, \alpha_{k-1})$ for $i + 1 \leq k \leq n - 1$, the same equation becomes

$$\begin{aligned}
\pi_i(b) &= \prod_{j=1; j \neq i}^{k-1} (1-p_j) \prod_{j=k}^n H_k(x) - b = \\
&= \prod_{j=1; j \neq i}^{k-1} (1-p_j) \left(\frac{\lambda+b}{\prod_{j=1}^{k-1} (1-p_j)} \right)^{\frac{n-k+1}{n-k}} - b = \\
&= \frac{\lambda+b}{1-p_i} \left(\frac{\lambda+b}{\prod_{j=1}^{k-1} (1-p_j)} \right)^{\frac{1}{n-k}} - b
\end{aligned}$$

Since $b \in [\alpha_k, \alpha_{k-1})$, then $b < \alpha_{k-1} = (1-p_{k-1})^{n-k+1} \prod_{j=1}^{k-2} (1-p_j) - \lambda = (1-p_{k-1})^{n-k} \prod_{j=1}^{k-1} (1-p_j) - \lambda$, and hence $\lambda+b < (1-p_{k-1})^{n-k} \prod_{j=1}^{k-1} (1-p_j)$. Therefore:

$$\begin{aligned}
\pi_i(b) &< \frac{\lambda+b}{1-p_i} \left(\frac{(1-p_{k-1})^{n-k} \prod_{j=1}^{k-1} (1-p_j)}{\prod_{j=1}^{k-1} (1-p_j)} \right)^{\frac{1}{n-k}} - b = \\
&= \frac{\lambda+b}{1-p_i} (1-p_{k-1}) - b = (\lambda+b) \frac{p_i - p_{k-1}}{1-p_i} + \lambda
\end{aligned}$$

Finally, as $i+1 \leq k$, $p_i \leq p_{k-1}$, hence $p_i - p_{k-1} \leq 0$, and $\pi_i(b) < \lambda$.

Note that if the valuation of the bidder is $\ell < 1$, the first equation becomes $\ell(\lambda+b) - b$, which is maximized for $b=0$, and the second equation is still smaller than $\ell(\lambda+b) - b$, so still maximized for $b=0$. \square

Appendix C

Proofs of Part III

C.1 Chapter 10

Theorem 10.1. *Computing the value $v(C)$ of a coalition C in a weakest link game can be done in polynomial time.*

Proof. Due to Observation 10.1, $v(C)$ takes one of the values in W or 0. For each of the possible edge weights $\tau \in W$, we can test whether there exists an s - t path that is comprised solely of the edges in C whose weight is at least τ , as follows: Let C^τ be the set of edges in C with weight at least τ . Denote by $G'(V, C^\tau)$ the subgraph with vertex set V and edge set C^τ . The graph $G'(V, C^\tau)$ can easily be computed in polynomial time, by iterating through the edges and eliminating those that have a weight lower than τ .

Given $G'(V, C^\tau)$ we can check whether there exists any path connecting s and t in it using a depth-first search (DFS), which again requires polynomial time. If such a path exists we say the test was positive for τ , which indicates that $v(C) \geq \tau$, and if such a path does not exist we say the test was negative, indicating that $v(C) < \tau$.

After iterating over all possible values $\tau \in W$ we return the maximal τ for which the test was positive. Since $|W| \leq |E|$ the entire procedure requires polynomial time. \square

Theorem 10.2. *Testing whether an imputation $p = (p_1, \dots, p_n)$ is in the core of a weakest link game can be done in polynomial time.*

Proof. We provide a polynomial time algorithm that takes a WLGE over a graph $G(V, E)$ and an imputation $p = (p_1, p_2, \dots, p_n)$ and either finds a coalition B that blocks the imputation or verifies that no such blocking coalition exists, so p is in the core.

To check if there exists a coalition B with $v(B) > p(B)$, we iterate over all possible values that $v(B)$ can take. By Observation 10.1 it suffices to use a procedure that searches for blocking coalitions with value exactly τ , and run it for all possible values $\tau \in W$. If no blocking coalition is found whose value is exactly τ for any τ in the set W , no blocking coalition exists.

