Social Choice Theory

A Modern Approach with Computational Aspects

Felix Brandt

This manuscript is a work in progress and will be revised and expanded regularly. The latest version is available at `www.social-choice-theory.com`.

Please cite as:
Felix Brandt. *Social Choice Theory: A Modern Approach with Computational Aspects*. Manuscript available at `www.social-choice-theory.com`, 2026.

Contents

Preface vii

Acknowledgements ix

I Basics of Social Choice 1

1 Introduction 3
1.1 Framework and Key Questions 4
1.2 Voting Rules 6
1.3 Desirable Properties 9
1.4 Key Takeaways 10
1.5 Further Reading 10
1.6 Exercises 12

2 Choice Theory 13
2.1 Rational Decision-Making 14
2.2 Preference and Maximality 14
2.3 Degrees of Rationality 15
2.4 Rationalizable Choice 17
2.5 Consistent Choice 18
2.6 Strong Expansion Consistency 21
2.7 Key Takeaways 22
2.8 Further Reading 22
2.9 Exercises 24

3 From Individual Choice to Social Choice 27
3.1 Basic Axioms 28
3.2 Resolute Social Choice Functions 30
3.3 May's Theorem 31
3.4 The Condorcet Paradox 35
3.5 Key Takeaways 37
3.6 Further Reading 37
3.7 Exercises 38

4 Arrow's Theorem 39
4.1 Social Choice Functions 39
4.2 Social Welfare Functions 41
4.3 Illustrations and Consequences 42
4.4 Proof of Arrow's Theorem 43
4.5 Variants of Arrow's Theorem 46
4.6 Escape Routes 52
4.7 A Word on Cardinal Preferences 52
4.8 Key Takeaways 55
4.9 Further Reading 56

4.10 Exercises 58

II Social Choice Functions **61**

5 Restricted Domains **63**
5.1 Bypassing Impossibilities 63
5.2 Dichotomous Preferences 65
5.3 Single-Peaked Preferences 66
5.4 Checking Single-Peakedness 69
5.5 Value-Restricted Preferences 71
5.6 Key Takeaways 74
5.7 Further Reading 74
5.8 Exercises 77

6 Scoring Rules and Condorcet Extensions **79**
6.1 Scoring Rules 80
6.2 Characterization of Scoring Rules 82
6.3 Condorcet Extensions 84
6.4 McGarvey's Theorem 86
6.5 Borda versus Condorcet 88
6.6 Kemeny's rule 92
6.7 Key Takeaways 100
6.8 Further Reading 100
6.9 Exercises 102

7 Majoritarian SCFs **105**
7.1 The Top Cycle 106
7.2 Tournaments: Disregarding Majority Ties 110
7.3 The Uncovered Set 113
7.4 The Banks Set 118
7.5 Set-Rationalizability and Stability 124
7.6 The Bipartisan Set 127
7.7 Key Takeaways 136
7.8 Further Reading 136
7.9 Exercises 138

III Strategic Manipulation **141**

8 Resolute SCFs **143**
8.1 Two Alternatives 144
8.2 Restricted Domains 145
8.3 The Gibbard-Satterthwaite Theorem 146
8.4 Abstention 150

8.5 Hardness of Manipulation . . . 153
8.6 Key Takeaways . . . 154
8.7 Further Reading . . . 154
8.8 Exercises . . . 155

9 Irresolute SCFs 157
9.1 Kelly's Extension . . . 158
9.2 Fishburn's Extension . . . 160
9.3 Even-Chance Extension . . . 162
9.4 Optimist and Pessimist Extension . . . 162
9.5 Complete Preferences over Sets . . . 162
9.6 Further Reading . . . 162
9.7 Exercises . . . 163

IV Convex Sets of Alternatives 165

10 Probabilistic Social Choice 167

11 Budget Aggregation 169

Appendices 173

A Notation and Terminology 173

B Algorithms and Computational Complexity 179

Preface

. . .

There are already many excellent books on social choice theory (e.g., Sen, 1970, 2017; Fishburn, 1973; Moulin, 1988a; Austen-Smith and Banks, 1999; Arrow et al., 2002, 2011; Gaertner, 2009; Nitzan, 2010). The purpose of this one is to provide a contemporary perspective on classic problems of social choice. By "contemporary", I mean several things: the inclusion of computational considerations, an emphasis on positive results (alongside the extensive discussion of Arrovian impossibilities), an in-depth exploration of the rich landscape of Condorcet extensions, a thorough treatment of strategic manipulation and abstention, and the consideration of convex sets of alternatives, as in probabilistic social choice. Throughout, special emphasis is placed on social choice functions with attractive axiomatic characterizations. Despite the word "modern" in the subtitle, some recent developments—driven by practical applications such as multi-winner elections and participatory budgeting—are only briefly touched upon. These topics are well covered in the growing literature cited in Section 1.5.

. . .

The reader should be familiar with mathematical concepts—such as sets, relations, functions, and directed graphs—and standard proof techniques. Basic familiarity with some concepts from theoretical computer science, such as polynomial-time algorithms and NP-hardness, will also be useful. A quick introduction to these concepts is given in Chapter B.

If a theorem only holds under certain conditions, such as a lower bound on the number of alternatives or a restricted domain of preferences, these conditions are given in the upper right corner of the box surrounding the theorem.

Exercises marked with a star (☆) are more challenging than the other exercises.

Acknowledgements

This book is based on the course “Computational Social Choice” that I have taught at the Technical University of Munich since 2010. I would like to express my deep gratitude to all the teaching assistants who have supported me over the years: Patrick Becker, Florian Brandl, Markus Brill, Chris Dong, Fabian Frank, Christian Geist, Paul Harrenstein, Patrick Lederer, René Romen, and Christian Stricker. Not only did they contribute some of the exercises included in this book, but we also engaged in countless enlightening discussions about exercises, notation, definitions, and proofs. The latter is also true for other current and former members of my research group: Haris Aziz, Martin Bullinger, Felix Fischer, Matthias Greger, Johannes Hofbauer, Evangelia Pyrga, Satyanand Rammohan, Christian Saile, Alexander Schlenga, Hans Georg Seedig, and Anaëlle Wilczynski. I am also grateful to former students Florian Grundbacher, Keyvan Kardel, Matija Novaković, and Paul Stursberg for their helpful feedback.

. . .

Part I

Basics of Social Choice

1 Introduction

How can we aggregate the preferences of multiple individuals into a collective choice?

This question arises in a wide variety of contexts, such as political elections, collective decision-making, and resource allocation, and has occupied great minds from various disciplines. It is formally studied in social choice theory, which goes back to the Age of Enlightenment, particularly during the French Revolution in the late 18th Century, and to the important contributions by Jean-Charles de Borda and Marie Jean Antoine Nicolas de Caritat, better known as the Marquis de Condorcet. The field gained recognition as a distinct area of scientific research with the discovery of Kenneth Arrow's seminal impossibility theorem, published in 1950. Since then, other Nobel Laureates, such as Amartya Sen and Eric Maskin, along with influential figures like Charles Dodgson (the author of "Alice in Wonderland", better known by his pen name Lewis Carroll) and mathematician and computer scientist John Kemeny, have made important contributions to social choice theory. Rooted in economic theory, political science, and mathematics, social choice theory has also attracted attention in philosophy, computer science, law, sociology, and operations research.

A skeptical reader might wonder whether preference aggregation is a topic substantial enough to merit an entire book. After all, when electing a political leader, what's wrong with letting each voter select their preferred candidate and then picking the candidate who received the most votes?

This rule, known as plurality (aka first-past-the-post), is indeed the most widespread voting rule in the world. Yet, it is severely flawed. To illustrate this, let us consider an example with three candidates, a, b, and c. Suppose

$$40\% \text{ vote for } a, \quad 35\% \text{ vote for } b, \quad \text{and } 25\% \text{ vote for } c.$$

Plurality will select candidate a, because he is supported by the largest fraction of voters.

It is safe to assume that voters not only have a most preferred candidate but also hold preferences between the two remaining candidates. Suppose the complete preference rankings of the voters are described in the following table, where each column represents the ranking of the corresponding group of voters.

40%	35%	25%
a	b	c
b	c	b
c	a	a

These preferences could be the result of extreme positions adopted by candidate a, which render him objectionable to supporters of the other two candidates.[1]

Given these preferences, picking candidate a does not seem like a good idea after all. A majority of voters (60%) think candidate a is the *worst* candidate. These voters would prefer any other candidate to a. An even larger majority of voters (75%) prefers candidate b to candidate c. Hence, candidate b wins all pairwise majority comparisons while candidate a loses all of them!

This is not a contrived example. Eric Maskin and Amartya Sen, two economists at Harvard University, argue that precisely this happened in the primaries of the 2016 US presidential election, which were won by Donald Trump (Maskin and Sen, 2016).[2] As a matter of fact, there is overwhelming consensus among experts that plurality is *not* a good voting rule because it only takes into account the voters' favorite candidates. In 2010, a poll was conducted among 22 leading social choice theorists at Chateau du Baffy in France to identify attractive voting rules (Laslier, 2011). Each of the experts could either approve or disapprove any one of 18 voting rules. Plurality received no support at all!

More than 200 years earlier, the Chevalier de Borda and the Marquis de Condorcet already agreed that plurality has serious shortcomings. Yet they sharply disagreed on which voting rule should replace it — a debate that remains very much alive in contemporary social choice. Alongside plurality's undeniable simplicity and the reluctance of politicians to change a voting system that helped elect them, this is one of the main reasons why plurality is still so popular today.

1.1 Framework and Key Questions

The ingredients for the formal analysis of social choice are

- a set of voters,
- a set of alternatives,
- the voters' preferences over alternatives, and
- an aggregation function.

[1] The example can be made more realistic by slightly adjusting the percentages or incorporating additional types of voters. As long as its core remains unchanged, the criticism of plurality still stands.

[2] Igersheim et al. (2022) compare alternative voting methods based on data from a survey leading up to the 2016 US presidential elections.

The voters are autonomous in the sense that they can contemplate the alternatives, form their own preferences, and communicate these preferences (truthfully or not). Voters need not be human; they could be artificial entities. Depending on the application, the set of alternatives can vary in type and size. It could, for example, consist of candidates standing for election, joint plans to pursue, budget allocations, or assignments of students to courses. The preferences of the voters are typically represented by binary relations that specify whether a given voter prefers one alternative to another. In some cases, it can be convenient to represent preferences by an ordinal utility function, which assigns numerical values to alternatives. There are different types of aggregation functions. They all have in common that they take the voters' preferences as inputs and return a collective outcome. Depending on the type of function, this outcome could be a single alternative, a set of alternatives, a ranking of the alternatives, a probability distribution over the alternatives, and so forth.

2	2	1
a	b	a
b	c	c
c	a	b

Let us illustrate the different types of aggregation functions using the example given in the margin. As in the previous example, each column represents a group of voters with identical preferences. The number of voters with these preferences is given in the column header. We thus have a total of five voters. Tables such as the one in the margin are called preference profiles. If we look for a collectively most preferred alternative, a good candidate could be alternative a, which is top-ranked by three of the five voters. At second glance, alternative b looks quite good as well; it is only top-ranked by two voters, but two other voters rank it second, while a is never second-ranked. Thus, one could argue that a and b are equally good and the set $\{a, b\}$ should be returned, with tie-breaking postponed to a later stage. Yet another type of aggregation function returns a collective ranking of the alternatives. These functions are at the center of Arrow's famous impossibility theorem. A suitable ranking in our example could be the ranking given by the pairwise majority relation: first a, then b, then c, as three out of five voters prefer a to both b and c, and four out of five voters prefer b to c. Later in this book, we will consider functions that return probability distributions over alternatives. One way to do this is to assign probabilities that are proportional to how often an alternative is ranked first. In the example, this would result in probability $3/5$ for a and probability $2/5$ for b.

For the most part, we will be studying functions that return sets of alternatives, allowing for the possibility that multiple alternatives are tied as the best options. Among the key questions studied in this book are the following.

- What does it mean to make rational choices?
- Which formal properties should an aggregation function satisfy?
- Which of these properties can be satisfied simultaneously?
- How difficult is it to compute collective choices?
- What can be done to prevent voters from lying about their preferences?

We will address all these questions formally by means of mathematical definitions and theorems. A key role will be played by the axiomatic method.

political elections
online polls
public goods provision
participatory budgeting
budget aggregation
college admission
resource allocation

Unsurprisingly, a framework as broad as the one sketched above allows for countless applications, including *political elections*, *online polls*, *public goods provision*, *participatory budgeting*, *budget aggregation*, *college admission*, and *resource allocation*. Note that in some of these applications, the set of alternatives is highly structured. For example, in budget aggregation, the set of alternatives consists of all distributions of a fixed amount of money to various projects or institutions. In resource allocation, on the other hand, the set of alternatives is the set of all possible allocations of objects to voters. This usually comes with a natural restriction of admissible preferences: each voter is assumed to be indifferent between all allocations in which he receives the same objects.

recommender systems
meta-search engines
machine learning
AI alignment

There are also a number of unexpected applications of social choice theory in computer science. *Recommender systems* aggregate user data, such as ratings, clicks, purchases, or preferences, to suggest items to similar-minded people. *Meta-search engines* query multiple search engines (such as Google, Bing, Yahoo, and DuckDuckGo) and then aggregate the rankings into a single ranking for the user. Here, the search engines are the voters, the websites are the alternatives, and the meta-search engine computes an aggregation function. In *machine learning*, researchers have used methods from social choice theory to evaluate artificially generated players that play games. The learning algorithms are based on frequently ranking the players based on their performance and then generating refined offspring. Here, the games are the voters, and the players are the alternatives. Other work in machine learning uses social choice for *AI alignment*, which is concerned with ensuring that AI systems operate in ways that reflect human values and preferences. Since humans have diverse and often conflicting preferences, aggregation functions from social choice theory can be utilized to generate representative collective preferences.

1.2 Voting Rules

The first recorded use of elections dates back to Ancient Greece (around 500 BC). These elections were limited to yes/no majority decisions by a show of hands. This marks an important milestone in the history of human civilization, but most of the characteristic challenges and issues in collective decision-making only arise when there are at least three alternatives.

The first elections with more than two alternatives occurred in medieval England. Over the years, various voting rules have been proposed. Below, we list five common voting rules that have all been used for real-world elections.

Plurality (ca. 1290) has already been discussed before, and is used in most democratic countries. Alternatives that are ranked first by most voters win. Its first documented use is for the English House of Commons in the late 13th century. It is most commonly used for legislative elections in single-member districts, such as the British House of Commons, the US House of Representatives, or the Lok Sabha in India.

Borda's rule (1435) is a scoring rule, where each voter assigns 0 points to his lowest-ranked alternative, 1 point to his second lowest-ranked alternative, and so on. The alternatives with the highest accumulated score win. The rule is named after the French mathematician, physicist, and naval officer Jean-Charles, Chevalier de Borda, who proposed this rule in 1770. However, the German Catholic bishop and polymath Nicholas of Cusa already described this method in 1435 (Emerson, 2013). It is used in Slovenia (to elect two of the ninety members of the National Assembly) as well as in some academic institutions, professional bodies, and for awarding sports awards such as the Major League Baseball Most Valuable Player Award.

Plurality with runoff (1832) Two alternatives with the highest plurality scores (i.e., that are ranked first by most voters) face off in a majority runoff. An alternative that is preferred by a majority over the other one wins.[3] It was first used in France in 1832, where it is still used for presidential elections. It has become the most common method for electing heads of state in various countries throughout the world, such as Indonesia, Brazil, and Poland.

Instant-runoff (1857) is a multi-round elimination method based on plurality scores. In each round, one alternative with minimal plurality score (i.e., an alternative that is ranked first by the lowest number of voters) is eliminated. Plurality scores are then recomputed for the remaining alternatives. The process stops when only a single alternative remains. The method was briefly discussed and dismissed by the French philosopher, mathematician, and politician Marie Jean Antoine Nicolas de Caritat, Marquis de Condorcet, in 1788. It is, however, usually associated with the British lawyer Thomas Hare, who invented the *single transferable vote* system, a proportional representation method based on the same principle. Instant-runoff (which is also known as *alternative vote* and *ranked choice voting*) has been used for the Australian House of Representatives since 1918 and to elect the heads of state in India, Ireland, and Sri Lanka.

single transferable vote

alternative vote

ranked choice voting

Nanson's rule (1882) is a multi-round elimination method based on Borda scores. In each round, all alternatives whose Borda score is not strictly larger than the average Borda score are eliminated. Borda scores are then recomputed for the remaining alternatives. When all Borda scores are identical, the remaining alternatives win. The rule was invented by the British mathematician Edward J. Nanson, who immigrated to Australia (Nanson, 1883). Nanson's rule was used by some Australian universities and associations. The city of Marquette in Michigan, USA, used it for city elections in the 1920s.

To illustrate the definitions of these rules, consider the following preference profile,

[3] In reality, the runoff election is typically held a few weeks after the first round, and the votes cast in the runoff are independent of those in the first round. Voters can thus strategize and submit different preferences in each round. If an alternative is top-ranked by a majority of voters, no runoff is held. We ignore these effects and simply compute the outcome based on the voters' original preferences.

representing the preferences of 14 voters over five alternatives.

5	4	3	2
a	c	d	b
e	b	e	d
b	e	b	c
d	d	c	e
c	a	a	a

The only relevant information for identifying plurality winners is the top row. Alternative a, which is top-ranked by five voters, is the plurality winner.

Borda scores are computed as follows.

	5	4	3	2
4	a	c	d	b
3	e	b	e	d
2	b	e	b	c
1	d	d	c	e
0	c	a	a	a

Borda scores		
a:	$5 \cdot 4$	$= 20$
b:	$5 \cdot 2 + 4 \cdot 3 + 3 \cdot 2 + 2 \cdot 4$	$= 36$
c:	$4 \cdot 4 + 3 + 2 \cdot 2$	$= 23$
d:	$5 + 4 + 3 \cdot 4 + 2 \cdot 3$	$= 27$
e:	$5 \cdot 3 + 4 \cdot 2 + 3 \cdot 3 + 2$	$= 34$

Hence, alternative b is the Borda winner.

The two alternatives with the highest plurality scores are a and c. A majority of voters (9 of 14) prefer c to a. Hence, alternative c wins plurality with runoff.

The sequence of eliminations for instant-runoff is visualized below. The eliminated alternatives with minimal plurality score are marked in gray.

5	4	3	2
a	c	d	b
e	b	e	d
b	e	b	c
d	d	c	e
c	a	a	a

5	4	3	2
a	c	d	b
b	b	b	d
d	d	c	c
c	a	a	a

5	4	5
a	c	d
d	d	c
c	a	a

5	9
a	d
d	a

Thus, the winner of instant-runoff is alternative d.

For Nanson's rule, we need to compare Borda scores with the average Borda score. The average Borda score assigned by each voter is 2, and the total average Borda score is $14 \cdot 2 = 28$. Thus, all alternatives, except b and e, are eliminated. A majority of the voters (8 of 14) prefer e to b. Hence, the Borda score of b in the reduced profile is below average, and it will be deleted. The Nanson winner is alternative e.

In summary, we have the following winners.

Voting rule	Winner
Plurality	a
Borda's rule	b
Plurality with runoff	c
Instant-runoff	d
Nanson's rule	e

Each of these rules yields a different winner. In other words, depending on whether we are electing the US House of Representatives, members of the National Assembly in Slovenia, the president of France, the Australian House of Representatives, or the city council of Marquette in Michigan, the outcome for the *very same* preference profile differs.

A critical problem is that laypersons can be easily convinced that *each* of these rules is reasonable, natural, or possibly even optimal. In order to shed more light on the advantages and disadvantages of specific rules, academics study which rules satisfy certain desirable properties known as *axioms*.

1.3 Desirable Properties

In this section, we will briefly discuss three desirable properties and assess whether they are satisfied by the rules introduced so far. To ensure that the rules return single winners, we assume that ties (arising at any stage of computing the outcome) are broken lexicographically. We will use the profile discussed in Section 1.2 (reproduced in the margin) as a running example.

5	4	3	2
a	c	d	b
e	b	e	d
b	e	b	c
d	d	c	e
c	a	a	a

No majority loser. In the profile in the margin (and the introductory example of this chapter), plurality selects an alternative that is ranked last by a majority of voters. This is undesirable. It is not difficult to prove that the other four rules will never do this (Exercise 1.2).

No manipulation. A voting rule can be strategically manipulated if a voter can misrepresent his preferences such that the rule returns an alternative that he prefers to the original winner. For example, in the profile in the margin, any of the voters in the first column can change the Borda winner from b to e by moving e to the top and b to the bottom of the ranking. This will increase the Borda score of e by 1 and decrease that of b by 2. As we will see in Section 8.3, *every* voting rule is manipulable when there are at least three alternatives! Borda's rule is among the most manipulable voting rules. Instant-runoff and Nanson's rule, on the other hand, are rarely manipulable, and finding beneficial manipulations is an NP-hard problem.

No strategic abstention. A voting rule can be manipulated by strategic abstention if a group of voters can obtain a more preferred alternative by abstaining from an election. Let us illustrate this using plurality with runoff and the profile in the margin. If at least three of the voters in the first column abstain, the runoff will be between c and d, and d will win the majority comparison with c. All these voters are happier with d than with c. It turns out that, of our five rules, only plurality and Borda's rule do not suffer from strategic abstention (Exercise 1.3). Nanson's rule is immune to strategic abstention when there are at most three alternatives.

The table below summarizes the previous observations.

Voting rule	No majority loser	No manipulation	No strategic abstention
Plurality	–	–	✓
Borda's rule	✓	–	✓
Plurality w/ runoff	✓	–	–
Instant-runoff	✓	–	–
Nanson's rule	✓	–	✓ (for 3 alternatives)

In Section 6.3, we will see that, of the five rules discussed in this chapter, only Nanson's rule satisfies an influential axiom called *Condorcet-consistency*, which demands that a candidate who wins all pairwise majority comparisons has to be elected. In the profile analyzed in Section 1.2, alternative e beats all other alternatives in pairwise majority comparisons, but only Nanson's rule selects e. However, social choice theory is full of tradeoffs. The main theorem of Section 8.4, for example, establishes that all Condorcet-consistent rules suffer from strategic abstention.

Other axioms studied in this book concern how a rule reacts to the unavailability of alternatives (e.g., by demanding that the set of winners should be unaffected by the removal of losing alternatives, see Section 7.5) and how choices by different electorates should be related to each other (e.g., by demanding that alternatives that are selected by two different electorates should also be selected by the union of both electorates). The axiomatic method has produced far-reaching impossibility theorems, which show that certain combinations of axioms cannot be satisfied simultaneously.

We will also see that voter uncertainty about how ties are broken, such as through randomization, can be leveraged to let voting rules satisfy desirable properties that are otherwise incompatible with each other (see Chapter 9 and Chapter 10).

1.4 Key Takeaways

Introduction

- The most widespread voting rule, plurality, is flawed.
- The outcomes of voting rules can vary significantly.
- Voting rules are compared with each other by defining desirable properties.

1.5 Further Reading

Black (1958) and McLean and Urken (1995) provide excellent overviews of the early contributions to voting theory by Borda, Condorcet, Dodgson, Nanson, and others. Additionally, McLean and Hewitt (1994) offer an in-depth discussion of Condorcet's groundbreaking contributions. Nicholas of Cusa's work on Borda's rule is mentioned by Emerson (2013). The impact of Nanson's work on voting rules has been analyzed in greater detail by McLean (1996).

The properties of common voting rules have been studied extensively in political science (see, e.g., Nurmi, 1999; Felsenthal, 2012; Felsenthal and Nurmi, 2018; Brandt et al., 2022b). The example discussed in Section 1.2 is due to Brandt et al. (2022b) and was generated with the help of a computer using integer programming. No example of this kind requires fewer types of voters, and no profile of this kind with four types of voters requires fewer voters.[4] It is inspired by similar examples constructed by Nurmi (1987) and Balinski (2002).

Fishburn and Brams (1983) showed that instant-runoff and plurality with runoff can be manipulated by abstention. Moulin (1988b) showed that plurality and Borda's rule are resistant to this kind of manipulation. He also proved that maximin, which is almost identical to Nanson's rule when there are at most three alternatives, cannot be manipulated by abstention. Lepelley and Smaoui (2019) pointed out that this also holds for Nanson's rule (see also Brandt et al., 2025). Bartholdi, III and Orlin (1991) and Davies et al. (2014) proved that manipulating instant-runoff and Nanson's rule, respectively, is NP-hard. There is a large body of work studying the frequency of manipulability and other election phenomena (see, e.g., Aleskerov et al., 2012; Green-Armytage et al., 2015; Gehrlein and Lepelley, 2017; Diss and Merlin, 2021; Durand, 2023, 2025).

For some specialized applications of social choice theory, comprehensive overviews are available. The progress on participatory budgeting is surveyed by Aziz and Shah (2021), De Vries et al. (2022), and Rey et al. (2025). Suksompong and Teh (2026) reports on recent findings in the area of voting in divisible settings, such as budget aggregation. College admission, kidney exchange, and other applications of matching markets are thoroughly covered by Manlove (2013), Klaus et al. (2016), and Echenique et al. (2023). A recent hot topic in social choice theory is multi-winner elections, which are concerned with rules that select a committee consisting of a fixed number of candidates from a larger pool of candidates. An overview of the rich special case of approval ballots is provided by Lackner and Skowron (2023). For an overview of recommender systems and collaborative filtering, see Ricci et al. (2022). Meta search-engine aggregation is, for example, discussed by Dwork et al. (2001) and Wang et al. (2024). For recent papers highlighting the connection between evaluation in machine learning and social choice theory, see Balduzzi et al. (2018); Vinyals et al. (2019); Lanctot et al. (2025); Brandt (2026); Khalaf et al. (2026). As to AI alignment and social choice, the reader is referred to recent papers by Conitzer et al. (2024), Munos et al. (2024), Maura-Rivero et al. (2025), and Gölz et al. (2025).

[4] The same phenomenon can be achieved with only 12 voters, but 7 different types of voters.

1.6 Exercises

1.1 *Computing winners*

Consider the profile below and compute the winners according to plurality, Borda's rule, plurality with runoff, instant-runoff, and Nanson's rule.

1	4	1	3	2
a	c	d	e	a
b	b	e	a	d
d	e	a	b	c
c	d	b	d	b
e	a	c	c	e

Tie-breaking in plurality with runoff is performed lexicographically, whereas in instant-runoff voting, ties are resolved in reverse lexicographic order at each elimination round.

1.2 *Avoiding majority losers*

The winners, according to the *anti-plurality* rule, are those alternatives that are ranked *last* by the fewest voters.

Show that the anti-plurality rule, Borda's rule, plurality with runoff, instant-runoff, and Nanson's rule never elect a majority loser when there are at least two alternatives.

1.3 *Strategic manipulation and abstention*

Prove the following statements.

(a) Instant-runoff suffers from strategic abstention.

(b) Instant-runoff can be strategically manipulated.

(c) Plurality and Borda's rule are immune to strategic abstention.

☆ **1.4** *Monotonicity and manipulation*

A voting rule is *monotonic* if a selected alternative will still be selected when it rises in someone's preference ranking while leaving everything else unchanged (see also Section 3.3).

(a) Consider plurality, Borda's rule, plurality with runoff, instant-runoff, and Nanson's rule. Which of these rules satisfies monotonicity?

(b) Show that whenever a rule violates monotonicity, it can be strategically manipulated.

The origin of action is choice, and that of choice is desire and reasoning with a view to an end.

Aristotle, 350 BC

2 Choice Theory

Learning Outcomes

- How do we model rational decision-making?
- Which properties should be satisfied by preference relations?
- What is the relationship between rational and consistent choice?

A prerequisite for analyzing collective choice is to understand individual choice. This chapter explores *rational choice theory*. We will define a simple yet versatile model, which is concerned with selecting elements from sets, and provide formal definitions of concepts like preference, rationality, and consistency. The main theorem of this chapter establishes the equivalence of rationality and two choice consistency conditions.

Let U be a finite set of $m \geqslant 2$ alternatives, called the *universe*, which consists of all potentially available alternatives. We then consider concrete situations where some alternatives from the universe are offered in the form of menus. For example, a guest in a restaurant may choose from a dessert menu consisting of apple pie (a), brownies (b), and crème caramel (c). The model requires that at least one alternative is chosen. If several alternatives are chosen, this denotes indifference on the part of the chooser, i.e., she is equally happy with any one of the alternatives. A menu like $\{a, b, c\}$ will be referred to as a feasible set. $\mathcal{F}$ denotes the *set of feasible sets*. Unless noted otherwise, we will assume that every finite and nonempty set is feasible, i.e. $\mathcal{F} = \mathcal{P}^*(U)$ where $\mathcal{P}^*(U) = 2^U \setminus \{\emptyset\}$ denotes the set of nonempty subsets of U. For example, if $U = \{a, b, c\}$, $\mathcal{F} = \mathcal{P}^*(U) = \{\{a\}, \{b\}, \{c\}, \{a, b\}, \{b, c\}, \{a, c\}, \{a, b, c\}\}$.[1]

universe U

set of feasible sets $\mathcal{F}$

$\mathcal{P}^*$

A *choice function* associates every feasible set with a subset of the available alternatives. Formally, it is a function

choice function f

$$f\colon \mathcal{F} \to \mathcal{P}^*(U) \quad \text{such that } f(A) \subseteq A \text{ for all } A \in \mathcal{F}.$$

The table in the margin shows an example of a choice function when $U = \{a, b, c\}$. It suffices to only specify choices from feasible sets that contain at least two alternatives. Choices from singleton sets are trivial because we always need to select at least one alternative.

A	$f(A)$
ab	a
bc	bc
ac	a
abc	a

[1] Not choosing anything could be accounted for by adding a dummy alternative that represents a null choice. Similarly, alternatives could represent bundles of items. Note, however, that such interpretations of alternatives would have consequences on the set of feasible sets.

Not every choice function complies with our intuitive understanding of rationality. Certain patterns of choice from varying feasible sets may be deemed inconsistent. This would, for example, be the case when the guest in the restaurant would pick apple pie from the full menu but switch to brownies if crème caramel was unavailable, i.e., $f(\{a,b,c\}) = \{a\}$ and $f(\{a,b\}) = \{b\}$. The first choice indicates that the guest prefers a to b. This preference is reversed in the absence of the seemingly irrelevant alternative c.

2.1 Rational Decision-Making

Let us now move on to a key assumption of rational decision-making in economic theory. According to this assumption, an agent follows a specific thought process when making choices.

1. What is desirable?
2. What is feasible?
3. Choose the most desirable from among the feasible.

A critical aspect of this procedure is the sequence of these steps. The agent clarifies his preferences *before* considering the availability of alternatives. Although this might seem inefficient—why deliberate over options that may ultimately be unavailable?—the core principle of this model is that

desirability should be independent of feasibility.

In other words, the preferences of an agent should be independent of the feasible set.

2.2 Preference and Maximality

preference relation

We have already referred to preferences on an intuitive level. Formally, preferences are modeled as binary relations. By $\succsim$ we denote a binary *preference relation* on U that represents the preferences of an agent. The interpretation of $a \succsim b$ is that "a is at least as good as b." Individual preferences are assumed to be complete, i.e., for all $x, y \in U$, $x \succsim y$ or $y \succsim x$. Completeness implies reflexivity, i.e., for all $x \in U$, $x \succsim x$.

strict preference relation

The *strict preference relation* $\succ$ is defined by taking the asymmetric part of $\succsim$. For any $x, y \in U$,

$$x \succ y \quad \Leftrightarrow \quad x \succsim y \text{ and not } y \succsim x.$$

indifference relation

The relation $\succ$ is asymmetric, i.e., $x \succ y$ implies not $y \succ x$ for all $x, y \in U$. We say that x *dominates* y when $x \succ y$ and that x *weakly dominates* y when $x \succsim y$. The *indifference relation* $\sim$ is given by the symmetric part of $\succsim$. For any $x, y \in U$,

$$x \sim y \quad \Leftrightarrow \quad x \succsim y \text{ and } y \succsim x.$$

The relation $\sim$ is symmetric, i.e., $x \sim y$ implies $y \sim x$ for all $x, y \in U$. The restriction of $\succsim$ to some feasible set A is denoted by $\succsim|_A = \succsim \cap (A \times A)$.

Rationality intuitively entails that something is maximized. To formally capture this, we define the *set of maximal elements* set of maximal elements

$$\text{Max}(A, \succsim) = \{x \in A \colon \neg(y \succ x) \text{ for all } y \in A\}.$$

In later chapters, we will also apply this notion of maximality to relations that are not complete. For complete relations, the set of maximal elements is identical to $\text{Max}(A, \succsim) = \{x \in A \colon x \succsim y \text{ for all } y \in A\}$. To see this, let $x \in A$ and $\succsim$ be a complete relation. We then have

$$x \succsim y \text{ for all } y \in A \quad \Leftrightarrow \quad \neg(y \succ x) \text{ for all } y \in A.$$

The direction from left to right follows immediately from the definition of the strict part of $\succsim$. The direction from right to left requires completeness.

2.3 Degrees of Rationality

Without making further assumptions on $\succsim$, the set of maximal elements can be empty. Imagine an agent whose preferences over apple pie, brownies, and crème caramel are represented by the following directed graph. An arrow from x to y denotes that x is weakly preferred to y.

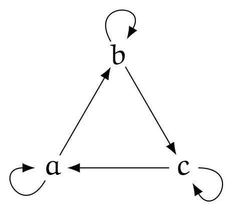

No alternative from the menu of all three desserts maximizes his satisfaction; there is always a better alternative.

To avoid this seemingly paradoxical phenomenon, it is typically assumed that preference relations of rational agents satisfy some notion of transitivity. Three such notions will be important in the following.

Let $\succsim$ be a binary relation on U.

- $\succsim$ is *transitive* if for all $x, y, z \in U$, $\quad x \succsim y$ and $y \succsim z \quad \Rightarrow \quad x \succsim z$. transitivity
- $\succsim$ is *quasi-transitive* if for all $x, y, z \in U$, $\quad x \succ y$ and $y \succ z \quad \Rightarrow \quad x \succ z$. quasi-transitivity
- $\succsim$ is *acyclic* if for all $x_1, \dots, x_n \in U$, $\quad x_1 \succ x_2 \succ \cdots \succ x_n \quad \Rightarrow \quad \neg(x_n \succ x_1)$. acyclicity

It can be proved that transitivity implies quasi-transitivity, and quasi-transitivity implies acyclicity (see Exercise 2.4). We thus have a hierarchy of transitivity notions. The following figure shows a preference relation $\succsim$ that is quasi-transitive but not transitive, since we have $b \succsim c$ and $c \succsim a$, but not $b \succsim a$.

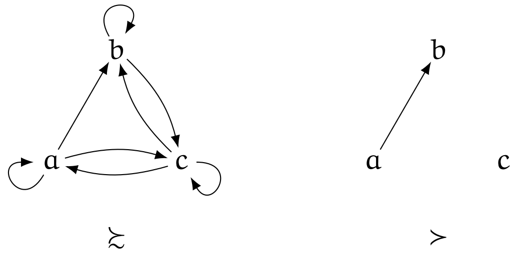

Here, $\mathrm{Max}(A, \succsim) = \{a, c\}$.

The following figure shows a preference relation $\succsim$ that is acyclic but not quasi-transitive, since we have $a \succ b$ and $b \succ c$, but not $a \succ c$.

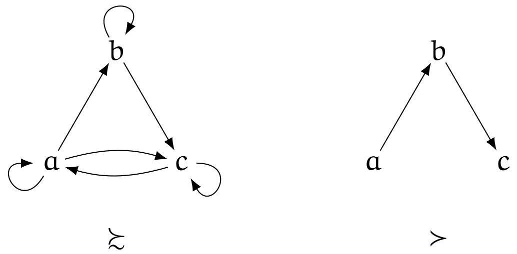

Here, $\mathrm{Max}(A, \succsim) = \{a\}$.

The two preceding examples show that all three notions are different from each other. When the relation $\succsim$ is antisymmetric, i.e., $x \succsim y$ and $y \succsim x$ imply $x = y$, then all three notions coincide.

Most readers are probably familiar with the strongest of the three transitivity notions. This is because transitive and complete relations are weak orders, which we, for example, obtain when assigning numeric utilities to alternatives and ranking them in order of decreasing utility.

Perhaps surprisingly, while transitivity is sufficient for the existence of maximal elements, it is not necessary. Acyclicity of the underlying preference relation is necessary and sufficient for the existence of maximal elements.

Lemma 2.1

Let $\succsim$ be a preference relation. Then,

$$\mathrm{Max}(A, \succsim) \neq \emptyset \text{ for all } A \in \mathcal{F} \quad \Leftrightarrow \quad \succsim \text{ is acyclic.}$$

Proof.

$\Rightarrow$ Let $A = \{x_1, \ldots, x_n\} \in \mathcal{F}$ with $x_1 \succ x_2 \succ \cdots \succ x_n$. By assumption, $\mathrm{Max}(A, \succsim) \neq \emptyset$. Since all alternatives except x_1 are dominated according to $\succsim$, this only leaves $\mathrm{Max}(A, \succsim) = \{x_1\}$. Hence, x_1 has to be undominated in A. In particular, $\neg(x_n \succ x_1)$.

$\Leftarrow$ Let $A \in \mathcal{F}$, $\succsim$ be an acyclic relation on U, and $x_1 \succ x_2 \succ \cdots \succ x_k$ a maximal sequence of distinct elements in A. Such a sequence must exist because A is finite. Since this sequence is of maximal length, x_1 cannot be dominated by any $y \in A \setminus \{x_1, \dots, x_k\}$. Moreover, x_1 cannot be dominated by any $y \in \{x_1, \dots, x_k\}$ because of the acyclicity of $\succsim$. Hence, $x_1 \in \text{Max}(A, \succsim)$.

□

A constructive proof of the direction from right to left of the preceding lemma is obtained by the following simple algorithm for finding an element of $\text{Max}(A, \succsim)$: let $x_1 \in A$. If $x_1 \notin \text{Max}(A, \succsim)$, then there has to be some $x_2 \in A$ with $x_2 \succ x_1$. Similarly, if $x_2 \notin \text{Max}(A, \succsim)$, then there has to be some $x_3 \in A$ with $x_3 \succ x_2$, and so on. We thus obtain the sequence $x_1 \prec x_2 \prec \dots \prec x_k$. In each step, the acyclicity of $\succsim$ implies that none of the previously considered alternatives can dominate x_k. By finiteness of A, there eventually has to be some $k \in \mathbb{N}$ such that $x_k \in \text{Max}(A, \succsim)$.

2.4 Rationalizable Choice

In reality, we cannot observe preferences, but we can observe choices. A central question in choice theory is, therefore:

Given some choice function f, can we find a preference relation that "rationalizes" f?

A choice function f is *rationalizable* if there exists a binary relation $\succsim$ on U such that rationalizability

$$f(A) = \text{Max}(A, \succsim) \text{ for all } A \in \mathcal{F}.$$

Given the large number of complete binary relations on U ($3^{m(m-1)/2}$), finding a rationalizing relation $\succsim$ appears to be a rather tedious task. However, after realizing that the relation has to rationalize the choice function in particular for all two-element sets, we obtain a very natural candidate for such a relation, the *base relation* $\succsim_f$ of f, which is base relation
defined by letting for all $x, y \in U$,

$$x \succsim_f y \quad \Leftrightarrow \quad x \in f(\{x, y\}).$$

Lemma 2.2

A choice function f is rationalizable iff it is rationalized by its base relation $\succsim_f$, i.e.,

$$f(A) = \text{Max}(A, \succsim_f) \text{ for all feasible sets } A.$$

Proof.

$\Rightarrow$ Let f be a choice function that is rationalized by $\succsim$. Then, for any $x, y \in U$,

$$x \succsim y \overset{\text{Def. of Max}}{\Leftrightarrow} x \in \text{Max}(\{x, y\}, \succsim) \overset{\text{Rat. of } f}{\Leftrightarrow} x \in f(\{x, y\}) \overset{\text{Def. of } \succsim_f}{\Leftrightarrow} x \succsim_f y.$$

We then have that $\succsim = \succsim_f$. So not only does $\succsim_f$ rationalize f, but it is also the *only* rationalizing relation.

$\Leftarrow$ If a choice function is rationalized by its base relation, it is rationalizable.

□

Lemma 2.2 has immediate algorithmic consequences. Without this insight, a naive algorithm for checking whether a choice function is rationalizable would enumerate all preference relations, all feasible sets, and then check whether the given relation rationalizes the choice from the given feasible set. The runtime of such an algorithm is in $O(3^{m(m-1)/2} \cdot 2^m \cdot m^2)$. Using Lemma 2.2, one can construct the base relation and then check whether it rationalizes the choice function, giving a runtime of $O(m^2 + 2^m \cdot m^2)$.