If a coalition B has value τ , it must contain a path P connecting s and t consisting solely of edges with weight at least τ . The value of path P as a coalition is also $v(P) = \tau$. Thus if B is a blocking coalition, P is also a blocking coalition. Therefore, to find a blocking coalition B where $v(B) = \tau$ it suffices to examine all the paths P where $v(P) = \tau$. If there are several such paths P where $v(P) = \tau$, it suffices to examine the path Q with minimal payoff $p(Q) = \sum_{i \in Q} p_i$: if $p(Q) < v(Q) = \tau$ then we have a blocking coalition Q , and if $p(Q) \geq v(Q) = \tau$ then for any path Q' where $v(Q') = \tau$ we have $p(Q') \geq p(Q) \geq v(Q) = \tau$ so Q' cannot be a blocking coalition.

Therefore, to seek a blocking coalition B where $v(B) = \tau$ it suffices to examine only the minimal payoff path P where $v(P) = \tau$ (i.e., an s - t path Q where $v(Q) = \tau$ that minimizes $p(Q) = \sum_{i \in Q} p_i$ of all such paths with value τ). If this path does not constitute a blocking coalition then there are no blocking coalitions with value τ .

To search for such a path, we construct a weighted graph G_τ with the same vertices as G , while dropping all edges where $w(e) < \tau$, retaining only edges with weight of τ or more. However, we change the weights of the retained edges – we replace the weight of an edge $e \in E$ with its payoff under the imputation, so $w'(e) = p_e$ (by $w'(e)$ we denote the new weight). In the generated graph G_τ we can find the “shortest” s - t path S_τ , under the new

weights, using Dijkstra's algorithm.

The payoff of S_τ under the imputation p is its total length in G_τ , under the new weights. If $p(S_\tau) < \tau$ then S_τ is a blocking coalition with value at least τ , and if $p(S_\tau) > \tau$ then no blocking coalition with value τ exists.¹ Since the above procedure takes polynomial time, and is repeated $|W| < |E|$ times (for each possible value of τ), the entire algorithm has a polynomial running time. \square

Theorem 10.3. *Testing core emptiness, finding an ϵ -core imputation and finding the least core value take polynomial time for WLGs.*

Proof. The algorithm of Theorem 10.2 can serve as a separation oracle for the core LP 10.1. It takes a proposed imputation $p = (p_1, \dots, p_n)$ and either returns a blocking coalition yielding a violating constraint, or verifies that no such coalition exists, in which case all LP constraints are satisfied. Thus it is possible to solve the core LP 10.1 in polynomial time, and either find a core imputation or verify that the core is empty.

We note that it is easy to adapt the algorithm in Theorem 10.2 to serve as a separation oracle for the ϵ -core LP 10.1. Rather than checking whether a path forms a blocking coalition for a given value of τ , we can perform a relaxed test: checking whether it is blocking by a margin of at least ϵ by constructing G_τ as in Theorem 10.2, finding the shortest path S_τ , and checking if $\tau = v(S_\tau) < p(S_\tau) - \epsilon$. Since we have a separation oracle for the ϵ -core LP 10.1, it can be solved in polynomial time, allowing us to either find an ϵ -core imputation or verify that the ϵ -core is empty.

To find the least core value we can perform a binary search on the minimal value of ϵ , at each step solving the ϵ -core LP 10.1 for the current value of ϵ . We repeat this as many times as needed to compute the least-core value up

¹Note that decreasing weights of some edges potentially reduces the values of some coalitions; thus the procedure might “miss” a blocking coalition, when the true value of the coalition under the new weights is lower than under the true weights. However, this is not a coalition whose value is τ , but rather one whose value is $\tau' > \tau$. This would be found later, when examining the value τ' .

to any required degree of numerical accuracy. \square

Corollary 10.1. *Calculating the Cost of Stability of a weakest link game can be done in polynomial time.*

Proof. Solving LP 10.2 allows finding the CoS. Again, we use the algorithm of Theorem 10.2 as a separation oracle, but make the appropriate changes so as to solve LP 10.2 rather than LP 10.1.

When testing whether an imputation is stable we use the constraint $\sum_{i=1}^n p_i = v(N)$. To switch from LP 10.1 to LP 10.2, we replace this constraint with the constraint $\sum_{i=1}^n p_i = v(N) + \Delta$, which tests whether there exists a stable Δ super-imputation (i.e., a payoff vector allocating a total of $v(N) + \Delta$). Changing the target function to be $\min \Delta$ (rather than just the feasibility goal of LP 10.1) results in the CoS formulation LP 10.2, which we solve in polynomial time using the same separation oracle. \square

Theorem 10.4. *If graph G is a parallel composition of graphs G_i , the CoS of G is $(\sum_{G_i} CoS(G_i) + v_i(G_i)) - \max_{G_i}(v_i(G_i))$.*

Proof. First, we show the CoS is not larger than the theorem's value. Since this is a WLG, the value of the grand coalition of the composition cannot be greater than the one of the maximal G_i , and as this is a parallel composition, it must be equal to the maximal G_i . Examine the super-imputation with the minimal sum of each G_i when it is considered on its own. For each graph, the sum of this super-imputation is $CoS(G_i) + v_i(G_i)$, which, when summed over and subtracting the value of the grand coalition, the CoS is $(\sum_{G_i} CoS(G_i) + v_i(G_i)) - \max_{G_i}(v_i(G_i))$, as we wanted.