Lemma 2.2 also implies that the unintuitive example function mentioned earlier, where $f(\{a,b,c\}) = \{a\}$ and $f(\{a,b\}) = \{b\}$, cannot be rationalized. The second choice indicates that $b \succ_f a$, which contradicts $a \in Max(\{a,b,c\}, \succsim_f)$.

2.5 Consistent Choice

It would be nice if irrationality (i.e., the non-existence of a rationalizing relation) could be pointed out directly by identifying “inconsistencies” in the choice function. Consistency conditions are conditions that relate choices from different feasible sets to each other. One typically distinguishes between contraction-consistency conditions and expansion-consistency conditions.

Perhaps the simplest consistency condition is defined as follows: if an alternative x is chosen from some feasible set, it should also be chosen from all feasible subsets that contain x. This condition is known as contraction.

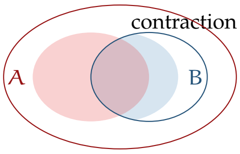

A choice function f satisfies *contraction* if for all $A, B \in \mathcal{F}$,

$$B \subseteq A \quad \Rightarrow \quad f(A) \cap B \subseteq f(B).$$

In the figure in the margin, A is represented by a red ellipse and $f(A)$ by a shaded red ellipse while B is represented by a blue ellipse and $f(B)$ by a shaded blue ellipse. Typical illustrations of contraction are of the following kind:

> *If a top European tennis player happens to be from Switzerland,*
> *then he should also be among the top players in Switzerland.*

3	2	2
a	b	c
b	c	b
c	a	a

In the following chapters, we will apply the consistency conditions defined here to collective decisions, that is, we will treat a group of agents as if it were a single individual and investigate to what extent group choices are consistent. To make this more concrete, consider plurality rule as discussed in Chapter 1 and the preference profile in the margin. Define plurality for variable feasible sets by first removing infeasible alternatives from the profile and then returning the plurality winners from the reduced profile. For the profile in the margin, we then have $f(\{a,b,c\}) = \{a\}$ and $f(\{a,b\}) = \{b\}$. These two choices are exactly as in the irrational dessert choice example, which means that the plurality choice function violates contraction.[2]

[2] This interpretation of plurality rule will be referred to as “narrow” plurality in Section 4.3. By contrast, a “broad” interpretation of plurality is defined by first computing plurality scores for all alternatives in the universe and then returning, for each feasible set, the feasible alternatives with maximal score. This choice function satisfies contraction. However, choices depend on the ranking of infeasible alternatives (see Section 4.3).

A necessary condition for rationalizability in light of Lemma 2.1 and Lemma 2.2 is that the base relation is acyclic. This condition is satisfied by any choice function that satisfies contraction.

Lemma 2.3

Let f be a choice function. Then,

$$f \text{ satisfies contraction} \quad \Rightarrow \quad \succsim_f \text{ is acyclic.}$$

Proof. We show the contrapositive: if $\succsim_f$ is cyclic, then f violates contraction. Let $A = \{x_1, \dots, x_n\}$, $x_1 \succ_f x_2 \succ_f \cdots \succ_f x_n \succ_f x_1$, and define $x_0 = x_n$. Since $f(A) \neq \emptyset$, there is some i with $1 \leqslant i \leqslant n$ such that $x_i \in f(A)$. If f satisfies contraction, $x_i \in f(\{x_{i-1}, x_i\})$. However, $x_{i-1} \succ_f x_i$ means that $f(\{x_{i-1}, x_i\}) = \{x_{i-1}\} \not\ni x_i$, a contradiction. □

Acyclicity of the base relation is not sufficient for rationalizability because the base relation fails to rationalize the choice function for feasible sets with more than two elements. As the example given in the margin shows, even contraction is insufficient for rationalizability. In this example, $\succsim_f$ is acyclic but fails to rationalize f because $f(\{a, b, c\}) = \{a\} \neq \{a, b\} = \text{Max}(\{a, b, c\}, \succsim_f)$.

A	$f(A)$
ab	ab
bc	b
ac	a
abc	a

We now introduce another consistency condition that—in conjunction with contraction—will be sufficient for rationalizability. This condition is called expansion.

A choice function f satisfies *expansion* if for all $A, B \in \mathcal{F}$,

expansion

$$f(A) \cap f(B) \subseteq f(A \cup B).$$

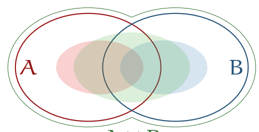

While contraction is a condition that specifies under which condition an alternative that is chosen in a feasible set has to be chosen in a feasible *subset*, expansion specifies under which conditions the alternative has to be chosen in a *superset*. A typical illustration of this property is:

> *A tennis player who is among the best female players and among the best teenagers is also a top player among all players who are female or teenagers.*

The choice function given in the margin above violates expansion as alternative b is chosen from every feasible set in which it is contained, except $\{a, b, c\}$. We thus have that contraction does not imply expansion. Similarly, one can easily construct examples showing that expansion does not imply contraction (see Exercise 2.2).

3	2	2
a	b	c
b	c	b
c	a	a

In the context of *social* choice, the example given in the margin not only shows that plurality rule violates contraction but also expansion. We have $f(\{a, b\}) = \{b\}$ and $f(\{b, c\}) = \{b\}$, but $f(\{a, b, c\}) = \{a\}$.

The following theorem establishes that contraction and expansion are necessary and sufficient for rationalizability.

Theorem 2.1 (Sen, 1971)

Let f be a choice function.

$$f \text{ is rationalizable} \quad \Leftrightarrow \quad f \text{ satisfies contraction and expansion.}$$

Proof.

$\Rightarrow$ Let f be a rationalizable choice function and $A, B \in \mathcal{F}$ with $B \subseteq A$. To prove that f satisfies contraction, we need to show that any $x \in f(A) \cap B$ is contained in $f(B)$.

$$x \in f(A) \overset{\text{Lem. 2.2}}{\Longrightarrow} x \in \mathrm{Max}(A, \succsim_f) \overset{x \in B \subseteq A}{\Longrightarrow} x \in \mathrm{Max}(B, \succsim_f) \overset{\text{Rat. of } f}{\Longrightarrow} x \in f(B)$$

For expansion, we take arbitrary $A, B \in \mathcal{F}$ and show that any $x \in f(A) \cap f(B)$ is contained in $f(A \cup B)$.

$$\begin{aligned} x \in f(A) \cap f(B) &\overset{\text{Lem. 2.2}}{\Longrightarrow} x \in \mathrm{Max}(A, \succsim_f) \cap \mathrm{Max}(B, \succsim_f) \\ &\overset{\text{Def. of Max}}{\Longrightarrow} \forall y \in A \cup B \colon \neg(y \succ_f x) \\ &\overset{\text{Def. of Max}}{\Longrightarrow} x \in \mathrm{Max}(A \cup B, \succsim_f) \overset{\text{Rat. of } f}{\Longrightarrow} x \in f(A \cup B) \end{aligned}$$

$\Leftarrow$ Let f be a choice function that satisfies contraction and expansion. We need to show that for all $A \in \mathcal{F}$, $\mathrm{Max}(A, \succsim_f) = f(A)$. First observe that

$$x \in \mathrm{Max}(A, \succsim_f) \overset{\text{Def. of Max}}{\Longleftrightarrow} \forall y \in A \colon \neg(y \succ_f x) \overset{\text{Def. of } \succsim_f}{\Longleftrightarrow} \forall y \in A \colon x \in f(\{x, y\}).$$

If $x \in \mathrm{Max}(A, \succsim_f)$ is chosen in all two-element subsets of A, we can repeatedly apply expansion (for exactly $|A| - 1$ times) to obtain $x \in f(A)$. Similarly, if $x \in f(A)$, $|A| - 1$ applications of contraction show that x is chosen in all two-element subsets of A and is thus contained in $\mathrm{Max}(A, \succsim_f)$.

□

Theorem 2.1 further reduces the effort to point out the *non-existence* of a rationalizing relation. It suffices to identify two feasible sets that yield a violation of contraction or expansion.

Schwartz (1976) gave an illuminating alternative characterization of contraction, expansion, and rationalizability. Let A, B be two feasible sets and $x \in A \cap B$.

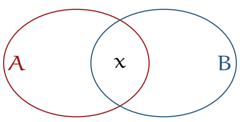

Contraction:	$x \in f(A \cup B)$	$\Rightarrow$	$x \in f(A) \cap f(B)$
Expansion:	$x \in f(A \cup B)$	$\Leftarrow$	$x \in f(A) \cap f(B)$
Rationalizability:	$x \in f(A \cup B)$	$\Leftrightarrow$	$x \in f(A) \cap f(B)$

This elegantly illustrates the duality of contraction and expansion, as well as the equivalence stated in Theorem 2.1.[3] This insight can be leveraged to rewrite rationalizability

[3]Note that, for contraction, $x \in f(B)$ can be dropped on the right-hand side. The given implication uses

by demanding that for all feasible sets A, B,

$$f(A \cup B) \cap A \cap B = f(A) \cap f(B). \qquad \text{(rationalizability)}$$

2.6 Strong Expansion Consistency

As we have seen, acyclicity is only a minimal rationality condition on preference relations. It is barely strong enough to guarantee nonempty choices from every feasible set. The choice function f given in the margin, for example, satisfies contraction and expansion, and, hence, by Theorem 2.1, rationalizability. However, it cannot be rationalized by a quasi-transitive relation. This can be seen by observing that the base relation $\succsim_f$ of f, which by Lemma 2.2 is the only candidate for a rationalizing relation and is depicted in the margin, fails to be quasi-transitive.

A	$f(A)$
ab	a
bc	b
ac	ac
abc	a

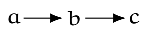

We now define a stronger expansion-consistency condition than expansion, which—in conjunction with contraction—can be used to characterize transitively rationalizable choice functions.

A choice function f satisfies *strong expansion* if for all feasible sets A, B,

$$B \subseteq A \text{ and } f(A) \cap B \neq \emptyset \quad \Rightarrow \quad f(B) \subseteq f(A).$$

strong expansion

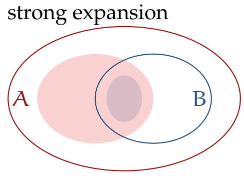

The condition can be illustrated as follows:

> *If there is a top European tennis player from Spain, then all players who are among the best in Spain are also among the best in Europe.*

This axiom is commonly seen as the strongest non-trivial expansion-consistency condition.[4]

Contraction demands that the choice set can only become larger when alternatives are removed. It is possible to phrase strong expansion in similar terms: *the choice set can only become smaller when removing alternatives (unless all chosen alternatives are removed)*. This may be confusing at first because strong expansion is an *expansion*-consistency condition, which states that alternatives that are chosen in some feasible set should also be chosen in certain supersets. However, equivalently, expansion-consistency conditions can be formulated as conditions that describe that *unchosen* alternatives remain unchosen when *reducing* the feasible set.

Theorem 2.2 (Bordes, 1976)

Let f be a choice function.

$$f \text{ is transitively rationalizable} \quad \Leftrightarrow \quad f \text{ satisfies contraction and strong expansion.}$$

---

contraction twice.

[4] A strengthening of strong expansion could be obtained by omitting $f(A) \cap B$ from the antecedent of the implication: an alternative that is chosen from some feasible set is chosen from all its supersets. However, it is fairly obvious that this property is only satisfied by the trivial choice function $f(A) = A$.

weak axiom of revealed preference (WARP)

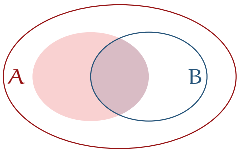

The conjunction of contraction and strong expansion is known as Arrow's choice axiom or the *weak axiom of revealed preference (WARP)* (Samuelson, 1938).

Let feasible sets A, B such that $B \subseteq A$ and $f(A) \cap B \neq \emptyset$.

Contraction:	$f(B)$	$\supseteq$	$f(A) \cap B$
Strong expansion:	$f(B)$	$\subseteq$	$f(A) \cap B$
WARP:	$f(B)$	$=$	$f(A) \cap B$

The natural appeal of choice functions that satisfy WARP comes from the fact that choices are made according to a (weak) ranking of the alternatives.

2.7 Key Takeaways

Choice Theory

- Choice functions map sets of alternatives to subsets.
- Rationalizable choice functions return maximal elements.
- Different degrees of rationality can be related to choice consistency conditions.

The relationships between rationalizability and choice consistency conditions are illustrated in the diagram below.

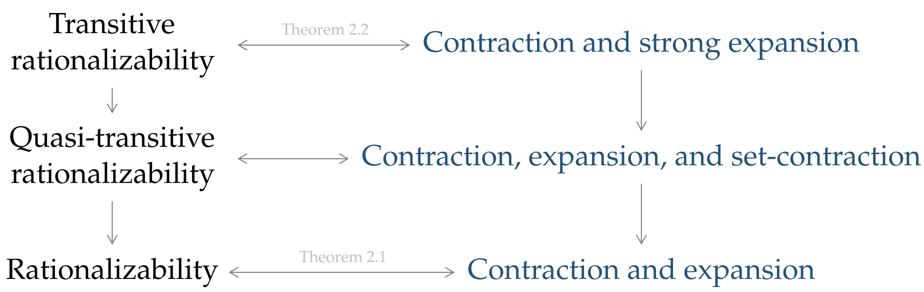

We will get back to choice theory in Section 7.5. There, we will also introduce the set-contraction axiom, which appears in the diagram above.

2.8 Further Reading

Austen-Smith and Banks (1999, Chapter 1), Sen (1977, 1986), and Moulin (1985a) provide excellent overviews of choice theory in the context of social choice. An entertaining and light-hearted introduction to choice theory is given by Allingham (2002). Gilboa (2010) masterfully introduces readers to various topics in microeconomic theory via the umbrella of rational choice.

The three-step rational decision-making process in Section 2.1 is taken from Rubinstein (2006).

Lemma 2.1 is due to von Neumann and Morgenstern (1944). The concept of rationalizability was introduced by Samuelson (1938). The idea of the base relation goes back to Herzberger (1973). Contraction (aka α) is due to Chernoff (1954). Expansion (aka γ) is due to Sen (1971). Strong expansion (aka β^+) is due to Bordes (1976).

Another influential choice consistency condition is *path independence* (Plott, 1973), which demands that $f(A \cup B) = f(f(A) \cup f(B))$ for all $A, B \in \mathcal{F}$. It is weaker than quasi-transitive rationalizability and equivalent to contraction and set-contraction. path independence

2.9 Exercises

2.1 *Properties of choice functions*

Consider the following choice function f.

A	$f(A)$
$\{a,b\}$	$\{b\}$
$\{a,c\}$	$\{a\}$
$\{b,c\}$	$\{b,c\}$
$\{a,b,c\}$	$\{b\}$

(a) Draw the strict part $\succ_f$ of the base relation $\succsim_f$.

(b) Does f satisfy contraction and expansion? Is f rationalizable? Why?

2.2 *Independence of contraction and expansion*

Let $U = \{a,b,c\}$.

(a) Give a choice function that satisfies contraction but not expansion.

(b) Give a choice function that satisfies expansion but not contraction.

2.3 *Disjunctive expansion (Brandt and Dong, 2025)*

A choice function f satisfies *disjunctive expansion* if for all $A, B \in \mathcal{F}$, $f(A) \subseteq f(A \cup B)$ or $f(B) \subseteq f(A \cup B)$. Prove the following statements.

☆ (a) Every choice function that satisfies *strong expansion* also satisfies *disjunctive expansion*.

(b) Every choice function that satisfies *disjunctive expansion* also satisfies *expansion*.

(c) There is a choice function on $U = \{a,b,c\}$ that satisfies *disjunctive expansion* but not *strong expansion*.

2.4 *Notions of transitivity*

Show that for binary relations, transitivity implies quasi-transitivity, and quasi-transitivity implies acyclicity.

2.5 *Negative transitivity*

A binary relation $\succsim \subseteq U \times U$ is *negatively transitive* if for all $x, y, z \in U$,

$$x \succ y \quad \Rightarrow \quad x \succ z \text{ or } z \succ y.$$

Show that a complete relation $\succsim$ is transitive iff it is negatively transitive.

☆ **2.6** *Resolute choice functions*

A choice function f is *resolute* if $|f(A)| = 1$ for all $A \in \mathcal{F}$. Show that a resolute choice function f is rationalizable if and only if it satisfies contraction. In that case, f is rationalizable by a *strict* and *transitive* relation.

For, when any number of men have, by the consent of every individual, made a community, they have thereby made that community one body, with a power to act as one body, which is only by the will and determination of the majority.

John Locke, 1690

3 From Individual Choice to Social Choice

Learning Outcomes

- Under which conditions is impartial single-valued social choice possible?
- How to choose between two alternatives?
- Is there a natural way to extend majority decisions beyond two alternatives?

In this chapter, we start to investigate collective choice, i.e., choices based on the preferences of multiple agents. As we will see, majority rule takes a central role when deciding between only two alternatives. Problems arise when we consider more than two alternatives.

Let $N = \{1, \ldots, n\}$ be a finite set of voters (the electorate) and $\mathcal{W}$ the set of all transitive and complete relations over U. The preferences of agents are assumed to satisfy the strongest of the three rationality notions introduced in the previous chapter: they can be conveniently represented as weak rankings. Throughout this and the next chapter, the *domain* $\mathcal{D}$ of admissible *preference profiles* is

$N = \{1, \ldots, n\}$

$\mathcal{W}$

domain $\mathcal{D}$

preference profile

$$\mathcal{D} = \mathcal{W}^N$$

unless noted otherwise.[1] Preference profiles are written as $P = (\succsim_1, \ldots, \succsim_n)$.

The main objects of study in this book are social choice functions. A *social choice function (SCF)* is a function

social choice function (SCF)

$$f : \mathcal{F} \times \mathcal{D} \to \mathcal{P}^*(U)$$

such that $f(A, P) \subseteq A$ for all $A \in \mathcal{F}$. In other words, an SCF maps a preference profile and a feasible set of alternatives to a nonempty subset of the feasible set. Every SCF induces a choice function $f(\cdot, P)$ for a fixed preference profile $P \in \mathcal{D}$, and rationalizability and consistency conditions carry over. An SCF is rationalizable (or consistent) if, for each $P \in \mathcal{D}$, the choice function $f(\cdot, P)$ is rationalizable (or consistent). Whenever A or P are clear from the context, we will sometimes just write $f(A)$ or $f(P)$. The SCF f is *trivial* if $f(A, P) = A$ for all $A \in \mathcal{F}$ and $P \in \mathcal{D}$. For two SCFs f and f', f is a *refinement* of f' if $f(A, P) \subseteq f'(A, P)$ for all $A \in \mathcal{F}$ and $P \in \mathcal{D}$. We will also say that f is *finer* than f', f' is *coarser* than f, and f' is a *coarsening* of f. Obviously, every SCF is a refinement of the trivial SCF.

trivial SCF

refinement

[1] For two sets X and Y, Y^X denotes the set of all functions from X to Y.

Later in this book, we will also consider preference profiles with varying sets of voters.
N_P To account for this, N_P denotes the set of voters present in profile P, and $n_P = |N_P|$. As long as $\mathcal{D} = \mathcal{W}^N$, N_P will invariably be N.

$P|_A$ The restriction of profile P with respect to a feasible set A is denoted by $P|_A = (\succsim_{i_1}|_A, \dots, \succsim_{i_n}|_A)$.

3.1 Basic Axioms

Two central fairness properties of SCFs are anonymity and neutrality. The underlying idea of both properties is that both the identities of the voters and the labels of the alternatives should be irrelevant. This can be formalized using bijections between sets of voters and sets of alternatives, respectively.

Let $P, P' \in \mathcal{D}$ be profiles such that there is a bijection

$$\pi\colon N_P \to N_{P'} \text{ with } \succsim_i = \succsim'_{\pi(i)} \text{ for all } i \in N_P.$$

anonymity Then, an SCF f is *anonymous* if $f(A, P) = f(A, P')$ for all $A \in \mathcal{F}$.

Anonymity ensures that all voters are treated equally, with no individual receiving special consideration. Renaming the voters has no effect on the outcome. It is a relatively mild symmetry property that all voting rules introduced in Chapter 1 satisfy. A simple example of an SCF violating anonymity when $n \geqslant 2$ is given below.

$$f(\{a, b\}, P) = \begin{cases} \{a\} & \text{if } a \succ_1 b \\ \{b\} & \text{otherwise.} \end{cases}$$

To see that f violates anonymity, consider a profile P in which $a \succ_1 b$ but $b \succ_i a$ for all $i \in N \setminus \{1\}$. Then, $f(P) = \{a\}$. For any non-trivial permutation $\pi\colon N \to N$, we obtain another profile P' by letting $\succsim_i = \succsim'_{\pi(i)}$ for all $i \in N$. However, $f(P') = \{b\} \neq \{a\}$.

For another example, consider US presidential elections. Due to the electoral college, the influence that individual citizens have on the final outcome depends on their location. For example, the coordinated relocation of voters between states (from "blue" states such as California or "red" states such as Texas to "swing states" such as Arizona, Pennsylvania, or North Carolina) can easily alter the outcome of tight elections. As a matter of fact, in 2000 and 2016, more American citizens voted for the losing candidate than for the winning candidate.

N_{xy} A useful notation for the set of voters who weakly prefer x to y is $N_{xy} = \{i \in N\colon x \succsim_i y\}$.
n_{xy} The number of voters in N_{xy} will be denoted by $n_{xy} = |N_{xy}|$. When $U = \{a, b\}$, the outcome of any anonymous SCF only depends on n_{ab} and n_{ba}.

Neutrality is defined similarly using a bijection between two feasible sets. Let A, B be feasible sets and P, P' profiles with $N_P = N_{P'}$ such that there is a bijection

$$\sigma\colon A \to B \text{ with } (x \succsim_i y) \Leftrightarrow (\sigma(x) \succsim'_i \sigma(y)) \text{ for all } i \in N_P \text{ and } x, y \in A.$$

neutrality Then, an SCF f is *neutral* if $f(B, P') = \{\sigma(x)\colon x \in f(A, P)\}$. Neutrality ensures that all alternatives are treated equally. However, in contrast to anonymity, the outcome of the SCF should not be invariant under relabeling alternatives, but the chosen alternatives need to be relabeled accordingly. Note that A and B must have the same size for a bijection between them to exist.

Neutrality, as defined above, also implies that SCFs are independent of preferences over alternatives that are not contained in the feasible set. This can be seen by letting $A = B$ and $\sigma(x) = x$. The preferences between alternatives not in A can be rearranged arbitrarily without affecting the outcome. This implication of neutrality is known as *independence of infeasible alternatives* and will be discussed in detail in Chapter 4.

Constant functions, such as $f(\{a, b\}, P) = \{a\}$ for all $P \in \mathcal{D}$, are the simplest examples of non-neutral SCFs. Neutrality is not imperative for reasonable collective decision-making. Consider an SCF on two alternatives a, b where a is the status quo, which will only be replaced with b if more than two-thirds of the voters strictly prefer b:

$$f(\{a,b\}, P) = \begin{cases} \{a\} & \text{if } n_{ab} \geqslant \frac{n}{3} \\ \{b\} & \text{otherwise.} \end{cases}$$

To see that f violates neutrality, let P be a profile such that $n_{ab} = n_{ba}$. Then, $f(P) = \{a\}$. If we swap a and b in P ($\sigma = (ab)$), we obtain profile P' for which again $n_{ab} = n_{ba}$. Hence, $f(P') = \{a\}$ but $\sigma(f(P)) = \{b\}$. It is not difficult to come up with more subtle examples, such as reducing the threshold in the previous SCF to 51%, that fail to be neutral but "almost" satisfy neutrality. All the rules introduced in Chapter 1 are neutral. However, lexicographically breaking ties between winners violates neutrality.

The weakest uncontroversial notion of social optimality is Pareto-optimality, named after the Italian economist Vilfredo Pareto.

Given $P \in \mathcal{D}$ and $x, y \in U$, x *Pareto-dominates* y if (Pareto-dominance)

$$x \succ_i y \text{ for all } i \in N_P.$$

Clearly, the Pareto dominance relation is transitive because the individual preferences are. Alternative $x \in A$ is *Pareto-optimal* in $P|_A$ if it is not Pareto-dominated by any $y \in A$.[2] (Pareto-optimality)

Social choice typically involves tradeoffs in the sense that an alternative might be preferred by some (but not all) to the status quo. By contrast, if the status quo is Pareto-dominated, *everybody* can be made strictly better off.

The *Pareto SCF* $\mathsf{PO}(A, P)$ returns all Pareto-optimal alternatives in $P|_A$. The Pareto SCF is anonymous and neutral. An SCF f is *Pareto-optimal* if $f(A, P) \subseteq \mathsf{PO}(A, P)$ for all $A \in \mathcal{F}$ and $P \in \mathcal{D}$. In other words, a Pareto-optimal SCF only returns Pareto-optimal alternatives. (Pareto SCF; Pareto-optimal SCF)

A technical property that is even weaker than Pareto-optimality is non-imposition. An SCF f is *non-imposing* if for every $A \in \mathcal{F}$ and $x \in A$, there is $P \in \mathcal{D}$ such that $f(A, P) = \{x\}$. In other words, every alternative will be uniquely selected in some situation. (non-imposition)

[2] This notion of Pareto-dominance is sometimes called strong Pareto-dominance. Similarly, the corresponding notion of Pareto-optimality is sometimes called weak Pareto-optimality.

3.2 Resolute Social Choice Functions

resoluteness

Ideally, we would like SCFs to always return a single alternative. Such SCFs are called resolute. Formally, an SCF f is *resolute* if $|f(A, P)| = 1$ for all $A \in \mathcal{F}$ and $P \in \mathcal{D}$. Resoluteness is in conflict with our basic fairness conditions: if all voters are indifferent between all alternatives, choosing a single alternative is at variance with neutrality for any m and n. Whether anonymous, neutral, and resolute SCFs exist for certain combinations for m and n becomes a more interesting question when restricting attention to *strict preferences*. The set of all strict (or antisymmetric) preference relations is denoted by

strict preferences $\mathcal{S}$

$$\mathcal{S} = \{\succsim \in \mathcal{W} \colon \forall x, y \in U \colon x \sim y \Rightarrow x = y\}$$

and the domain of strict (or linear) preference profiles is $\mathcal{S}^N$. Technically, $\mathcal{S}$ consists of all complete, transitive, and antisymmetric relations over U.

Consider the preference profile P depicted below on the left-hand side and let f be a resolute SCF.

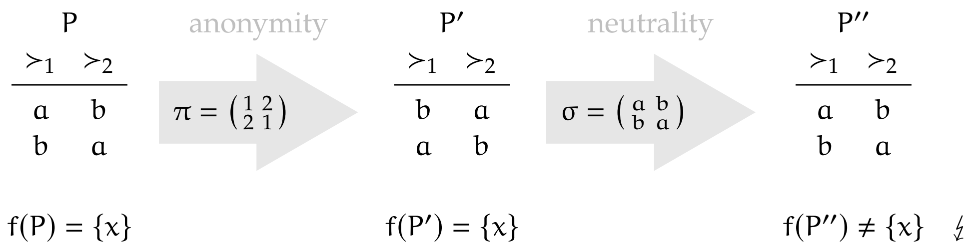

f returns a single alternative x (which is either a or b). When permuting voters according to $\pi = \left(\begin{smallmatrix}1 & 2\\ 2 & 1\end{smallmatrix}\right)$, one obtains profile P', depicted in the middle. Anonymity requires that $f(P')$ is still $\{x\}$. Now, permuting the alternatives using $\sigma = \left(\begin{smallmatrix}a & b\\ b & a\end{smallmatrix}\right)$ gives P''. Neutrality demands that $f(P'') \neq \{x\}$. However, $P'' = P$ and $f(P) = \{x\}$, a contradiction. As a consequence, no anonymous, neutral, and resolute SCF exists when $m = 2$ and $n = 2$. In general, such SCFs exist iff n has no prime factor less than or equal to m.

Theorem 3.1 (Moulin, 1983) $\mathcal{D} = \mathcal{S}^N$

There is an anonymous, neutral, and resolute SCF iff n cannot be divided by any $q \in \mathbb{N}$ with $2 \leqslant q \leqslant m$.

Proof.

$\Rightarrow$ Let $n = q \cdot k$ for some $q, k \in \mathbb{N}$ with $2 \leqslant q \leqslant m$ and assume for contradiction that f is anonymous, neutral, and resolute. Now, pick some $A = \{x_1, \dots, x_q\} \in \mathcal{F}$ with $|A| = q$ and let $P \in \mathcal{D}$ such that $P|_A$ looks as follows.

$$\begin{array}{cccc} k & k & \dots & k \\ \hline x_1 & x_2 & \dots & x_q \\ x_2 & x_3 & \dots & x_1 \\ \vdots & \vdots & \ddots & \vdots \\ x_q & x_1 & \dots & x_{q-1} \end{array}$$

Then, similar to the $m = n = 2$ example discussed above, anonymity and neutrality imply that $f(A, P) = A$. Since $|A| > 1$, this contradicts resoluteness.

$\Leftarrow$ Let n be indivisible by any $q \in \mathbb{N}$ with $2 \leqslant q \leqslant m$. We will show that the *iterated plurality rule* is anonymous, neutral, and resolute. The iterated plurality rule f breaks ties between plurality winners by re-computing plurality scores of winning alternatives in the reduced profile, which consists only of these alternatives. This is repeated until all remaining alternatives are plurality winners. It is easily seen that f satisfies anonymity and neutrality. Now, assume for contradiction that f is irresolute and $f(A, P) = X$ with $2 \leqslant |X| \leqslant |A|$ for some $A \in \mathcal{F}$ and $P \in \mathcal{D}$. This is only possible when all alternatives in X are top-ranked by the same number of voters in $P|_X$. Hence, $n = |X| \cdot k$ for some $k \in \mathbb{N}$, which implies that n is divisible by $|X|$ with $2 \leqslant |X| \leqslant m$, a contradiction.

□

When insisting on anonymity, neutrality, and universal SCFs, defined for any combination of m and n, allowing sets as outcomes of SCFs is inevitable.

3.3 May's Theorem

Let us first restrict attention to feasible sets of size two. Even when demanding anonymity and neutrality, there is a wide variety of SCFs. We have already defined PO, which, in the case of two alternatives, collapses to the following rule.

$$\text{PO}(\{a, b\}, P) = \begin{cases} \{a\} & \text{if } n_{ba} = 0 \\ \{b\} & \text{if } n_{ab} = 0 \\ \{a, b\} & \text{otherwise} \end{cases} \qquad (\textit{Pareto rule})$$

Here are two further examples, majority rule and a silly rule.

$$f(\{a, b\}, P) = \begin{cases} \{a\} & \text{if } n_{ab} > n_{ba} \\ \{b\} & \text{if } n_{ba} > n_{ab} \\ \{a, b\} & \text{otherwise} \end{cases} \qquad (\textit{majority rule})$$

$$f(\{a, b\}, P) = \begin{cases} \{a\} & \text{if } n_{ab} \text{ odd and } n_{ba} \text{ even} \\ \{b\} & \text{if } n_{ba} \text{ odd and } n_{ab} \text{ even} \\ \{a, b\} & \text{otherwise} \end{cases} \qquad (\textit{silly rule})$$

These SCFs are anonymous because their outcomes only depend on n_{ab} and n_{ba}. The neutrality of these SCFs can be verified by observing that the case distinctions are completely symmetric. Both rules are resolute when $\mathcal{D} = \mathcal{S}^N$ and n is odd, which is in line with Theorem 3.1.

The silly rule lives up to its name, even though it is anonymous and neutral. It even satisfies Pareto-optimality when n is odd. Moreover, it has the interesting property that in every profile, every voter can affect the outcome by changing his preferences. In particular, every voter who is unhappy with the outcome (because his less preferred alternative is selected) can change his preferences such that his more preferred alternative will be selected. This instability makes the silly rule a strategic nightmare. Predicting the outcome of an election using the silly rule appears to be an impossible task. In fact, protocols similar to the silly rule are used in cryptography to let multiple agents jointly generate random numbers. What is particularly troublesome is that there are attractive rules for two alternatives that cannot be manipulated. Majority rule is one of them (see Section 8.1).

What the silly rule lacks is some form of monotonicity: increasing the support for a winning alternative should not turn it into a losing alternative. This is precisely the axiom we introduce next.

P_{-i} For a given profile $P \in \mathcal{D}$, we write P_{-i} as a shorthand for the profile in which voter i's preferences have been removed: $P_{-i} = P \setminus \{(i, \succsim_i)\}$. Let $P, P' \in \mathcal{D}$, $a \in U$, and $i \in N$ such that $P_{-i} = P'_{-i}$ and for all $x, y \in U \setminus \{a\}$,

$$\begin{aligned} x \succsim_i y \quad &\Leftrightarrow \quad x \succsim'_i y \\ a \sim_i y \quad &\Rightarrow \quad a \succsim'_i y, \text{ and} \\ a \succ_i y \quad &\Rightarrow \quad a \succ'_i y. \end{aligned}$$

The equivalence in the first line states that the preferences between all alternatives other than a are unchanged. The two implications state that alternative a rises in voter i's preference ranking. This is a bit tedious to write down formally because there are three ways how alternative a can be reinforced against alternative y when preferences are weak: *(i)* from $y \succ_i a$ to $a \sim'_i y$, *(ii)* from $y \succ_i a$ to $a \succ'_i y$, and *(iii)* from $a \sim_i y$ to $a \succ'_i y$.

monotonicity f is *monotonic* if for all $A \in \mathcal{F}$,

$$a \in f(A, P) \quad \Rightarrow \quad a \in f(A, P').$$

This can be strengthened to a condition called positive responsiveness.

positive responsiveness f is *positively responsive* if for all $A \in \mathcal{F}$,

$$a \in f(A, P) \text{ and } P|_A \neq P'|_A \quad \Rightarrow \quad \{a\} = f(A, P').$$

Monotonicity requires that a chosen alternative remains chosen even when it is reinforced. Positive responsiveness requires that a chosen alternative is chosen *uniquely* when it is reinforced. For resolute SCFs, monotonicity and positive responsiveness are equivalent. Sometimes, the contrapositive of monotonicity turns out to be useful: an unchosen alternative remains unchosen when it is weakened.

The silly rule clearly violates monotonicity. The Pareto rule is monotonic but not positively responsive. Majority rule satisfies positive responsiveness on top of anonymity and neutrality. The following theorem shows that it is the only rule for which this is the case.

Theorem 3.2 (May, 1952) $m = 2$

Majority rule is the *only* SCF that satisfies anonymity, neutrality, and positive responsiveness.

Proof. The fact that majority rule satisfies all three properties is obvious from its definition on page 31. To see that it is the only such rule, we will use the three axioms to assign unique outcomes to one profile after another until the SCF is fully specified and identical to majority rule. Let $U = \{a, b\}$. First, observe that the outcome that any anonymous rule returns for a given profile only depends on how many voters have $a \succ b$ $(n - n_{ba})$, $b \succ a$ $(n - n_{ab})$, and $a \sim b$ $(n_{ba} + n_{ab} - n)$. Each combination of these three parameters can be viewed as a point in three-dimensional space, representing an equivalence class of profiles that are identical up to anonymity. The resulting two-dimensional discrete simplex is shown in the figure below for $n = 3$.

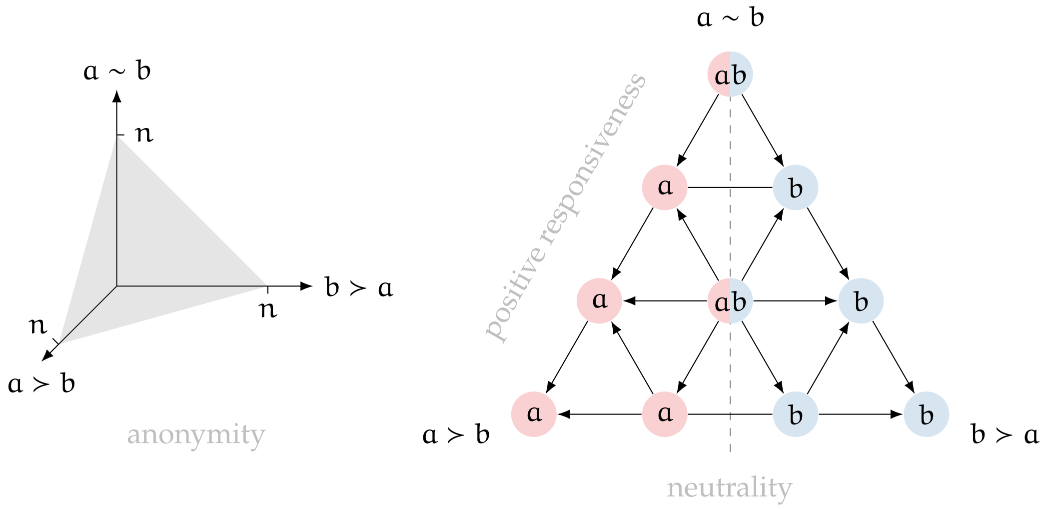

We have a total of 10 equivalence classes of profiles, which, for simplicity, will be referred to as profiles in the following. The profile in the middle of the simplex, for example, has one voter for each strict preference and one indifferent voter. In the profile on the far left, all three voters prefer a to b. For two alternatives, there is only one non-trivial permutation of the alternatives: $\sigma = \begin{pmatrix} a & b \\ b & a \end{pmatrix}$. In the diagram, this boils down to the vertical symmetry axis, denoted by the dashed line. This has immediate consequences for profiles that lie on the axis, for which we can only return the outcome $\{a, b\}$. All the remaining profiles can be processed by repeatedly applying positive responsiveness. Each of the arrows facing left (upwards, straight, or downwards) represents the modification of a single voter's preference in favor of a. For example, an upward-left arrow represents a change from $b \succ a$ to $a \sim b$. Positive responsiveness thus implies that a has to be returned for all profiles to the left of the symmetry axis. Analogously, b has to be returned for all profiles to the right of the axis. Alternatively, this also follows from neutrality. The SCF is now fully specified and identical to majority rule. □

Since Theorem 3.2 is the first axiomatic characterization in this book, let us appreciate a simple consequence of this result: *any* SCF on two alternatives other than majority rule violates at least one of the three axioms. In other words, if we insist on anonymity, neutrality, and positive responsiveness, we need look no further than majority rule. When inspecting the proof of Theorem 3.2, it becomes apparent that similar proofs also work for many subdomains of $\mathcal{W}^N$, most importantly $\mathcal{S}^N$ (see Exercise 3.3).

We will see later in this book that when there are more than two alternatives, positive responsiveness is a rather restrictive property satisfied by only a few SCFs (see Section 9.1). There is an alternative characterization of majority rule which replaces positive responsiveness with the weaker property of monotonicity, at the expense of restricting attention to most decisive SCFs.

Theorem 3.3 $m = 2$

Majority rule is the finest SCF that satisfies anonymity, neutrality, and monotonicity.

The three axioms characterize so-called "threshold rules." The six possible SCFs satisfying anonymity, neutrality, and monotonicity when $n = 3$ are shown below.

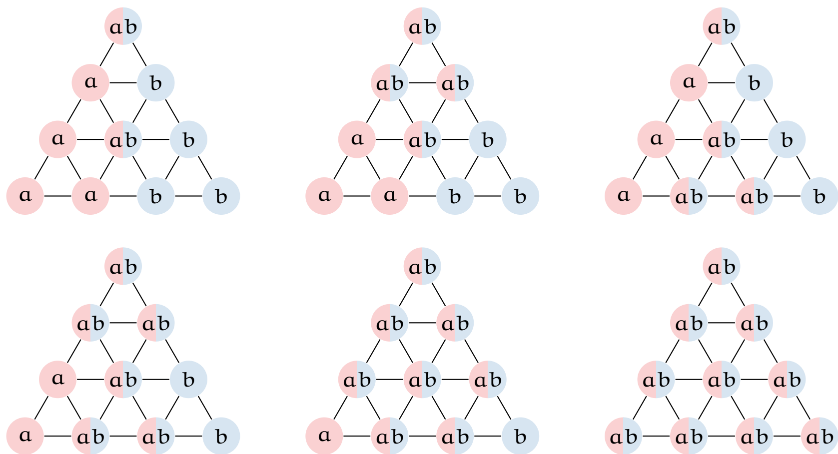

The first one is (relative) majority rule. The second one is absolute majority rule, which only returns a single alternative if there is a majority of voters with a *strict* preference. The latter two are the Pareto rule and the trivial rule, respectively.