Now, suppose there is a blocking coalition C , for which $v(C) > \sum_{j \in C} p_j$. As there are no edges connecting the separate G_i s, every route between s and t passes through only a single G_i , so there is an i for which $v(C \cap G_i) > \sum_{j \in C \cap G_i} p_j$. However, since $p \cap G_i$ is a super-imputation over G_i , that is impossible.

A CoS smaller than our lower bound state above is not possible either. Suppose there is a super-imputation with a smaller sum, so there is a G_i for which $\sum_{j \in G_i} p_j < v_i(G_i) + CoS(G_i)$. This contradicts the very definition of the CoS. \square

Theorem 10.5. *If G is a series composition of the graphs G_i , the CoS of G is $\min_i CoS(G_i^{\min_{j \neq i}(v(G_j))})$, where $G_i^{\min_{j \neq i}(v(G_j))}$ is G_i in which all edges with weight above $\min_{j \neq i}(v(G_j))$ are lowered to that value.*

Proof. We first show the CoS cannot be larger. Note that every path from s to t has a maximal value of $\min_j(v(G_j))$, so no path from s_i to t_i can have a larger value.

A valid super-imputation is a super-imputation of $G_i^{\min_{j \neq i}(v(G_j))}$, giving 0 to everyone else. As all routes from s to t pass through G_i (with the capacity limit), the super-imputation does not induce any coalitions which do not receive their value – if there is such a coalition, it is particularly also a coalition from s_i to t_i with the same value, and that is what the super-imputation of $G_i^{\min_{j \neq i}(v(G_j))}$ deals with.

Now, we prove a smaller CoS is not possible by induction. Given graphs G_1 and G_2 , suppose the CoS is smaller than $CoS(G_1^{v(G_2)})$ and $CoS(G_2^{v(G_1)})$. Then construct a path made of a single path from s_1 to t_1 with value of $v(G_1)$ and from t_1 , all of G_2 . This is actually $G_2^{v(G_1)}$ (due to the constraints of the first path), and we know that the smaller imputation does not satisfy it, i.e., there is a coalition of edges from s_2 to t_2 which have an incentive to leave the grand coalition (with the single path from s_1 to t_1 , of course).

For any n graphs, we look at the first $n - 1$ graphs as a single graph G' , hence $CoS(G) = \min(CoS(G'^{v(G_n)}), CoS(G_n^{v(G')}))$. Since $v(G') = \min_{i \neq n} v_i(G_i)$ and from the induction definition

$$CoS(G'^{v(G_n)}) = \min_{i \neq n} (G_i^{\min(v(G_n), \min_{j \neq i, j \neq n}(v(G_j)))}) = \min_{i \neq n} (G_i^{\min_{j \neq i}(v(G_j))})$$

as required. \square

Theorem 10.6. *It is NP-complete to determine whether the value of the optimal coalition structure in a weakest link game exceeds an input k .*

Proof. We use a reduction from the NP-complete problem Disjoint Paths Problem (DPP).² In the DPP problem we are given an undirected graph $G(V, E)$ and k pairs of source-target vertex pairs $\{(s_i, t_i)\}_{i=1}^k$, and are asked whether there are k edge-disjoint paths in G such that the i 'th path connects s_i and t_i .

We reduce a DPP to finding the optimal coalition structure in a WLG. We take the original graph $G(V, E)$ and add two special vertices: a meta-source s and a meta-target t . We add k edges from s to the k sources $\{s_i\}_{i=1}^k$ with weight $1 - \epsilon_i$ for an arbitrary set of k distinct values $\{\epsilon_i\}_{i=1}^k$ in range $(0, 1)$ (by distinct we mean that $\epsilon_i \neq \epsilon_j$ for any $i \neq j$). Similarly we add an edge from each t_i to t with weight $1 - \epsilon_i$ for any $1 \leq i \leq k$. We set the weights of all edges in G to be 1.

In the optimal coalition structure problem we search for disjoint paths between s and t maximizing the sum of the values of the paths. At best, for each $1 \leq i \leq k$, there is a path from s_i to t_i , and one can use the edges (s, s_i) and (t_i, t) , each with weight $1 - \epsilon_i$, to complete it to an s - t path. Thus $\sum_{i=1}^k (1 - \epsilon_i)$ is an upper bound for the optimal coalition structure's value in the reduced instance.

This upper bound is achieved only if the weights of the two end-edges of our s, t paths match: if one of our paths starts with weight $1 - \epsilon_i$ and ends with weight $1 - \epsilon_j$ for some $i \neq j$, there is no way to complete this solution with total value $\sum_{i=1}^k 1 - \epsilon_i$. In this case, we only get $\min\{w((s, s_i)), w((t_j, t))\}$ for this part of the partition, failing to achieve a total value of $\sum_{i=1}^k (1 - \epsilon_i)$.

Therefore, the generated instance of the WLG optimal coalition structure input allows a solution of total value of $\sum_{i=1}^k (1 - \epsilon_i)$ if and only if the DPP instance is a positive instance (i.e., if there are k edge-disjoint paths connecting the pairs $\{s_i, t_i\}_{i=1}^k$).

NP membership as given an optimal coalition, calculating its value is done in polynomial time, and can than be compared to k . \square

Theorem 10.7. *A polynomial time $\mathcal{O}(\log n)$ -approximation exists for the*

²NP-completeness for DPP shown in Karp [126].

optimal coalition structure problem in weakest link games.

Proof. We first consider the following problem: given a weighted graph $G(V, E)$ with designated source vertex $s \in V$, target $t \in V$, and threshold τ , find the maximal number of edge-disjoint s - t paths that only use edges whose weight is at least τ . We present a polynomial time algorithm to solve this problem.

First, remove all edges that weigh below the threshold τ , and set the weights of the remaining edges to be 1 (unit weight), to obtain the graph G_τ . Note that every path in G that only uses edges whose weight is at least τ is equivalent to a path in G_τ .

Thus, it suffices to find the maximal number of edge-disjoint s - t paths in G_τ , which can be done by finding the maximal flow between s - t (for example by using the Edmonds-Karp maximal-flow algorithm).

The value of this flow is the maximal number of edge-disjoint s - t paths in G_τ , since due to unit capacity no edge is used twice (a partition into paths can be obtained by keeping track of augmenting paths found during the run).

Let w' be the value of the coalition of all agents, i.e., $v(N)$. Define n_i to be the maximum number of disjoint s - t paths in G that only use the edges with weight at least $\frac{w'}{2^i}$. The value of the optimal coalition structure is upper-bounded by $\sum_{i=1}^{\infty} n_i \frac{w'}{2^{i-1}}$. Because the number of coalitions in the optimal solution with value in the range $[\frac{w'}{2^i}, \frac{w'}{2^{i-1}}]$ does not exceed n_i , and for each of them we get value at most $\frac{w'}{2^{i-1}}$.

To find an $\mathcal{O}(\log(n))$ approximation of the optimal coalition structure, we perform the following procedure. For all possible thresholds τ in the set W , we find the maximum number of disjoint paths in G_τ . We then find the value $\tau = \tau^*$ that maximizes the product of τ and the number of disjoint paths in G_τ . We claim that these disjoint paths in G_{τ^*} form an $\mathcal{O}(\log(n))$ approximation solution.

The analysis is similar to the $\log(n)$ -competitive algorithms for the matroid secretary problem [21]. We prove that in the sum $\sum_{i=1}^{\infty} n_i \frac{w'}{2^{i-1}}$, the sum of terms for $i > 2 \log(n)$ is not more than $2 \frac{w'}{n}$ which is at most $\frac{2}{n}$ fraction of

the whole sum.

We know that n_i is at most n , the number of agents, for every i . Thus the sum of those terms is not more than $nw' \sum_{i=2 \log(n)+1}^{\infty} \frac{1}{2^{i-1}} = n \frac{w'}{2^{2 \log(n)-1}} = 2 \frac{w'}{n}$. We conclude that more than $1 - \frac{2}{n}$ fraction of the sum is concentrated in the first $2 \log(n)$ terms, and consequently there exists an i for which $n_i \frac{w'}{2^{i-1}}$ is at least $\frac{1-\frac{2}{n}}{2 \log(n)}$ fraction of the sum.