The above results underline that majority rule is uncontroversial when there are only two alternatives. This is also evident from the fact that all voting rules mentioned in Chapter 1 (plurality, Borda, plurality with runoff, instant-runoff, and Nanson's rule) coincide with majority rule on two alternatives. Majority rule is also transitively rationalizable, but this is a trivial observation that holds for every SCF on two alternatives.

We will denote weakenings of axioms (such as positive responsiveness, monotonicity, and Pareto-optimality) that only hold for feasible sets of size two by positive responsiveness$_2$, monotonicity$_2$, and Pareto-optimality$_2$.

3.4 The Condorcet Paradox

As seen in the previous section, majority rule unsurprisingly plays a central role when deciding between two alternatives. Therefore, an important concept is the *pairwise majority relation*, denoted by $\succsim_M$, which is defined by

pairwise majority relation $\succsim_M$

$$x \succsim_M y \quad \Leftrightarrow \quad n_{xy} \geqslant n_{yx}$$

for all $x, y \in U$.

We have already seen in Section 2.5 that (narrow) plurality is not rationalizable because it violates both contraction and expansion. This statement can be generalized to a much larger class of SCFs.

Theorem 3.4 (Condorcet, 1785; May, 1952) $m \geqslant 3, n \geqslant 3$

No anonymous, neutral, and positively responsive$_2$ SCF is rationalizable.

Proof. Assume without loss of generality that $a, b, c \in U$ and let f be an SCF with the desired properties. In particular, f is rationalizable, and we denote its rationalizing relation by $\succsim$. Now, consider a preference profile P such that $P|_{\{a,b,c\}}$ is given as in the left table below. If $n > 3$, the additional voters are indifferent between a, b, and c.

1	1	1
a	b	c
b	c	a
c	a	b

$P|_{\{a,b,c\}}$

A	$f(A, P)$
ab	a
bc	b
ac	c

$f(\cdot, P)$

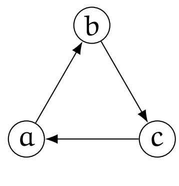

$\succsim_f = \succsim_M$

Lemma 2.2 implies that $\succsim = \succsim_f$. May's theorem implies that f coincides with majority rule for feasible sets of size 2 (see the table in the middle). In other words, $\succsim_f = \succsim_M$. However, $\succsim_M$ is cyclic (see the graph on the right) and therefore cannot rationalize f. □

Theorem 3.4 can be seen as a poor man's version of Arrow's impossibility theorem. It is much easier to prove and provides an immediate insight into the main difficulty of social choice theory: the majority relation can be cyclic. This phenomenon is also known as the Condorcet paradox. An immediate consequence of this insight is that—even without referencing formal properties such as rationalizability—selecting any one of the three alternatives in the Condorcet paradox profile will lead to dissatisfaction among a majority of voters. Furthermore, this majority is in agreement about which alternative should replace the chosen one. However, this replacement will not provide any relief, as another majority will propose to replace the chosen alternative with yet another one. This cycle of replacements can continue indefinitely. In order to arrive at Arrow's impossibility theorem in Chapter 4, we are going to significantly weaken anonymity, neutrality, and positive responsiveness$_2$ while slightly strengthening rationalizability to transitive rationalizability.

When disregarding majority ties, there are no majority cycles exactly when every feasible set admits an alternative that strictly majority-dominates every other alternative in the set. Such an alternative is called a Condorcet winner. Formally, alternative x is a

Condorcet winner

Condorcet winner in $A \in \mathcal{F}$ if

$$x \succ_M y \text{ for all } y \in A \setminus \{x\}.$$

As we have seen, Condorcet winners may not exist. Whenever they do exist, they are

weak Condorcet winner

unique. Alternative x is a *weak Condorcet winner* in $A \in \mathcal{F}$ if $x \succsim_M y$ for all $y \in A$. Weak Condorcet winners need not be unique. In fact, it is possible that all alternatives are weak Condorcet winners. The set of weak Condorcet winners for given $A \in \mathcal{F}$ and $P \in \mathcal{D}$

Cond

is defined as

$$\mathsf{Cond}(A, P) = \text{Max}(A, \succsim_M).$$

How likely is the existence of a Condorcet winner? To address this question, let $\mathcal{D} = \mathcal{S}^N$ and define $p_{\text{Cond}}(m, n)$ as the proportion of preference profiles for m alternatives and n voters that admit a Condorcet winner. This proportion can be interpreted as the probability of the existence of a Condorcet winner under the simplifying assumption that every preference profile is equally likely (this is known as the "impartial culture" model). For example, for the case of three alternatives and three voters, $p_{\text{Cond}}(3, 3) = {}^{17}/_{18}$. To see this, observe that when $m = n = 3$, only "latin square" profiles—in which each alternative appears exactly once in each row and column—admit no Condorcet winner: if an alternative is ranked first twice, it is obviously a Condorcet winner; given that the three first-ranked alternatives are different from each other, any alternative that is ranked second twice, also has to be a Condorcet winner. This only leaves the latin square profiles. Now fix, without loss of generality, the preferences of the first voter to $a \succ_1 b \succ_1 c$. Since there are $3! = 6$ preference relations on three alternatives, there are $6 \cdot 6 = 36$ combinations of preference relations for the remaining two voters. There can be no majority ties with three voters. Hence, if there is *no* Condorcet winner, $\succsim_M$ has to be cyclic. This is only possible when $b \succ_i c \succ_i a$ and $c \succ_j a \succ_j b$ for $i = 2$, $j = 3$ or $i = 3$, $j = 2$. Thus, $\succsim_M$ is cyclic in two out of 36 cases. Hence, $p_{\text{Cond}}(3, 3) = 1 - {}^{2}/_{36} = {}^{17}/_{18}$ (about 94%).

There have been significant efforts to compute p_{Cond} for other combinations of m and n. Of particular interest are the limit values when n goes to infinity, which allows us to abstract away from the subtle issue of majority ties. The following table shows some limit proportions $p_{\text{Cond}}(m, n)$ for varying m when $n \to \infty$ (Gehrlein and Fishburn, 1979).

m	1	3	5	7	9	11	13	15	17	19
$p_{\text{Cond}}(m, \infty)$	100%	91%	75%	63%	54%	48%	43%	39%	36%	33%

These values are surprisingly difficult to obtain. For example,

$$p_{\text{Cond}}(5, \infty) = \frac{5}{16} + \frac{15}{4\pi} \sin^{-1}\left(\frac{1}{3}\right) + \frac{15}{2\pi^2} \int_0^{\frac{1}{3}} \frac{\sin^{-1}(\alpha/(1+2\alpha))}{\sqrt{1-\alpha^2}} d\alpha.$$

Moreover, except for special cases, such as $p_{\text{Cond}}(3, 3)$ or $p_{\text{Cond}}(5, \infty)$, no closed-form expression for $p_{\text{Cond}}(m, n)$ is known and these values can only be approximated. It is not even known whether $p_{\text{Cond}}(m, n)$ decreases when m or n increase (n being odd)! Both statements ($p_{\text{Cond}}(m, n) > p_{\text{Cond}}(m + 1, n)$ and $p_{\text{Cond}}(m, n) > p_{\text{Cond}}(m, n + 2)$ for $m, n \geqslant 3$ and n odd) were conjectured by Kelly (1974).[3]

Even though intuition and numerical computations suggest that both conjectures are true, a proof remains elusive. The only noteworthy progress was made by Fishburn et al. (1979), who proved the conjectures for the special cases when either $m = 3$ or $n = 3$. They conclude that "the more general case will be extremely difficult to resolve" and may require "entirely new techniques that we cannot now foresee." Kelly (1988) writes that "obtaining a proof of this claim or finding a counterexample is one of the oldest unsolved problems in social choice theory."

3.5 Key Takeaways

From Individual Choice to Social Choice

- Resolute social choice often conflicts with fairness.
- Majority rule is essential when deciding between two alternatives.
- Condorcet winners do not always exist.

3.6 Further Reading

The original version of Theorem 3.1 by Moulin (1983) also needs Pareto-optimality because Moulin does not consider a framework with variable feasible sets. The two-alternative case of social choice is extensively covered by Fishburn (1973, Part I). Further information on the proportion of profiles that admit Condorcet winners has been compiled by Gehrlein (1983, 2006) and Fishburn (1984a, Section 3). More recently, Sauermann (2022) gives an asymptotic bound for $p_{\text{Cond}}(m, n)$ when m is large.

[3] It is known that $p_{\text{Cond}}(\infty, n) = 0$ for $n \geqslant 3$ (May, 1971).

3.7 Exercises

3.1 *Non-imposition*

Show that Pareto-optimality implies non-imposition.

3.2 *Independence of May's axioms*

Show that the axioms used in Theorem 3.2—neutrality, anonymity, and positive responsiveness—are independent from each other when $m = 2$, i.e., none of them is implied by the other two.

Hint: For each of these axioms, construct an SCF that does not satisfy it but satisfies all of the remaining axioms.

3.3 *May's theorem in subdomains*

Let $U = \{a, b\}$. Prove that majority rule is the only anonymous, neutral, and positively responsive SCF in each of the following domains.

(a) $\mathcal{D} = \mathcal{S}^N$

(b) $\mathcal{D} = \{P \in \mathcal{S}^N : n_{ab} \notin \{3, 7\}\}$ and $n = 8$

☆ **3.4** *A variant of Condorcet's Paradox*

Show that there is *no* anonymous, neutral, positively responsive$_2$, and Pareto-optimal SCF that satisfies *strong expansion* when $m \geqslant 3$ and n is even.

Hint: First, show that f is majority rule on feasible sets of size 2. Then show a contradiction by constructing a preference profile for $m = 3$ and $n = 2$ for which no alternative is selected.

I believe that Kenneth Arrow has made a first-rate contribution to man's body of knowledge.

Paul Samuelson, 1967

4 Arrow's Theorem

Learning Outcomes

- What are the limitations of rationalizable social choice functions?
- How can we aggregate preference relations into a collective preference relation?
- What are preference intensities, and should they be taken into account?

Arrow's impossibility theorem is the most central and influential result in social choice theory. We will arrive at this seminal statement by replacing the four assumptions of the Condorcet-May impossibility (Theorem 3.4) with similar properties. Neutrality, positive responsiveness$_2$, and anonymity will be replaced with conditions that are weaker under mild assumptions. Rationalizability, on the other hand, will be slightly strengthened to transitive rationalizability. We will then introduce an equivalent formulation of Arrow's theorem for *social welfare functions*, which map a preference profile to a collective preference relation. In this formulation, transitive rationalizability is inherently included in the definition of the aggregation function, making the result even cleaner and more compelling. We also discuss various variations of Arrow's theorem, which arise from weakening some axioms while strengthening others.

The reason we discuss these negative results in such detail is that they clearly define the limits of what can be achieved in principle. These results indicate that we do not need to waste time attempting to construct an SCF that simultaneously satisfies a number of seemingly mild properties. Instead, they highlight which property should be weakened or omitted altogether to obtain positive results. We conclude the chapter by pondering various ways to circumvent Arrovian impossibilities. Three such escape routes are explored in detail in Part II of this book.

4.1 Social Choice Functions

We start by defining an important axiom that we already encountered when discussing neutrality (see Section 3.1). This axiom demands that choices from any feasible set must only depend on the preferences of the voters among alternatives contained in the feasible set. An SCF f satisfies *independence of infeasible alternatives (IIA)* if for all $P, P' \in \mathcal{D}$ and $A \in \mathcal{F}$, IIA (for SCFs)

$$P|_A = P'|_A \quad \Rightarrow \quad f(A, P) = f(A, P').$$

This condition can be understood as a framework assumption that provides meaning to feasible sets. By definition, IIA is weaker than neutrality. To see this, let $A = B$ and σ be the identity function in the definition of neutrality on page 28.

Pareto-optimality (for SCFs)

For the sake of completeness, let us repeat the definition of *Pareto-optimality*: for all $A \in \mathcal{F}$, profiles $P \in \mathcal{D}$, and $x, y \in A$,

$$x \succ_i y \text{ for all } i \in N \quad \Rightarrow \quad y \notin f(A, P).$$

As mentioned earlier, Pareto-optimality is perhaps the weakest notion of social optimality: unanimously disliked alternatives should not be selected. Pareto-optimality$_2$ can be understood as a weakening of monotonicity$_2$ when also assuming non-imposition$_2$ and IIA$_2$ (see Exercise 4.1).

Finally, we demand that there should be no voter who always gets one of his most preferred alternatives, *no matter* what the preferences of all the other voters are. An SCF

dictatorial SCF

f is *dictatorial* if there exists $i \in N$ such that for all profiles $P \in \mathcal{D}$ and $A \in \mathcal{F}$,

$$f(A, P) \subseteq \text{Max}(A, \succsim_i).$$

The crucial part of the definition is the order of the existential and the universal quantifier. In the opposite order, this condition would be much less harmful—perhaps even desirable—as it would only require that for any profile and feasible set, there happens to be a voter whose favorite alternatives are chosen. Dictatorship, on the other hand, represents the most severe failure of anonymity: there is a fixed voter who can single-handedly decide the social choice by submitting preferences with a unique top choice, independently of the preferences of everyone else! Since there can be at most one dictator, non-dictatorship is weaker than anonymity whenever $n \geqslant 2$.

Note that a dictatorial SCF f with dictator i may consider the preferences of other voters when making selections from within the set $\text{Max}(A, \succsim_i)$. However, if voter i submits preferences such that $|\text{Max}(A, \succsim_i)| = 1$, then the outcome of f is completely determined by i.

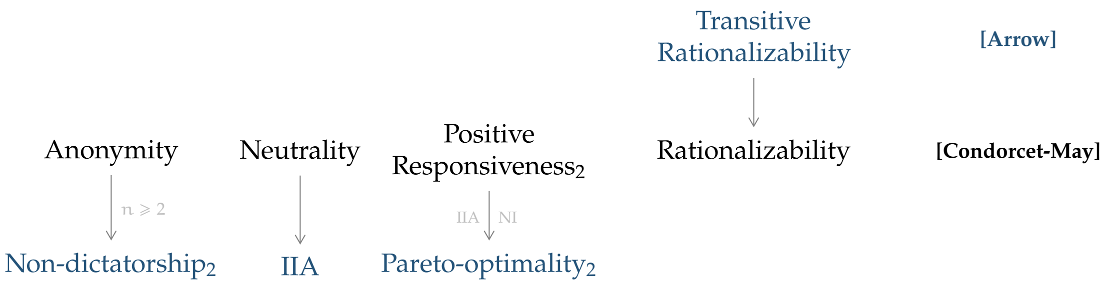

We are now ready to state Arrow's theorem for SCFs.

Theorem 4.1 (Arrow, 1951) $m \geqslant 3$

There is no transitively rationalizable SCF that satisfies non-dictatorship$_2$, IIA$_2$, and Pareto-optimality$_2$.

Note that none of the axioms argues over feasible sets that contain more than two alternatives. We still assume that $\mathcal{F} = \mathcal{P}^*(U)$, but Theorem 4.1 holds for any $\mathcal{F}$ that contains all two-element subsets of U.

4.2 Social Welfare Functions

Arrow's theorem is best known in its formulation for so-called social welfare functions. A *social welfare function (SWF)* is a function $g : \mathcal{D} \to \mathcal{W}$. An SWF aggregates individual preference relations into a collective or social preference relation.

social welfare function (SWF)

There is a simple one-to-one correspondence between transitively rationalizable SCFs and SWFs. As we have seen in Lemma 2.2, rationalizable SCFs only depend on their choices from two-element feasible sets. Hence, every SCF f admits an equivalent SWF g such that for all $x, y \in U$,

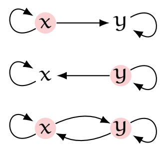

$$x \in f(\{x,y\}, P) \quad \Leftrightarrow \quad x \succsim_{f(\cdot,P)} y \quad \Leftrightarrow \quad x\; g(P)\; y.$$

One can think of the collective preference relation returned by an SWF as the rationalizing relation of the corresponding SCF and *vice versa*. The three possible choices from a two-element set correspond to the three possible complete relations between two alternatives shown in the margin.

IIA$_2$, Pareto-optimality$_2$, and non-dictatorship$_2$ can thus be appropriately redefined for SWFs by considering the SCF's base relation. Let g be an SWF and denote the collective preference relation of some profile $P \in \mathcal{D}$ by $\succsim = g(P)$. Similarly, we let $\succsim' = g(P')$ for $P' \in \mathcal{D}$.

g satisfies *independence of irrelevant alternatives (IIA)* if for all $P, P' \in \mathcal{D}$ and $x, y \in U$,

IIA (for SWFs)

$$P|_{\{x,y\}} = P'|_{\{x,y\}} \quad \Rightarrow \quad \succsim|_{\{x,y\}} = \succsim'|_{\{x,y\}}.$$

g is *Pareto-optimal* if for all $P \in \mathcal{D}$ and $x, y \in U$,

Pareto-optimality (for SWFs)

$$x \succ_i y \text{ for all } i \in N \quad \Rightarrow \quad x \succ y.$$

g is *dictatorial* if there exists $i \in N$ such that for all $P \in \mathcal{D}$ and $x, y \in U$,

dictatorship (for SWFs)

$$x \succ_i y \quad \Rightarrow \quad x \succ y.$$

Since the transitivity of collective preferences is already incorporated into the definition of SWFs, we obtain the following compelling formulation of Arrow's theorem.

Theorem 4.2 (Arrow, 1951) $m \geqslant 3$

Every SWF that satisfies IIA and Pareto-optimality is dictatorial.

Whether an SWF satisfies Pareto-optimality and non-dictatorship is usually quite obvious, and SWFs that are typically used to aggregate the preferences of multiple voters in real-world applications satisfy these conditions. It follows from Theorem 4.2 that any such SWF has to violate IIA.

4.3 Illustrations and Consequences

To grasp the practical consequences of IIA failures, consider a university that lets students rank courses and then publishes a collective ranking based on the courses' Borda scores. For simplicity, let us assume there are three courses (a, b, and c) and 5 students (you are welcome to multiply this number by any number of your choice to make the example more realistic). Courses a and b are quite popular while Course c is not (see profile P below).

	3	2
2	b	a
1	a	c
0	c	b

P

$\underset{7}{a} \succ \underset{6}{b} \succ \underset{2}{c}$

	3	2
2	b	a
1	a	b
0	c	c

P'

$\underset{8}{b} \succ \underset{7}{a} \succ \underset{0}{c}$

	3	2
3	b	a
2	b'	c
1	a	b
0	c	b'

$\hat{P}$

$\underset{11}{b} \succ \underset{9}{a} \succ \underset{6}{b'} \succ \underset{4}{c}$

In the collective ranking for profile P, Course a is ranked before Course b (grey numbers denote Borda scores). Now, in the next year, Course c becomes even worse, while a and b remain unchanged from the previous year (see P'). As a consequence, Course b will be ranked before Course a in the collective Borda ranking. This situation represents a failure of IIA because the individual preferences between a and b have not changed ($P|_{\{a,b\}} = P'|_{\{a,b\}}$); course c is an "irrelevant alternative" for the comparison of a and b. To illustrate a similar phenomenon without affecting the position of c, let us go back to the original profile P. The lecturer of Course b is unhappy about his narrow defeat in profile P and, therefore, decides to offer an additional variant of Course b, which is just like Course b, only a little bit worse. The resulting profile $\hat{P}$ leads to a Borda ranking, in which b is ranked before a. Again, the individual preferences between a and b in P and $\hat{P}$ are identical, resulting in a violation of IIA.

	3	2
2	a	c
1	b	a
0	c	b

For another example that frequently occurs in the real world, think of a five-member hiring committee that ranks three job candidates according to their Borda scores, as shown in the margin. The Borda scores for Candidates a, b, and c are 8, 3, and 4. Hence, the first offer is made to Candidate a. Now suppose Candidate a politely declines the offer, and the committee then issues an offer to their second choice, Candidate c. This does not seem like a good idea because, once Candidate a is out, the choice is between Candidates b and c, and a majority of the committee members prefer b to c.[1]

Theorem 4.2 shows that these phenomena are not peculiarities of Borda scores: *every* reasonable SWF will run into these problems in some situations!

Since all conditions in Theorem 4.2 are universal and intuitive, the implications of Arrow's theorem extend far beyond the analysis of voting procedures. Arrow raised the question of whether a reasonable notion of *collective preference* can be defined and showed that such a notion cannot be sustained under seemingly mild and innocuous conditions. This had a deep impact on welfare economics, which is concerned with the well-being of societies of agents. Simply put, Theorem 4.2 showed that it is unclear what is best for a society of agents and that societies cannot be treated like single rational decision-makers.

In the context of voting, the message of Arrow's theorem is sometimes incorrectly reduced to the discouraging conclusion that "only dictatorships or two-party systems can work." When having a closer look at Arrow's theorem for social *choice* functions (Theorem 4.1), it becomes evident that Arrow pointed out a conflict between the independence of infeasible alternatives and collective rationality (or, equivalently, choice consistency). Before we discuss these issues in Section 4.6, let us prove that Theorem 4.2 holds.

4.4 Proof of Arrow's Theorem

To prove Arrow's theorem, we introduce two properties of groups of agents: decisiveness and semi-decisiveness. To this end, let g be an SWF, $G \subseteq N$ a group of voters, and $a, b \in U$ with $a \neq b$ two distinct alternatives.

We say that G is *decisive for* a *against* b (denoted by $a\ D_G\ b$) if for all $P \in \mathcal{D}$, decisiveness $a\ D_G\ b$

$$a \succ_i b \text{ for all } i \in G \quad \Rightarrow \quad a \succ b.$$

Clearly, if $G \subseteq N$ is decisive for some pair, then every $G' \subseteq N$ with $G \subseteq G'$ is also decisive for this pair. Whenever G is decisive for all pairs $a, b \in U$, we simply say that G is *decisive*. The existence of a decisive group is not undesirable. In fact, Pareto-optimality implies (and is implied by) the decisiveness of N. Moreover, since supersets of decisive groups are decisive, demanding Pareto-optimality is equivalent to demanding the existence of a decisive group. If $m = 2$ and g is majority rule, then G is decisive iff $|G| > n/2$. What is problematic is the decisiveness of small groups of voters, in particular, singleton decisive groups, which correspond to dictators.

We say that G is *semi-decisive for* a *against* b (denoted by $a\ \tilde{D}_G\ b$) if for all $P \in \mathcal{D}$, semi-decisiveness $a\ \tilde{D}_G\ b$

[1] The same problem appears if we had used plurality scores instead of Borda scores. What seems like a better idea is to recompute all scores once a candidate rejects an offer. In Section 6.6, we will explore Kemeny's rule, which is particularly well-suited for this type of application, even though it still does not satisfy IIA.

$$a \succ_i b \text{ for all } i \in G \text{ and } b \succ_j a \text{ for all } j \notin G \quad \Rightarrow \quad a \succ b.$$

Clearly, $a\ D_G\ b$ implies $a\ \tilde{D}_G\ b$ (but not vice versa). The notions of decisiveness and semi-decisiveness can also be used for a rationalizable SCF f. The relation $\succsim$ in the consequence of the conditions then refers to the base relation of f.

The following lemma shows that if some group is semi-decisive over some pair, then it is decisive (over every pair).

Lemma 4.1 (Field Expansion)

Let $m \geqslant 3$, g be an SWF that satisfies IIA and Pareto-optimality, and $G \subseteq N$. Then,

$$(\exists a, b \in U\colon a\ \tilde{D}_G\ b) \quad \Rightarrow \quad (\forall x, y \in U\colon x\ D_G\ y)$$

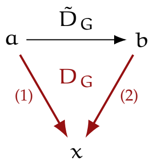

Proof. Let $a\ \tilde{D}_G\ b$ and $x \neq a, b$. We now show that (1) $a\ D_G\ x$ and (2) $b\ D_G\ x$.

(1) Consider an arbitrary $P \in \mathcal{D}$ with $a \succ_i x$ for all $i \in G$. Then, define P' with $P|_{\{a,x\}} = P'|_{\{a,x\}}$ but $a \succ'_i b \succ'_i x$ for all $i \in G$ and $b \succ'_j a, x$ for all $j \notin G$, as shown in the margin.

G	N \ G
a	b
b	·
x	·

$P'|_{\{a,b,x\}}$

$$\left.\begin{array}{rcl} a\ \tilde{D}_G\ b & \Rightarrow & a \succ' b \\ \forall i \in N\colon b \succ'_i x & \overset{\text{Pareto}}{\Rightarrow} & b \succ' x \end{array}\right\} \overset{\text{Trans.}}{\Rightarrow} a \succ' x \overset{\text{IIA}}{\Rightarrow} a \succ x \overset{\text{Def.}}{\Rightarrow} a\ D_G\ x$$

(2) Consider an arbitrary $P \in \mathcal{D}$ with $b \succ_i x$ for all $i \in G$. Then, define P' with $P|_{\{b,x\}} = P'|_{\{b,x\}}$ but $b \succ'_i a \succ'_i x$ for all $i \in G$ and $b, x \succ'_j a$ for all $j \notin G$, as shown in the margin.

G	N \ G
b	·
a	·
x	a

$P'|_{\{a,b,x\}}$

$$\left.\begin{array}{rcl} a\ D_G\ x & \Rightarrow & a \succ' x \\ \forall i \in N\colon b \succ'_i a & \overset{\text{Pareto}}{\Rightarrow} & b \succ' a \end{array}\right\} \overset{\text{Trans.}}{\Rightarrow} b \succ' x \overset{\text{IIA}}{\Rightarrow} b \succ x \overset{\text{Def.}}{\Rightarrow} b\ D_G\ x$$

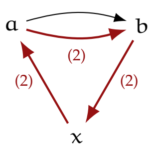

Repeated application of (2) implies that the relation D_G is *complete*, i.e., for all $x, y \in U$, $x\ D_G\ y$ or $y\ D_G\ x$. First, for $x \neq a, b$, we have

$$a\ \tilde{D}_G\ b \overset{(2)}{\Rightarrow} b\ D_G\ x \overset{(2)}{\Rightarrow} x\ D_G\ a \overset{(2)}{\Rightarrow} a\ D_G\ b,$$

as illustrated in the margin. For distinct $x, y \neq a, b$, we have

$$a\ \tilde{D}_G\ b \overset{(2)}{\Rightarrow} b\ D_G\ x \overset{(2)}{\Rightarrow} x\ D_G\ y.$$

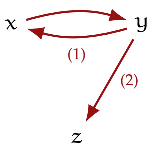

Moreover, D_G is *symmetric*, i.e., for all $x, y \in U$, $x\ D_G\ y$ implies $y\ D_G\ x$. To see this, let $x, y, z \in U$. Then,

$$x\ D_G\ y \overset{(2)}{\Rightarrow} y\ D_G\ z \overset{(1)}{\Rightarrow} y\ D_G\ x,$$

as illustrated in the margin. (Note that, when $m \geqslant 4$, symmetry of D_G follows directly from repeated application of (2).) □

Note that Lemma 4.1 only requires *quasi*-transitivity of the collective preference relation. We will reuse it for the proof of Theorem 4.3.

The next lemma establishes that every decisive group, consisting of at least two voters, contains a strictly smaller decisive subgroup. This will enable us to show that Pareto-optimality implies the existence of a dictator.

Lemma 4.2 (Group Contraction)

Let $m \geqslant 3$, g be an SWF that satisfies IIA and Pareto-optimality. If $G \subseteq N$ is decisive and $|G| \geqslant 2$, then there is a decisive group $G' \subset G$.

Proof. Since $|G| \geqslant 2$, we can partition G into two nonempty subgroups G_1, G_2 such that $G = G_1 \cup G_2$ and $G_1 \cap G_2 = \emptyset$. Now consider the following familiar profile P, which assigns a single preference relation to all voters in G_1, another one to all voters in G_2, and a third one to all remaining voters. (The preferences over alternatives other than a, b, and c are irrelevant.)

G_1	G_2	$N \setminus G$
a	b	c
b	c	a
c	a	b

$P|_{\{a,b,c\}}$

Note that $N \setminus G$ can be empty, which means that there are no voters associated with the preferences of the third column. As always, $\succsim = g(P)$. First, $b \; D_G \; c$ implies that $b \succ c$. We now distinguish between two cases.

First, assume that $a \succ c$. For all voters $i \in G_1$, we have $a \succ_i c$, while for all voters $j \notin G_1$, we have $c \succ_j a$. IIA implies that the collective preference $a \succ c$ has to remain intact for all profiles P' in which the above is true. This means that $a \; \tilde{D}_{G_1} \; c$. By Lemma 4.1, G_1 is decisive.

The other case is that $c \succsim a$. We then know from $b \succ c$ and transitivity of $\succsim$ that $b \succ a$. (Otherwise, $c \succsim a$ and $a \succsim b$ would imply $c \succsim b$, a contradiction.) For all voters $i \in G_2$, we have $b \succ_i a$, while for all voters $j \notin G_2$, we have $a \succ_j b$. Analogously to the previous case, we therefore have that $b \; \tilde{D}_{G_2} \; a$ and, by Lemma 4.1, that G_2 is decisive. □

This lemma allows us to easily prove Arrow's theorem. Let $m \geqslant 3$ and g be an SWF that satisfies IIA and Pareto-optimality. The latter condition implies that N is decisive. Now let G be an inclusion-minimal decisive group, which must exist by the finiteness of N. If $|G| \geqslant 2$, Lemma 4.2 implies that a strict subset of G is decisive, contradicting the inclusion-minimality of G. Hence, $G = \{i\}$ for some $i \in N$, which means that voter i is a dictator.

Arrow's theorem for SCFs (Theorem 4.1) shows that every non-dictatorial and Pareto-optimal SCF *(i)* cannot be rationalized by a transitive collective preference relation or *(ii)* violates independence of irrelevant alternatives. All conditions of Theorem 4.1 are required for the impossibility to hold (see Exercise 4.2).

Sen (1977, pp. 78f) provides an illuminating discussion of two different interpretations of Borda's rule for variable feasible sets that highlight the tension between transitive rationalizability and IIA. While IIA demands that the choice set only depends on preferences over elements contained in the feasible set, rationalizability requires that choices from any nonempty feasible set $A \subseteq U$ can be rationalized by a single collective preference relation ranging over all alternatives in U. Now, the *broad* Borda rule first assigns Borda scores to all alternatives in U and then returns the alternatives with maximal scores within A. By contrast, the *narrow* Borda rule directly assigns Borda scores to alternatives in A and then returns those with maximal score.

The broad Borda rule can be rationalized by a transitive collective preference relation (the ranking of all alternatives by their Borda score), but violates IIA, while the narrow Borda rule satisfies IIA, but cannot be rationalized by any preference relation (it violates contraction). Arrow's theorem shows that this tradeoff concerns *all* non-dictatorial and Pareto-optimal social choice functions. Moreover, as we will see in Chapter 7, Sen observed that transitive rationalizability can be replaced with contraction in Arrow's theorem and many related results. While voters could justifiably complain that, under the broad Borda rule, the social choice from feasible set $\{a, b, c\}$ depends on their preferences over other unrelated alternatives, say, d or e (violating IIA), they could be similarly concerned about the narrow Borda rule, under which it is possible that alternative a is chosen from $\{a, b, c\}$ but not from $\{a, b\}$. Both phenomena are troubling: the lack of IIA because seemingly irrelevant information is taken into account for the social choice, and the lack of contraction because introducing or removing, say, clearly inferior alternatives can influence the social choice.

4.5 Variants of Arrow's Theorem

A natural question is which rationalizable SCFs satisfy IIA_2 and Pareto-optimality$_2$ when only requiring that the base relation is quasi-transitive or acyclic, rather than transitive. As we have seen in Lemma 2.1, the weakest form of rationalizability we consider suffices for the existence of maximal elements in every feasible set. The Pareto dominance relation is an asymmetric and transitive relation. As a consequence, PO is quasi-transitively rationalizable. (The profile in the margin shows that PO fails *transitive* rationalizability.) Furthermore, it is easily seen that PO also satisfies IIA_2 and Pareto-optimality$_2$ and thus satisfies two of Arrow's axioms plus quasi-transitive rationalizability. PO is obviously very indecisive, as it almost always returns all alternatives. As we will see in Theorem 4.3 below, quasi-transitive rationalizability generally only allows for SCFs that exhibit a similarly severe lack of decisiveness or undesirably high concentration of power, such as in dictatorships. Arrow's theorem is remarkably robust in the sense that the impossibility still holds when weakening one axiom and strengthening another.

1	1
a	c
b	a
c	b

To state variants of Arrow's theorem, we will define four groups of voters that have less power than dictators: weak dictators, quasi-dictators, oligarchies, and collegia. We are interested in the *non-existence* of these groups. Hence, requiring the non-existence of dictators is the weakest desideratum in the diagram below.

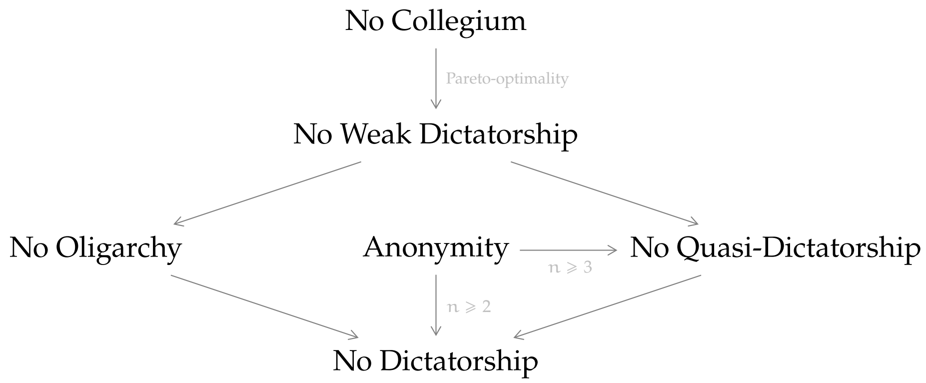

Recall the definition of a dictator. A dictator is a voter i such that for all $P \in \mathcal{D}$ and A, $f(A, P) \subseteq \text{Max}(A, \succsim_i)$. In particular, when $A = \{x, y\}$, $x \succ_i y$ implies that $\{x\} = f(\{x, y\}, P)$, or for SWFs, $x \succ_i y$ implies $x \succ y$. A dictator is a decisive group that consists of a single element. When a voter is a dictator when reversing his preferences, he is called a *reverse dictator*. In other words, when i is a reverse dictator, then for all $P \in \mathcal{D}$ and A, $f(A, P) \subseteq \text{Max}(A, \precsim_i)$. reverse dictator

A *weak dictator* is a voter i such that for all $P \in \mathcal{D}$ and $A \in \mathcal{F}$, $f(A, P) \cap \text{Max}(A, \succsim_i) \neq \emptyset$. In particular, when $A = \{x, y\}$, $x \succ_i y \Rightarrow x \in f(\{x, y\}, P)$, or for SWFs, $x \succ_i y$ implies $x \succsim y$. A weak dictator for SCFs can force an alternative into the choice set while a weak dictator for SWFs can veto a strict collective preference. Weak dictators are, therefore, sometimes called *vetoers*. weak dictator

A *quasi-dictator* is a weak dictator i that forms a decisive group when he joins forces with any other voter, i.e., every $G \subseteq N$ with $i \in G$ and $|G| \geqslant 2$ is decisive. quasi-dictator

An *oligarchy* is a decisive group consisting of weak dictators. An oligarchy can single-handedly decide on strict collective preferences, and, moreover, every member can veto a strict collective preference. As a consequence, whenever two oligarchs have conflicting strict preferences, there has to be collective indifference. The preferences of non-oligarchy members can only be taken into account when oligarchs have non-conflicting preferences and at least one oligarch is indifferent. oligarchy

A *collegium* is the nonempty intersection of all decisive groups. A collegium is neither decisive nor weakly decisive. However, if a voter is contained in a collegium, even all remaining $n - 1$ voters taken together do not have enough power to enforce a strict collective preference. collegium

Whenever there is a dictator, this voter forms a singleton oligarchy. Moreover, the existence of an oligarchy implies the existence of a weak dictator, as every oligarch, by definition, is a weak dictator. Whenever a decisive group exists (for example, because of Pareto-optimality), any weak dictator has to be contained in all decisive groups and thus in the collegium. Otherwise, a decisive group could enforce a strict preference while the weak dictator could prevent the same strict preference. As already pointed out earlier, non-dictatorship is weaker than anonymity when $n \geqslant 2$. When $n \geqslant 3$, anonymity also prevents the existence of quasi-dictators. Suppose there are two quasi-dictators, i and j, and a third voter, k. i and k form a decisive group that can dictate strict preferences. j, however, is a weak dictator and thus can veto strict preferences, yielding a contradiction. The implications are completed by the simple observations that every dictator is a quasi-dictator and every quasi-dictator is a weak dictator. There is no logical relationship between oligarchies and quasi-dictators.

The table below shows eight Arrovian impossibility theorems. Each row corresponds to a theorem: no SCF satisfies all properties listed in the same row.

Anonymity$_2$	Neutrality	Positive Responsiveness$_2$	Rationalizability	$m \geqslant 3$ $n \geqslant 3$	Condorcet (1785) May (1952)
No Dictator$_2$	IIA$_2$	Pareto-Optimality$_2$	Transitive Rationalizability	$m \geqslant 3$	Arrow (1951)
No (reverse) Dictator$_2$	IIA$_2$	Non-Imposition$_2$	Transitive Rationalizability	$m \geqslant 3$	Wilson (1972)
No Oligarchy$_2$	IIA$_2$	Pareto-Optimality$_2$	Quasi-Transitive Rationalizability	$m \geqslant 3$	Gibbard (1969)
No Quasi-Dictator$_2$	IIA$_2$	Positive Responsiveness$_2$	Rationalizability	$m \geqslant 3$ $n \geqslant 4$	Mas-Colell and Sonnenschein (1972)
No Collegium$_2$	—	Pareto-Optimality$_2$	Rationalizability	$m \geqslant n$	Brown (1975) Banks (1995a)
No Weak Dictator$_2$	Neutrality	Monotonicity$_2$	Rationalizability	$m \geqslant n$	Blau and Deb (1977)
No Weak Dictator$_2$	Neutrality	Pareto-Optimality$_2$	Rationalizability	$m > n$	Blair and Pollak (1982)

We have already discussed the Condorcet-May impossibility (Theorem 3.4) and Arrow's theorem (Theorem 4.1). Wilson (1972) showed that if we replace Pareto-optimality$_2$ with the much weaker property of non-imposition$_2$, we only get reverse dictatorships on top of dictatorships.

Retaining Pareto-optimality$_2$ but weakening transitive rationalizability to quasi-transitive rationalizability implies the existence of an oligarchy. To prove this, we define weaker variants of decisiveness and semi-decisiveness.

Voter $i \in N$ is *weakly semi-decisive* for a against b (denoted by $a \tilde{W}_i b$) if for all $P \in \mathcal{D}$,

$$a \succ_i b \text{ and } b \succ_j a \text{ for all } j \neq i \quad \Rightarrow \quad a \succsim b.$$

Similarly, voter i is *weakly decisive* for a against b (denoted by $a\, W_i\, b$) if for all $P \in \mathcal{D}$,

$$a \succ_i b \quad \Rightarrow \quad a \succsim b.$$

Theorem 4.3 (Gibbard, 1969) $m \geqslant 3$

Every quasi-transitively rationalizable SCF that satisfies IIA$_2$ and Pareto-optimality$_2$ is oligarchic.

Proof. Since every rationalizable SCF f is completely determined by its base relation, it suffices to argue about the base relations induced by f. The base relations of $f(\cdot, P)$ and $f(\cdot, P')$ for profiles $P, P' \in \mathcal{D}$ will simply be denoted by $\succsim$ and $\succsim'$.
We first show the implication below. The statement and the proof closely resemble those of Lemma 4.1. Let $m \geqslant 3$, f be a quasi-transitively rationalizable SCF that satisfies IIA and Pareto-optimality, and $i \in N$. Then,

$$(\exists a, b \in U\colon a \tilde{W}_i b) \quad \Rightarrow \quad \text{Voter } i \text{ is a weak dictator}$$

Let $a \tilde{W}_i b$ and $x \neq a, b$. We now show that (1) $a W_i x$ and (2) $b W_i x$.

(1) Consider an arbitrary $P \in \mathcal{D}$ with $a \succ_i x$. Then, define P' with $P|_{\{a,x\}} = P'|_{\{a,x\}}$ but $a \succ'_i b \succ'_i x$ and $b \succ'_j a, x$ for all $j \neq i$, as shown in the margin.