By the definition of τ^* , we know the solution we get has value of at least $n_i \frac{w'}{2^i}$, which proves that our solution has at least $\Omega(\log(n))$ fraction of the above sum, and therefore it is an $\mathcal{O}(\log(n))$ approximation. \square

C.2 Chapter 11

Theorem 11.1. *Axioms 1–6 and axioms 1–3, 7–9 (the sets we deal with) are all independent of one another.*

Proof. We list some odd mechanisms that are consistent with each 5 of our axioms, demonstrating the necessity of each (we present here the main points of the odd mechanism – the complete version can be constructed by using the other axioms). First, we begin with axioms 1–6.

- *Anonymity:* A mechanism that recommends + for every singleton.
- *Positive response:* A system for which if the group contains a nonvoter which is influenced by both + and - nodes (of whatever number) from outside the group, it is recommended 0.
- *IIS:* Except for star groups, all nodes outside a group are considered to be influencing all nodes inside it.
- *α -centripetal:* All star groups are always recommended - (including those consisting of arbitrarily large M of + voters, and one nonvoter connected to a sole - voter).




Figure C.1: Internal consistency problem

- *(β, r) -centrifugal*: All star groups are always recommended + (including those consisting of one + voter, arbitrarily large M_1 of nonvoters, each connected to arbitrarily large M_2 - voters).
- *Internal consistency*: Taking 3 nonvoters, 3 + voters, and a single - voter. Connecting each nonvoter to a single + voter, and two of the nonvoters are also connected to the - voter (see Figure C.1). Our groups are made of the pairs of + voters with the nonvoters connected to them. When we take a single pair, one that is also connected to the - voter, the recommendation is +. Adding another pair that is not connected to the - voter, the recommendation is now -.

We now continue with axioms 1–3 and 7–9:

All examples given in appendix 1 of [12] for axioms 1–3 work. Hence, we just need to show for 7–9:

- *Trust propagation*: A group is recommended by summing for each vertex over number of + and - nodes it is directly connected to.
- *Scale invariance*: A group is recommended by a random walk from each vertex in the group, but to the relative influence of each edge is added the number of outgoing edges from the influencing vertex.

- *Proportional inclusiveness*: A group ignores all outside influences – performs simple majority on + and - nodes inside it.

□

Theorem 11.2. *No recommendation system satisfies axioms 1 – 6, i.e., is anonymous, positive responsive, IIS, is internally consistent and is α – centripetal and (β, r) – centrifugal for $1 < \alpha, \beta < \infty$.*

Proof. First, we note that if the (β, r) – centrifugal axiom uses an $r \notin \mathbb{N}$, then we shall use as $r := \lceil r \rceil$.

Our proof is constructed using 3 steps in which we build a graph, find a specific group in it, and show that the axioms require it to both be recommended + and -, creating a contradiction.

Step 1: Build a graph

Since $\alpha > 1$, there is some $\ell \in \mathbb{N}$ for which $\ell\alpha - (\ell + 2) \geq 0$. We define k as $\lceil \frac{\beta-1}{\alpha-1} \rceil + \ell$ and s as $\lfloor k\alpha \rfloor$. We now build the following graph, consisting of a star group containing k + voters (v_1, \dots, v_k) and $k \cdot r$ non voters (u_1, \dots, u_{kr}) , and outside the star group are $s \cdot k \cdot r$ - voters (t_1, \dots, t_{skr}) , with each nonvoter connected to $s \cdot r$ - voters (i.e., every nonvoter u_i has the edges (u_i, t_h) for $h \in \mathbb{N}, i \leq h \leq (i - 1) + s$).

Step 2: Build an indivisible positive set

We now wish to construct a set C which has no partition for which each part is recommended +. We call our star group \tilde{C} , (with k + voters and kr nonvoters). According to the α – centripetal assumption, since $s \leq k \cdot \alpha$ for each nonvoter, \tilde{C} 's recommendation is +. We now seek to find the minimal + part of \tilde{C} – suppose \tilde{C} has a partition for which every part is recommended +. In at least one of these parts the number of nonvoters exceeds (or is equal to) r times its number of voters, and we call it \tilde{C}' , and continue the process. This process ends with a set C with a nonvoters and b + voters ($b \leq rb \leq a \leq rk$), for which the recommendation system recommends +.

We shall now show that according to the (β, r) -centrifugal assumption and internal consistency, it needs to be recommended -, causing a contradiction.

Step 3: Build a contradictory partition

We need to show a partition which results in every part being recommended -, and show that there is no partition for which each part get recommended +. The latter is trivial according to our minimization process – if there is such a partition, then the process has not ended yet. However, we can partition C into a sets of one + voter and r nonvoters (possibly, some nonvoters end without any voter to group them with – they are grouped apart). We shall now show that these sets need to be recommended -, causing the contradiction.

Since + voters in the set are more heavily weighted that outside it, we can focus our proof just for sets with r nonvoters and one voter, and that will suffice for the case of a lone nonvoter. Note that in these sets our nonvoter is connected to one + voter in its set, $k-1$ + and s - voters outside it. According to the positive response axiom (number 2), we can remove one + and one - voters (when both do not belong to the set) from each nonvoter without changing the recommendation. We are left with one + voter connected to the nonvoter in the set, and $s - (k - 1)$ - voters connected to the nonvoter. we now wish to prove that $s - k + 1 \geq \beta$.

Due to k 's definition, we know

$$\begin{aligned} \frac{\beta - 1 + l\alpha - l}{\alpha - 1} \leq k \leq \frac{\beta - 1 + (\ell + 1)\alpha - (\ell + 1)}{\alpha - 1} &= \\ &= \frac{\beta + (\ell + 1)\alpha - (\ell + 2)}{\alpha - 1} \end{aligned}$$

Similarly, we know

$$s \geq \alpha \frac{\beta - 1 + l\alpha - l}{\alpha - 1} - 1 = \frac{\beta\alpha + l\alpha^2 - (\ell + 2)\alpha + 1}{\alpha - 1}$$

Hence:

$$\begin{aligned}
s - k + 1 &\geq \\
&\frac{\beta\alpha + \ell\alpha^2 - (\ell + 2)\alpha + 1 - \beta - (\ell + 1)\alpha + (\ell + 2) + \alpha - 1}{\alpha - 1} = \\
&= \frac{\beta\alpha + \ell\alpha^2 - (2\ell + 2)\alpha - \beta + \ell + 2}{\alpha - 1} = \\
&= \frac{(\alpha - 1)(\beta + \ell\alpha - (\ell + 2))}{\alpha - 1} = \beta + \ell\alpha - (\ell + 2)
\end{aligned}$$

Thanks to our definition of ℓ , this means $s - k + 1 \geq \beta$, hence the set is recommended -, reaching a contradiction. \square

Theorem 11.3. *The group random-walk recommendation system is the only one which satisfies axioms 1 – 3 and 7 – 9, i.e., is anonymous, positive responsive, IIS, has trust propagation, scale invariance and proportional inclusiveness.*

Proof. We shall prove the theorem by looking at a voting system $G(N, V_+, V_-, E)$, and looking at a specific group $C \subseteq N$. We will show what the axioms would force its recommendation to be – and that the same recommendation would be made by the group random-walk recommendation system.

Part 1: Using the axioms

Suppose there is an undecided voter $u \in N \setminus C$ which is influenced by $v_1, \dots, v_m \in N$ (possibly for some $i \neq j$ $v_i = v_j$) and influences w_1, \dots, w_t (again, possibly for some $i \neq j$ $w_i = w_j$). According to scale invariance, recommendations do not change if w_i multiplies its connections by m . Now, using trust propagation, w_i connects directly to v_1, \dots, v_m , and is no longer connected at all to u . As we do this for all w_1, \dots, w_t , when we finish, node u no longer influences any other vertex.

Performing these steps for every undecided voter $u \in N \setminus C$ which is influencing and being influenced, we end up with the nodes in the group C either directly connected by an edge – and being influenced – to voters or nonvoter sinks (i.e., nonvoters which are not influenced by others). Thanks to

the IIS axiom, we can ignore all vertices which are not in the same connected component as C . Note that this axiom also means we can ignore all voters or nonvoter sinks which do not directly influence (with a single edge) any node in the group.

Now, using proportional inclusiveness, we eliminate from C all nonvoters which are influenced by voters, leaving in the group, at most, nonvoter sinks (all other members are voters), i.e., there are now y_- voters for $-$, y_+ voters for $+$ and y_o nonvoter sinks. Suppose $y_o = 0$ – from anonymity axiom we know if $y_- = y_+$, recommendation is 0, and hence from the positive response axiom, the recommendation is $type(\max(y_-, y_+))$ (from the same axiom, that is also the recommendation if $y_o > 0$).