$\{i\}$	$N \setminus \{i\}$
a	b
b	$\cdot$
x	$\cdot$

$P'|_{\{a,b,x\}}$

$$\left.\begin{array}{rcl} a \tilde{W}_i b & \Rightarrow & a \succsim' b \\ \forall i \in N\colon b \succ'_i x & \overset{\text{Pareto}}{\Rightarrow} & b \succ' x \end{array}\right\} \overset{\text{Q.–Trans.}}{\Rightarrow} a \succsim' x \overset{\text{IIA}}{\Rightarrow} a \succsim x \overset{\text{Def.}}{\Rightarrow} a W_i x$$

(2) Consider an arbitrary $P \in \mathcal{D}$ with $b \succ_i x$. Then, define P' with $P|_{\{b,x\}} = P'|_{\{b,x\}}$ but $b \succ'_i a \succ'_i x$ and $b, x \succ'_j a$ for all $j \neq i$, as shown in the margin.

$\{i\}$	$N \setminus \{i\}$
b	$\cdot$
a	$\cdot$
x	a

$P'|_{\{a,b,x\}}$

$$\left.\begin{array}{rcl} a W_i x & \Rightarrow & a \succsim' x \\ \forall i \in N\colon b \succ'_i a & \overset{\text{Pareto}}{\Rightarrow} & b \succ' a \end{array}\right\} \overset{\text{Q.–Trans.}}{\Rightarrow} b \succsim' x \overset{\text{IIA}}{\Rightarrow} b \succsim x \overset{\text{Def.}}{\Rightarrow} b W_i x$$

Just like in the proof of Lemma 4.1, repeated application of (1) and (2) shows that the relation W_i is complete and symmetric. Hence, $x W_i y$ for all $x, y \in U$, which means that voter i is a weak dictator. To show that a voter is a weak dictator, it thus suffices to show that this voter is weakly semi-decisive over one pair of alternatives.
The rest of the proof resembles that of Lemma 4.2. Let G be an inclusion-minimal decisive group, i.e., G is decisive and does not contain another decisive group. The existence of such a group is guaranteed as the group N is decisive due to Pareto-optimality$_2$ and n is finite. Now, if $|G| = 1$, we are done since a decisive group of size 1 is an oligarchy. We hence assume that $|G| > 1$ and aim to show that every voter in G is a weak dictator. To this end, let $i \in G$ be an arbitrary voter in this group and consider the following profile P.

$\{i\}$	$G \setminus \{i\}$	$N \setminus G$
a	b	c
b	c	a
c	a	b

$P|_{\{a,b,c\}}$

First, $b D_G c$ implies that $b \succ c$.
Assume for contradiction that $c \succ a$. We then know from $b \succ c$ and quasi-transitivity of $\succsim$ that $b \succ a$. For all voters $j \in G \setminus \{i\}$, we have $b \succ_j a$, while for all voters $k \notin G \setminus \{i\}$,

we have $a \succ_k b$. Hence, IIA implies that $b\ \tilde{D}_{G\setminus\{i\}}\ a$. It follows from Lemma 4.1 (which also holds for quasi-transitively rationalizable SCFs) that $G \setminus \{i\}$ is decisive, which contradicts the minimality of G.
As a consequence, $a \succsim c$. Since $a \succ_i c$ and for all voters $j \neq i$, $c \succ_j a$, IIA implies that $a\ \tilde{W}_i\ c$. Invoking the statement shown at the beginning of this proof, we thus obtain that i is a weak dictator. Since i was picked arbitrarily, all voters in G are weak dictators, and G is an oligarchy. □

When further weakening quasi-transitive rationalizability to just rationalizability and replacing Pareto-optimality$_2$ with positive responsiveness$_2$, Mas-Colell and Sonnenschein (1972) have proved the existence of a quasi-dictator.

Positive responsiveness$_2$ is a comparatively strong assumption.[2] Based on earlier work by Brown (1975), Banks (1995a) has shown that for acyclic preference aggregation, Pareto-optimality$_2$ and $m \geqslant n$ already imply the existence of a collegium, even without demanding IIA$_2$.

Theorem 4.4 (Brown, 1975; Banks, 1995a) $m \geqslant n$

Every rationalizable SCF that satisfies Pareto-optimality$_2$ admits a collegium.

Proof. Let f be a Pareto-optimal$_2$ SCF. Pareto-optimality$_2$ implies that N is decisive. Now, assume for contradiction that f does not admit a collegium. Then, for every $i \in N$, there exists a decisive group $G \subseteq N$ with $i \notin G$. Since all supersets of decisive groups are also decisive, we even have that $N \setminus \{i\}$ is decisive for every $i \in N$. Since $m \geqslant n$, let $x_1, \dots, x_n \in U$ be n distinct alternatives and consider a preference profile $P \in \mathcal{D}$ whose restriction to $\{x_1, \dots, x_n\}$ looks as follows.

1	1	…	1
x_1	x_2	…	x_n
x_2	x_3	…	x_1
⋮	⋮	⋱	⋮
x_n	x_1	…	x_{n-1}

$P|_{\{x_1,\dots,x_n\}}$

Then, let $x_0 = x_n$ and observe that for all $i \in N$, $N_{x_{i-1}x_i} = N \setminus \{i\}$. Since all of these groups are decisive, the base relation $\succsim$ of f satisfies

$$x_1 \succ x_2 \succ \dots \succ x_n \succ x_1,$$

violating acyclicity of $\succsim$. □

The proof of Theorem 4.4 strongly hinges on the assumption that $m \geqslant n$. It is quite different from Arrow's proof, as the statement does not even require IIA$_2$. As a

[2] Duggan (2016) has shown that the theorem by Mas-Colell and Sonnenschein (1972) remains intact for certain weakenings of positive responsiveness$_2$ (see Section 4.9).

consequence, any Pareto-optimal SCF that selects alternatives from feasible sets based on a single global ranking—such as the broad Borda rule—admits a collegium. In the case of the broad Borda rule, N is the only decisive group when $m \geqslant n$. A collegium consisting of *all* voters indicates a lack of decisiveness rather than a high concentration of power. The narrow Borda rule avoids the existence of a collegium (when $n \geqslant 3$) but fails to be rationalizable.

When further increasing m, such that $m > n$, and assuming neutrality (which implies IIA), Blair and Pollak (1982) have proved the existence of a weak dictator.

Theorem 4.5 (Blair and Pollak, 1982) $m > n$

Every rationalizable SCF that satisfies neutrality and Pareto-optimality$_2$ is weakly dictatorial.

Proof. Let f be an SCF with the stated properties. Further, let $x_1, \dots, x_{n+1} \in U$ be $n+1$ distinct alternatives and assume for contradiction that no voter $i \in N$ is weakly decisive for x_{i+1} against x_i. Hence, by IIA, for every $i \in N$, there is a profile $P^i \in \mathcal{D}$ such that $x_{i+1} \succ_i^i x_i$ but $x_i \succ^i x_{i+1}$. Now select a profile $P \in \mathcal{D}$ such that for all $i \in N$,

$$P|_{\{x_i, x_{i+1}\}} = P^i|_{\{x_i, x_{i+1}\}} \text{ and } x_{n+1} \succ_i x_1.$$

To see that such a profile P (consisting of transitive individual preference relations) exists, observe that for each voter $i \in N$, we specify the pairwise preferences between x_j and x_{j+1} (for all $j \in N$) and $x_{n+1} \succ_i x_1$. These $n+1$ pairwise constraints admit a transitive completion unless they form a cycle $x_1 \succ_i x_2 \succ_i \cdots \succ_i x_{n+1} \succ_i x_1$. However, this cycle requires $x_i \succ_i x_{i+1}$, which contradicts $x_{i+1} \succ_i^i x_i$.
We now obtain the following consequences for the base relation of f by IIA,

$$x_1 \succ x_2 \succ \cdots \succ x_n \succ x_{n+1}.$$

Moreover, by Pareto-optimality$_2$, $x_{n+1} \succ x_1$, which contradicts the acyclicity of $\succsim$.
We have shown that there is some $i \in N$ such that i is weakly decisive for x_{i+1} against x_i. By neutrality, voter i is weakly decisive for *all* pairs of alternatives and hence a weak dictator. □

When interpreting the lower bounds on m in Theorem 4.4 and Theorem 4.5, note that $m = |U|$ merely refers to the number of *potential* alternatives and the theorems hold even when we are never choosing from feasible sets that contain more than two alternatives (e.g., when $\mathcal{F} = \{\{x, y\} : x, y \in U\}$).

A theorem by Blau and Deb (1977) shows that the consequence of Theorem 4.5 even holds for $m \geqslant n$ when replacing Pareto-optimality$_2$ with monotonicity$_2$.

4.6 Escape Routes

Despite the persuasiveness of Arrow's theorem and its variants, it does not justify giving up and declaring fair and consistent social choice to be impossible. Rather, it precisely shows us which desiderata need to be weakened or replaced to achieve positive results. Part II of this book explores three escape routes from Arrow's theorem. Each of these routes is based on weakening or replacing one of the critical assumptions in Arrow's theorem. The first escape route is based on weakening an *implicit* assumption in Arrow's theorem, namely that the voters' preferences are unrestricted.

Restricted domains Restricting the domain of admissible preference profiles to dichotomous and single-peaked preferences, respectively, will lead us to two simple and attractive SCFs: *approval voting* and *median voting* (Chapter 5).

Omitting some of the axioms explicitly stated in Theorem 4.1 leaves very little room for positive results. Non-dictatorship, for example, seems non-negotiable since we are not interested in dictatorial SCFs. Similarly, going against the will of *all* people by giving up Pareto-optimality is undesirable. Pareto-optimality can be seen as a principle of minimal social consensus, and Wilson's theorem shows that, even if we were to give it up, no interesting possibilities arise. That leaves IIA and transitive rationalizability. For SCFs (not SWFs!), IIA is merely a framework assumption that gives meaning to feasible sets. We will, therefore, either abandon variable feasible sets and IIA altogether or weaken transitive rationalizability.

Variable electorate consistency Abandoning IIA and consistency with respect to variable feasible sets opens up countless possibilities. When then invoking a natural consistency condition with respect to variable electorates called reinforcement, we obtain axiomatic characterizations of scoring rules (such as *Borda's rule*) and *Kemeny's rule* (Chapter 6).

Weakening consistency Dropping contraction but retaining strong expansion, expansion, and a further weakening leads to axiomatic characterizations of the *top cycle*, the *uncovered set*, and the *Banks set*. Weakening the conjunction of contraction and strong expansion to a new condition called stability leads to the *top cycle* (again), the *minimal covering set*, and the *bipartisan set* (Chapter 7).

4.7 A Word on Cardinal Preferences

Throughout this book, we assume that the preferences of agents are represented by binary relations. At first glance, it may seem as if these relations miss an essential feature of preferences: *intensity*. In everyday language, we often claim to "strongly" prefer one thing to another, or perhaps even that we like one thing "twice as much." It is, therefore, tempting to replace preference relations with cardinal preferences or, more specifically, with *utility functions* that assign numerical values to each alternative, with higher numbers indicating more happiness. A natural solution to the problem of

utility functions

social choice could then be to select an alternative that maximizes the sum of individual utilities— a principle known as *utilitarianism*.

utilitarianism

The difficulty lies in giving cardinal preferences a coherent interpretation. With ordinal preferences, the meaning is straightforward and behaviorally observable: if an agent chooses a over b when both of them are available, he reveals a preference for a. By contrast, cardinal utility requires that numerical differences between utilities represent meaningful differences in preference intensity. This assumption is problematic because there is no convincing method for observing or measuring such magnitudes. Proponents of cardinal preferences sometimes argue that many humans are comfortable assigning scores to alternatives, even if they do not understand the exact meaning of these numbers. However, these numbers will likely be inconsistent among single individuals and certainly not comparable across different individuals. For many years, the editorial interface of Elsevier journals asked editors to (optionally) rate a manuscript on a scale from 0 to 100. Of course, an editor might quickly enter a 73 without hesitation to mark a good but not groundbreaking manuscript, but if pressed to justify why it is a 73 and not a 79, the choice would be difficult to defend. In modern microeconomic theory, utility is only assumed to be purely ordinal because there is no justification to assume that happiness, pleasure, or satisfaction are *measurable* quantities. During the so-called ordinal revolution in the 1930s, economists began to question the utilitarian approach going back to Bentham (1789) (see, e.g., Robbins, 1932; Hicks and Allen, 1934). This development can be traced back to Pareto (1906) and eventually led to Arrow's impossibility. Arrow himself addresses the measurability and interpersonal comparison of utility in the second chapter of his monograph:

immeasurability of cardinal utilities

> *The viewpoint will be taken here that interpersonal comparison of utilities has no meaning and, in fact, that there is no meaning relevant to welfare comparisons in the measurability of individual utility. The controversy is well-known and hardly need be recited here. During the entire controversy, the proponents of measurable utility have been unable to produce any proposition of economic behavior which could be explained by their hypothesis and not by those of the indifference-curve theorists. Indeed, the only meaning the concepts of utility can be said to have is their indications of actual behavior, and, if any course of behavior can be explained by a given utility function, it has been amply demonstrated that such a course of behavior can be equally well explained by any other utility function which is a strictly increasing function of the first. If we cannot have measurable utility, in this sense, we cannot have interpersonal comparability of utilities a fortiori.*
>
> *(Arrow, 1951, pp. 9–10)*

Cardinal utilities are particularly troublesome if we engage in their *interpersonal comparison*. Even if there were a way to assign numbers with cardinal meaning to alternatives, these numbers might strongly depend on the agent's subjective perception. Furthermore, since we cannot tell whether voters exaggerate their utilities, utilitarianism with unrestricted utility functions risks degenerating into a contest of "who can name the bigger number." Even when utilities are normalized to a fixed range, comparing utilities

interpersonal comparison of utilities

across individuals remains problematic. Imagine that, instead of forming a line, people who need to use a public restroom are admitted based on their subjective perception of the urgency of their need. This seemingly absurd example illustrates several key issues: incomparability of individual sensations, vulnerability to strategic manipulation, and the lack of mechanisms that reliably measure preference intensities. Similar problems arise when aggregating grades or scores assigned by different instructors. A C+ from one instructor might correspond to a B- from another.

In some cases, however, in particular when the set of alternatives is structured and infinite, it may be more convenient to work with (ordinal) utility functions than with preference relations. As long as no cardinal meaning is associated with these numbers, this is perfectly fine. For example, in consumer theory, it is usually assumed that preferences over elements of Euclidean space can be represented by continuous and quasi-concave utility functions. This narrows down the set of potential preference relations to those that seem reasonable given the structural restriction of the outcome space, allowing for a compact representation of preferences by, say, a polynomial function. When the set of alternatives is the probability simplex, it is common to consider linear

expected utility

(*expected utility*) functions. Von Neumann and Morgenstern (1947) have shown that a preference relation over probability distributions can be represented by a linear utility function iff it satisfies completeness, transitivity, continuity, and independence. It should be kept in mind that, even here, the utility function merely serves as a compact representation of ordinal preferences. As Hylland (1980a) puts it: "From the individual's point of view, only preferences exist; the utility scale is merely the theorist's abstraction." An SCF or SWF may work with the numerical representation of those preferences—as utilitarianism does—as long as these rules are evaluated in terms of ordinal axioms (such as Pareto-optimality or strategyproofness). Because linearity implies that the specification of utilities for degenerate probability distributions suffices (as these values can be extended by expectation), it is particularly tempting to misinterpret the differences in utility as intensities. Von Neumann and Morgenstern (1953, p. 16) themselves cautioned against such cardinal readings, a warning reiterated by Luce and Raiffa (1957, p. 32, "Fallacy 3"), Schoemaker (1982), and Fishburn (1989). Classes of utility functions—such as linear utility functions—are better understood as restricted domains of preferences. Apart from probability distributions, linear utilities arise naturally in settings with divisible goods, bundles, or a tradable numeraire. Whether Arrow's impossibility still holds in the restricted domain of preferences that can be represented by linear utility functions is an interesting question that will be addressed in Chapter 10. In the context of elections, this is mostly relevant when allowing for randomization, as the set of alternatives does not possess the kind of structure that justifies the use of linear utility representations (candidates are neither divisible nor multipliable nor commensurable with money).

Linear utility functions that represent the same preference relation are unique up to positive affine transformations. Two common ways, already mentioned by Sen (1970), to normalize the utility values are *unit range* (where the lowest utility is 0 and the highest is 1) and *unit sum* (where the sum of all utility values is 1). Each of these methods gives rise to unique utility representations of linear utility functions. They are closely connected to two cardinal voting rules: *score voting* (aka range voting) and *cumulative voting*. In score voting, each voter can grade each alternative (say, on a scale from 1 to 10). In cumulative voting, each voter can allocate a fixed number of votes among the alternatives. It is permissible to allocate multiple votes to a single alternative. Under both rules, the alternatives with the highest accumulated scores/votes win. Although these rules face the conceptual challenges associated with cardinal utilities, there have been efforts to understand their properties both theoretically (see, e.g., Pivato, 2014) and experimentally (see, e.g., Baujard et al., 2014, 2018). Score-voting is essentially equivalent to a cardinal SWF called *relative utilitarianism* (see, e.g., Dhillon and Mertens, 1999; Börgers and Choo, 2017; Brandl, 2021).

score voting

cumulative voting

relative utilitarianism

It is sometimes claimed that these methods offer an escape from Arrow's impossibility. Even when only ordinally reading the scores, both rules clearly satisfy Pareto-optimality and non-dictatorship. The accumulated scores also yield a transitive ranking according to which we can pick maximal elements from any feasible set. However, the following example shows that score voting violates IIA. Say there are two voters and $U = \{a, b, c\}$. Voter 1 assigns a score of 10 to a and 0 to both b and c, while Voter 2 assigns a score of 3 to a, 10 to b, and 0 to c. Then, alternative a will be chosen from the feasible set $\{a, b\}$ because its accumulated score is 13, which is higher than b's score of 10. Now, consider a second preference profile, which is identical to the first one, except that Voter 1 increases his score for b to 5. The relative rankings of a and b are identical in both profiles (Voter 1 prefers a while Voter 2 prefers b). However, now alternative b will be chosen from $\{a, b\}$ because its score of 15 exceeds a's score of 13. Voter 1's changed preference between b and c affected the choice between a and b. An obvious way around this problem is to define a cardinal version of IIA, which merely demands that choices from a feasible set should only depend on the *scores* of alternatives within the feasible set. However, such an interpretation of IIA runs into the problems of immeasurability of cardinal utility and interpersonal comparison of utilities.

4.8 Key Takeaways

Arrow's Theorem

- Rationalizable social choice functions suffer from indecisiveness or high concentration of power.
- Common escape routes include restricting the domain of preferences and abandoning or weakening rationalizability.
- Cardinal utilities raise two fundamental problems: measurability and interpersonal comparisons.

4.9 Further Reading

An early formulation of Arrow's impossibility theorem appeared in a research paper (Arrow, 1950) as well as Arrow's PhD thesis, which was also published as a monograph (Arrow, 1951). A revised second edition, which presents the result as we know it today and corrects a mistake in the domain assumption discovered by Blau (1957), was published twelve years later (Arrow, 1963). The theorem, as well as other contributions by Arrow, are extensively discussed by Feiwel (1987a,b) and Maskin and Sen (2014).

Tang and Lin (2009) gave an induction-based proof of Arrow's theorem for $\mathcal{D} = \mathcal{S}^N$ where they manually proved induction steps for m and n and let a SAT solver—a computer program that decides the satisfiability of Boolean formulas in propositional logic—"prove" the base case ($m = 3$, $n = 2$). Exhaustive search is not an option, as there are about 10^{28} SWFs for two voters and three alternatives. Similar techniques have recently been used to prove new theorems in social choice theory.

Sen (1969) initiated the study of Pareto-optimal SWFs that satisfy IIA and weakenings of transitive rationalizability. Comprehensive overviews of the extensive literature on Arrovian impossibility theorems are given by Sen (1977, 1986), Schwartz (1986), Fishburn (1987), Austen-Smith and Banks (1999), and Campbell and Kelly (2002). Gibbard's oligarchy theorem, shown in the table on page 48, and variants thereof were established independently by Guha (1972), Mas-Colell and Sonnenschein (1972), and Schwartz (1972). Particularly noteworthy are results about acyclic collective preference relations (e.g., Mas-Colell and Sonnenschein, 1972; Brown, 1975; Blau and Deb, 1977; Blair and Pollak, 1982, 1983; Banks, 1995a; Duggan, 2016) because acyclicity is necessary and sufficient for the existence of maximal elements when there is a finite number of alternatives. Duggan (2016) strengthens the impossibility by Mas-Colell and Sonnenschein (1972) by weakening positive responsiveness$_2$ to a condition that only requires that ties are broken when roughly one third of the voters change their preference. Nakamura (1975) and Moulin (1985b) characterize rationalizable SCFs that are neutral and anonymous, respectively. Sen (1995) aptly summarizes this stream of research by writing that "the arbitrariness of power of which Arrow's case of dictatorship is an extreme example, lingers in one form or another even when transitivity is dropped, so long as *some* regularity is demanded (such as the absence of cycles)."

For more background and discussion on cardinal utilities and utilitarianism, the reader is referred to Plott (1976), Mandler (1999), and Fleurbaey et al. (2008). Sen (1970) has initiated the study of so-called *social welfare functionals (SWFLs)*, which map a profile of cardinal utilities to a transitive and complete collective preference relation (see also d'Aspremont and Gevers, 2002). The corresponding notion of IIA is the one mentioned at the end of Section 4.7 and takes into account the absolute values of utilities (rather than only ordinal comparisons between these values). Hence, utilitarianism (as well as relative utilitarianism and score voting) satisfies this notion of IIA. To account for the fact that linear utilities are invariant under positive affine transformations, Sen introduced the axiom of *cardinality and non-comparability*, which prescribes that collective preferences returned by the SWFL are invariant under positive affine transformations of the individual utility functions. However, this assumption effectively transforms the

social welfare functionals (SWFLs)

problem into one of ordinal preference aggregation, as the utility values assigned to two different alternatives in two different utility profiles can be made identical across profiles by applying a positive affine transformation. Hence, the cardinal variant of IIA implies standard IIA, and Arrow's theorem holds (Sen, 1970, Theorem 8*2).

Dhillon and Mertens (1999) introduce a weakening of IIA called *independence of redundant alternatives*, which only considers feasible sets for which every infeasible outcome is *unanimously* indifferent to some feasible outcome, and show that relative utilitarianism satisfies this condition.

In cumulative voting, the scores assigned by the voters must add up to a fixed value. A variant of this method, called *quadratic voting*, is obtained by demanding that the sum of squares of each score add up to a fixed value (Lalley and Weyl, 2018).

quadratic voting

Majority judgment is another voting rule based on cardinal preferences, which returns an alternative with the highest median score (Balinski and Laraki, 2011). It represents a radical departure from the traditional social choice setting and already diverges from majority rule on two alternatives.

majority judgment

The properties of these so-called "evaluative" voting rules (which also include score voting and cumulative voting) have, for example, been studied by Laslier (2012, 2019), Baujard et al. (2014, 2018, 2025), Pivato (2014), Macé (2018), and El Ouafdi et al. (2020).

4.10 Exercises

4.1 *Monotonicity and Pareto-optimality*

Let f be an SCF that satisfies non-imposition$_2$ and IIA$_2$.
Show that when f is monotonic$_2$, it is also Pareto-optimal$_2$.

4.2 *Independence of Arrow's axioms*

Show that the axioms used in Theorem 4.1—transitive rationalizability, IIA, Pareto-optimality, and non-dictatorship—are independent from each other, i.e., none of them is implied by the other three.

Hint: For each of these axioms, construct an SCF that does not satisfy it but satisfies all of the remaining axioms.

4.3 *The lonely dictator*

Let $n = 1$.

(a) Show that Pareto-optimality implies dictatorship.

(b) Show that every transitively rationalizable SCF that satisfies IIA$_2$ and non-imposition is dictatorial or reversely dictatorial (Wilson's theorem).

4.4 *The majority dictator*

Fix $i \in N$ and define $f(A, P) = \text{Max}(A, \succ_M \cap \succ_i)$ for all $A \in \mathcal{F}$ and $P \in \mathcal{D}$.

(a) Show that f satisfies IIA, rationalizability, and Pareto-optimality.

(b) Is f dictatorial, quasi-dictatorial, or weakly dictatorial?

4.5 *UN Security Council* *(Austen-Smith and Banks, 1999, pp. 41–42)*

Prior to 1965, the United Nations Security Council consisted of five permanent members (China, France, Great Britain, Soviet Union, United States) and six other members. The passage of a motion required seven "yes" votes, as well as the concurring votes of all five permanent members. If we interpret a "yes" vote as strictly preferring adoption of a motion, interpret a concurring vote as weakly preferring the motion, and let individuals $1, \dots, 5$ denote the permanent members, the social choice function of the Security Council (restricted to the case of two alternatives $a, b \in U$) is then

$$f(\{a, b\}, P) = \begin{cases} \{a\} & \text{if } |\{i\colon a \succ_i b\}| \geqslant 7 \text{ and } a \succsim_i b \text{ for all } i \in \{1, \dots, 5\}, \\ \{b\} & \text{if } |\{i\colon b \succ_i a\}| \geqslant 7 \text{ and } b \succsim_i a \text{ for all } i \in \{1, \dots, 5\}, \\ \{a, b\} & \text{otherwise.} \end{cases}$$

(a) Which groups are decisive? Are there any weak dictators? Are there any quasi-dictators? Is there an oligarchy or a collegium?

(b) Show that $\succsim_f$ is acyclic.

Hint: Focus on the preference relation of a permanent member. Even though f is only defined for feasible sets of size 2, U can consist of more than two alternatives.

☆ **4.6** *Extremal alternatives (Geanakoplos, 2005)*

Let $m \geqslant 3$, $n \geqslant 2$, and g denote an SWF that satisfies IIA and Pareto-optimality.

An alternative x is *extremal* for $\succsim \subseteq U \times U$ if $x \succ y$ for all $y \in U \setminus \{x\}$ or $y \succ x$ for all $y \in U \setminus \{x\}$. Moreover, an alternative x is *extremal* for $P \in \mathcal{D}$ if it is extremal for all $\succsim_i$ with $i \in N$. A voter i is *extremally pivotal* for $x \in U$ if there are $P, P' \in \mathcal{D}$ such that *(i)* x is extremal for P, *(ii)* $P_{-i} = P'_{-i}$, *(iii)* $y \succ_i x$ and $y \succ x$ for all $y \in U \setminus \{x\}$, *(iv)* $x \succ'_i y$ and $x \succ' y$ for all $y \in U \setminus \{x\}$ (where $\succsim = g(P)$ and $\succsim' = g(P')$).

Prove the following claims without using Theorem 4.2 or Lemma 4.1.

(a) If $x \in U$ is extremal for $P \in \mathcal{D}$, it is also extremal for $\succsim = g(P)$.

(b) For every $x \in U$, there is an extremally pivotal voter $i \in N$.

(c) A voter i who is extremally pivotal for $x \in U$ is decisive for every $y, z \in U \setminus \{x\}$.

(d) There is a voter $i \in N$ who is extremally pivotal for all alternatives $x \in U$.

4.7 *Cumulative voting*

Show that cumulative voting violates IIA.

Part II

Social Choice Functions

Individual preferences are determined not by turning a roulette wheel over all possible alternatives, but by certain specific social, economic, political, and cultural forces. This may easily produce some patterns in the set of individual preferences.

Amartya Sen, 1970

5 Restricted Domains

Learning Outcomes

- Which subdomains of preferences allow for rationalizable social choice functions?
- How can one identify weak Condorcet winners in these subdomains?
- Can it be efficiently verified whether a preference profile belongs to a subdomain?

An implicit assumption in Arrow's impossibility theorem is that all preference profiles are admissible ($\mathcal{D} = \mathcal{W}^N$). Restricting the domain of admissible preference profiles $\mathcal{D}$ can allow for more positive results.

We will distinguish between two types of domains: Cartesian domains and non-Cartesian domains. A domain $\mathcal{D}$ is *Cartesian* if $\mathcal{D} = \mathcal{R}^N$ for some $\mathcal{R} \subseteq \mathcal{W}$.[1] Two examples of Cartesian domains are $\mathcal{W}^N$ and $\mathcal{S}^N$. The domain where two voters cannot strictly disagree, i.e., $x \succ_1 y$ implies $x \succsim_2 y$ for all $x, y \in U$, is an example of a non-Cartesian domain. Cartesian domain

Verifying that the proof of Arrow's theorem never uses indifferences shows that the theorem also holds when $\mathcal{D} = \mathcal{S}^N$. Note, however, that in contrast to many other statements, Arrow's theorem is *not* stronger for strict preferences than for weak preferences. The reason for this is that non-dictatorship is not inherited to subdomains: an SWF that is non-dictatorial for $\mathcal{W}^N$ may be dictatorial in the subdomain $\mathcal{S}^N$. For a concrete example, consider the SWF g that returns a fixed (say, lexicographic) ranking of all alternatives if at least one preference relation is not strict. For all remaining profiles, Voter 1 is a dictator. There are profiles in $\mathcal{W}^N$ where g is not dictatorial, but within $\mathcal{S}^N$, g is dictatorial. When $\mathcal{D} = \mathcal{S}^N$, dictatorial SWFs are projections, i.e., the collective preference relation is always completely identical to the dictator's preference relation.

5.1 Bypassing Impossibilities

When accepting that pairwise choices should be based on majority rule (for example, because of May's theorem), there is only a single candidate for a rationalizable SCF: the SCF that returns the maximal elements according to the pairwise majority relation: Cond.

[1] This definition of Cartesian domains includes that the domain is symmetric in the sense that $\mathcal{R}$ is identical for all voters.

We know from Theorem 3.4 that $\mathsf{Cond}(A, P)$ may be empty for some $A \in \mathcal{F}$ and $P \in \mathcal{S}^N$ and that Cond is not a well-defined SCF in $\mathcal{S}^N$. In fact, Cond is a well-defined SCF iff $\succsim_M$ is acyclic for all $P \in \mathcal{D}$ (see Lemma 2.1). Moreover, if $\succsim_M$ is transitive for all profiles $P \in \mathcal{D}$, Cond satisfies all conditions of Arrow's theorem, including transitive rationalizability. To this end, define

$$\mathcal{D}_{\mathsf{TRANS}} = \{P \in \mathcal{W}^N : \succsim_M \text{ is transitive}\}.$$

This domain is *not* Cartesian (see Exercise 5.2). Clearly, when $\mathcal{D} \subseteq \mathcal{D}_{\mathsf{TRANS}}$, Cond is the *only* SCF rationalized by $\succsim_M$. It is thus possible to bypass Arrow's theorem by shrinking the domain $\mathcal{W}^N$ to $\mathcal{D}_{\mathsf{TRANS}}$.

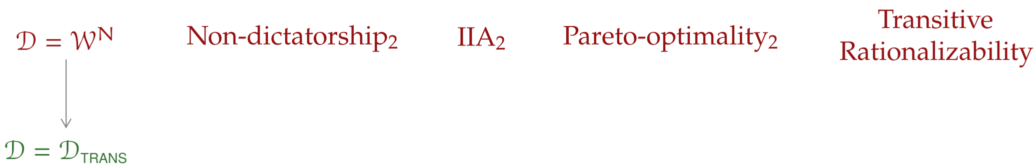

When using May's axioms to enforce that pairwise choices are based on majority rule, this immediately gives a simple axiomatic characterization: Cond is the *only* SCF satisfying anonymity, neutrality, positive responsiveness$_2$, and transitive rationalizability when $\mathcal{D} = \mathcal{D}_{\mathsf{TRANS}}$.[2] As we will see below, this holds even when relaxing positive responsiveness$_2$ to Pareto-optimality$_2$. By leveraging parts of the proof of Arrow's theorem, it can be shown that *only* refinements of Cond satisfy Arrow's axioms when $\mathcal{D} = \mathcal{D}_{\mathsf{TRANS}}$. In the absence of majority ties or when assuming anonymity and neutrality, Arrow's theorem can even be turned into a complete axiomatic characterization of Cond.

Theorem 5.1 $\mathcal{D} = \mathcal{D}_{\mathsf{TRANS}}, m \geqslant 3, n \geqslant 2$

Let f be an SCF that satisfies Pareto-optimality$_2$, IIA$_2$, non-dictatorship$_2$, and transitive rationalizability. Then, the following statements hold:

(i) $f \subseteq \mathsf{Cond}$,

(ii) $f = \mathsf{Cond}$ if f is anonymous and neutral, and

(iii) $f = \mathsf{Cond}$ if $\mathcal{D} = \{P \in \mathcal{D}_{\mathsf{TRANS}} : n_{xy} \neq n_{yx} \text{ for all } x, y \in U\}$.

Proof. Since every rationalizable SCF f is completely determined by its base relation, it suffices to argue about the base relation $\succsim$ induced by f for a given $P \in \mathcal{D}_{\mathsf{TRANS}}$.
To prove *(i)*, we extend the notion of weak semi-decisiveness to groups of voters by defining that G is *weakly semi-decisive for* a *against* b (denoted by $a\ \tilde{W}_G\ b$) if for all $P \in \mathcal{D}$,

$$a \succ_i b \text{ for all } i \in G \text{ and } b \succ_j a \text{ for all } j \notin G \quad\Rightarrow\quad a \succsim b.$$

[2] This also applies to all subdomains of $\mathcal{D}_{\mathsf{TRANS}}$ in which May's theorem holds, including all subdomains that permit arbitrary weak or strict preferences between pairs of alternatives, such as $\mathcal{W}^N_{\mathsf{DICH}}$, $\mathcal{S}^N_{\mathsf{SP}(\geqslant)}$, and $\mathcal{S}^N_{\mathsf{SC}(\geqslant)}$, which are introduced later in this chapter.

First, observe that Lemma 4.1 only requires $a\ \tilde{W}_G\ b$ instead of $a\ \tilde{D}_G\ b$ and also holds when $\mathcal{D} = \mathcal{D}_{\text{TRANS}}$ and $|G| < n/2$. This follows because all the profiles that appear in the proof of Lemma 4.1 are contained in $\mathcal{D}_{\text{TRANS}}$: for the profiles P' in (1), we have $b \succ_M a$ and $b \succ_M x$, and for the profiles P' in (2), we have $b \succ'_M a$ and $x \succ'_M a$. Hence, $\succsim'_M|_{\{a,b,x\}}$ is transitive. Any additional alternatives can be arranged in an arbitrary fixed order such that $P' \in \mathcal{D}_{\text{TRANS}}$. Hence, weak semi-decisiveness of $G \subseteq N$ with $|G| < n/2$ for one pair of alternatives implies decisiveness of G for all pairs of alternatives.

We now show that $x \succ_M y$ implies $x \succ y$ for all $P \in \mathcal{D}_{\text{TRANS}}$ and $x, y \in U$, which implies that f has to be a refinement of Cond. Assume for contradiction that there is a profile $P \in \mathcal{D}_{\text{TRANS}}$ and $x, y \in U$ such that $x \succ_M y$ but $y \succsim x$. By IIA, $G = \{i \in N : y \succ_i x\}$ with $|G| < n/2$ is weakly semi-decisive for y against x, and our variant of the field expansion lemma implies that G is decisive. The rest of the proof follows from Lemma 4.2 for the special case where $|G| < n/2$. The profile P used in the proof of Lemma 4.2 is contained in $\mathcal{D}_{\text{TRANS}}$ because $|N \setminus G| > n/2$ and a majority of agents have the same preferences over a, b, c. The proof thus yields that f is dictatorial, contradicting the assumption of non-dictatorship$_2$.

As to *(ii)*, observe that for any $A \in \mathcal{F}$, $P \in \mathcal{D}$, and $x, y \in \mathsf{Cond}(A, P)$, we have $x \sim_M y$. By anonymity and neutrality, $x \sim y$, which implies that $f = \mathsf{Cond}$.

As to *(iii)*, note that, in the absence of majority ties, Cond is resolute and thus does not admit a strict refinement. It can also be verified that the profiles appearing in the proof of *(i)* do not have majority ties. □

Clearly, all refinements of Cond satisfy Pareto-optimality and non-dictatorship. Not all refinements of Cond satisfy IIA$_2$ and transitive rationalizability, for example, when breaking ties using the broad or narrow Borda rule. Strict refinements that do satisfy all axioms must violate anonymity or neutrality, such as when using dictatorial or lexicographic tie-breaking. As the proof of Arrow's theorem, the proof of Theorem 5.1 also works if $\mathcal{D} = \mathcal{S}^N$. When weakening transitive rationalizability to quasi-transitive rationalizability, the proof of Theorem 4.3 can also be adapted to characterize Cond.

In the following sections, we will study natural Cartesian subdomains of $\mathcal{D}_{\text{TRANS}}$, i.e., domains that arise from natural restrictions on preference relations that lead to transitive (or quasi-transitive) majority relations.

5.2 Dichotomous Preferences

A very simple subdomain of $\mathcal{W}^N$ is the one where each preference relation may admit at most two indifference classes, i.e., voters can only approve or disapprove each alternative. Formally, the domain of *dichotomous preferences* is defined as

dichotomous preferences

$$\mathcal{W}^N_{\text{DICH}} = \{\succsim \in \mathcal{W} : \forall x, y, z \in U : x \succ y \Rightarrow z \sim x \vee z \sim y\}^N.$$

An example of a dichotomous preference profile is shown in the margin. Alternatives in the same indifference class are printed next to each other. The first voter approves a and b while the second one approves b and c.

1	1
$a\,b$	$b\,c$
c	a

Dichotomous preference profiles cannot induce majority cycles. They even guarantee that the majority relation is transitive. This is due to the fact that $x \succsim_M y$ iff x is approved at least as often as y.

Theorem 5.2 (Inada, 1964)

$\succsim_M$ is transitive for all dichotomous profiles, i.e., $\mathcal{W}^N_{\text{DICH}} \subseteq \mathcal{D}_{\text{TRANS}}$.

Proof. For a profile $P \in \mathcal{W}^N_{\text{DICH}}$ and $x \in U$, let $n(x) = |\{i \in N : x \in \text{Max}(U, \succsim_i)\}|$ be the number of voters who top-rank x. We then have for every $P \in \mathcal{W}^N_{\text{DICH}}$ and $x, y \in U$,

$$x \succsim_M y \quad \overset{\text{Def. of } \succsim_M}{\Leftrightarrow} \quad n_{xy} \geqslant n_{yx} \quad \overset{P \in \mathcal{W}^N_{\text{DICH}}}{\Leftrightarrow} \quad n(x) \geqslant n(y)$$

The last equivalence holds because for every voter $i \in N$,

$$x \succ_i y \quad \Leftrightarrow \quad x \in \text{Max}(U, \succsim_i) \not\ni y.$$

Now take $x, y, z \in U$ such that $x \succsim_M y$ and $y \succsim_M z$. According to the equivalence above, this means that $n(x) \geqslant n(y)$ and $n(y) \geqslant n(z)$, and hence that $n(x) \geqslant n(z)$, which is equivalent to $x \succsim_M z$. □

approval voting

Within domain $\mathcal{W}^N_{\text{DICH}}$, Cond is known as *approval voting*: each voter approves some subset of alternatives, and the alternatives with the highest numbers of approvals are selected. In other words, approval voting is the SCF that is rationalized by the majority relation for dichotomous preferences. In Section 8.2, we will see that approval voting cannot be manipulated when preferences are dichotomous.

Of course, the restriction to dichotomous preferences is quite severe. Nevertheless, approval voting can be (and is!) used, even when preferences are not dichotomous. A crucial issue for voters is which alternatives they approve, given their non-dichotomous preference relations. One would expect that voters approve all alternatives that meet some individual threshold for approval. Both theoretical and empirical studies have shown that approval voting is much more likely to return a Condorcet winner than plurality. Approval voting is used by a number of academic organizations, e.g., the Society of Social Choice and Welfare, the Game Theory Society, the Mathematical Association of America, the Institute for Operations Research and the Management Sciences, and the American Statistical Association. Attempts to adopt approval voting for local elections in the US have been met with limited success.