Part 2: Using the recommendation system

Now we need to show that the procedure described above reaches the same recommendation as a group random walk would recommend. This recommendation system, in effect, gives all members of C the same weight (say, 1), and while voters put all their weight on their vote, nonvoters divide their weight according to the random walk. Therefore, we shall show that the contribution of each voter and nonvoter to the final tally is maintained by the changes we do to C in the procedure described above using the axioms.

Furthermore, we note that scale invariance and trust propagation do not affect the result of a random walk, hence we can examine the graph as it looks following our multiple applications of these two axioms (just before we begin to apply proportional inclusiveness). Therefore, we need to show that applying proportional inclusiveness does not change the weight in the group. If we manage to show proportional inclusiveness does not change the recommendation of the group random-walk, since the axiom's application leaves us with a group consisting of voters only (and un-influencable nonvoters) we can conduct a simple plurality between the votes, as both group random-walk and the procedure above indicate should happen.

Let the nonvoter we apply proportional inclusiveness to be u , which is

connected to the voters $v_1, \dots, v_m \notin C$ and to the nodes $v_{m+1}, \dots, v_t \in C$ to each with k_i connections (we define $s = \sum_{i=1}^t k_i$). There are also the nodes $v_{t+1}, \dots, v_r \in C$ which are not connected to u . Notice that u 's "vote" in the group random-walk system gives $\frac{k_i}{s}$ weight to v_i , for $1 \leq i \leq t$.

Following an application of proportional inclusiveness, we now have s copies of v_{t+1}, \dots, v_r in the new C , and $s + k_i$ copies of v_i , $m < i \leq t$. We also have s copies of v_i , $1 \leq i \leq m$, of which k_i copies are in C' , and we have s copies of any nodes which are not v_i ($1 \leq i \leq r$), i.e., any node which was not connected to u and not in C .

Let us focus on nonvoters in $C \setminus \{u\}$. Each copy of the nonvoter is connected to the same nodes as it was before, except for those which were connected to u . Each nonvoter random-walk recommendation before proportional inclusiveness was giving equal weight to each of its connection, so that if a nonvoter had w outgoing edges, hence $\frac{1}{w}$ weight was given to each node connected to it, including u . Therefore, $\frac{k_i}{s} \frac{1}{w}$ weight for each v_i $1 \leq i \leq t$ (the nodes connected to u). Following proportional inclusiveness, the nonvoter has $w \cdot s$ outgoing edges, the weight of each node it is connected to is $\frac{s}{ws} = \frac{1}{w}$, except the nodes v_i ($1 \leq i \leq t$), the weight of which is $\frac{k_i}{ws}$. Hence the random-walk recommendation for each nonvoter remains the same.

The recommendation for the group C remains the same – prior to the proportional inclusiveness each node received a weight of $\frac{1}{|C|}$. Following the application of the axiom, $|C'| = s(|C| - 1) + s = s|C|$, and each node $c \in C$ such that $c \neq v_i$ ($m + 1 \leq i \leq t$) now has the s copies, and as each copy has the same recommendation of each single one, it contributes the weight of $\frac{s}{s|C|} = \frac{1}{|C|}$. Each node connected to nonvoter u contributed $\frac{k_i}{s} \frac{1}{|C|}$, and now each $v_i \notin C$ contributes $\frac{k_i}{s|C|}$, as there are k_i copies of v_i in C' for $1 \leq i \leq m$. $v_i \in C$ now contribute $\frac{s+k_i}{s|C|}$, which is the same of v_i 's contribution as a node of C and as a component of u 's recommendation.

Therefore, each node which in the group random-walk would have an influence over the recommendation maintains that level of influence after applying proportional inclusiveness. Multiple applications of this axiom leave

us with a group consisting of purely voters and nonvoter sinks, which according to the axioms leads to a plurality votes among voters, which is exactly the procedure followed by the group random-walk in this case as well. \square

Appendix D

Information on Local Dominance Simulations

We simulated voting in an iterative setting, where voters start from a particular state, and then iteratively make strategic moves until convergence. We constructed a simulator that enables us to control the following features of the simulation. First, determine the parameters of the preference profile:

- Number of voters. We used $n \in \{10, 20, 50, 100\}$.
- Number of candidates. We used $m \in \{3, \dots, 8\}$.
- Distribution of preferences. We used all the 6 distributions described below.

The simulator allows us to choose different distance metric; different r values; different k values (for truth/lazy-biased voters); and different types of voters in the simulations. We could also set the initial starting point to be truthful or randomly assigned, and a scheduler that allows concurrent updates (both properties results not examined here).