5.3 Single-Peaked Preferences

In this section, we consider domains that are based on a linear order of the alternatives. As an introductory example, consider preferences over the maximal income tax rate. Say we have five alternatives: 20%, 30%, 40%, 50%, and 60% and the following preference relations:

$\succsim_1$	$\succsim_2$	$\succsim_3$
40 %	50 %	20 %
30 %	60 %	30 %
50 %	40 %	40 %
20 %	30 %	50 %
60 %	20 %	60 %

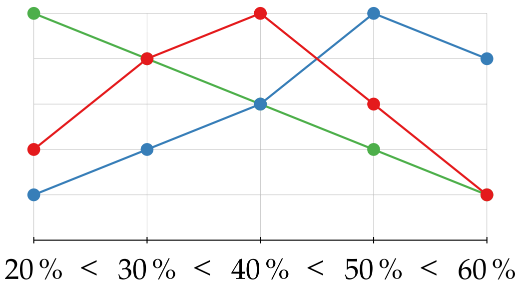

Plotting these preferences in a graph where the alternatives are ordered according to the linear order on the x-axis and the individual ranking on the y-axis results in the graph on the right-hand side. The colored lines highlight that each voter has a unique top choice—his *peak*—with declining affinity as one moves away from the peak.[3] Preferences such as 20 % ≻ 60 % ≻ 40 % would violate this condition, and, indeed, such preferences over tax rates appear to be unreasonable. Formally, the domain of *single-peaked preferences* with respect to a linear order $\geqslant \in \mathcal{S}$ is

single-peaked preferences

$$\mathcal{S}^N_{\mathrm{SP}(\geqslant)} = \{\succsim \in \mathcal{S} \colon \forall x,y,z \in U \colon (x > y > z) \vee (z > y > x) \Rightarrow (x \succ y \Rightarrow y \succ z)\}^N.$$

Note that a preference relation is not completely specified by its peak. The preference relation 40 % ≻ 50 % ≻ 30 % ≻ 20 % ≻ 60 % has the same peak as that of Voter 1 but differs in the ranking of 30 % and 50 %.

The original motivation for single-peaked preferences stems from political science, where candidates can sometimes be located on a one-dimensional left-right spectrum, ranging from very liberal to very conservative. Other applications where single-peaked preferences seem natural include the location of a desirable facility on a road, setting the temperature for a joint thermostat, deciding on the size of a public park, allocating budget to a library, and aggregating school grades.

For fixed $\geqslant \in \mathcal{S}$, $\mathcal{S}^N_{\mathrm{SP}(\geqslant)}$ is a Cartesian domain. For example, when $a > b > c$, there are only the four possible strict preference relations shown in the margin. Note, however, that the domain of all profiles that are single-peaked according to *some* linear order $\bigcup_{\geqslant \in \mathcal{S}} \mathcal{S}^N_{\mathrm{SP}(\geqslant)}$ is *not* Cartesian (see Exercise 5.2).

a	b	b	c
b	a	c	b
c	c	a	a

The majority relation of any single-peaked preference profile is quasi-transitive. Let $\mathcal{D}_{\mathrm{QTRANS}} = \{P \in \mathcal{W}^N : \succsim_M \text{ is quasi-transitive}\}$.

Theorem 5.3 (Black, 1948)

$\succsim_M$ is quasi-transitive for all single-peaked profiles, i.e. $\mathcal{S}^N_{\mathrm{SP}(\geqslant)} \subseteq \mathcal{D}_{\mathrm{QTRANS}}$ for all $\geqslant \in \mathcal{S}$.

Proof. Let $x, y, z \in U$. We need to show that $x \succ_M y$ and $y \succ_M z$ imply that $x \succ_M z$. There are $3! = 6$ possible linear orders on three alternatives. We will distinguish three cases,

[3]Preferences are still ordinal rather than cardinal. The y-value of an alternative only represents its rank in the preference relation.

each of which covers two linear orders of x, y, and z.

Case 1. $x > y > z$ (or $z > y > x$):

$$N_{xy} \overset{\text{SP}}{\subseteq} N_{yz} \overset{\text{Trans.}}{\Rightarrow} N_{xy} \subseteq N_{xz} \Rightarrow n_{xz} \geqslant n_{xy} > \frac{n}{2}$$

Case 2. $z > x > y$ (or $y > x > z$):

$$N_{zx} \overset{\text{SP}}{\subseteq} N_{xy} \overset{\text{Trans.}}{\Rightarrow} N_{zx} \subseteq N_{zy} \Rightarrow \frac{n}{2} > n_{zy} \geqslant n_{zx} \Rightarrow n_{xz} > \frac{n}{2}$$

Case 3. $y > z > x$ (or $x > z > y$):

$$N_{yz} \overset{\text{SP}}{\subseteq} N_{zx} \overset{\text{Trans.}}{\Rightarrow} N_{yz} \subseteq N_{yx} \Rightarrow \frac{n}{2} < n_{yz} \leqslant n_{yx} \Rightarrow n_{xy} < \frac{n}{2} \quad \lightning \qquad \square$$

single-caved

An analogous theorem holds for *single-caved* (sometimes also called single-troughed or single-dipped) preferences with respect to a linear order $\geqslant \in \mathcal{S}$:

$$\mathcal{S}^N_{\text{SC}(\geqslant)} = \{\succsim \in \mathcal{S} \colon \forall x, y, z \in U \colon (x > y > z) \vee (z > y > x) \Rightarrow (y \succ x \Rightarrow z \succ y)\}^N$$

These preferences can, for example, arise when placing an undesirable facility, such as a garbage dump, on a one-dimensional road. A profile is single-caved iff its reverse profile (consisting of all reversed preference relations) is single-peaked.

Quasi-transitivity of $\succsim_M$ implies the existence of weak Condorcet winners. It turns out that weak Condorcet winners can be identified by considering only the peaks of the voters. To this end, let p_i be the peak of voter $i \in N$ ($\text{Max}(A, \succsim_i)$). Further, let $x, y \in A$ such that $x < y$ and there is no $z \in A$ with $x < z < y$. Then,

$$x \succsim_M y \text{ iff } |\{i \in N \colon p_i \leqslant x\}| \geqslant |\{i \in N \colon y \leqslant p_i\}|.$$

Hence, x is a weak Condorcet winner iff at least half of the peaks lie at or to the left of x and at least half of the peaks lie at or to the right of x, i.e.,

$$|\{i \in N \colon p_i \leqslant x\}| \geqslant n/2 \quad \text{and} \quad |\{i \in N \colon p_i \geqslant x\}| \geqslant n/2.$$

median voter

When n is odd, there are no majority ties, and every single-peaked profile admits a unique Condorcet winner. It follows from our previous observations that the Condorcet winner is the peak of a *median voter*: if we sort voters by $\geqslant$ according to their peaks, breaking ties arbitrarily, the $((n+1)/2)$th voter is a median voter.

2	2	3
a	b	c
b	a	b
c	c	a

Despite the fact that the peaks of the voters suffice to compute the Condorcet winner, plurality (which famously only takes into account peaks) can select an alternative different from the Condorcet winner, even when preferences are single-peaked. To see this, consider the preference profile in the margin with linear order $a > b > c$. The voters in the middle column are the median voters. Hence, alternative b is the Condorcet winner. However, the plurality winner is alternative c.

median voting

In domain $\mathcal{S}^N_{\text{SP}(\geqslant)}$, Cond is known as *median voting*. When n is even, there can be median voters with two different peaks, and $\mathsf{Cond}(A, P)$ consists of all alternatives in between the two median alternatives. When n is even, but m is odd, we can add a "phantom voter" whose peak is invariably the alternative in the middle and then return the Condorcet winner. It is thus possible to have a fair and resolute SCF for single-peaked preferences whenever m or n is odd.

In Section 8.2, we will see that median voting cannot be manipulated when preferences are single-peaked.

Unfortunately, in the context of political parties or candidates, single-peakedness is a restrictive—and often unrealistic—assumption. Sophisticated political models describing the positions of parties require more than one policy dimension, and the majority relation may be cyclic again (see also Section 5.7).

With single-caved preferences, Cond takes an even simpler form than with single-peaked preferences. The outcome only depends on the top choices of the voters, which can only be one of the two extreme alternatives with respect to the linear order. Hence, identifying weak Condorcet winners just boils down to a majority vote between the two extreme alternatives.

5.4 Checking Single-Peakedness

Can we efficiently check whether a given preference profile $P \in \mathcal{D}$ is single-peaked? When $\geqslant \in \mathcal{S}$ is given, checking whether $P \in \mathcal{S}^N_{\mathsf{SP}(\geqslant)}$ is straightforward. It is much less obvious whether one can efficiently check whether a given preference profile $P \in \mathcal{D}$ is single-peaked according to *some* linear order: is there $\geqslant \in \mathcal{S}$ such that $P \in \mathcal{S}^N_{\mathsf{SP}(\geqslant)}$?

1	1	1
a	b	c
b	c	a
c	a	b

In light of Theorem 5.3, we can discard all profiles P for which $\succ_M$ is not transitive, such as the classic Condorcet paradox profile, shown in the margin. However, the domains $\mathcal{D}_{\mathsf{QTRANS}}$ and $\bigcup_{\geqslant \in \mathcal{S}} \mathcal{S}^N_{\mathsf{SP}(\geqslant)}$ do not coincide. There are profiles that admit a quasi-transitive majority relation but that are not single-peaked. For an example, consider the profile given in the margin. How can we verify that this profile is not single-peaked? Of course, we can enumerate all linear orders and check single-peakedness for each order. However, this approach is not feasible in polynomial time because the number of orders is $m!$. It turns out that there is a very simple necessary condition for single-peakedness that can be easily checked. Any last-ranked alternative has to lie at the leftmost or rightmost position according to the linear order. Hence, there can be at most two different last-ranked alternatives. This insight can be leveraged to devise an algorithm that arranges alternatives according to the linear order $\geqslant$ from the outside to the inside by repeatedly considering pairs of bottom-ranked alternatives. The naive algorithm that examines all possible arrangements of pairs of last-ranked alternatives must check $2^{\frac{m}{2}}$ linear orders. While this represents a notable improvement over $m!$, the number still grows exponentially with m. However, the following observation enables us to fully determine a linear order without the need to explore numerous possibilities.

1	1	1
a	a	b
b	c	c
c	d	d
d	b	a

Let l and r be the current left innermost and right innermost alternatives, respectively. If x and y still need to be arranged and there is a voter i such that $l \succ_i x \succ_i r$ and $y \succ_i x$, then x has to be placed next to r. For suppose otherwise, x will be placed next to l, which implies that voter i's peak is to the left of x because $l \succ_i x$. However, alternative y with $y \succ_i x$ will then be placed to the right of x, contradicting the single-peakedness of $\succ_i$. By symmetry, the same rule holds when swapping l and r, i.e., if $r \succ_i x \succ_i l$, then x has to be placed next to l. We thus obtain the following linear-time algorithm for checking whether P is single-peaked. Let $z_l, z_r \notin U$ be two markers for the leftmost and rightmost positions on the linear spectrum.

```
procedure SINGLE-PEAKED ORDER(P)
    Put z_l in the leftmost slot of ⩾
    Put z_r in the rightmost slot of ⩾
    A ← U
    while |A| ⩾ 2 do
        l ← current left innermost alternative of ⩾
        r ← current right innermost alternative of ⩾
        B  ← {x ∈ A : ∃i ∈ N such that ∀y ∈ A : y ≿_i x}
        B_l ← {x ∈ B : ∃i ∈ N such that r ≻_i x ≻_i l and ∃y ∈ A : y ≻_i x}
        B_r ← {x ∈ B : ∃i ∈ N such that l ≻_i x ≻_i r and ∃y ∈ A : y ≻_i x}
        if |B| ⩽ 2 ∧ |B_l| ⩽ 1 ∧ |B_r| ⩽ 1 ∧ B_l ∩ B_r = ∅ then
            Put the alternative in B_l (if any) next to l
            Put the alternative in B_r (if any) next to r
            Put the alternatives in B \ (B_l ∪ B_r) (if any) either next to l or next to r
        else
            return ∅                                        ▷ P is not single-peaked.
        A ← A \ B
    if |A| = 1 then put x ∈ A in between l and r
    return ⩾                                                ▷ P is single-peaked.
```

1	1	1
b	c	d
d	d	c
c	b	a
a	a	b

For a correctness proof of this algorithm, see the references given in Section 5.7. Let us illustrate the algorithm by running it on the profile in the margin. We start by placing the markers z_l and z_r and initializing $A = \{a, b, c, d\}$, $l = z_l$, and $r = z_r$. There are two last-ranked alternatives, $B = \{a, b\}$. In the first round, B_l and B_r are always empty because $z_l, z_r \notin U$. We can thus place a and b arbitrarily, for example, as below.

$$z_l \quad a \quad \_ \quad \_ \quad b \quad z_r$$

Both alternatives in $B = \{a, b\}$ have been placed and $A = A \setminus \{a, b\}$. In the next round, $l = a$, $r = b$, and the two last-ranked alternatives are $B = \{c, d\}$. Since $r \succ_1 c \succ_1 l$ and $d \succ_1 c$, we have $c \in B_l$. No such restrictions apply to d and $B_r = \emptyset$. We thus obtain the following linear order $\geqslant$.

$$z_l \quad a \quad c \quad d \quad b \quad z_r$$

By arranging *highest*-ranked alternatives from the outside to the inside, we obtain an analogous algorithm for single-caved preferences.

In the spirit of Kuratowski's famous characterization of planar graphs in terms of excluded minors, it is possible to characterize the domain $\bigcup_{\geqslant \in \mathcal{S}} \mathcal{S}^N_{\mathrm{SP}(\geqslant)}$ by excluding the five minimal profiles that fail to be single-peaked.

Theorem 5.4 (Ballester and Haeringer, 2011)

A profile $P \in \mathcal{D}$ is contained in $\bigcup_{\geqslant \in \mathcal{S}} \mathcal{S}^N_{\mathrm{SP}(\geqslant)}$ iff there are no $i, j, k \in N$ and $x, y, z, w \in U$ such that $P|_{\{x,y,z\}}$ or $P|_{\{x,y,z,w\}}$ restricted to voters i, j, and k looks as follows.

$\succsim_i$	$\succsim_j$	$\succsim_k$
x	y	z
y	z	x
z	x	y

$\succsim_i$	$\succsim_j$	$\succsim_k$
x	z	x
y	y	z
z	x	y

$\succsim_i$	$\succsim_j$
w	w
x	z
y	y
z	x

$\succsim_i$	$\succsim_j$
w	z
x	w
y	y
z	x

$\succsim_i$	$\succsim_j$
x	z
w	w
y	y
z	x

Note that the first two subprofiles represent all cases with more than two last-ranked alternatives, while the other three profiles represent all cases where an alternative (y) simultaneously needs to be placed to the left and to the right.

This observation provides an alternative polynomial-time algorithm for recognizing single-peaked preferences: one can simply go over all subprofiles of a given size, and check that none of them is isomorphic to one of the five excluded profiles. Theorem 5.4 can be adapted to single-caved preferences by inverting each excluded subprofile.

5.5 Value-Restricted Preferences

We have seen several examples of Cartesian subdomains of $\mathcal{D}_{\mathrm{QTRANS}}$, such as all domains of single-peaked preferences $\mathcal{S}^N_{\mathrm{SP}(\geqslant)}$ and all domains of single-caved preferences $\mathcal{S}^N_{\mathrm{SC}(\geqslant)}$.[4]

1	1	1
a	c	d
b	a	a
c	b	b
d	d	c

Which other Cartesian subdomains of $\mathcal{D}_{\mathrm{QTRANS}}$ are there? To address this question, consider the preference profile in the margin. This profile is neither single-peaked nor single-caved. It is not single-caved because there are three different top-ranked alternatives. To see that it is not single-peaked, we can either run the algorithm from Section 5.4 (which rejects the profile in the second round because alternative b needs to be placed both to the left and to the right) or observe that the profile restricted to the last two voters is isomorphic to the last excluded profile in Theorem 5.4. Still, it can be verified that $\succsim_M$ is transitive. This is because the profile satisfies single-peakedness *for every triple* of alternatives. Here, we can choose a different linear order $\geqslant$ for every triple. This train of thought leads to new Cartesian subdomains of $\mathcal{D}_{\mathrm{QTRANS}}$: fix a linear order $\geqslant_{xyz}$ for every triple of distinct alternatives $x, y, z \in A$ and demand that every profile in this domain consists of preference relations $\succsim$ that are triple single-peaked, i.e., for every triple $x, y, z \in A$, $\succsim|_{\{x,y,z\}} \in \mathcal{D}_{\mathrm{SP}(\geqslant_{xyz})}$. These domains admit quasi-transitive majority relations because of the proof of Theorem 5.3. Quasi-transitivity is a property of triples, and the proof of Theorem 5.3 only considers the linear order restricted to three arbitrary alternatives. It is possible to define triple-single-caved subdomains of $\mathcal{D}_{\mathrm{QTRANS}}$ analogously. However, triple single-peakedness and triple single-cavedness still do not capture all Cartesian subdomains of $\mathcal{D}_{\mathrm{QTRANS}} \cap \mathcal{S}^N$.

[4]Note that $\mathcal{D}_{\mathrm{TRANS}} \cap \mathcal{S}^N = \mathcal{D}_{\mathrm{QTRANS}} \cap \mathcal{S}^N$ when n is odd.

value-restricted domain

We will now give a complete characterization of these domains. A Cartesian domain $\mathcal{D} = \mathcal{R}^N$ with $\mathcal{R} \subseteq \mathcal{S}$ is *value-restricted* if for each distinct $x, y, z \in U$, there is some alternative, say x, such that for all $\succsim \in \mathcal{R}$,

$$(x \succ y) \vee (x \succ z) \text{ or} \qquad \text{(x is never at the bottom)}$$
$$(y \succ x) \vee (z \succ x), \text{ or} \qquad \text{(x is never at the top)}$$
$$((x \succ y) \wedge (x \succ z)) \vee ((y \succ x) \wedge (z \succ x)). \qquad \text{(x is never in the middle)}$$

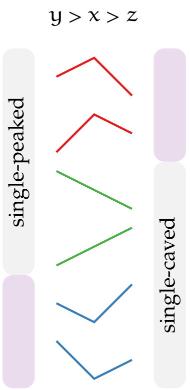

The conditions in each of these lines correspond to a domain restriction on triples (see the illustration in the margin). The first condition demands that $\mathcal{D}$ is not triple single-peaked with respect to the linear order $y > x > z$, leading to the red (x is top-ranked) and green (x is ranked in the middle) curves in the margin. The second condition demands that $\mathcal{D}$ is not triple single-caved with respect to the linear order $y > x > z$, leading to the green and blue (x is ranked last) curves in the margin. The last condition corresponds to a new type of domain restriction mixing aspects of triple single-peakedness and triple single-cavedness (highlighted in purple in the margin).

Theorem 5.5 (Sen and Pattanaik, 1969) $n \geqslant 3$

Let $\mathcal{R} \subseteq \mathcal{S}$. Then,

$$\mathcal{R}^N \subseteq \mathcal{D}_{\text{QTRANS}} \quad \Leftrightarrow \quad \mathcal{R}^N \text{ is value-restricted.}$$

Proof. For $a, b, c \in U$, define $N_{abc} = \{i \in N : a \succ_i b \succ_i c\}$.

$\Leftarrow$ Let $\mathcal{R}^N$ be a value-restricted domain and let $P \in \mathcal{R}^N$. To show that $\mathcal{R}^N \subseteq \mathcal{D}_{\text{QTRANS}}$, pick arbitrary $x, y, z \in U$ with $x \succ_M y$ and $y \succ_M z$. Then, $n_{xy} > \frac{n}{2}$ and $n_{yz} > \frac{n}{2}$, and we need to show that $n_{xz} > \frac{n}{2}$. We distinguish nine cases.

1. *x is never top-ranked:* $N_{xy} = N_{zxy} \subseteq N_{zy}$. $n_{xy} > \frac{n}{2} \Rightarrow n_{zy} > \frac{n}{2}$ $\lightning$
2. *y is never top-ranked:* $N_{yz} = N_{xyz} \subseteq N_{xz}$. $n_{yz} > \frac{n}{2} \Rightarrow n_{xz} > \frac{n}{2}$
3. *z is never top-ranked:* $N_{xy} = N_{xzy} \cup N_{xyz} \subseteq N_{xz}$. $n_{xy} > \frac{n}{2} \Rightarrow n_{xz} > \frac{n}{2}$
4. *x is never middle-ranked:* $N_{xy} = N_{xzy} \cup N_{xyz} \subseteq N_{xz}$. $n_{xy} > \frac{n}{2} \Rightarrow n_{xz} > \frac{n}{2}$
5. *y is never middle-ranked:* $N_{xy} = N_{zxy} \cup N_{xzy} \subseteq N_{zy}$. $n_{xy} > \frac{n}{2} \Rightarrow n_{zy} > \frac{n}{2}$ $\lightning$
6. *z is never middle-ranked:* $N_{yz} = N_{xyz} \cup N_{yxz} \subseteq N_{xz}$. $n_{yz} > \frac{n}{2} \Rightarrow n_{xz} > \frac{n}{2}$
7. *x is never bottom-ranked:* $N_{yz} = N_{xyz} \cup N_{yxz} \subseteq N_{xz}$. $n_{yz} > \frac{n}{2} \Rightarrow n_{xz} > \frac{n}{2}$
8. *y is never bottom-ranked:* $N_{xy} = N_{xyz} \subseteq N_{xz}$. $n_{xy} > \frac{n}{2} \Rightarrow n_{xz} > \frac{n}{2}$
9. *z is never bottom-ranked:* $N_{yz} = N_{yzx} \subseteq N_{yx}$. $n_{yz} > \frac{n}{2} \Rightarrow n_{yx} > \frac{n}{2}$ $\lightning$

$\Rightarrow$ Suppose that $\mathcal{R}^N$ is *not* value-restricted. This implies that $m \geqslant 3$. We first show that there are three relations $\succsim_1, \succsim_2, \succsim_3 \in \mathcal{R}$ and three alternatives $x, y, z \in U$ such that $\succsim_1|_{\{x,y,z\}}$, $\succsim_2|_{\{x,y,z\}}$, and $\succsim_3|_{\{x,y,z\}}$ form a latin square.

$\succsim_1\|_{\{x,y,z\}}$	$\succsim_2\|_{\{x,y,z\}}$	$\succsim_3\|_{\{x,y,z\}}$
x	y	z
y	z	x
z	x	y

Assume for contradiction, one of the relations is missing; without loss of generality, $x \succ y \succ z$. Then, there has to be one relation that top-ranks x, and it cannot be $x \succ y \succ z$. So, it is $x \succ_1 z \succ_1 y$. There has to be a relation that middle-ranks y, and it cannot be $x \succ y \succ z$. So, it has to be $z \succ_2 y \succ_2 x$. And there has to be a relation that bottom-ranks z, and it cannot be $x \succ y \succ z$. So, it has to be $y \succ_3 x \succ_3 z$. In that case, we can just rename the variables for alternatives according to $\sigma = \left(\begin{smallmatrix} x & y & z \\ z & y & x \end{smallmatrix}\right)$ and the indices for relations according to $\pi = \left(\begin{smallmatrix} 1 & 2 & 3 \\ 3 & 1 & 2 \end{smallmatrix}\right)$ to obtain exactly the configuration shown above.

We can now construct a preference profile using $\succsim_1$, $\succsim_2$, and $\succsim_3$ that is not contained in $\mathcal{D}_{\mathsf{QTRANS}}$.

$\lceil \frac{n-1}{2} \rceil$	$\lfloor \frac{n-1}{2} \rfloor$	1
x	y	z
y	z	x
z	x	y

Observe that $n_{xy} = \lceil \frac{n+1}{2} \rceil > \frac{n}{2}$, $n_{yz} = n - 1 > \frac{n}{2}$, but $n_{xz} = \lceil \frac{n-1}{2} \rceil \leqslant \frac{n}{2}$. As a consequence, $\succsim_M$ fails to be quasi-transitive.

□

$\mathcal{S}^N_{\mathsf{SP}(\geqslant)}$ and $\mathcal{S}^N_{\mathsf{SC}(\geqslant)}$ are value-restricted for all linear orders $\geqslant$. Note that the domain in which preference relations are single-peaked or single-caved according to the same linear order ($\mathcal{S}^N_{\mathsf{SP}(\geqslant)} \cup \mathcal{S}^N_{\mathsf{SC}(\geqslant)}$) is *not* value-restricted (see Exercise 5.3). A domain is value-restricted if for every triple of alternatives, the domain is *either* single-peaked *or* single-caved (or of the third type).

$\succsim_i$	$\succsim_j$	$\succsim_k$
x	y	z
y	z	x
z	x	y

Value-restriction rules out that a certain substructure, defined on three alternatives x, y, and z, can appear in preference profiles. This substructure is the latin square depicted in the margin. The domain must not contain a preference profile $P \in \mathcal{D}$ such that, there are $i, j, k \in N$ and $x, y, z \in U$ such that $P|_{\{x,y,z\}}$ restricted to voters i, j, and k looks like the profile in the margin. Note that this substructure is precisely the core of the Condorcet paradox. It is also one of the five subprofiles that appear in Theorem 5.4.

Whether a Cartesian domain $\mathcal{R}^N$ is value-restricted can be checked in polynomial time by enumerating all triples and checking the conditions from the definition of value restriction for each of the preference relations in $\mathcal{R}^N$.

5.6 Key Takeaways

Restricted Domains

- Approval voting always returns weak Condorcet winners when preferences are dichotomous.
- Median voting always returns weak Condorcet winners when preferences are single-peaked.
- Single-peaked preference profiles can be recognized in polynomial time.
- There is a characterization of all Cartesian domains of strict preferences that admit quasi-transitive majority relations.

The set inclusions between the various domains introduced in this chapter are illustrated below. While for both $\mathcal{W}^N$ and $\mathcal{S}^N$, we have negative results in the form of Arrow's impossibility and its variants, we obtained positive results for the non-Cartesian domain of profiles that admit acyclic majority relations, in particular, its subdomains of profiles with quasi-transitive and transitive majority relations, all domains that are single-peaked with respect to some linear ordering, and the domain of dichotomous preferences.

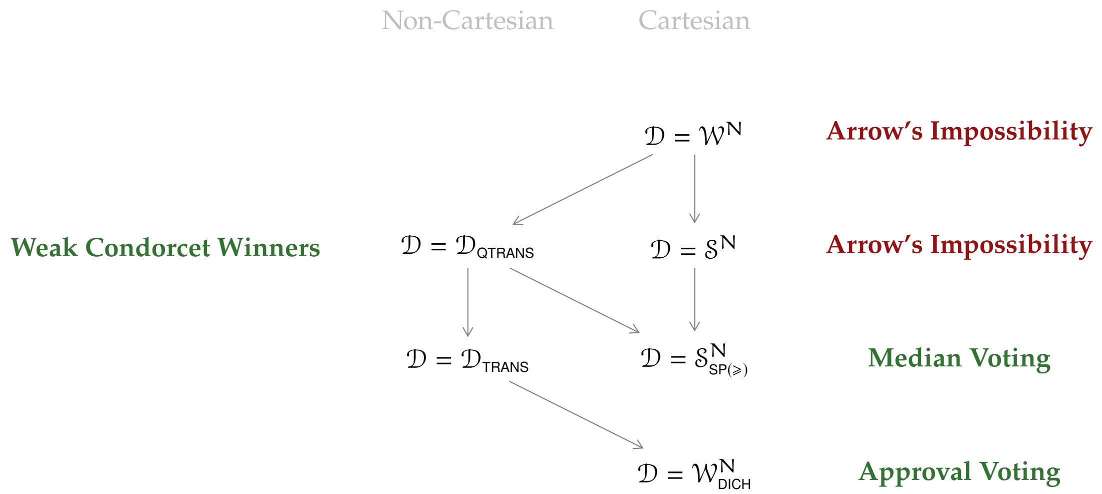

5.7 Further Reading

Various aspects of restricted domains of preferences, such as the possibility of Arrovian aggregation and non-manipulable SCFs as well as structural and computational properties, have been thoroughly examined (see, e.g., Gaertner, 2001; Le Breton and Weymark, 2011; Felsenthal and Nurmi, 2019; Elkind et al., 2017, 2025; Barberà et al., 2020, for overviews).

Two comprehensive references on approval voting are the books by Brams and Fishburn (2007) and Laslier and Sanver (2010). Its formal study was initiated by Brams and Fishburn (1978). Maniquet and Mongin (2015) characterized approval voting using Arrow's axioms within the domain of dichotomous preferences. Brandl and Peters (2022) provide an overview of axiomatic characterizations of approval voting using reinforcement (which will be introduced in Chapter 6).

An important question concerns the generalization of single-peakedness to weak preferences (see, e.g., Barberà, 2007). A natural definition for weak preferences is to have a set of consecutive most-preferred alternatives with weakly decreasing preference as one moves away from this set. (This amounts to demanding that all weak upper contour sets are intervals with respect to the linear ordering.) However, it turns out that Theorem 5.3 does not hold for this generalization. Fishburn (1973, Chapter 9) shows that the theorem can be salvaged for *single-plateaued preferences*, which allow for ties at the top—the plateau—and between pairs of alternatives residing on different sides of the plateau.

single-plateaued preferences

Single-peaked preferences are often considered for infinite sets of alternatives, such as real intervals. Theorem 5.3 and the subsequent observations regarding median voters also hold for an infinite number of alternatives.

The algorithm presented in Section 5.4 is based on the algorithm by Escoffier et al. (2008), which was predated by an earlier algorithm by Doignon and Falmagne (1994). Bartholdi, III and Trick (1986) reduced the problem of recognizing whether a profile is single-peaked to the consecutive ones problem, which first established that single-peakedness can be verified in polynomial—but not linear—time. A rigorous correctness proof of the algorithm has been given by Elkind et al. (2025, Section 4.1).

Puppe (2018) has shown that every connected Cartesian subdomain of $\mathcal{D}_{\mathrm{QTRANS}}$ such that each alternative is top-ranked in some preference relation and that contains two completely opposed preference relations has to be single-peaked for some linear order.

A common extension of single-peaked preferences to multiple dimensions is *spatial voting*, where both alternatives and voters are represented by points in Euclidean space and the preferences of voters are given by their distances to the alternatives (e.g., Austen-Smith and Banks, 1999, 2005). McKelvey (1976) and Schofield (1978) proved that when letting all points of some multi-dimensional Euclidean space be alternatives, Condorcet winners are unlikely to exist, and there are usually majority paths between any pair of alternatives. This result is known as the *McKelvey-Schofield chaos theorem* (see also Banks, 1995b).

spatial voting

McKelvey-Schofield chaos theorem

Sen (1966) proposed the condition of value restriction (based on earlier work by Ward (1965)). Value-restricted (Cartesian) domains are also called *Condorcet domains* (Puppe and Slinko, 2026) or *acyclic sets of linear orders* (Fishburn, 1997). (It can be shown that Cartesian domains guarantee quasi-transitive majority relations iff they guarantee acyclic majority relations.) A particularly challenging question concerns the maximum number of different preference relations that can be contained in a Condorcet domain with m alternatives. According to Fishburn (2002), this is "one of the most fascinating and intractable combinatorial problems in social choice theory." The number of preference relations available in $\mathcal{S}^N_{\mathrm{SP}(\geqslant)}$ is 2^{m-1}. When $m \geqslant 4$, there are larger Condorcet

Condorcet domains

domains. Leedham-Green et al. (2024) showed with the help of computers that the largest Condorcet domain for $m = 8$ consists of 224 relations and is unique up to isomorphism.

single-crossing domains

A natural class of value-restricted domains not considered in this chapter consists of *single-crossing domains*. These domains are based on arranging voters (rather than alternatives) according to a linear order such that for every pair of alternatives $x, y \in U$, the voters in N_{xy} form a consecutive interval. It can be shown that the top choices of the two median voters are weak Condorcet winners (see, e.g., Elkind et al., 2025; Puppe and Slinko, 2026).

5.8 Exercises

5.1 *Two voters*

Show that $\mathcal{W}^{\{1,2\}} \subseteq \mathcal{D}_{\text{QTRANS}}$.

5.2 *Non-Cartesian Domains*

Show that the following domains are *not* Cartesian, even when $m = 3$ and $n = 2$.

(a) $\mathcal{D}_{\text{TRANS}}$

(b) $\bigcup_{\geqslant \in \mathcal{S}} \mathcal{S}^N_{\text{SP}(\geqslant)}$

5.3 *Single-peaked and single-caved preferences*

Prove or disprove the following statements.

(a) When $m = 3$ and $n = 2$, every preference profile is single-peaked.

(b) $\succsim_M$ is quasi-transitive if *each* voter has single-peaked or single-caved preferences with respect to the same linear order $\geqslant$.

(c) Show that the profile shown on the right is single-peaked.

1	1	1	1
a	d	c	c
e	c	a	d
c	b	d	b
d	a	e	a
b	e	b	e

5.4 *Number of orders*

Show that the number of orders with respect to which a profile is single-peaked can be exponential in m.

5.5 *More on single-peakedness and single-cavedness*

Prove or disprove the following statements.

(a) Inverting a single-peaked preference profile always yields a preference profile that is single-caved.

(b) Plurality and Condorcet winners coincide for single-caved preferences.

(c) Borda and Condorcet winners coincide for single-peaked preferences.

5.6 *Recognizing single-peakedness*

Determine which of the following preference profiles are single-peaked.

(a)

1	1	1	1
f	d	a	c
c	a	f	f
a	f	c	a
d	b	d	e
b	c	e	d
e	e	b	b

(b)

1	1	1
a	a	b
b	d	d
c	b	a
f	c	e
d	f	c
e	e	f

Democracy is the recurrent suspicion that more than half the people are right more than half the time.

E. B. White, 1943

6 Scoring Rules and Condorcet Extensions

Learning Outcomes

- What classes of social choice functions are there?
- Are there Condorcet-consistent scoring rules?
- How can the ideas of Borda and Condorcet be reconciled?

The concept of representative democracy as we know it today arose largely from ideas and institutions that developed during the Age of Enlightenment, especially in the American and French Revolutions. While there were some earlier precursors, this period also marks the emergence of social choice theory as a formal discipline. The two key figures of this development were Jean-Charles, Chevalier de Borda (1733–1799), and Marie Jean Antoine Nicolas de Caritat, Marquis de Condorcet (1743–1794), both members of the French Academy of Sciences.

Borda, a member of the applied fraction of the academy, was a physicist, mathematician, and Navy officer who made important contributions to hydraulics, mechanics, and optics. He also participated in the construction of the standard meter (defined as 1/10 000 000 of the distance between the north pole and the equator) by embarking on a 7-year expedition to measure the longitude line between Dunkirk and Barcelona (through Paris) using triangulation. During the American Revolutionary War, Borda was in charge of several French Navy vessels. His name is one of 72 names inscribed on the Eiffel Tower.

Condorcet, on the other hand, was a member of the theoretical fraction of the academy. A true rationalist in the spirit of the Enlightenment, he made major contributions to philosophy, political economy, and mathematics (collaborating with Leonhard Euler, among others). Condorcet was a passionate advocate for equal rights, supporting public education, free markets, and universal suffrage—including women's suffrage, which was radical at the time.[1] Condorcet, a humanist and outspoken opponent of the death penalty, died in a prison cell after being captured by the French Revolutionary authorities.

While Borda and Condorcet agreed that plurality has serious shortcomings, they had differing opinions on which voting rule should be used instead. Borda advocated for the score-based system that bears his name, while Condorcet highlighted the importance of pairwise majority comparisons between alternatives. The conflict between these two

[1] Although France was one of the earliest European countries to embrace modern democratic ideals, it did not grant women the right to vote until 1944.

principles (rankings vs. pairwise comparisons) prevails in contemporary social choice and is the main theme of this chapter.

On the technical side, we make two important changes to our model in this chapter. First, since we will not consider consistency or rationalizability axioms, we do not vary feasible sets and simply let $\mathcal{F} = \{U\}$. This implies that the only feasible set $A \in \mathcal{F}$ is U itself. When working with SCFs, which are defined as functions of a feasible set A and a preference profile P, we usually omit A and focus on P.

Second, we extend the domain $\mathcal{S}^N$ by considering variable sets of voters. To this end, we consider a universal set of voters $N \subseteq \mathbb{N}$ and let

$\mathcal{S}^{\subseteq N}$

$$\mathcal{S}^{\subseteq N} = \bigcup_{N' \in \mathcal{P}^*(N),\, N' \text{ is finite}} \mathcal{S}^{N'}$$

be the set of all profiles of strict preferences formed by a subset of voters of N. Each preference profile still uses a finite number of voters. Recall that N_P is the set of voters present in P. Whenever N is finite, $n = |N|$. When $N = \mathbb{N}$, we write $n = \infty$. Throughout this chapter, $\mathcal{D} = \mathcal{S}^{\subseteq N}$.

6.1 Scoring Rules

score vector

scoring rule

Scoring rules are defined by associating each rank with a score. A *score vector* is a vector $s = (s_1, \dots, s_m)$ of real numbers. If a voter ranks an alternative at the ith position, it gets s_i points. For a given score vector s, the corresponding *scoring rule* f_s is defined as

$$f_s(P) = \arg\max_{x \in A} \sum_{i \in N_P} s_{|y \in A \colon y \succsim_i x|}.$$

The scoring rule chooses those alternatives for which the accumulated score is maximal. When the profile P is clear from the context, we will abuse notation by writing $s(x)$ for the score of alternative $x \in U$.

Three common examples of scoring rules are

- Borda's rule where $s = (m-1, m-2, \dots, 0)$,
- plurality rule where $s = (1, 0, \dots, 0)$, and
- anti-plurality where $s = (1, \dots, 1, 0)$.

Scoring rules are invariant under positive affine transformations of the score vectors. Two score vectors s and t with $t_i = \alpha s_i + \beta$ for all i represent exactly the same SCF when $\alpha > 0$. Scoring rules can be efficiently computed.

Obviously, one usually would want to assign higher scores to higher ranks. This implies that the scoring rule is monotonic. It is not difficult to show that every monotonic scoring rule has to satisfy this property.

Proposition 6.1 $\mathcal{F} = \{U\}, \mathcal{D} = \mathcal{S}^{\subseteq N}$

Let f_s be a scoring rule with score vector s. Then,

$$f_s \text{ is monotonic} \quad \Leftrightarrow \quad s_1 \geqslant \cdots \geqslant s_m.$$

Proof.

$\Leftarrow$ The direction from right to left is straightforward. When s satisfies $s_1 \geqslant \cdots \geqslant s_m$, reinforcing an alternative can only increase its score while the score of other alternatives may decrease.

$\Rightarrow$ The direction from left to right can be shown by contraposition. If it is not the case that $s_1 \geqslant \cdots \geqslant s_m$, there has to be some k such that $s_{k+1} > s_k$. Now consider the profile P on the right.

In this profile, each alternative is first-ranked once, second-ranked once, and so on. As consequence, $f(x) = \sum_{i=1}^{m} s_i$ for all $x \in U$ and $f(P) = U$. Now, let the first voter swap alternatives a_k and a_{k+1} in his preference relation such that a_{k+1} is now ranked higher than a_k. The resulting profile is called P'. Then, the score of a_k increases while the score of a_{k+1} decreases. As a result, $f(P') = \{a_k\}$. This constitutes a monotonicity violation because alternative a_{k+1} was reinforced but has dropped out of the choice set.

1	1 ...	1
a_1	$a_2 \dots$	a_m
a_2	$a_3 \dots$	a_1
$\vdots$		
a_k	$\vdots$	$\vdots$
a_{k+1}		
$\vdots$		
a_m	$a_1 \dots$	a_{m-1}

□

A monotonic scoring rule is *non-trivial* if $s_1 > s_m$. This merely implies that the scoring rule does not always return all alternatives.

If several alternatives receive the same score, they will all be returned by the scoring rule. One way to break these ties is to use another score vector and compare the corresponding scores of the tied alternatives. If some of the alternatives are still tied, one could use another score vector, and so on. By following this train of thought, we arrive at the definition of composite scoring rules. An SCF f is a *composite scoring rule* if there are $k \in \mathbb{N}$ and score vectors $s^1, \dots, s^k$ such that $f = f_k$ and

composite scoring rule

$$f_j(P) = \begin{cases} \arg\max_{x \in f_{j-1}(P)} \sum_{i \in N_P} s^j_{|\{y \in A \colon y \succsim_i x\}|} & \text{if } j > 1 \\ f_{s^1} & \text{otherwise.} \end{cases}$$

In other words, composite scoring rules are scoring rules with score-based tie-breaking.

The rule that returns those Borda winners that are top-ranked by most voters is a composite scoring rule where s^1 is the Borda score vector and s^2 is the plurality score vector. Another example is the rule that returns those plurality winners that are last-ranked least often. Composite scoring rules are compositions of *broad* (rather than narrow) scoring rules.

6.2 Characterization of Scoring Rules

A remarkable theorem by Peyton Young shows that scoring rules are the *only* anonymous and neutral SCFs that satisfy an intuitive variable-electorate condition called reinforcement. An SCF f satisfies *reinforcement* if for all $P, P' \in \mathcal{D}$ with $N_P \cap N_{P'} = \emptyset$,

reinforcement

$$f(P) \cap f(P') \neq \emptyset \quad \Rightarrow \quad f(P \cup P') = f(P) \cap f(P').$$

In other words, the alternatives that are chosen simultaneously by two disjoint electorates are precisely the alternatives chosen by the union of both electorates. This can be seen as an analog of contraction and expansion for variable electorates!

Contraction and expansion :	$x \in f(A) \cap f(A')$	$\Leftrightarrow$	$x \in f(A \cup A')$	$[x \in A \cap A']$
Reinforcement :	$x \in f(P) \cap f(P')$	$\Leftrightarrow$	$x \in f(P \cup P')$	$[f(P) \cap f(P') \neq \emptyset]$

Hence, on this escape route from Arrovian impossibilities, choice consistency is replaced with consistency with respect to a variable electorate.

Theorem 6.1 (Young, 1975) $\mathcal{F} = \{\mathcal{U}\}, \mathcal{D} = \mathcal{S}^{\subseteq N}, n = \infty$

Let f be an anonymous and neutral SCF.

$$f \text{ is a composite scoring rule} \quad \Leftrightarrow \quad f \text{ satisfies reinforcement.}$$

Proving the direction from right to left of Theorem 6.1 is rather difficult and usually achieved by applying the separating hyperplane theorem. Even though scoring rules are defined for finite electorates, the theorem requires an infinite set of potential voters. Note that the class of (composite) scoring rules includes trivial and non-monotonic scoring rules.

Let us call the scoring rule with score vector $(0, 1, \ldots, m-1)$ the *reverse Borda rule*. With the class of scoring rules, Borda's rule and the reverse Borda rule can be singled out by additionally demanding a technical axiom called cancellation. An SCF f satisfies *cancellation* if for all $P \in \mathcal{D}$,

cancellation

$$(\forall x, y \in A\colon n_{xy} = n_{yx}) \quad \Rightarrow \quad f(P) = A.$$

Clearly, cancellation is only meaningful for strict preferences when the number of voters is even.

Proposition 6.2

Borda's rule and the reverse Borda rule are the only non-trivial composite scoring rules satisfying cancellation.

Proof. It is easily seen that Borda's rule and the reverse Borda rule satisfy cancellation. On page 89, we will point out that the Borda score of alternative $x \in A$ can be rewritten as $\sum_{y \in A \setminus \{x\}} n_{xy}$. Similarly, the reverse Borda score of x is $\sum_{y \in A \setminus \{x\}} n_{yx}$. Hence, if all margins are identical, all Borda scores and all reverse Borda scores are identical as well.

Let us first show the rest of the statement for *non-composite* scoring rules. Let f_s be a scoring rule with score vector s that is different from both Borda's rule and the reverse Borda rule. Since f is non-trivial, there are $k, \ell \in \{1, \ldots, m-1\}$ with $k \neq \ell$ such that $s_k - s_{k+1} \neq s_\ell - s_{\ell+1}$.

Now consider a profile $P \in \mathcal{D}$ with $m!$ voters such that every voter has a different preference relation. Clearly, $n_{xy} - n_{yx} = 0$ for all $x, y \in A$, and every scoring rule (including f) will return all alternatives in A. Fix two alternatives a, b and let $i, j \in N$ be such that voter i has a in position k and b in position $k+1$ and voter j has b in position ℓ and a in position $\ell + 1$, as shown below. In particular, a and b are adjacent in the preference relations of both voter i and j.

	$\cdots$	$\succsim_i$	$\cdots$	$\succsim_j$	$\cdots$
		$\vdots$		$\vdots$	
$s_k \longrightarrow$		a		$\cdot$	
$s_{k+1} \longrightarrow$		b		$\cdot$	
		$\vdots$		$\vdots$	
$s_\ell \longrightarrow$		$\cdot$		b	
$s_{\ell+1} \longrightarrow$		$\cdot$		a	
		$\vdots$		$\vdots$	

P

$\cdots$	$\succsim_i$	$\cdots$	$\succsim_j$	$\cdots$
	$\vdots$		$\vdots$	
	b		$\cdot$	
	a		$\cdot$	
	$\vdots$		$\vdots$	
	$\cdot$		a	
	$\cdot$		b	
	$\vdots$		$\vdots$	

P'

We now swap alternatives a and b in the preference relations of voters i and j to obtain profile P'. Still, $n'_{xy} - n'_{yx} = 0$ for all $x, y \in A$. However, the scores of alternatives a and b are no longer equal. Hence, f will not return all alternatives, violating cancellation.

The previous argument shows that the only non-trivial *composite* scoring rules satisfying cancellation consist of sequences of Borda score vectors and reverse Borda score vectors. It turns out that all these sequences coincide with either Borda's rule or the reverse Borda rule, as all alternatives that have the same Borda score w have the same reverse Borda score $n(m-1) - w$, and vice versa. □

By combining Theorem 6.1 and Proposition 6.2, one immediately obtains that Borda's rule and the reverse Borda rule are the only anonymous and neutral SCFs that satisfy reinforcement and cancellation. It turns out that anonymity is not required because it is already implied by the conjunction of reinforcement and cancellation. By adding Pareto-optimality (or just Pareto-optimality in 1-voter profiles, which Young calls "faithfulness"), the reverse Borda rule is excluded.

Corollary 6.1 (Young, 1974) $\mathcal{F} = \{\mathcal{U}\}, \mathcal{D} = \mathcal{S}^{\subseteq N}, n = \infty$

Borda's rule is the only SCF satisfying neutrality, Pareto-optimality, reinforcement, and cancellation.

Proof. Let f be an SCF that satisfies reinforcement and cancellation. We show that f is anonymous. To this end, let $P, P' \in \mathcal{D}$ such that there is a bijection $\pi\colon N_P \to N_{P'}$ with $\succsim_i = \succsim'_{\pi(i)}$ for all $i \in N_P$. Let $\hat{P}$ be a copy of profile P on a disjoint electorate, i.e. $N_P \cap N_{\hat{P}} = \emptyset$ and there is a bijection $\hat{\pi}\colon N_P \to N_{\hat{P}}$ with $\succsim_i = \hat{\succsim}_{\hat{\pi}(i)}$ for all $i \in N_P$. Similarly, let $\bar{P}$ be a profile on a disjoint electorate in which all preference relations of profile P are reversed, i.e., $N_{\bar{P}} \cap (N_P \cup N_{\hat{P}}) = \emptyset$ and there is a bijection $\bar{\pi}\colon N_P \to N_{\bar{P}}$ with $\succsim_i = (\bar{\succsim}_{\bar{\pi}(i)})^{-1}$ for all $i \in N_P$. We then have the following equalities due to reinforcement and cancellation.

$$\begin{aligned} f(P) &= f(P) \cap A = f(P \cup \hat{P} \cup \bar{P}) = A \cap f(\hat{P}) = f(\hat{P}) \\ &= f(\hat{P} \cap A) = f(\hat{P} \cup \bar{P} \cup P') = A \cap f(P') = f(P') \end{aligned}$$

Now that we have established anonymity, we can invoke Theorem 6.1 and Proposition 6.2 to obtain a characterization of Borda's rule, the reverse Borda rule, and the trivial rule. The latter two do not satisfy Pareto-optimality, which completes the proof. □

Borda proposed his rule to the French Academy of Sciences in 1770. It was then used for 20 years until it was abolished by Napoleon Bonaparte, who replaced it with his own rule. It remains unknown which rule that was.

6.3 Condorcet Extensions

While scoring rules are exclusively concerned with the *ranks* of alternatives, Condorcet extensions focus on the *pairwise comparisons* between alternatives.

Condorcet extension
Condorcet-consistency

An SCF f is a *Condorcet extension* (or *Condorcet-consistent*) if $f(A, P) = \{x\}$ whenever x is a Condorcet winner in A according to P.

Condorcet-consistency on its own leaves a lot of freedom to define SCFs. Let us consider some examples of Condorcet extensions, all of which are based on the idea of selecting alternatives that are "as close as possible to being Condorcet winners" when no actual Condorcet winner exists. A common method to categorize SCFs and, in particular, Condorcet extensions is to distinguish them by their informational basis.

Copeland's rule

Condorcet winners majority-dominate all other alternatives. *Copeland's rule* returns those alternatives that dominate most other alternatives (with majority ties counting as "half" dominations), i.e.,

$$CO(A, P) = \arg\max_{x \in A} (|y \in A : x \succ_M y| + 1/2 \cdot |y \in A : x \sim_M y|)$$

In Chapter 7, we still study several other Condorcet extensions that, like Copeland's rule, only depend on the pairwise majority relation $\succsim_M$. We say that an SCF f is *majoritarian* if it is neutral and for all $P, P' \in \mathcal{D}$ and $A \in \mathcal{F}$, $\succsim_M = \succsim_f$, and $\succsim_M = \succsim'_M$ implies $f(A,P) = f(A,P')$. majoritarian SCFs

The rationale for the *maximin* rule is to return those alternatives that are less preferred than any other alternative by the fewest number of voters, i.e., maximin

$$\mathsf{MM}(A,P) = \arg\max_{x \in A} \min_{y \in A \setminus \{x\}} n_{xy}.$$

Note that if x is a Condorcet winner, x is the only alternative for which $\min_{y \in A\setminus\{x\}} n_{xy}$ is positive and will, therefore, be selected uniquely. Maximin is not majoritarian (see Exercise 6.8) as its outcome is based on the pairwise supports. Moreover, like many other Condorcet extensions, it only needs the *differences* of pairwise supports, the so-called majority margins. We say that an SCF is *margin-based* if for all $P, P' \in \mathcal{D}$ and $A \in \mathcal{F}$, margin-based SCFs

$$\forall x, y \in A\colon n_{xy} - n_{yx} = n'_{xy} - n'_{yx} \quad \Rightarrow \quad f(A,P) = f(A,P').$$

Since the pairwise majority relation can be reconstructed from the majority margins, every majoritarian SCF is also margin-based.

Finally, as an example of a Condorcet extension that is neither majoritarian nor margin-based, consider *Young's rule* YO, which returns those alternatives that can be turned into a weak Condorcet winner by removing as few voters as possible from the profile. Clearly, if P admits a Condorcet winner, it will be the only alternative returned by YO, as no voter needs to be removed. Young's rule

Note that all three of these SCFs are based on maximizing some measure (majority-dominated alternatives, worst pairwise defeats, number of included voters) to approximate Condorcet winners. CO is also similar in spirit to YO because, in the absence of majority ties, it selects those alternatives that can be turned into weak Condorcet winners by removing the smallest number of alternatives (rather than voters).

To illustrate these definitions, consider the example profile given below. Given a profile $P \in \mathcal{D}$, the *support matrix* is defined as $(n_{xy})_{x,y \in A}$ and the *margin matrix* is defined as $(n_{xy} - n_{yx})_{x,y\in A}$. Both matrices can be conveniently represented as weighted directed graphs, which help us to identify the outcomes of margin-based SCFs. Majority ties correspond to edges with weight 0 in the margin graph and are usually omitted to improve readability. Minority edges are greyed out. support matrix margin matrix

2	3	2	1
a	d	c	c
b	a	b	d
c	b	d	b
d	c	a	a

P

$$(n_{xy})_{x,y\in U} = \begin{pmatrix} & a & b & c & d \\ a & 8 & 5 & 5 & 2 \\ b & 3 & 8 & 5 & 4 \\ c & 3 & 3 & 8 & 5 \\ d & 6 & 4 & 3 & 8 \end{pmatrix}$$

support matrix

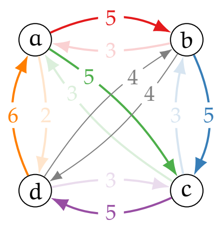

support graph

$$(n_{xy} - n_{yx})_{x,y\in U} = \begin{pmatrix} & a & b & c & d \\ a & 0 & 2 & 2 & -4 \\ b & -2 & 0 & 2 & 0 \\ c & -2 & -2 & 0 & 2 \\ d & 4 & 0 & -2 & 0 \end{pmatrix}$$

margin matrix

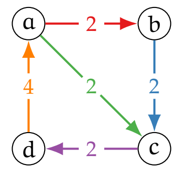

margin graph

For example, $\mathrm{CO}(P) = \{a\}$, as the Copeland scores for a, b, c, and d are 2, 1.5, 1, and 1.5, respectively. $\mathrm{MM}(P) = \{b, c, d\}$, as the maximin scores are 2, 3, 3, and 3, respectively. CO and MM can be efficiently computed in general.

Identifying the alternatives returned by Young's rule is more difficult. The margin graph reveals that at least two voters need to be removed to have a weak Condorcet winner and that only alternatives b, c, and d can be turned into weak Condorcet winners in this manner. By removing the first two voters, both c and d become weak Condorcet winners. Similarly, by removing two voters from the second column, b becomes a weak Condorcet winner. Hence, $\mathrm{YO}(A, P) = \{b, c, d\}$. In general, the problem of deciding whether a given alternative is a Young winner has been shown to be Θ_2^p-complete. Θ_2^p is a superclass of NP. The existence of a polynomial-time algorithm for computing YO would thus imply that the polynomial hierarchy collapses, which is considered extremely unlikely by most complexity theorists.

6.4 McGarvey's Theorem

Since most common Condorcet extensions are margin-based (and a significant fraction of them are majoritarian), it would be useful to know which margin matrices (and which majority graphs) can be induced by preference profiles when preferences are strict. Clearly, when n is odd, all majority margins are odd as well. Similarly, when n is even, all majority margins are even. It turns out that this is the only structural restriction imposed on majority margins.

The following theorem was shown by David McGarvey as part of his MSc thesis, written under the supervision of Kenneth May to whom we owe Theorem 3.2. A matrix M is skew-symmetric if $M = -M^T$.

Theorem 6.2 (McGarvey, 1953; Debord, 1987) $\mathcal{F} = \{U\}, \mathcal{D} = \mathcal{S}^{\subseteq N}, n = \infty$

Every skew-symmetric $U \times U$ matrix, in which all off-diagonal entries have the same parity, is the margin matrix of some preference profile $P \in \mathcal{S}^{\subseteq N}$.

Proof. Let $M = (m_{xy})_{x,y \in U}$ be a skew-symmetric matrix. First, consider the case in which all off-diagonal entries are even. We now construct a preference profile $P \in \mathcal{S}^{\subseteq N}$. For all $x, y \in U$ with $m_{xy} > 0$, add $m_{xy}/2$ voters with the preference relation

$$x \succ y \succ z_1 \succ z_2 \succ \cdots \succ z_{m-2}$$

and $m_{xy}/2$ voters with the preference relation

$$z_{m-2} \succ z_{m-3} \succ \cdots \succ z_1 \succ x \succ y$$

where $z_1, \dots, z_{m-2}$ is an arbitrary enumeration of $U \setminus \{x, y\}$. In other words, all voters prefer x to y while there are majority ties between all other pairs of alternatives. This allows us to deal with one pairwise comparison at a time. When all pairs of alternatives have been taken care of, the margin matrix of P is exactly M.
In case all off-diagonal entries of M are odd, we initiate P by adding a voter with arbitrary preferences $\succsim_1 \in \mathcal{S}$ and "subtract" this voter's preferences from M. More precisely, we construct a matrix $M' = (m'_{xy})_{x,y \in U}$ where for all $x, y \in U$ with $x \succ_1 y$, $m'_{xy} = m_{xy} - 1$ and $m'_{yx} = m_{yx} + 1$. The off-diagonal entries of M' are all even, and we can proceed as in the first case.

□

Let us illustrate the construction described in the proof of Theorem 6.2 by again considering the example from page 86. Each edge of the margin graph, except the heavy orange edge from d to a, can be represented using two voters. The orange edge requires four voters.

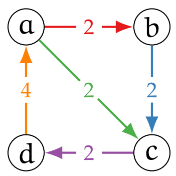

margin graph

1	1	1	1	1	1	1	1	2	2
a	d	b	b	a	d	c	b	d	c
b	c	c	a	c	b	d	a	a	b
c	a	a	b	b	a	a	c	b	d
d	b	b	c	d	c	b	d	c	a

preference profile

Note that the resulting profile has twelve voters while the profile given on page 86 induces the same margin graph but only requires eight voters. In fact, the margin graph cannot be induced by fewer than eight voters.

When also allowing weak preferences, *every* skew-symmetric matrix is induced by some profile in $\mathcal{W}^{\subseteq \mathbb{N}}$ (see Exercise 6.7).

An important consequence of Theorem 6.2 is that *every* directed graph can be induced by some preference profile.

Corollary 6.2

For every directed graph $G = (V, E)$, there exists a preference profile $P \in \mathcal{S}^{\subseteq N}$ such that $\succ_M = E$.

Since McGarvey's construction adds two voters for every edge in the majority graph, we have that $n_P \leqslant 2 \cdot \binom{m}{2} = m(m-1)$. Hence, $n_P \in O(m^2)$. Exercise 6.6 is concerned with an optimized construction such that $n_P \in O(m)$. Erdős and Moser (1964) have shown that the number of required voters is in $\Theta(m/\log m)$.

Corollary 6.2 is of crucial importance when studying majoritarian SCFs (see Chapter 7). It tells us that we can assume that the set of all possible inputs for a majoritarian SCF is the set of all directed graphs, rather than the set of all possible preference profiles. When furthermore assuming that there are no majority ties (either because they are unlikely for large electorates or because n is odd), the majority graph is a complete oriented graph, a *tournament* (see Section 7.2).

tournament

6.5 Borda versus Condorcet

When $m = 2$, majority rule is the only non-trivial monotonic scoring rule and is identical to Cond. However, once there are more than two alternatives, irreconcilable differences between the ideas of Borda and Condorcet emerge.

3	2
a	b
b	c
c	a

First, Condorcet (1785) gave a convoluted example, which proves that Borda's rule is not Condorcet-consistent.[2] The simplest profile to show this is given in the margin. Alternative a is a Condorcet winner because three out of five voters rank it at the top. The Borda scores are $s(a) = 3 \cdot 2 = 6$, $s(b) = 3 \cdot 1 + 2 \cdot 2 = 7$, and $s(c) = 2 \cdot 1 = 2$. Hence, alternative b is the Borda winner.

This observation can be significantly strengthened by showing that *no* composite scoring rule is Condorcet-consistent.

Theorem 6.3 $\mathcal{F} = \{U\}, \mathcal{D} = \mathcal{S}^{\subseteq N}, m \geqslant 3, n \geqslant 5$

No (composite) scoring rule is Condorcet-consistent.

Proof. Let (s_1, s_2, s_3) be an arbitrary score vector and consider the following profile.

2	1	1	1
a	b	b	c
b	a	c	a
c	c	a	b

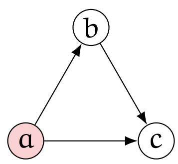

[2]Interestingly, Borda was aware of Condorcet's criterion, as he implicitly uses it in his work. In particular, he motivated his rule by pointing out that it does not select Condorcet losers. It is possible that he believed that his rule was Condorcet-consistent before Condorcet pointed out the opposite (see, e.g., McLean and Hewitt, 1994; Nurmi, 1999).

Alternative a is a Condorcet winner, but $s(a) = s(b) = 2s_1 + 2s_2 + s_3$. Hence, every scoring rule that returns a will also return b, which contradicts Condorcet-consistency. The statement extends to larger m by letting all voters rank additional alternatives below a, b, and c. Any odd $n > 5$ can be accommodated by adding pairs of voters with opposed preferences and c ranked in between a and b. Even n requires case distinctions, left to the avid reader. □

The bound of $n \geqslant 5$ is tight as the scoring rule with score vector $(3, 1, 0)$ is Condorcet-consistent for up to four voters. This rule is also the only scoring rule that returns the Condorcet winner (possibly among other alternatives) when $n \leqslant 6$ (see Exercise 6.1). A variation of Theorem 6.3 shows that no scoring rule always returns Condorcet winners (possibly among other alternatives) when $n \geqslant 7$.

Theorem 6.3 entails that Young's characterization of scoring rules (Theorem 6.1) does not hold in all subdomains of $\mathcal{S}^{\subseteq N}$. For example, when only allowing for profiles with transitive majority relations and odd numbers of at least five voters, Cond satisfies reinforcement but is no composite scoring rule.

Since Borda scores are equidistant, Borda's rule is margin-based. To see this, we can rewrite the Borda score of alternative $x \in A$ as

$$\sum_{i \in N} \left|\{y \in A \colon x \succ_i y\}\right| = \sum_{y \in A \setminus \{x\}} |N_{xy}| = \sum_{y \in A \setminus \{x\}} n_{xy}.$$

Majority margins and majority supports are just linear transformations of each other ($n_{xy} - n_{yx} = n - 2n_{yx}$). Hence, the margin matrix (or the margin graph) of a profile can be used to identify Borda winners by summing up the entries in each row (or the weights of all incident edges). In the example given on page 86, the unique Borda winner is alternative d.

Borda's rule chooses the alternatives with the highest *average rank*. Similarly, the alternatives with minimal Borda score are those with the lowest average rank. Since Condorcet winners have to have an above-average Borda score, they can never have the lowest Borda score.

Theorem 6.4 (Smith, 1973) $\mathcal{F} = \{U\}, \mathcal{D} = \mathcal{S}^N$

A Condorcet winner is never the alternative with the lowest Borda score. Borda's rule is the only scoring rule for which this is the case.

Proof. Recall that the Borda score of alternative x can be rewritten as $\sum_{y \in A \setminus \{x\}} n_{xy}$. When x is a Condorcet winner,

$$\sum_{y \in U \setminus \{x\}} n_{xy} > \sum_{y \in U \setminus \{x\}} \frac{n}{2} = (m-1) \cdot \frac{n}{2}.$$

The sum of all Borda scores over all alternatives and all voters is $n \cdot \frac{m \cdot (m-1)}{2}$. To get the average Borda score per alternative, we simply divide by m to obtain $\frac{n \cdot (m-1)}{2}$. Hence, the Condorcet winner has a Borda score above average.

For the second statement, consider a scoring rule f_s with score vector s different from Borda's rule. If f_s is the trivial rule or the reverse Borda rule, every single-voter profile has a Condorcet winner with minimal score. If f_s is neither Borda's rule nor the reverse Borda rule nor the trivial rule, there are $k, \ell \in \{1, \dots, m-1\}$ such that $d = (s_k - s_{k+1}) - (s_\ell - s_{\ell+1}) > 0$. Let $\nu = (\max_{k\in\{1,\dots,m\}} s_k - \min_{k\in\{1,\dots,m\}} s_k)/d + 1$ and construct a profile P as follows. For each strict preference relation on m, there are ν voters with these preferences, resulting in identical majority margins and identical scores for all alternatives. Now consider a subset of ν voters who rank a in kth position and b in $(k+1)$st position, and let them swap a and b. Similarly, consider a subset of ν voters who rank b in ℓth position and a in $(\ell+1)$st position, and let them swap b and a. This decreases the score of a by $\nu \cdot d$ and increases the score of b by $\nu \cdot d$ while all other scores remain unchanged. Finally, add a single voter who top-ranks a, completing P with a total of $\nu \cdot m! + 1$ voters. Alternative a is a Condorcet winner in P, but a is the unique alternative with minimal score. □

Condorcet loser

Equivalently, a Condorcet loser is never the alternative with the highest Borda score. x is a *Condorcet loser* in A if $y >_M x$ for all $y \in A \setminus \{x\}$.

Theorem 6.5 (van Newenhizen, 1992) $\mathcal{F} = \{U\}, \mathcal{D} = \mathcal{S}^N, n \to \infty$

When all preference relations are equally likely, Borda's rule maximizes the probability over all scoring rules that a Condorcet winner is chosen whenever one exists.

Is there an SCF that combines the ideas of Borda and Condorcet?

Black's rule

There have been attempts to reconcile the ideas of Borda and Condorcet. *Black's rule*, for example, returns the Condorcet winner if one exists and the Borda winner otherwise. Black's rule is one of the few anonymous and neutral SCFs that satisfy positive responsiveness. Despite the piecewise definition, Black's rule also satisfies monotonicity. It is also obviously Condorcet-consistent by definition. Hence, Theorem 6.3 implies that it is no composite scoring rule and, by Theorem 6.1 and the anonymity and neutrality of Black's rule, it cannot satisfy reinforcement.

Nanson's rule

Baldwin's rule

There is a clever way to use Borda scores to define Condorcet-consistent SCFs. *Nanson's rule*, defined in Chapter 1, is a multi-round elimination method where, in each round, all alternatives whose Borda score is not strictly larger than the average Borda score ($(m-1)n/2$) are eliminated. A variant of this rule, due to Baldwin (1926), only eliminates the alternatives with minimal Borda score. When all Borda scores are identical, the remaining alternatives win. Since Condorcet winners always have above-average Borda scores, both rules are Condorcet extensions. Both rules fail to be monotonic (Exercise 6.9). Moreover, by the same arguments as above, both rules violate reinforcement.[3]

Theorem 6.3 and Theorem 6.1 entail that no anonymous and neutral Condorcet extension satisfies reinforcement when the number of potential voters is unbounded. The following theorem shows that this incompatibility persists even when dropping

[3]In contrast to Baldwin's rule, Nanson's rule satisfies participation (to be introduced in Section 8.4) when $m \leqslant 3$ (Brandt et al., 2022b, 2025).

anonymity and neutrality and assuming at least nine voters. The theorem is of central importance because it establishes, two centuries after Borda and Condorcet, that the rationales between both ideas are inherently irreconcilable.

Theorem 6.6 (Young and Levenglick, 1978) $\mathcal{D} = \mathcal{S}^{\subseteq N}, m \geqslant 3, n \geqslant 9$

No Condorcet extension satisfies reinforcement.

Proof. We give an alternative, simpler proof than the one by Young and Levenglick (1978), which requires $n \geqslant 13$. A computer analysis by Brandt et al. (2025) has shown that the statement holds iff $n \geqslant 8$ (see Section 6.8).

Let f be a Condorcet extension satisfying reinforcement and consider the following profile P, which induces a Condorcet cycle with weight 2 on each edge:

2	2	2
a	b	c
b	c	a
c	a	b

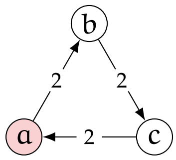

Since $f(P) \neq \emptyset$, we may assume without loss of generality that $a \in f(P)$ (otherwise relabel the alternatives). Now consider the following profile P', defined on a disjoint set of voters, i.e., $N_P \cap N_{P'} = \emptyset$:

1	2
c	a
a	c
b	b

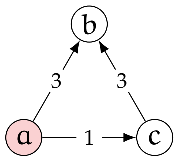

In this profile, a is the Condorcet winner, so $f(P') = \{a\}$. Since f satisfies reinforcement, we have $f(P \cup P') = f(P) \cap f(P') = \{a\}$. However, in the combined profile $P \cup P'$, c is the Condorcet winner:

2	2	3	2
a	b	c	a
b	c	a	c
c	a	b	b

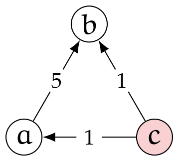

Thus, since f is Condorcet-consistent, $f(P \cup P') = \{c\}$, a contradiction.

When $m > 3$, place all additional alternatives below a, b, c in an arbitrary —but identical—order in profile P above. It remains to show that $f(P) \subseteq \{a, b, c\}$. Assume for contradiction that there is $x \in f(P) \setminus \{a, b, c\}$. Then, consider the 3-voter profile P'' where all voters have the preferences $x \succ a \succ b \succ c \dots$. Clearly, $f(P'') = \{x\}$ because of Condorcet-consistency. Reinforcement implies that $f(P \cup P'') = \{x\}$, which contradicts the fact that alternative a is a Condorcet winner in $P \cup P''$. □

Note that the same proof also holds when weakening reinforcement by replacing the consequence with either $f(P \cup P') \subseteq f(P) \cap f(P')$ or $f(P \cup P') \supseteq f(P) \cap f(P')$.

Theorem 6.6 establishes a fundamental dilemma in social choice theory. The set of Condorcet extensions and that of all SCFs that satisfy reinforcement are disjoint:

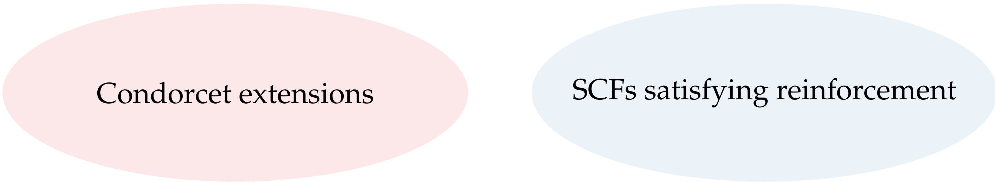

When generalizing Borda's proposal to demanding that rules should satisfy at least reinforcement, we obtain that the principles underlying Condorcet's and Borda's proposals are inherently incompatible. While there are Condorcet extensions that are based on scores, such as Black's rule and Nanson's rule, any such rule has to violate reinforcement.

Theorem 6.6 is often understood as a criticism of Condorcet extensions. In particular, the proof of Theorem 6.6 shows that the Condorcet winner can change when adding a completely symmetric Condorcet cycle profile, in which—by anonymity and neutrality—all alternatives have to be returned. Note, however, that the direction of the Condorcet cycle is crucial for this phenomenon to appear.

99	1
a	b
b	c_1
c_1	$\vdots$
$\vdots$	c_{100}
c_{100}	a

There is a similarly devastating criticism of Borda's rule. Consider the profile in the margin. The Borda score of alternative a is $s(a) = 99 \cdot 101 = 9\,999$, which is less than $s(b) = 99 \cdot 100 + 101 = 10\,001$. Hence, Borda's rule will return alternative b, even though 99% of the voters think a is the best alternative.

In the next section, we will consider aggregation functions that map each preference profile to a set of preference relations. It then turns out that reinforcement, a strong interpretation of Condorcet's principle, and neutrality characterize exactly one function: Kemeny's rule!

6.6 Kemeny's rule

Let us begin this section by examining a probabilistic model that dates back to the Marquis de Condorcet. This model, which has its origins in the Enlightenment era, assumes that there is a "true" objective ranking of all alternatives. When voters form their opinions, each voter selects a "true" pairwise comparison with a probability of $0.5 < p < 1$. The voters' preferences are imperfect estimates of the "truth." They tend to be right more often than wrong ($p > 0.5$). The canonical motivating example is a jury, where each juror forms his own opinion on whether a defendant is guilty or not.

maximum likelihood SCF

Condorcet initially focused on the case of two alternatives and used a statistical method known today as *maximum likelihood estimation*. For this, he leveraged the, at the time, new theory of probability calculus. An SCF is a *maximum likelihood SCF* for a given p if it yields all alternatives that are most likely to be top-ranked in the "true" ranking. Let us stick with this informal definition to avoid defining a sophisticated, concrete model, which would require introducing additional notation.

Theorem 6.7 (Condorcet Jury Theorem, 1785) $\mathcal{D} = \mathcal{S}^N, m = 2$

Majority rule is the maximum likelihood SCF for $0.5 < p < 1$. As n goes to infinity, majority rule converges to the "truth" with probability 1.

Proof. For each profile P, one can compute how likely it is that this profile was sampled according to the ground truth $a > b$ and compare this to the probability for the ground truth $b > a$. The first probability is $p^{n_{ab}} \cdot (1-p)^{n_{ba}}$ while the second one is $p^{n_{ba}} \cdot (1-p)^{n_{ab}}$. Performing some arithmetic conversions results in the following equivalences.

$$\begin{aligned}
& p^{n_{ab}} \cdot (1-p)^{n_{ba}} > p^{n_{ba}} \cdot (1-p)^{n_{ab}} \\
\Leftrightarrow \quad & p^{n_{ab}-n_{ba}} \cdot (1-p)^{n_{ba}-n_{ab}} > 1 \\
\Leftrightarrow \quad & \underbrace{\left(\frac{p}{1-p}\right)}_{>1}^{n_{ab}-n_{ba}} > 1 \\
\Leftrightarrow \quad & n_{ab} - n_{ba} > 0
\end{aligned}$$

Hence, the ground truth $a > b$ is more likely iff $a >_M b$. The probability with which majority rule is correct when there are n voters is

$$\sum_{i=\lceil (n+1)/2 \rceil}^{n} \binom{n}{i} \cdot p^i \cdot (1-p)^{n-i}.$$

This probability converges to 1, as n increases. However, it is not strictly increasing; in fact, it is lower for even numbers of voters because of majority ties. □

For example, when $p = 2/3$ (meaning voters are right twice as often as they are wrong), majority rule is correct with probability 0.9 when there are 15 voters and with probability 0.99 when there are 47 voters.

Condorcet's maximum likelihood analysis of the case of three or more alternatives is vague and equivocal. Young (1988) has picked up Condorcet's original ideas and proved the following two theorems about maximum likelihood estimation in social choice.

Theorem 6.8 (Young, 1988) $\mathcal{D} = \mathcal{S}^N$

Borda's rule is the maximum likelihood SCF if p is sufficiently close to 0.5.

This essentially follows from the fact that Borda winners are precisely those alternatives that receive most "pairwise votes", where a voter gives a pairwise vote to an alternative if he prefers it to another alternative. Note that, when considering $m > 2$, sampling each pairwise comparison independently from each other can lead to intransitive individual preference relations.

Young (1988) was able to make sense of Condorcet's writings by considering social preference functions. A *social preference function (SPF)* is a function $h\colon \mathcal{D} \to \mathcal{P}^*(\mathcal{S})$. The mathematician and computer scientist John Kemeny has proposed an SPF that is now known as Kemeny's rule.

social preference function (SPF)

$$\text{KE}(P) = \arg\max_{\succsim \in \mathcal{S}} \sum_{i \in N_P} |\succ \cap \succ_i|$$

Kemeny's rule yields all rankings that maximize pairwise agreements, i.e., a ranking of the alternatives that agrees with as many pairwise preferences of the voters as possible. Kemeny rankings are sometimes also called "consensus rankings."

In the definition above, we take the intersection of relations. This is possible because relations are just sets consisting of ordered pairs of alternatives. For example, when $\succ_1 = \{(a,b),(b,c),(a,c)\}$ and $\succ_2 = \{(c,a),(a,b),(c,b)\}$. Then, $\succ_1 \cap \succ_2 = \{(a,b)\}$. The cardinality of the symmetrical difference $|\succ_1 \triangle \succ_2| = (\succ_1 \setminus \succ_2) \cup (\succ_2 \setminus \succ_1)$ is called the *Kendall-Tau distance* between two relations. Kemeny's rule returns all rankings that minimize the sum of Kendall-Tau distances to each individual ranking.

Kendall-Tau distance

Just as in the case of SCFs, we can define maximum likelihood SPFs. A *maximum likelihood SPF* for given p returns those rankings that are most likely "correct."

maximum likelihood SPF

Theorem 6.9 (Condorcet, 1785; Young, 1988) $\mathcal{D} = \mathcal{S}^N$

Kemeny's rule is the maximum-likelihood SPF for $0.5 < p < 1$.

Kemeny score

As an example, consider the preference profile P below. The *Kemeny score* $\sum_{i \in N_P} |\succ \cap \succ_i|$ of each of the possible six rankings $\succsim \in \mathcal{S}$ of three alternatives is given on the right-hand side.

5	4	2	2
a	b	c	c
b	c	a	b
c	a	b	a

P

Ranking	Kemeny score
$a \succ b \succ c$	$7+9+5=21$
$a \succ c \succ b$	$5+4+7=16$
$b \succ a \succ c$	$6+5+9=20$
$b \succ c \succ a$	$9+8+6=23$
$c \succ a \succ b$	$8+7+4=19$
$c \succ b \succ a$	$4+6+8=18$

The unique Kemeny ranking is $b \succ c \succ a$.

To get more insights into how Kemeny rankings can be computed, consider the support graph and the margin graph of profile P.

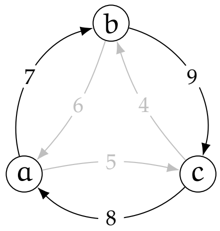

support graph

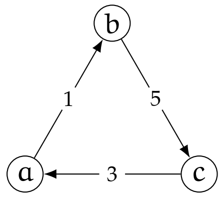

margin graph

A Kemeny ranking is an acyclic subgraph with maximum weight. If cyclic rankings were allowed, $\succsim_M$ would have a maximal Kemeny score of $7 + 9 + 8 = 24$. In comparison to majority rule, every Kemeny edge (y, x) that does not coincide with the corresponding majority edge x, y invokes a penalty of $n_{xy} - n_{yx}$.

Now consider the margin graph with weights $n_{xy} - n_{yx}$. In this graph, a Kemeny ranking can be found by finding a set of edges with minimal accumulated weight whose inversion makes the graph acyclic and then inverting precisely these edges.

The following lemma implies that any set of edges of minimal weight whose *deletion* makes the graph acyclic also has to be a minimal set whose *inversion* makes the graph acyclic.

Lemma 6.1

Let $G = (V, E)$ be a directed graph and $E' \subseteq E$. Then, G can be made acyclic by inverting a subset of edges in E' iff $(V, E \setminus E')$ is acyclic.

Proof.

- $\Rightarrow$ Let E' be a set of edges whose inversion makes G acyclic. Then, removing E' from G also makes G acyclic.
- $\Leftarrow$ Let E' be a set of edges whose removal makes G acyclic. Now, insert edges from E' one after another, orienting each one such that no cycle forms; if an edge cannot be oriented without having a cycle, the previous graph was cyclic, a contradiction. Hence, E' contains a subset of edges whose inversion makes G acyclic.

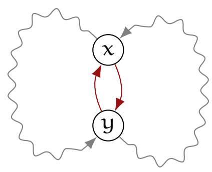

□

Hence, the inclusion-minimal sets of edges whose removal make G acyclic and the inclusion-minimal sets of edges whose inversion makes G acyclic coincide. Note that referring to a subset of edges in E' in the second part of the statement is crucial for the correctness of the lemma. This can be seen by considering a 3-cycle. Inverting all three edges retains the cyclicity of the graph, whereas removing all three edges results in an acyclic graph. Computing a Kemeny ranking thus boils down to identifying a set of edges—at least one from each cycle—with minimal accumulated weight. Cycles may overlap and, therefore, cannot be treated separately.

Recall the preference profile from page 86, which results in the margin graph shown below.

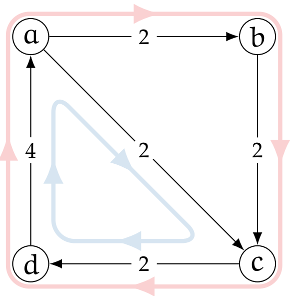

This margin graph contains two cycles: (a, c, d), highlighted in blue, and (a, b, c, d), highlighted in red. These cycles share two edges, (d, a) with weight 4 and (c, d) with weight 2. Deleting any of these makes the margin graph acyclic but deleting (c, d) is "cheaper." In fact, deleting (c, d) is the *only* possibility to get rid of all cycles by deleting edges with accumulated weight of 2 or less. (There are several ways to remove all cycles by deleting edges with accumulated weight of 4, e.g., by deleting the single edge (d, a) or by deleting edges (a, c) and (a, b).) Thus, the unique Kemeny ranking is $d \succ a \succ b \succ c$.

Peyton Young and Arthur Levenglick have given two very appealing axiomatic characterizations of Kemeny's rule, both of which are based on the reinforcement axiom introduced in Section 6.2.

reinforcement (for SPFs)

An SPF h satisfies *reinforcement (for SPFs)* if for all $P, P' \in \mathcal{D}$ with $N_P \cap N_{P'} = \emptyset$,

$$h(P) \cap h(P') \neq \emptyset \quad \Rightarrow \quad h(P \cup P') = h(P) \cap h(P').$$

Since there are many more rankings than there are alternatives, reinforcement for SPFs applies less often than for SCFs and, hence, seems easier to satisfy.

majority-consistency (for SPFs)

Two alternatives x and y are *adjacent* in a preference relation $\succsim$ if there is no z such that $x \succ z \succ y$ or $y \succ z \succ x$. An SPF h is *majority-consistent* if for all $P \in \mathcal{D}$, and $\succsim \in h(P)$, and $x \succ y$ adjacent in $\succsim$, $x \succsim_M y$. Moreover, $x \sim_M y$ implies $\succsim \setminus \{(x,y)\} \cup \{(y,x)\} \in h(P)$. In other words, adjacent pairwise comparisons are consistent with the majority relation. Note that majority-consistency implies that Condorcet winners are ranked at the top of every returned ranking. Moreover, each of the returned rankings represents a Hamiltonian path in the weak majority graph.

neutrality (for SPFs)

Finally, an SPF h is *neutral* if $\sigma(h(P)) = h(P')$ for all $P, P' \in \mathcal{D}$ with $N_P = N_{P'}$ such that there is a permutation $\sigma\colon U \to U$ with $P' = \sigma(P)$.