As using different metrics did not result in qualitatively different results, and when starting from a truthful position biases are not significant (though lazy-bias takes longer to converge), we focus here on showing results using ℓ_1 additive metric. We generated profiles from all distribution types for various

numbers of voters and candidates. From each distribution we sampled 200 instances.¹ Then, we simulated strategic voting on each instance varying the distance metric (ℓ_1 , multiplicative ℓ_1), the voters' types (basic, truth-biased, lazy) and the uncertainty parameters r and k . We repeated each simulation 100 times (as the scheduler may pick a different path each time), and collected multiple statistics on the equilibria outcomes.

D.1 Distributions

We included in our simulations several different distributions.²

Uniform Also known as the *impartial culture* distribution, this is the simplest distribution to study. While people's votes are rarely distributed at random, the uniform distribution allows more confidence that our results are not particular and specific to the distributions analyzed, and is thus often used in simulations of voting [174].

Single-peaked This distribution assigns each candidate a point on the interval $[0, 1]$, and each voter is randomly assigned a point on the interval, which defines its preferences — it prefers candidates closer to its point. This distribution has been long used in sociological and political research (as resembling the common right-left political axis) [129], but has also been widely examined in game theoretic scenarios. A particular interesting property is that for single-peaked preferences can be aggregated using strategy-proof mechanisms. The most prominent such mechanism is the *median vote* [220].³

¹We also used three datasets from German pre-election polls, with 100 voters, 3 candidates, and no sampling, as well as all 225 currently available full preferences from PrefLib (<http://preflib.org>) [152], but results did not display anything at variance with the simulations.

²Note that these are extensible within the code and researchers can simply add additional distributions easily to the existing framework.

³We also tried simulations with *single-dipped* preferences, where the voter's most pre-

Polya-Eggenberger urn This model was developed and used to model the grouping of much of society to major homogenous groups [58, 229, 201]. In a k -urn model, k preference orders are chosen, and an urn is built to let voters choose preference orders from it. Each of the k chosen preferences gets $\frac{1}{k+1}$ of the preference orders in the urn, with the remaining $\frac{1}{k+1}$ of the urn filled by all preference orders in uniform. Preferences chosen from this urn have significant likelihood to be of the k main groups. In this work, we used the 2-urn and 3-urn model.

Riffle In a riffle model we get preferences of each voter by interleaving two separate preference orders on subsets of candidates in an independent manner. Huang and Guestrin [118] showed real-world elections which resembled this distribution.

Plackett-Luce In the Plackett-Luce model each candidate has an intrinsic cardinal value in the interval $[0, 1]$ (the “ground truth”). Each vote is then sampled from a particular distribution which adds noise to the true ranking [219].

D.2 Collected Variables

For every generated profile, we measured the following variables (all averaged over 100 simulations with a random singleton scheduler). We then averaged again over all 200 generated profiles of a given distribution.

Number of steps The number of steps from the initial (truthful) profile to convergence.

Number of stable states The number of distinct equilibrium outcomes (for the same preference profile), in terms of voting profiles.

ferred candidates are at the extreme. However, in such profiles the truthful vote has only two candidates with positive support (the extremes), and no voter ever has an incentive to move.

Number of distinct winners The number of distinct equilibrium outcomes (for the same preference profile), in terms of winner's identity.

Maximal ratio of equilibria with same winner The maximal fraction of simulations (for the same preference profile) that ended with the same winner.

Plurality agreement The fraction of simulations where the winner was the original Plurality winner.

Borda agreement The fraction of simulations where the winner was the Borda winner.

Copland agreement The fraction of simulations where the winner was the Copland winner.

Maximin agreement The fraction of simulations where the winner was the Maximin winner.

Condorcet agreement The fraction of simulations where the winner was the Condorcet winner, when one exists (not counted otherwise).

Social Welfare The relative rank of the winner, according to its Borda score (lower is better). Equivalently: the complement of the average social welfare of voters, assuming Borda utilities.

Gap 1-2 The ratio between the score of the winner and the score of the runner-up.

Gap 2-3 The ratio between the score of the second and the third candidates.

Total Duverger Fraction of simulations where at most two candidates received votes.

Relative Duverger The fraction of votes for the two leading candidates.

Furthermore, for the Plackett-Luce distribution, we also measured *Winner Ground Rank*, which is the rank of the winner according to the ground truth used to generate the profile (between 0 and $m - 1$).