Theorem 6.10 (Young and Levenglick, 1978) $\mathcal{D} = \mathcal{S}^{\subseteq N}, n = \infty$

Kemeny's rule is the only neutral SPF that satisfies reinforcement and majority-consistency.

Proving that Kemeny's rule satisfies these properties is easy. To see that Kemeny's rule satisfies reinforcement, it helps to view it as a scoring rule on rankings: When merging two electorates, the Kemeny scores of each ranking are just added up, and every ranking that had maximal Kemeny score in both original electorates also has maximal Kemeny score in the union of these electorates. To see that Kemeny's rule satisfies majority-consistency, assume for contradiction that two adjacent alternatives are not ordered according to the pairwise majority relation. Then, the Kemeny score of this ranking could be increased by swapping these alternatives. Moreover, if there is a majority tie between two adjacent alternatives, then those alternatives can be swapped without affecting the Kemeny score. Neutrality is obviously satisfied. Proving the converse is the difficult part of the theorem.

Since majority-consistency can be seen as an (admittedly strong) interpretation of Condorcet's principle, Theorem 6.10 can be compared to Theorem 6.6, showing that no Condorcet extension satisfies reinforcement. What we referred to as a fundamental dilemma in social choice theory then turns into a complete characterization of Kemeny's rule when moving from SCFs to SPFs! It should be emphasized that the SCF that returns all alternatives that are top-ranked in Kemeny rankings does *not* satisfy reinforcement. This follows immediately from Theorem 6.6.

For the second characterization of Kemeny's rule, we will consider a weakening of majority-consistency. An SPF h satisfies *local independence of irrelevant alternatives (LIIA)* if for all $P, P' \in \mathcal{D}$, and $\succsim \in h(P)$, $\succsim' \in h(P')$, and $x, y \in U$ adjacent in both $\succsim$ and $\succsim'$,

local independence of irrelevant alternatives (LIIA)

$$P|_{\{x,y\}} = P'|_{\{x,y\}} \quad \Rightarrow \quad \succsim \setminus \{(x,y),(y,x)\} \cup \succsim'|_{\{x,y\}} \in h(P).$$

Furthermore, let us extend Pareto-optimality and anonymity to SPFs. An SPF h is *Pareto-optimal* if for all $P \in \mathcal{D}$, $x, y \in U$, and $\succsim \in h(P)$, $(\forall i \in N_P : x \succ_i y) \Rightarrow x \succ y$. An SPF h is *anonymous* if $h(P) = h(P')$ for all $P, P' \in \mathcal{D}$ such that there is a bijection $\pi\colon N_P \to N_{P'}$ with $\succsim_i = \succsim'_{\pi(i)}$ for all $i \in N_P$.

Pareto-optimality (for SPFs)

anonymity (for SPFs)

Theorem 6.11 (Young, 1988) $\mathcal{D} = \mathcal{S}^{\subseteq N}, n = \infty$

Kemeny's rule is the only SPF that satisfies anonymity, neutrality, Pareto-optimality, reinforcement, and LIIA.

Proof. Exercise ◻

This theorem nicely illustrates how Kemeny's rule circumvents Arrow's impossibility (Theorem 4.2). Even though LIIA is defined for SPFs, rather than SWFs, it feels like a weakening of IIA: how two adjacent alternatives x and y are ranked only depends on the voters' preferences between x and y. Kemeny's rule satisfies an even stronger version of LIIA, which states that every Kemeny ranking, restricted to a consecutive interval of alternatives, only depends on the preferences of the voters over alternatives within this interval. Hence, when removing the top-ranked or the bottom-ranked alternative, the Kemeny ranking of the remaining alternatives is unaffected. This is a very useful property, for example, in the hiring committee example discussed in Section 4.3.

Borda's rule can be interpreted as an SPF by returning all rankings in which x is ranked above y when x has a higher Borda score than y. This SPF satisfies all conditions of Theorem 6.11 except LIIA.

On top of the given axiomatic characterizations, there are many ways to arrive at Kemeny's rule. Kemeny rankings are maximum-likelihood rankings (Theorem 6.9). They maximize the accumulated score given by pairwise votes and correspond to maximum weight acyclic majority subgraphs as well as median rankings, as they minimize the sum of distances to all individual preference rankings in terms of the Kendall tau metric. Moreover, one arrives at Kemeny rankings when studying which unanimous preference profile (in which all voters have identical preferences) is reached by successively swapping the lowest number of pairs of adjacent alternatives, starting with the actual preferences of the voters. Because of Young's important contributions towards a better understanding of Kemeny's rule, it is sometimes also called the "Kemeny-Young method." Young (1995) concludes his latest paper on Kemeny's rule by stating, "I predict that the time will come when [Kemeny's rule] is considered a standard tool for political and group decision making."

While the Kemeny score of a given ranking can be computed efficiently, choosing combinations of objects with minimal accumulated weight (edges from cycles) smells like NP-hardness to experienced theoretical computer scientists. Indeed, it can be shown that computing a Kemeny ranking is NP-hard.

Feedback Arc Set

As we will see, the problem can be reduced to *Feedback Arc Set*, one of Karp's original 21 NP-complete decision problems:

Is it possible to make a given directed graph acyclic by removing k *edges?*

Theorem 6.12 (Karp, 1972)

Feedback Arc Set is NP-complete.

Let us refer to the problem of deciding whether there exists a ranking whose Kemeny score is at least some fixed number $s \in \mathbb{N}$ as Kemeny Score.

Theorem 6.13 (Bartholdi, III et al., 1989; Bachmeier et al., 2019) $\mathcal{D} = \mathcal{S}^N, n \geqslant 7$

Kemeny Score is NP-complete.

Proof. We give the original proof by Bartholdi, III et al. (1989) for $\mathcal{D} = \mathcal{S}^{\subseteq N}$ and mention at the end how the result has been extended to $\mathcal{D} = \mathcal{S}^N$ with small n. The proof is a reduction from Feedback Arc Set, i.e., we translate any instance of Feedback Arc Set to an instance of Kemeny Score. A Feedback Arc Set instance consists of a directed graph $G = (V, E)$ and an integer k. We construct a preference profile $P \in \mathcal{S}^{\subseteq N}$ on alternatives $U = V$ with odd $n_P \leqslant 2|E|$ voters such that $>_M = E$ and $n_{xy} - n_{yx} \in \{-2, 2\}$ for all $x, y \in U$ using McGarvey's theorem (Theorem 6.2). Lemma 6.1 implies that G can be made acyclic by removing k edges iff P admits a ranking with Kemeny score at least

$$s = \frac{m(m-1)}{2}\frac{(n_P+2)}{2} - 2k.$$

The product is the product of the number of edges $\binom{m}{2}$ and n_{xy} for each majority edge (x, y).
Dwork et al. (2001) showed that the Kemeny Score remains NP-complete if $\mathcal{D} = \mathcal{S}^N$ and n is even and at least 4. Their reduction contained a small error that was fixed by Biedl et al. (2009). Proving that Feedback Arc Set remains hard for majority graphs induced by four voters is relatively easy because one can simply subdivide every edge by introducing a new alternative. This construction preserves the size of the minimal feedback arc set and ensures that the new graph can be induced by only four voters. Larger even numbers of voters can be accounted for by adding pairs of voters with opposed preferences. The bound of 4 for even n is tight as Kemeny's rule can be computed in polynomial time for only two voters (e.g., the preference rankings of the voters are Kemeny rankings).
Showing that Kemeny Score remains NP-complete for odd n is much trickier. When n is odd, there are no majority ties and the majority graph is a tournament. Whether Feedback Arc Set remains NP-complete for tournaments was an acknowledged open problem for several decades before it was settled independently by three sets of authors (Alon, 2006; Charbit et al., 2007; Conitzer, 2006). Bachmeier et al. (2019) analyzed the tournaments appearing in the reduction by Conitzer (2006) and proved that these tournaments can be induced using only seven voters. Hardness has thus been established for even $n \geqslant 4$ and odd $n \geqslant 7$. □

The proof shows that Kemeny Score is NP-complete, even when all majority margins are equal. Whether Kemeny Score is NP-hard for $n \in \{3, 5\}$ is open. It is thus unknown whether we can efficiently aggregate three rankings into a median ranking.

Theorem 6.13 has established that the decision problem Kemeny Score is at least as hard as Feedback Arc Set. This implies the NP-hardness of the important problem of *finding* a Kemeny ranking. To see this, consider the following Turing reduction. Suppose there was an oracle that efficiently computes a Kemeny ranking $\succsim$ for a given preference profile P. s^*, the Kemeny score of $\succsim$, is the highest possible score of all rankings according to P. Hence, we could use this oracle to efficiently solve Kemeny Score by answering "yes" iff $s \leqslant s^*$. As a consequence, finding a Kemeny ranking is NP-hard when $n \geqslant 7$.

It can also be shown that deciding whether a given alternative is a Kemeny winner is NP-hard (see Exercise 6.15). This implies that no algorithm for computing Kemeny winners whose runtime is polynomial in the number of alternatives is known, and it is unlikely that one exists. However, this does not necessarily mean that computing Kemeny winners in practice is hopeless. NP-hardness is a worst-case measure. Realistic distributions of preferences for political elections often admit Condorcet winners, and even when no Condorcet winner exists, these instances might have Kemeny winners that can be easily identified.

Given Young's 1988 reading of Condorcet (1785), which equates Condorcet's maximum likelihood method with Kemeny's rule, it is fascinating to read that in 1793 Condorcet wistfully looked forward to "the formation of arithmetic machines or methods which could be used to determine the results of a very large ballot." (McLean and Hewitt, 1994, p. 45).

6.7 Key Takeaways

Scoring Rules and Condorcet Extensions

- A defining characteristic of scoring rules is reinforcement.
- Every Condorcet-consistent social choice function violates reinforcement.
- Every margin graph with only odd or only even weights is induced by some preference profile.
- Kemeny's rule is the maximum likelihood social preference function.
- Kemeny's rule is Condorcet-consistent and satisfies reinforcement.
- Computing Kemeny's rule is NP-hard.

6.8 Further Reading

Theorem 6.1 can be shown using a separating hyperplane argument. Corollary 6.1 and characterizations of approval voting using reinforcement only require elementary mathematics (see, e.g., Hansson and Sahlquist, 1976; Fishburn, 1978; Brandl and Peters, 2022). Independently of Young (1975), Smith (1973) has characterized score-based *SWFs* using reinforcement, which allows for a simpler proof. Based on Young's findings, specific scoring rules can easily be characterized using reinforcement. There are various characterizations of plurality (Lepelley, 1992; Sekiguchi, 2012; Öztürk, 2020) and anti-plurality (Bossert and Suzumura, 2016; Kurihara, 2018).

Theorem 6.1 can be rephrased for scoring rules (rather than *composite* scoring rules) by adding a continuity axiom (Young, 1975). As Zwicker (2016) notes, any composite scoring rule is equivalent to some simple scoring rule if we restrict the domain by fixing an upper bound on the number of voters.

Theorem 6.5 is based on earlier work by Gehrlein and Fishburn (1978), who studied which scoring rules are most likely to select Condorcet winners when n is finite.

Overviews of the rich landscape of Condorcet extensions have been provided by Fishburn (1977), Brams and Fishburn (2002), Laslier (1997), Brandt et al. (2016a), Fischer et al. (2016), and Caragiannis et al. (2016). Fishburn (1977) and several subsequent studies distinguish between C1, C2, and C3 SCFs, which are similar to our definitions of majoritarian and margin-based SCFs. The set of majoritarian SCFs consists precisely of all neutral C1 SCFs, and the set of support-based SCFs consists precisely of all neutral C1 and C2 SCFs.

Copeland's rule is usually attributed to Copeland (1951), although this idea goes back to earlier work by Ramon Lull in the 13th century (Hägele and Pukelsheim, 2001; Fidora and Sierra, 2011). Maximin is sometimes also called *minimax* (since it minimizes the worst loss) or the *Simpson-Kramer* rule (Simpson, 1969; Kramer, 1977).

Condorcet's original profile showing that Borda's rule fails to be Condorcet-consistent uses 81 voters and works for every scoring rule with strictly monotonic score vectors. Fishburn (1984b) gives an example with 7 voters to the same effect. The profile used in Theorem 6.3 is due to Florian Grundbacher.

The original version of Theorem 6.6 by Young and Levenglick (1978) is only phrased in terms of *weak* Condorcet winners, but they briefly mention a variant of this result that can be turned into a 13-voter impossibility for Condorcet extensions (see, e.g., Moulin, 1988a, Theorem 9.2). The 9-voter proof of Theorem 6.6 is due to Keyvan Kardel. Brandt et al. (2025) showed that the minimal number of voters required for this statement is 8. Their proof was found with the help of a computer, and they were unable to extract a human-readable proof. However, they give an 8-voter proof for anonymous SCFs and a 5-voter proof for anonymous and neutral SCFs. They also show that both bounds are tight. The example profile, criticizing Borda's rule, on page 92 is due to Klaus Nehring.

Barthélemy and Monjardet (1981) provide an early survey of the properties of Kemeny's rule by connecting it to medians in cluster analysis. When ignoring the weights in the margin graph, one obtains a majoritarian version of Kemeny's rule called *Slater's rule* (Slater, 1961).

Slater's rule

Can and Storcken (2013) prove an alternative characterization of Kemeny's rule. Kemeny's rule has been rediscovered many times in equivalent formulations (see, e.g., Bowman and Colantoni, 1973; Blin and Whinston, 1975; Merchant and Rao, 1976; Adelsman and Whinston, 1977). For example, in optimization, computing Kemeny rankings of a margin graph is known as the *linear ordering problem*. Charon and Hudry (2007, 2010), Hudry (2008, 2012) and Ali and Meila (2012) survey the algorithmic aspects of computing Kemeny rankings and related problems. Deciding whether a given ranking is a Kemeny ranking is coNP-complete (Fitzsimmons and Hemaspaandra, 2021), and deciding whether a given alternative is a Kemeny winner is Θ_2^p-complete (Hemaspaandra et al., 2005), even when restricting attention to tournaments (Lampis, 2022). Feedback Arc Set cannot be approximated efficiently; it is APX-hard (Kann, 1992). There exists a polynomial-time approximation scheme (PTAS) for weighted *tournament* Feedback Arc Set (Kenyon-Mathieu and Schudy, 2007).

linear ordering problem

A number of interesting margin-based Condorcet extensions emerge when deleting edges in the margin graph to break cycles and then returning the maximal elements. These methods include *split cycle* and its refinements *ranked pairs*, *Schulze's rule*, and *river* (Tideman, 1987; Schulze, 2003; Holliday and Pacuit, 2023; Döring et al., 2026). Split cycle, for example, removes all edges with minimal margins from each cycle.

split cycle
ranked pairs
Schulze's rule
river

6.9 Exercises

☆ **6.1** *Scoring rules*

Show that for $m = 3$ and $n = 5$, the only scoring rule that always returns Condorcet winners (possibly among other alternatives) has the score vector $(3, 1, 0)$.

☆ **6.2** *Cancellation*

Show that Proposition 6.2 holds even when $n \leqslant 4$ (with no restrictions placed on m): Borda's rule and the reverse Borda rule are the only non-trivial composite scoring rules satisfying cancellation.

6.3 *Margin-based SCFs and cancellation*

Let f be a neutral SCF that satisfies reinforcement. Prove that f is margin-based iff it satisfies cancellation.

6.4 *Reinforcement in restricted domains*

Prove or disprove whether Cond satisfies reinforcement in the given domains.

(a) $\mathcal{D}_{\text{TRANS}}$

(b) $\mathcal{W}^N_{\text{DICH}}$

(c) $\mathcal{S}^N_{\text{SP}(\geqslant)}$

6.5 *McGarvey's Theorem*

Give a preference profile whose majority graph corresponds to the tournament below.

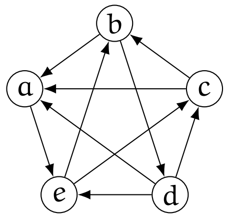

☆ **6.6** *McGarvey's Theorem with a linear number of voters*

Show that for every directed graph $G = (V, E)$, there exists a preference profile $P \in \mathcal{S}^{\subseteq N}$ such that $\succ_M = E$ and $n_P \leqslant 2m$.

6.7 *McGarvey's Theorem for weak preferences*

Show that every skew-symmetric $U \times U$ matrix is the margin matrix of some preference profile $P \in \mathcal{W}^{\subseteq \mathbb{N}}$.

6.8 *Margin-based SCFs*

Show that maximin (MM) and Borda's rule (BO) are not majoritarian.

☆ **6.9** *Nanson's rule and Baldwin's rule*

Show that Nanson's rule and Baldwin's rule are Condorcet extensions and violate monotonicity.

6.10 MI *and* MA

Consider the following two margin-based SCFS.

$$\mathrm{MI}(A, P) = \arg\max_{x \in A} \sum_{y \in A} \min(0, n_{xy} - n_{yx})$$
$$\mathrm{MA}(A, P) = \arg\max_{x \in A} \sum_{y \in A} \max(0, n_{xy} - n_{yx})$$

Check whether these SCFs satisfy Pareto-optimality, reinforcement, and Condorcet-consistency.

6.11 *Kemeny's rule*

Compute all Kemeny rankings of the following profiles.

(a)

1	1	1
a	b	c
b	c	a
c	a	b

(b)

2	2	1
a	b	c
b	c	a
c	a	b

(c)

3	3	2	1
a	b	c	d
b	d	d	c
c	a	a	b
d	c	b	a

6.12 *Hamiltonian path*

A *Hamiltonian path* is a path that visits each vertex of a directed graph exactly once.

(a) Let $G = (V, E)$ be a directed graph such that for all $x, y \in V$, $(x, y) \in E$ or $(y, x) \in E$. Prove that G contains a Hamiltonian path.

(b) Show that every Kemeny ranking constitutes a Hamiltonian path in $(A, \succsim_M)$.

☆ **6.13** *Pareto-optimality of Kemeny's rule*

Show that Kemeny's rule is Pareto-optimal.

☆ **6.14** *Theorem 6.11*

Let $n = \infty$. Show that Kemeny's rule is the only SPF that satisfies anonymity, neutrality, Pareto-optimality, reinforcement, and LIIA.

Hint: First, prove that majority rule is the only SPF on two alternatives that satisfy anonymity, neutrality, Pareto-optimality, and reinforcement. Then, show that, under the given conditions, LIIA is equivalent to majority-consistency. Finally, apply Theorem 6.10 and show that Kemeny's rule satisfies all the axioms.

6.15 *Finding Kemeny winners*

Top-ranked alternatives in Kemeny rankings are called *Kemeny winners*. Show that finding a Kemeny winner is NP-hard.

Hint: Give a reduction argument that repeatedly invokes an oracle that finds Kemeny winners. You may use the fact that computing Kemeny rankings is NP-hard.

One of the consequences of the assumptions of rational choice is that the choice in any environment can be determined by a knowledge of the choices in two-element environments.

Kenneth Arrow, 1951

7 Majoritarian SCFs

Learning Outcomes

- What are reasonable majoritarian social choice functions?
- How can we circumvent Arrow's impossibility by weakening rationalizability?
- How many alternatives can be excluded by consistent social choice functions?

Blair et al. (1976) and Sen (1977) observed that all the proofs of the Arrovian impossibilities listed in Section 4.5 are actually statements about the base relation. Moreover, every impossibility involving rationalizability (i.e., acyclicity of the base relation) can be turned into a stronger impossibility that only requires contraction. Consider, for example, the Condorcet-May impossibility (Theorem 3.4), where we eventually derived a contradiction because the base relation was acyclic. This implies a violation of contraction (see Lemma 2.3). Expansion is not needed.

Theorem 7.1 (Condorcet, 1785; May, 1952; Sen, 1977) $m \geqslant 3, n \geqslant 3$

There is no SCF satisfying anonymity$_2$, neutrality$_2$, positive responsiveness$_2$, and contraction.

1	1	1
a	b	c
b	c	a
c	a	b

Proof. Let f be an SCF with the desired properties and consider the Condorcet paradox profile P depicted in the margin. Based on the symmetry of P and the fact that f always returns at least one alternative, we can assume without loss of generality that $a \in f(A, P)$. Contraction then implies that $a \in f(\{a, c\}, P)$. However, anonymity$_2$, neutrality$_2$, and positive responsiveness$_2$ imply that pairwise choices are made according to majority rule by May's theorem (Theorem 3.2). For P, we have $c \succ_M a$, which means that $f(\{a, c\}, P) = \{c\}$, a contradiction. □

Similarly, one obtains a strong version of the Mas-Colell and Sonnenschein (1972) impossibility: every SCF satisfying IIA$_2$, positive responsiveness$_2$, and contraction admits a quasi-dictator. Sen (1977) pointed out that an even weaker variant of contraction suffices for these results: for all $A \in \mathcal{F}$ with $|A| = 3$, there is some $x \in f(A)$ such that $x \in f(\{x, y\})$ for all $y \in A$. This condition is equivalent to the absence of 3-cycles in the base relation.[1]

[1] Other Arrovian impossibility theorems, such as Theorem 4.4 and Theorem 4.5, can also be strengthened

A	$f(A)$
ab	a
bc	b
ac	c
abc	abc

Strong expansion, on the other hand, has no implications on the acyclicity of the base relation. The choice function f, given in the margin, satisfies strong expansion, but $\succsim_f$ is cyclic. In summary, contraction consistency, even in very weak forms, has devastating consequences on social choice. Expansion consistency, even in its strongest form, appears to be much less harmful. Hence, we will ignore contraction and explore a rich escape route from Arrow's theorem by looking for reasonable SCFs that satisfy strong expansion, expansion, or even weaker versions of expansion consistency. In particular, we will introduce three majoritarian SCFs that can be characterized using expansion consistency conditions. We start with the top cycle.

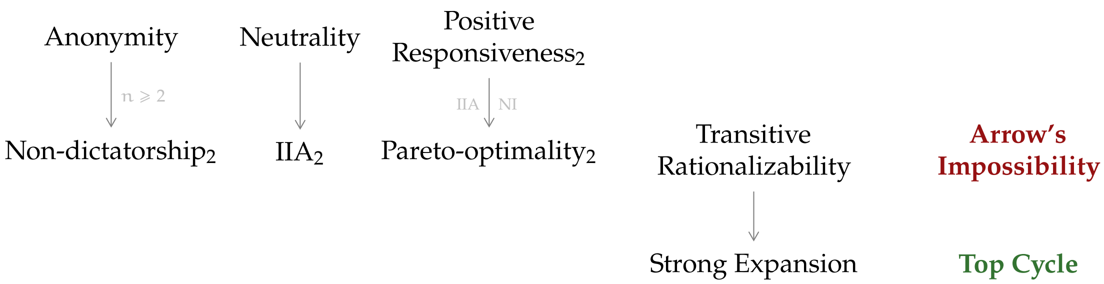

7.1 The Top Cycle

The first SCF we consider on this escape route is called the top cycle and is based on an auxiliary notion that generalizes the notion of Condorcet winners to sets of alternatives.

dominant set

A *dominant set* is a nonempty set of alternatives $B \subseteq A$ such that for all $x \in B$ and $y \in A \setminus B$, $x >_M y$. Let

$$\text{Dom}(A, >_M) = \{B \subseteq A : B \text{ is dominant}\}.$$

While Condorcet winners need not exist, every profile admits at least one dominant set, namely the set of all alternatives A. When there is a Condorcet winner, the singleton set consisting of the Condorcet winner is another dominant set. In general, we will be interested in small dominant sets. It turns out that for any pair of dominant sets, one has to be contained in the other.

Proposition 7.1

$\text{Dom}(A, >_M)$ is totally ordered by set inclusion.

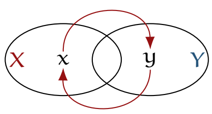

Proof. Assume for contradiction that there is a profile P, two dominant sets $X, Y \in \text{Dom}(A, >_M)$, and $x \in X \setminus Y$, $y \in Y \setminus X$. By dominance of X, $x >_M y$. Analogously, by dominance of Y, $y >_M x$, a contradiction. □

by replacing rationalizability with contraction. However, the described weakening of contraction for three alternatives does not suffice.

For every majority graph, dominant sets thus form a hierarchy as shown in the margin and admit a unique minimal dominant set (Ward, 1961). The minimal dominant set is called the *top cycle (TC)*. There are various equivalent ways to define TC, e.g., by taking the intersection of all dominant sets or by invoking the Max operator on the subset relation.

top cycle (TC)

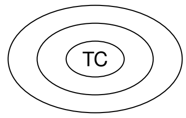

$$\mathrm{TC}(A, P) = \bigcap \mathrm{Dom}(A, \succ_M) = \mathrm{Max}(\mathrm{Dom}(A, \succ_M), \subseteq).$$

TC is a Condorcet extension, because, as mentioned above, any singleton set consisting of a Condorcet winner is a minimal dominant set. The top cycle is also known as the *majority set* (Ward, 1961), the *Good set* (Good, 1971), the *Smith set* (Smith, 1973), or *GETCHA* (Schwartz, 1986). The name is derived from the observation that the elements of the top cycle form a Hamiltonian cycle with respect to $\succsim_M$.

We now give a characterization of the top cycle by weakening one of Arrow's conditions and strengthening the other three, as shown on page 106. The core of the proof is based on a characterization by Bordes (1976), who used slightly different axioms.

Theorem 7.2 (Bordes, 1976)

TC is the finest SCF satisfying anonymity, neutrality, monotonicity$_2$, and strong expansion.

Proof. First, recall the variant of May's theorem (Theorem 3.3), which establishes that majority rule is the finest SCF satisfying anonymity, neutrality, and monotonicity$_2$ when $m = 2$. This implies the statement of the theorem for all $A \in \mathcal{F}$ with $|A| = 2$.

The rest of the proof consists of two parts. First, we show that any f that satisfies anonymity, neutrality, monotonicity$_2$, and strong expansion is a coarsening of TC, i.e., $\mathrm{TC} \subseteq f$. To see this, fix some profile $P \in \mathcal{D}$ and feasible set A and let $B = \{x, y\}$ with $x \in f(A)$ and $y \in A \setminus f(A)$. We can apply strong expansion since $B \subseteq A$ and $x \in f(A) \cap B \neq \emptyset$. We thus obtain that $f(B) = f(\{x, y\}) \subseteq f(A) \not\ni y$. This only leaves $f(B) = \{x\}$, which by our initial observation is equivalent to $x \succ_M y$. With $x \in f(A)$ and $y \in A \setminus f(A)$ being arbitrary, we have shown that $f(A) \in \mathrm{Dom}(A, \succ_M)$: f returns a dominant set. Since $\mathrm{TC}(A)$ is contained in all dominant sets, we get $\mathrm{TC} \subseteq f$.

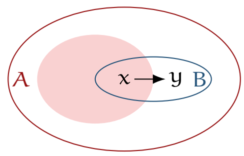

It remains to show that TC satisfies anonymity, neutrality, monotonicity$_2$, and strong expansion. The first three properties are straightforward. To see that TC satisfies strong expansion, take two arbitrary feasible sets A, B such that $\mathrm{TC}(A) \cap B \neq \emptyset$. A dominant set remains dominant when removing alternatives as long as at least one alternative of the dominant set remains. Hence, $\mathrm{TC}(A) \cap B$ is dominant in $(B, \succ_M)$. By definition, $\mathrm{TC}(B)$ is contained in all dominant sets of $(B, \succ_M)$. Hence $\mathrm{TC}(B) \subseteq \mathrm{TC}(A) \cap B \subseteq \mathrm{TC}(A)$. □

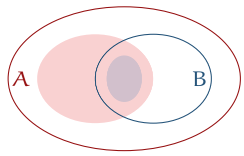

Note that there are other SCFs satisfying these axioms (such as the trivial SCF that always returns all alternatives or the SCF that returns all alternatives except alternatives that are unanimously last-ranked). However, TC is the finest among these. Theorem 7.2 can be interpreted as follows: we fix choices from two-element sets via majority rule and then leverage the strongest expansion consistency condition to extend these choices to large feasible sets. It is noteworthy that there is a *unique* finest SCF that satisfies the given axioms.

To get more insight into TC and how it could be computed, consider the following three examples. The alternatives contained in the top cycle are marked in red.

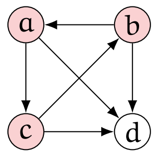

$\text{TC}(A, P) = \{a, b, c\}$

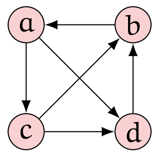

$\text{TC}(A, P) = \{a, b, c, d\}$

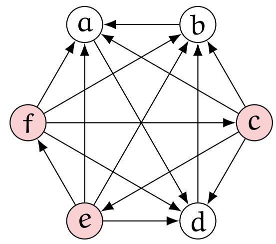

$\text{TC}(A, P) = \{c, e, f\}$

While it can be easily verified whether a set is dominant, it is not obvious how to check whether a set contains a dominant subset without checking every subset.

The minimal dominant set containing a given alternative x can be computed using a simple greedy algorithm: initialize working set B with $\{x\}$ and then iteratively add all alternatives that weakly majority-dominate an alternative in B until no more such alternatives can be found. This will result in

$$\bar{D}^*(x) = \{x_1 \in A \colon \exists x_2, x_3, \dots, x_k = x \text{ such that } x_1 \succsim_M x_2 \succsim_M \dots \succsim_M x_k\}.$$

Since dominant sets are totally ordered by set inclusion (Proposition 7.1), *every* dominant set is of the form $\bar{D}^*(x)$ for some $x \in A$. We thus have

$$\text{Dom}(A, \succ_M) = \{\bar{D}^*(x) \colon x \in A\}.$$

Using the algorithm described above, we can compute $\text{Dom}(A, \succ_M)$ in $O(|A|^3)$.[2] Once we have done this, we can identify $\text{TC}(A, P)$, the inclusion-minimal element of $\text{Dom}(A, \succ_M)$.

The previous algorithm can be improved by only computing $\bar{D}^*(x)$ for a single alternative x that is definitely contained in the top cycle. It turns out that any Copeland winner does the job, as the set of Copeland winners is a subset of the top cycle.

Proposition 7.2

$\text{CO} \subseteq \text{TC}$.

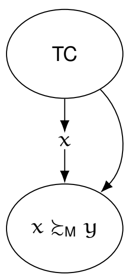

Proof.
Assume for contradiction that $x \in \text{CO}(A, P)$ but $x \notin \text{TC}(A, P)$. Recall that the Copeland score of alternative x is $|\{y \in A \colon x \succ_M y\}| + 1/2 \cdot |\{y \in A \colon x \sim_M y\}|$. It follows from the

[2] Since TC only depends on the majority graph, and the majority graph of a preference profile can be easily constructed in linear time ($O(m \cdot n)$), we consider the majority graph, restricted to the feasible set, to be the input of the algorithm.

definition of dominant sets that every $y \in A$ with $x \succ_M y$ or $x \sim_M y$ is not contained in $\mathsf{TC}(A, P)$. Hence, for an arbitrary $z \in \mathsf{TC}(A, P)$, we have $z \succ_M y$ and, moreover, $z \succ_M x$. As a consequence, the Copeland score of z is strictly greater than that of x, a contradiction. □

Since CO can be easily computed in linear time, we thus have a linear-time algorithm for computing TC. (Note that a linear-time algorithm is $O(|A|^2)$ because the size of the majority graph is quadratic in $|A|$.)

```
procedure TC(A, P)
    B ← CO(A, P)
    C ← B
    repeat
        C ← ⋃_{y∈C} {x ∈ A \ B : x ≿_M y}
        B ← B ∪ C
    until C = ∅
    return B
```

The proof of Proposition 7.2 actually shows that the Copeland score of any alternative in the top cycle is strictly greater than that of any alternative outside of the top cycle. This can be exploited algorithmically by first sorting all alternatives by decreasing Copeland score and returning the smallest dominant prefix of this sequence.

The top cycle is a simple and robust concept that sometimes appears when trying to circumvent Arrow's impossibility without drawing on strong expansion or dominant sets. The essence of Arrovian impossibility results is that the base relation fails to be transitive (or acyclic). A natural idea is to simply consider the *transitive closure* $\succsim_M^*$ of $\succsim_M$. Define $\succsim_M^*$ by letting

transitive closure

$$x \succsim_M^* y \quad \Leftrightarrow \quad x \in \bar{D}^*(y)$$

for all $x, y \in U$.

As it turns out, the maximal elements of the transitive closure of the majority relation are precisely the alternatives in the top cycle.

Theorem 7.3 (Deb, 1977)

$\mathsf{TC}(A, P) = \mathrm{Max}(A, \succsim_M^*)$.

Proof.

$$\begin{aligned} x \in \mathrm{Max}(A, \succsim_M^*) \quad &\Leftrightarrow \quad \forall y \in A \colon x \succsim_M^* y \\ &\Leftrightarrow \quad \forall y \in A \colon x \in \bar{D}^*(y) \\ &\Leftrightarrow \quad \forall B \in \mathrm{Dom}(A, \succsim_M) \colon x \in B \\ &\Leftrightarrow \quad x \in \mathsf{TC}(A, P) \end{aligned}$$

□

Incidentally, TC can be understood as an abbreviation of "top cycle" as well as "transitive closure." Theorem 7.3 leads to an alternative definition of the top cycle: It consists precisely of those alternatives x that, when weakly majority dominated by y, are the source of a weak majority path from x to y.

This insight allows for alternative linear-time TC algorithms based on the Kosaraju-Sharir algorithm or Tarjan's algorithm for finding strongly connected components.

At this stage, TC may appear to be an attractive solution to the issues raised by Condorcet and Arrow. However, there is a significant drawback to the top cycle: it is fairly large. In fact, it is so large that it may contain Pareto-dominated alternatives when there are more than two alternatives! A minimal profile that exhibits this phenomenon is shown in the margin. a Pareto-dominates b, but the top cycle contains all three alternatives. Similar examples without majority ties can be constructed with at least four alternatives, as shown in the margin. Since Pareto-optimality is an essential ingredient of all Arrovian impossibilities, this escape route is so far not entirely convincing. Although technically, these results only require Pareto-optimality$_2$ (which TC satisfies).

1	1
a	c
b	a
c	b

1	1	1
a	b	d
b	c	a
c	d	b
d	a	c

7.2 Tournaments: Disregarding Majority Ties

TC and CO are majoritarian, i.e., they are neutral SCFs whose outcome only depends on the majoritarian choices from two-element sets. Majoritarian SCFs satisfy three of Arrow's conditions (non-dictatorship, IIA, and Pareto-optimality$_2$) and form a surprisingly rich subclass of SCFs. We will restrict attention to majoritarian SCFs in this chapter. In fact, we are going to restrict the input of SCFs even further by ruling out majority ties. Majority ties are unlikely to arise for large electorates and, for example, can never appear when preferences are strict and there is an odd number of voters. The absence of majority ties allows for a much simpler exposition of majoritarian SCFs. For instance, the Copeland score of an alternative simply becomes $|\{y \in A \colon x \succ_M y\}|$. This assumption also allows for simpler axiomatic characterizations and more elegant results. Similar to the introduction of strict preferences, which rules out ties in the individual preferences, the absence of majority ties makes our lives a bit easier.

We thus define the following non-Cartesian subdomain of $\mathcal{S}^{\subseteq N}$:

$$\mathcal{D}_{\text{TOUR}} = \{P \in \mathcal{S}^{\subseteq N} \colon \forall x, y \in U \colon n_{xy} \neq n_{yx}\}.$$

For any $P \in \mathcal{D}_{\text{TOUR}}$, $\succsim_M$ is antisymmetric, and its irreflexive part coincides with $\succ_M$. A feasible set A and $\succ_M|_A$ define a *tournament* $(A, \succ_M)$, an oriented complete graph.[3]

When $\mathcal{D} = \mathcal{D}_{\text{TOUR}}$, we will denote majoritarian SCFs as functions of tournaments. For example, we write $\text{CO}(A, \succ_M)$ instead of $\text{CO}(A, P)$ and $\text{TC}(A, \succ_M)$ instead of $\text{TC}(A, P)$. Such SCFs are also known as *tournament solutions*. We just refer to dominance rather than majority-dominance when discussing $\succ_M$. In tournaments, an undominated (or equivalently, a dominant) alternative is a Condorcet winner. Furthermore, all three notions of transitivity coincide, and, in particular, a tournament is transitive iff it contains no 3-cycle. The Copeland score of an alternative is called the *degree* of an alternative.

degree

[3]We will simply write $(A, \succ_M)$ instead of $(A, \succ_M|_A)$ since majoritarian SCFs satisfy IIA by definition.

The following observation will be useful.

Lemma 7.1

Let f be a majoritarian SCF and $a \succ_M b \succ_M c \succ_M a$. Then, $f(\{a,b,c\},\succ_M) = \{a,b,c\}$.

Proof. Observe that f is neutral and that for any $x, y \in A = \{a,b,c\}$, there is a permutation $\pi : A \to A$ (either $\pi = \left(\begin{smallmatrix} a & b & c \\ b & c & a \end{smallmatrix}\right)$ or $\pi = \left(\begin{smallmatrix} a & b & c \\ c & a & b \end{smallmatrix}\right)$) such that $x \succ_M y$ iff $\pi(x) \succ_M \pi(y)$. Hence, if $x \in f(\{a,b,c\},\succ_M)$, then $y \in f(\{a,b,c\},\succ_M)$. The statement then immediately follows from the non-emptiness of choice sets. □

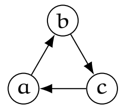

The 3-cycle is only one of many tournaments in which all alternatives have to be selected because of neutrality. Such tournaments are called vertex-homogeneous and are defined by demanding that for every $x, y \in A$ there is a tournament automorphism mapping x to y. Only tournaments with an odd number of alternatives can be vertex-homogeneous. For $m \in \{3, 5\}$, all vertex-homogeneous tournaments are "regular" in the sense that all vertices have the same degree. This is no longer true when $m \geqslant 7$. There are three regular tournaments with $m = 7$, only two of which are vertex-homogeneous.

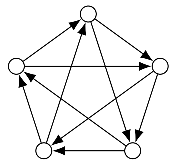

The following notation will simplify matters when working with tournaments. Let $(A, \succ_M)$ be a tournament and $x \in A$. The *dominion* of x is denoted by

dominion

$$D(x) = \{y \in A \colon x \succ_M y\}$$

and the *dominators* of x by

dominators

$$\bar{D}(x) = \{y \in A \colon y \succ_M x\}.$$

Note that, for a given alternative x, every tournament can be partitioned into $\bar{D}(x)$, $D(x)$, and $\{x\}$ itself. $D^k(x)$ returns the set of alternatives reachable from x in at most k steps and is defined inductively by letting $D^0(x) = \{x\}$ and

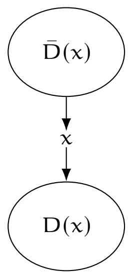

$$D^k(x) = D^{k-1}(x) \cup \bigcup_{y \in D^{k-1}(x)} D(y).$$

$D^*(x) = \bigcup_{k \geqslant 0} D^k(x)$ is the set of all alternatives reachable from x.

Similarly, we have

$$\bar{D}^k(x) = \bar{D}^{k-1}(x) \cup \bigcup_{y \in \bar{D}^{k-1}(x)} \bar{D}(y)$$

with $\bar{D}^0(x) = \{x\}$ and

$$\bar{D}^*(x) = \bigcup_{k \geqslant 0} \bar{D}^k(x)$$

for the set of alternatives from which x can be reached in at most k steps and the set of all alternatives from which x can be reached. Using this notation, we can rewrite TC by letting $\mathrm{TC} = \{x \in A \colon D^*(x) = A\}$.

Tournaments can be nicely visualized by omitting all edges that point downwards. The following three figures are graphical representations of the same tournament.

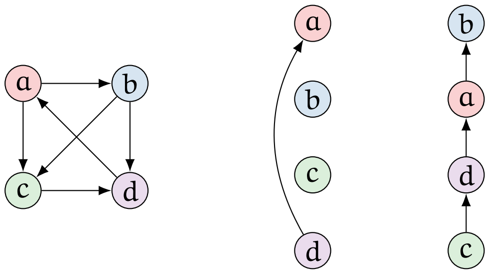

This tournament can also be used to illustrate the previous definitions: $D(c) = \{d\}$, $\bar{D}(c) = \{a, b\}$, $D^2(c) = \{a, c, d\}$, $\bar{D}^2(c) = \{a, b, c, d\}$, and $D^*(c) = \{a, b, c, d\}$.

A natural computational question concerns the optimal graphical representation of tournaments such that a maximum number of edges point downwards and can thus be omitted. Interestingly, the NP-hardness of this problem follows directly from the NP-completeness of Feedback Arc Set (see Section 6.6): the top-to-bottom ordering of alternatives in an optimal graphical representation directly corresponds to a Kemeny ranking of the underlying majority tournament under the assumption that all margins are equal.

The SCFs we introduce in the following will be characterized among the class of majoritarian SCFs. In order to relate these characterizations to Arrow's theorem, let us characterize the class of majoritarian SCFs using four axioms, each of which is logically related to an axiom of Arrow's theorem.

binariness An SCF f is *binary* if for all $A \in \mathcal{F}$ and $P, P' \in \mathcal{D}$,

$$\forall x, y \in A \colon f(\{x, y\}, P) = f(\{x, y\}, P') \quad \Rightarrow \quad f(A, P) = f(A, P').$$

Binariness demands that which alternatives are chosen from a feasible set only depends on the pairwise choices within this feasible set. It is thus weaker than rationalizability.[4]

Proposition 7.3

An SCF is majoritarian if it satisfies binariness, anonymity, neutrality, and positive responsiveness$_2$.

Proof. The statement immediately follows from May's theorem (Theorem 3.2) and the definition of binariness. □

The idea of Theorem 7.2 and the characterizations of the following SCFs is that we fix pairwise choices using majority rule and then use expansion consistency conditions to extend these choices in a meaningful way.

[4] It may seem that binariness also implies IIA. Note, however, that the pairwise choices can depend on preferences over infeasible alternatives.

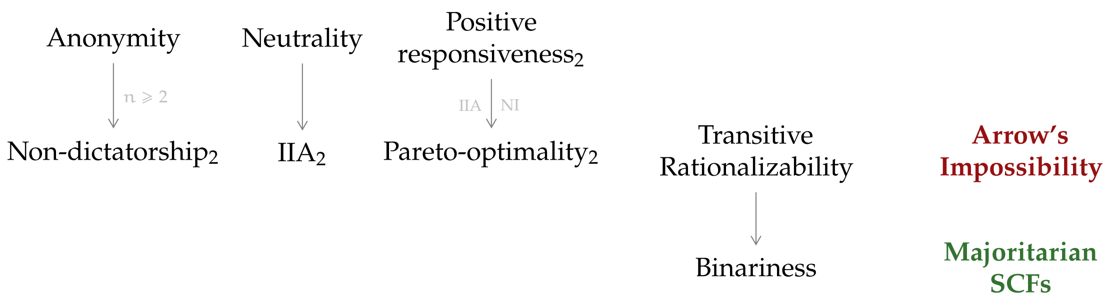

7.3 The Uncovered Set

Based on the top cycle's lack of decisiveness and, in particular, its lack of Pareto-optimality, we continue to escape Arrow's impossibility by looking for SCFs that can be characterized using expansion consistency conditions.

This section is concerned with the uncovered set, a refinement of the top cycle proposed independently by Fishburn (1977) and Miller (1980). While the top cycle was defined using the auxiliary notion of dominant sets, the uncovered set is based on a subrelation of the dominance relation: the covering relation, which goes back to a game-theoretic notion used by Gillies (1959). Given a tournament $(A, \succ_M)$, x covers y (denoted by $x \mathrel{C} y$), if $D(y) \subset D(x)$.

The covering relation is a transitive subrelation of the dominance relation. To see that

$$x \mathrel{C} y \;\Rightarrow\; x \succ_M y \quad \text{for all } x, y \in U,$$

observe that, in tournaments, $x \succ_M y$ is violated iff $y \succ_M x$, which contradicts $x \mathrel{C} y$ as x does not dominate itself. The transitivity of the covering relation,

$$x \mathrel{C} y \wedge y \mathrel{C} z \;\Rightarrow\; x \mathrel{C} z \quad \text{for all } x, y, z \in U,$$

follows immediately from the transitivity of the set inclusion relation.

The definition of the covering relation is very robust for tournaments in the sense that many similar definitions coincide. The following equivalences hold for all $x, y \in U$.

$$\begin{aligned} x \mathrel{C} y \quad &\Leftrightarrow \quad D(y) \subset D(x) \\ &\Leftrightarrow \quad \bar{D}(x) \cup \{x\} \subset \bar{D}(y) \cup \{y\} \\ &\Leftrightarrow \quad \bar{D}(x) \subset \bar{D}(y) \text{ and } x \succ_M y \\ &\Leftrightarrow \quad \bar{D}(x) \subset \bar{D}(y) \end{aligned}$$

It also does not matter whether one uses strict or weak inclusion in any of these definitions, as two dominions can never be identical.

The *uncovered set* (UC) consists of all uncovered alternatives, i.e., uncovered set (UC)

$$UC(A, \succ_M) = Max(A, C).$$

It is interesting to observe that the uncovered set contains the maximal elements of a transitive *subrelation* of the dominance relation, the top cycle contains the maximal elements of a transitive *superrelation* of the dominance relation (see Theorem 7.3).

Clearly, Condorcet winners are uncovered because they are undominated. Moreover, Condorcet winners cover all other alternatives, and are, hence, returned exclusively. UC is a Condorcet extension.

An alternative way to arrive at the definition of the uncovered set goes as follows. We face the problem that most tournaments do not admit Condorcet winners. However, every tournament contains *subtournaments* that admit Condorcet winners, for example, all subtournaments with two or fewer alternatives. It turns out that the Condorcet winners of all inclusion-maximal subtournaments that admit a Condorcet winner are precisely the alternatives in the uncovered set (Brandt, 2011).

Let us consider two example tournaments to get more intuition into the uncovered set. The covering relation is marked by blue edges, and the alternatives in the uncovered set are marked in red.

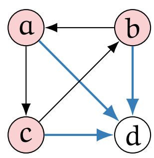

$\mathsf{UC}(A, \succ_M) = \{a, b, c\}$

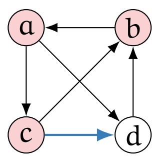

$\mathsf{UC}(A, \succ_M) = \{a, b, c\}$

In the tournament on the right-hand side, for example, c covers d because b, which is dominated by d, is also dominated by c. For comparison, $\mathsf{TC}(A, \succ_M) = \{a, b, c, d\}$ in both tournaments.

Similar to Theorem 7.2, which characterized TC using strong expansion, UC can be characterized using expansion.

trans. rat. → strong expansion → expansion
Arrow, TC, UC

Theorem 7.4 (Moulin, 1986) $\mathcal{D} = \mathcal{D}_{\mathsf{TOUR}}$

UC is the finest majoritarian SCF satisfying expansion.

Proof. Just like the proof of Theorem 7.2, the proof consists of two parts.
We first show that for any majoritarian SCF f that satisfies expansion, we have $\mathsf{UC} \subseteq f$. To show this, let $A \in \mathcal{F}$ and $a \in \mathsf{UC}(A)$. It needs to be shown that $a \in f(A)$. We are now covering the entire tournament $(A, \succ_M)$ with two-element and three-element subtournaments. In each of these tournaments, a is chosen by f.
For the two-element tournaments, we take an arbitrary alternative $x \in D(a)$. Since f is majoritarian and $a \succ_M x$, $f(\{a, x\}) = \{a\}$. In summary,

$$\forall x \in D(a)\colon\ a \overset{\text{Maj.}}{\in} f(\{a, x\}). \tag{1}$$

For the three-element tournaments, the following equivalence comes in handy.

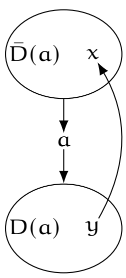

$$a \in \mathsf{UC}(A) \quad \Leftrightarrow \quad \nexists x \in A\colon x \mathrel{C} a \quad \Leftrightarrow \quad \forall x \in \bar{D}(a)\ \exists y \in D(a)\colon y \succ_M x \tag{*}$$

The above observation says that for each alternative $x \in \bar{D}(a)$, there is an $y \in D(a)$ with $y \succ_M x$. We thus have $x \succ_M a \succ_M y \succ_M x$, a three-cycle. Majoritarianness implies that $f(\{a, x, y\}) = \{a, x, y\}$. We thus have that

$$\forall x \in \bar{D}(a)\ \exists y \in D(a)\colon a \overset{\text{Maj.}}{\in} f(\{a, x, y\}). \tag{2}$$

Since all alternatives other than a are either in $D(a)$ or in $\bar{D}(a)$, these two- and three-element subtournaments cover the entire tournament $(A, \succ_M)$. Repeated application of expansion then gives $a \in f(A)$. As a consequence, $\mathsf{UC} \subseteq f$.

It remains to be shown that UC is majoritarian and satisfies expansion. Clearly, UC is majoritarian by definition. Regarding expansion, assume for contradiction that there is a profile $P \in \mathcal{D}$ and feasible sets A, B such that $x \in \mathsf{UC}(A) \cap \mathsf{UC}(B)$ but $x \notin \mathsf{UC}(A \cup B)$. Then there is some $y \in A \cup B$ such that $y \mathrel{C} x$ in $A \cup B$. Without loss of generality, let $y \in A$. Then $y \mathrel{C} x$ in A, which contradicts the assumption that $x \in \mathsf{UC}(A)$. □

Note that the proof of Theorem 7.4 only requires majoritarianness to conclude that all alternatives have to be selected if the majority relation between them forms a three-cycle. It is, therefore, possible to strengthen Theorem 7.4 by replacing majoritarianness with this property.

Since strong expansion implies expansion, $\mathsf{UC} \subseteq \mathsf{TC}$.[5] In contrast to TC, UC never returns Pareto-dominated alternatives. As in the case of Theorem 7.2, it is remarkable that a *unique* finest majoritarian SCF satisfying expansion exists.

Proposition 7.4 $\mathcal{D} = \mathcal{D}_{\mathsf{TOUR}}$

UC satisfies Pareto-optimality.

Proof. Consider a profile $P \in \mathcal{D}$ such that $a \succ_i b$ for all $i \in N$. We then have the following chain of implications.

$$x \in D(b) \quad \Leftrightarrow \quad b \succ_M x \quad \overset{\text{Trans.}}{\Rightarrow} \quad a \succ_M x \quad \Leftrightarrow \quad x \in D(a)$$

The implication in the middle follows from the transitivity of *individual* preferences: if all of the voters prefer a to b and more than half of them prefer b to x, then more than half of them also prefer a to x. Of course, we do not assume here that $\succ_M$ is transitive. We thus have $D(b) \subseteq D(a)$, which means that $a \mathrel{C} b$ and hence $b \notin \mathsf{UC}(A, \succ_M)$ for all $A \in \mathcal{F}$ that contain a. □

The previous proof shows that the Pareto relation is a subrelation of the covering relation. The covering relation is thus sandwiched between the majority relation and the Pareto relation.

If we only have access to the dominance relation, the *only* way to discard Pareto-dominated alternatives is to get rid of covered alternatives.

[5]This inclusion also immediately follows from the observation that all alternatives in the top cycle cover all alternatives not in the top cycle.

Theorem 7.5 (Brandt et al., 2016b) $\mathcal{D} = \mathcal{D}_{\text{TOUR}}$

UC is the coarsest majoritarian SCF satisfying Pareto-optimality.

Proof. Proposition 7.4 has already established that UC is Pareto-optimal. Let f be a Pareto-optimal majoritarian SCF f. It remains to show that $f \subseteq \text{UC}$. Assume for contradiction that there is a tournament $(A, \succ_M)$ and $a, b \in A$ such that $a \mathrel{C} b$ and $b \in f(A, \succ_M)$.

We will now construct a preference profile $P \in \mathcal{S}^{\subseteq N}$ with majority relation $\succ_M$ such that a Pareto-dominates b, contradicting Pareto-optimality of f. We start by adding a single voter to P with the preference relation

$$\bar{D}(a) \succ a \succ D(a) \cap \bar{D}(b) \succ b \succ D(b)$$

where the alternatives within each of the sets $\bar{D}(a)$, $D(a) \cap \bar{D}(b)$, and $D(b)$ can be ordered arbitrarily. This relation is complete because $a \mathrel{C} b$ and, thus, $\bar{D}(a) \cap D(b) = \emptyset$. Note that the preferences of this voter agree with all majority edges incident to a or b. Similar to McGarvey's construction (see Theorem 6.2), we now add pairs of voters until the majority relation of P is $\succ_M$. For every pair of alternatives $x, y \in A \setminus \{a, b\}$, add one voter with preferences

$$x \succ y \succ a \succ b \succ z_1 \succ \cdots \succ z_{m-4}$$

and one with preferences

$$z_{m-4} \succ \cdots \succ z_1 \succ a \succ b \succ x \succ y.$$

The majority relation of the resulting profile P is identical to $\succ_M$. Moreover, $a \succ_i b$ for all $i \in N$, contradicting Pareto-optimality of f. □

As an immediate consequence of Theorem 7.4 and Theorem 7.5, UC is the only majoritarian SCF that satisfies Pareto-optimality and expansion.

The uncovered set can be efficiently computed by first computing the covering relation for every pair of alternatives and then returning its maximal elements. The runtime of such an algorithm is in $O(|A|^3)$.

There is yet another elegant characterization of UC, which can be exploited to arrive at a more efficient algorithm.

Theorem 7.6 (Shepsle and Weingast, 1984) $\mathcal{D} = \mathcal{D}_{\text{TOUR}}$

UC returns precisely those alternatives that reach every other alternative in at most two steps, i.e.,

$$\text{UC}(A, \succ_M) = \{x \in A \colon D^2(x) = A\}.$$

Proof. The statement follows directly from the (*) statement in the proof of Theorem 7.4. □

Alternatives x for which $D^2(x) = A$ are called kings in graph theory. In general directed graphs, kings need not exist.

Theorem 7.6 can be leveraged to compute UC by recalling that the square of the adjacency matrix M tells us which alternative reaches another alternative in exactly two steps. Which alternatives can be reached in one step is encoded in the adjacency matrix, while the identity matrix I tells us which alternatives can be reached in zero steps. Hence, an alternative is uncovered iff the corresponding row in the matrix $M^2 + M + I$ does not contain a zero.

```
procedure UC(A, P)
    for all i, j ∈ A do
        if i >_M j then m_ij ← 1
        else m_ij ← 0
    M ← (m_ij)_{i,j∈A}
    U ← M^2 + M + I
    return {i ∈ A: ∀j ∈ A: u_ij ≠ 0}
```

The bottleneck of this algorithm is matrix multiplication (which does not become easier if both factors are identical). Matrix multiplication is a topic of ongoing research in computer science. The runtime of the fastest known algorithm is in $O(|A|^{2.371339})$ (Alman et al., 2025). It is based on the Coppersmith-Winograd algorithm (Coppersmith and Winograd, 1990), which in turn is based on Strassen's algorithm (Strassen, 1969), the first sub-cubic algorithm for this task. Aside from Strassen's algorithm, these methods are referred to as "galactic algorithms." This indicates that the constants hidden within the O-notation are extremely large, making these algorithms impractical for matrices encountered in real-world problems. There is no evidence that matrix multiplication is not feasible in linear time ($O(|A|^2)$).

To illustrate the algorithm, consider the following example.

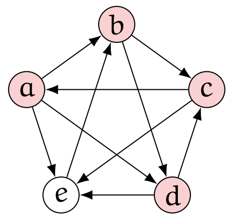

$$\begin{pmatrix} 0&1&0&1&1\\ 0&0&1&1&0\\ 1&0&0&0&1\\ 0&0&1&0&1\\ 0&1&0&0&0 \end{pmatrix}^2 + \begin{pmatrix} 0&1&0&1&1\\ 0&0&1&1&0\\ 1&0&0&0&1\\ 0&0&1&0&1\\ 0&1&0&0&0 \end{pmatrix} + \begin{pmatrix} 1&0&0&0&0\\ 0&1&0&0&0\\ 0&0&1&0&0\\ 0&0&0&1&0\\ 0&0&0&0&1 \end{pmatrix} = \begin{pmatrix} 1&2&2&2&2\\ 1&1&2&1&2\\ 1&2&1&1&2\\ 1&1&1&1&2\\ 0&1&1&1&1 \end{pmatrix}$$

The only row of the resulting matrix that contains a zero is the last one, which corresponds to alternative e. In fact, e does not reach a in at most two steps; it is covered by a. Hence, $\mathrm{UC}(A, >_M) = \{a, b, c, d\}$.

In general, it follows from the definition of the covering relation that for two alternatives $x, y \in U$, x covers y iff y does *not* reach x in at most two steps, i.e.,

$$x \mathrel{C} y \quad \Leftrightarrow \quad x \notin D^2(y) \quad \Leftrightarrow \quad y \notin \bar{D}^2(x).$$

This means that, in any majority graph, the covering edges are precisely those edges that do *not* lie on a 3-cycle. This effectively illustrates the relationship between TC and UC for majority tournaments (and why $\mathrm{UC} \subseteq \mathrm{TC}$):

- TC removes all edges that lie on a cycle and then returns the maximal elements.
- UC removes all edges that lie on a 3-cycle and then returns the maximal elements.
- The split cycle SCF (see Section 6.8) removes all edges with minimal margins from each cycle and then returns the maximal elements.

All three SCFs satisfy expansion. This is the case for every SCF that breaks majority cycles by removing edges, as long as the rule that specifies which edges ought to be removed from a cycle only depends on the preferences between alternatives contained in the cycle (Brandt and Dong, 2025). Kemeny's rule (and ranked pairs and Schulze's rule mentioned in Section 6.8) are also based on removing edges and then returning maximal elements. However, which edges will be removed is not decided locally for each cycle, and consequently, these SCFs violate expansion.

7.4 The Banks Set

In this section, we will dig deeper into the hierarchy of majoritarian SCFs by introducing yet another idea to get around the dominance relation's lack of transitivity. While a tournament may fail to be transitive, it always contains transitive subtournaments.

transitive subset

A *transitive subset* of a tournament $(A, \succ_M)$ is a set of alternatives $B \subseteq A$ such that $\succ_M|_B$ is transitive. Let $\text{Trans}(A, \succ_M) = \{B \subseteq A \colon B \text{ is transitive}\}$.

Banks set (BA)

Every alternative is the maximal element in some transitive subset (it could be just a singleton set). The *Banks set (BA)* consists of the maximal elements of all inclusion-maximal transitive subsets (Banks, 1985), i.e.,

$$\mathsf{BA}(A, \succ_M) = \{\text{Max}(B, \succ_M) \colon B \in \text{Max}(\text{Trans}(A, \succ_M), \supseteq)\}$$

When we have a transitive subset, adding an alternative that does not replace the maximal element is not critical. We can, therefore, replace inclusion-maximality with the property that a transitive set cannot be extended "from above."

$$x \in \mathsf{BA}(A, \succ_M) \quad \Leftrightarrow \quad \exists B \in \text{Trans}(A, \succ_M) \colon (x \in \text{Max}(B, \succ_M) \land (\nexists a \in A, \forall b \in B \colon a \succ_M b))$$

Whenever there is a Condorcet winner, it extends every transitive set. Hence, BA is a Condorcet extension.

Let us consider two example tournaments to get more intuition into the Banks set and to start thinking about algorithms to compute it.

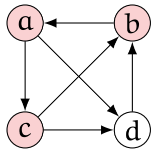

$\mathsf{BA}(A, \succ_M) = \{a, b, c\}$

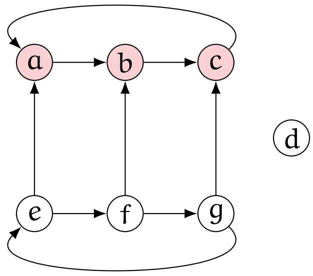

$\mathsf{BA}(A, \succ_M) = \{a, b, c\}$

All missing edges in the tournament on the right-hand side point downwards. This tournament has a non-trivial automorphism, $\sigma = \begin{pmatrix} a & b & c & d & e & f & g \\ b & c & a & d & f & g & e \end{pmatrix} = (abc)(efg)$. By relabeling the vertices according to σ, we obtain exactly the same tournament. Hence, any majoritarian SCF that returns a will also return b and c. An analogous relationship holds for alternatives e, f, and g. a is contained in the Banks set because it is the maximal element of several maximal transitive sets, such as $\{a, b, d, g\}$. d is not contained in the Banks set because every transitive set in which it is maximal can be extended. For example, $\{d, e, f\}$ can be extended by c. For comparison, in this tournament, $\mathsf{UC}(A, \succ_M) = \{a, b, c, d\}$ and $\mathsf{TC}(A, \succ_M) = \{a, b, c, d, e, f, g\}$. In fact, this is the smallest tournament in which the Banks set is different from the uncovered set.

BA can be characterized analogously to the characterization of TC and UC. To this end, we introduce a new consistency condition. A majoritarian SCF f satisfies *dominator expansion* if for all $A \in \mathcal{F}$ and $x \in A$ with $\bar{D}(x) \neq \emptyset$,

dominator expansion

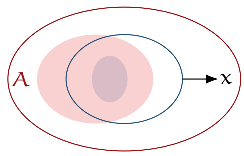

$$f(\bar{D}(x)) \subseteq f(A).$$

We have phrased this condition for majoritarian SCFs. It can be interpreted as a purely choice-theoretic condition when letting $\bar{D}(x) = \{y \in A \setminus \{x\} : y \in f(\{x, y\})\}$.

In other words, the best elements from all dominator sets have to be chosen. All alternatives that dominate x are better than x, but only the best of these need to be returned. Since choices from a feasible subset—a set of dominators—are extended to a feasible superset, this is an expansion consistency condition. For majoritarian SCFs, it is weaker than expansion.

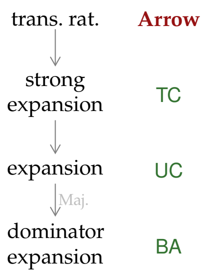

Proposition 7.5 $\mathcal{D} = \mathcal{D}_{\mathsf{TOUR}}$

A majoritarian SCF that satisfies expansion also satisfies dominator expansion.

Proof. Let f be a majoritarian SCF that satisfies expansion and $a \in A$. To show that $f(\bar{D}(a)) \subseteq f(A)$, consider an arbitrary

$$x \in f(\bar{D}(a)). \tag{1}$$

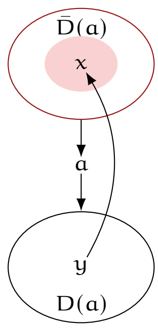

Now consider an arbitrary $y \in D(a)$. We know that $x \succ_M a$ and $a \succ_M y$. No matter whether $y \succ_M x$ or $x \succ_M y$, $x \in f(\{a, x, y\})$. In the first case, $\succ_M|_{\{a,x,y\}}$ forms a three-cycle, and all three alternatives have to be chosen by any majoritarian function. In the second case, x is a Condorcet winner in $\succ_M|_{\{a,x,y\}}$ and thus has to be chosen by expansion. In summary,

$$\forall y \in D(a)\colon x \in f(\{a, x, y\}). \tag{2}$$

Now observe that $\bar{D}(a) \cup \bigcup_{y \in D(a)} \{a, x, y\} = A$ and hence it follows from (1), (2), and repeated application of expansion that $x \in f(A)$. □

Note that Proposition 7.5 is not a purely choice-theoretic statement because it only holds for majoritarian SCFs. The argument for (2) requires that all three alternatives have to be chosen from a three-cycle. Proposition 7.5 could be turned into a choice-theoretic statement by demanding that $f(\{x, y, z\}) = \{x, y, z\}$ whenever $f(\{x, y\}) = \{x\}$, $f(\{y, z\}) = \{y\}$, and $f(\{x, z\}) = \{z\}$.

BA can be characterized using dominator expansion.

Theorem 7.7 (Brandt, 2011) $\mathcal{D} = \mathcal{D}_{\text{TOUR}}$

BA is the finest majoritarian SCF satisfying dominator expansion.

Proof. Analogous to the proofs of Theorems 7.2 and 7.4, the proof consists of two parts. We first show that for any majoritarian SCF f that satisfies dominator expansion, $\text{BA} \subseteq f$. Let $(A, \succ_M)$ be a tournament and $x \in \text{BA}(A)$. It needs to be shown that $x \in f(A)$. Since $x \in \text{BA}(A)$, there has to be some inclusion-maximal transitive subset $B \in \text{Max}(\text{Trans}(A, \succ_M), \supseteq)$ such that $\text{Max}(B, \succ_M) = \{x\}$. Let $B = \{x = x_0, x_1, \dots, x_k\}$ with $x_i \succ_M x_j$ for all $i < j$. Define $C \subseteq A$ by letting

x
x_1
x_2
$\vdots$
x_k

$$C = \bigcap_{i=1}^{k} \bar{D}(x_i) = \{x\}.$$

Now, consider the following chain of inclusions.

$$f(A) \supseteq f(\bar{D}(x_k)) \supseteq f(\bar{D}(x_k) \cap \bar{D}(x_{k-1})) \supseteq \dots \supseteq f(C) = f(\{x\}) = \{x\}.$$

The first inclusion is an immediate application of dominator expansion with respect to x_k. For the following inclusion, we consider the tournament $(\bar{D}(x_k), \succ_M|_{\bar{D}(x_k)})$ and apply dominator expansion with respect to x_{k-1}. We continue considering smaller and smaller subsets of alternatives until we reach $\bar{D}(x_k) \cap \bar{D}(x_{k-1}) \cap \dots \bar{D}(x_1) = C = \{x\}$. The sequence of inclusions thus yields $x \in f(A)$.

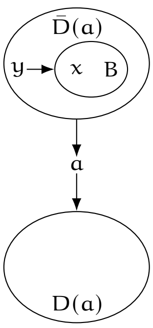

It remains to show that BA is majoritarian and satisfies dominator expansion. The former follows directly from the definition. For the latter statement, let $(A, \succ_M)$ be a tournament, $a \in A$, and $x \in \text{BA}(\bar{D}(a))$. It needs to be shown that $x \in \text{BA}(A)$. Since $x \in \text{BA}(\bar{D}(a))$, there is some inclusion-maximal transitive set $B \in \text{Max}(\text{Trans}(\bar{D}(a), \succ_M), \supseteq)$ of $\bar{D}(a)$ such that $\text{Max}(B, \succ_M) = \{x\}$. Now consider the set $C = B \cup \{a\}$. Since $B \subseteq \bar{D}(a)$, $\succ_M|_C$ is transitive. Any alternative y that dominates all alternatives in C has to be contained in $\bar{D}(a)$. Since $B \in \text{Max}(\text{Trans}(\bar{D}(a), \succ_M), \supseteq)$, no such alternative exists and $x \in \text{BA}(A)$. □

Since strong expansion implies expansion, which, in turn, implies dominator expansion,

$$\mathsf{BA} \subseteq \mathsf{UC} \subseteq \mathsf{TC}.$$

As a consequence, BA satisfies Pareto-optimality.

The question of whether the Banks set can be computed efficiently has two different answers, depending on how it is interpreted. Random alternatives in the Banks set can be found efficiently by a simple algorithm that constructs a transitive set that cannot be extended from above by iteratively adding alternatives from the intersection of dominator sets (Hudry, 2004).

```
procedure BA-Element(A, >_M)
    a ∈ A
    B ← ∅
    C ← {a}
    repeat
        a ∈ C
        B ← B ∪ {a}
        C ← ⋂_{b∈B} D̄(b)
    until C = ∅
    return a
```

This algorithm runs in linear time. Yet, computing the *entire* Banks set is NP-hard. The difficulty of this problem is rooted in the potentially exponential number of maximal transitive subtournaments.

Theorem 7.8 (Woeginger, 2003; Bachmeier et al., 2019) $\mathcal{D} = \mathcal{D}_{\mathsf{TOUR}}, n \geqslant 5$

Deciding whether a given alternative is contained in the Banks set is NP-complete.

Proof. Woeginger (2003) proved the statement by reduction from the 3-colorability of undirected graphs. We here give the arguably simpler reduction from Boolean satisfiability by Brandt et al. (2010). Whether a given subset is a maximal transitive set of a tournament can be checked in polynomial time. Hence, the problem is in NP.

Let $\varphi = \bigwedge_{i\in\{1,\dots,k\}} \bigvee_{x\in C_i} x$ be a propositional formula in conjunctive normal form, where each set C_i consists of literals, i.e., propositional variables and their negations. For example,

$$\varphi = (\neg p \vee s \vee q) \wedge (p \vee s \vee r) \wedge (p \vee q \vee \neg r).$$

For any literal x, we have $\bar{x} = \neg p$ if $x = p$, and $\bar{x} = p$ if $x = \neg p$, where p is a propositional variable. We may assume that x and $\bar{x}$ do not occur in the same clause and that, if x occurs in φ, so does $\bar{x}$.

We now construct a tournament $(A, \succ)$ with

$$A = \{c_1, \dots, c_{2k-1}\} \cup \{d\} \cup U_1 \cup \dots \cup U_{2k-1},$$

where for $i \in \{1,\dots,2k-1\}$,

$$U_i = \begin{cases} \{(i,x) \colon x \in C_{\frac{i+1}{2}}\} & \text{if } i \text{ is odd} \\ \{u_i\} & \text{if } i \text{ is even} \end{cases}.$$

The majority relation is defined such that $U_i \in \text{Trans}(A, \succ_M)$ for all $i \in \{1,\dots,2k-1\}$. Moreover, for $(i,x) \in U_i$ and $(j,y) \in U_j$ with $i < j$, we have $(i,x) \succ_M (j,y)$ iff $x \neq \bar{y}$. For the majority relation on the remaining alternatives, the reader is referred to the figure showing the tournament induced by the example formula φ. Omitted edges point downwards.

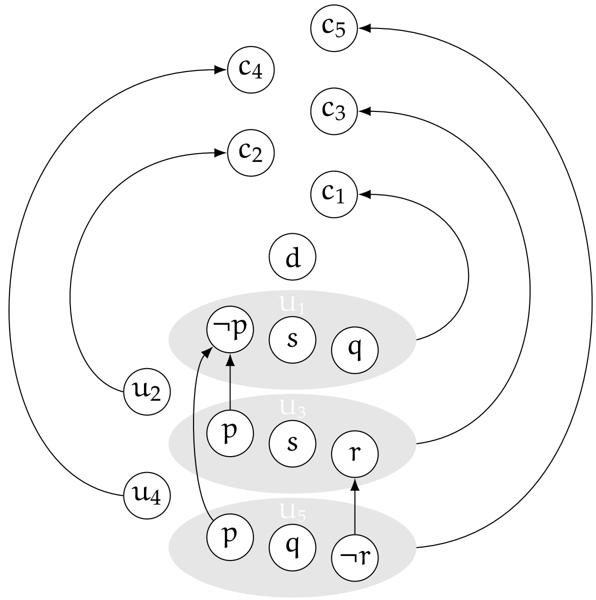

Observe that for every maximal transitive set $B \in \text{Max}(\text{Trans}(A, \succ_M), \supseteq)$ with $\text{Max}(B, \succ_M) = \{d\}$, it holds that

(i) $B \cap U_i \neq \emptyset$ for all $i \in \{1,\dots,2k-1\}$, i.e., B contains an alternative from each level below d, and

(ii) for all $(i,x),(j,y) \in B$, $x \neq \bar{y}$, i.e., B contains no upward edges.

For (i), assume that $B \cap U_i = \emptyset$. By assumption, $\text{Max}(B, \succ_M) = \{d\}$, which implies that $B \cap \{c_1,\dots,c_{2k-1}\} = \emptyset$. However, c_i dominates all alternatives in B, contradicting that $B \in \text{Max}(\text{Trans}(A, \succ_M), \supseteq)$. For (ii), assume $(i,x),(j,y) \in B$ with $x = \bar{y}$. Without loss of generality, we may assume that $i < j$. (i) implies that $u_{i+1} \in B$. Then, $(i,x) \succ_M u_{i+1} \succ_M (j,y) \succ_M (i,x)$ contradicts that $B \in \text{Trans}(A, \succ_M)$.
The previous insights imply that there is a one-to-one correspondence between sets $B \in \text{Max}(\text{Trans}(A, \succ_M), \supseteq)$ with $\text{Max}(B, \succ_M) = \{d\}$ and satisfying assignments of φ. Any such set B corresponds to an assignment in which literal x is set to true iff $(i,x) \in B$ for some $i \in \{1,\dots,2k-1\}$. By virtue of (ii), this assignment is well-defined (as every

variable is either true or false), and with (i) it follows that this assignment satisfies φ (as every clause contains at least one positive literal).
Hence,

$$\varphi \text{ is satisfiable} \quad \Leftrightarrow \quad d \in \mathsf{BA}(A, >_M).$$

Bachmeier et al. (2019) proved that all the tournaments appearing in the reduction for a reduced set of propositional formulas, for which deciding satisfiability is still NP-complete, can be induced by only five voters. □

It is open whether computing the Banks set remains NP-hard when $n = 3$. The computational state of affairs for the Banks set can be summarized as follows: Given a tournament, it is easy to find an alternative in the Banks set. But when somebody pinpoints a particular alternative and inquires whether it belongs to the Banks set, there is no known efficient algorithm to provide an answer to this question.

BA (and all its coarsenings) satisfy a strengthening of Condorcet-consistency. Not only do they select a Condorcet winner whenever it exists, but they also never select a single alternative in the absence of a Condorcet winner. An SCF f is a *strong Condorcet extension* if $f(A, P) = \{x\}$ iff x is a Condorcet winner.

strong Condorcet extension

Proposition 7.6 $\mathcal{D} = \mathcal{D}_{\mathsf{TOUR}}$

BA, UC, and TC are strong Condorcet extensions.

Proof. We have already shown that all three SCFs are Condorcet extensions. It remains to show that if $f(A, P) = \{x\}$, x has to be a Condorcet winner. When $\mathsf{BA}(A, >_M) = \{x\}$, every transitive subset has to be eventually extended by x. Hence, x has to dominate all other alternatives and thus be a Condorcet winner. When $f(A, P) = \{x\}$ for any coarsening of BA, $\{x\} = \mathsf{BA}(A, >_M)$ and x has to be a Condorcet winner as well. □

This concludes the escape route of characterizations using weakenings of strong expansion. The hierarchy of majoritarian SCFs and their corresponding expansion consistency conditions are shown below.

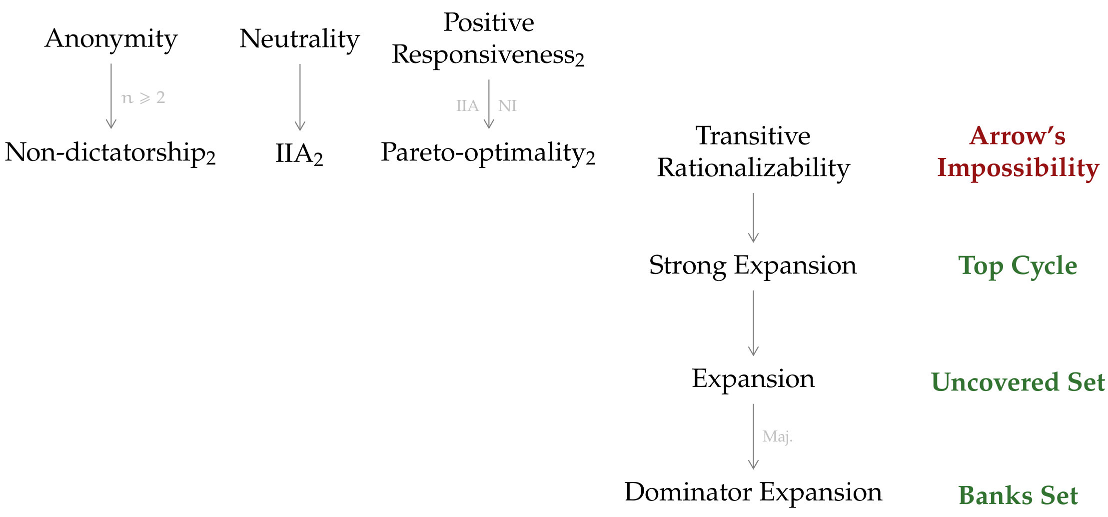

7.5 Set-Rationalizability and Stability

In this section, we will explore an alternative escape route from the Arrovian impossibilities by redefining the conditions for choice consistency. We will consider the entire set of chosen alternatives, rather than its individual elements. Choice functions yield sets of alternatives. Yet, rationality and consistency conditions are defined in terms of single alternatives, and rationalizing relations, such as the base relation, are defined on single alternatives.

We now define a notion of rationalizability for sets, which is based on a rationalizing relation, defined on sets, such that every choice set is undominated within its feasible set.

set-rationalizability

A choice function f is *set-rationalizable* if there exists $\succsim \subseteq \mathcal{P}^*(U) \times \mathcal{P}^*(U)$ such that for all $A, X \in \mathcal{P}^*(U)$,

$$X = f(A) \quad \Leftrightarrow \quad X \in \text{Max}(\mathcal{P}^*(A), \succsim).$$

Note that the choice set X is not only undominated within the feasible set A, but it is also the *unique* dominant element.

extended base relation

The base relation can be extended to sets by letting

$$X \succsim_f Y \quad \Leftrightarrow \quad X = f(X \cup Y),$$

i.e., X is weakly preferred to Y iff X is chosen in the presence of Y. Clearly, the extended base relation is antisymmetric but not complete. To illustrate the definitions, consider the following choice function f for $U = \{a, b, c\}$ and its extended base relation $\succsim_f$.

A	f(A)
ab	a
bc	b
ac	c
abc	abc

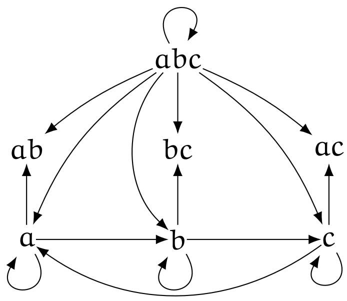

The previous example already shows the potential of the definition of set-rationalizability. The choice function corresponds to the classic Condorcet profile, in which the majority relation cycles. This choice function is not rationalizable, but it is set-rationalizable.

The contraction and expansion consistency conditions can also be redefined in terms of sets of alternatives. For comparison, we first rewrite the definitions of contraction and expansion from Chapter 2. Let $A, B \in \mathcal{F}$ and $x \in A \cap B$.

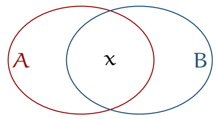

Contraction:	$x \in f(A \cup B)$	$\Rightarrow$	$x \in f(A)$ and $x \in f(B)$
Expansion:	$x \in f(A \cup B)$	$\Leftarrow$	$x \in f(A)$ and $x \in f(B)$

We now propose set-based versions of these axioms by syntactically replacing "$x \in$" with "$X =$."

Let $A, B \in \mathcal{F}$ and $X \subseteq A \cap B$.

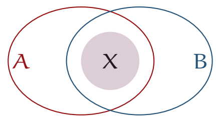

Set-contraction:	$X = f(A \cup B)$	$\Rightarrow$	$X = f(A)$ and $X = f(B)$
Set-expansion:	$X = f(A \cup B)$	$\Leftarrow$	$X = f(A)$ and $X = f(B)$

Note that the set-based versions coincide with their traditional counterparts when restricting attention to resolute choice functions.

Set-contraction can be rewritten as a choice condition that requires that removing unchosen alternatives does not affect the choice set.

Lemma 7.2

Let f be a choice function. Then,

$$f \text{ satisfies set-contraction} \quad \Leftrightarrow \quad f(V) = f(W) \text{ for all } V, W \in \mathcal{F} \text{ with } f(V) \subseteq W \subseteq V.$$

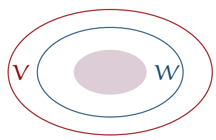

Proof.

$\Rightarrow$ Let $f(V) \subseteq W \subseteq V$. Then, $X = f(V \cup W) = f(V)$ and $X = f(W)$.

$\Leftarrow$ Let $X \subseteq A \cap B$ and $X = f(A \cup B)$. Then,

$$f(A \cup B) \subseteq A \subseteq A \cup B \quad \Rightarrow \quad f(A \cup B) = f(A) \text{ and}$$
$$f(A \cup B) \subseteq B \subseteq A \cup B \quad \Rightarrow \quad f(A \cup B) = f(B).$$

□

We thus arrive at the following definitions.

set-contraction A choice function f satisfies *set-contraction* if for all $A, B \in \mathcal{F}$,

$$f(A) \subseteq B \subseteq A \quad \Rightarrow \quad f(A) = f(B).$$

In voting, one often refers to the "spoiler effect" as the phenomenon that occurs when a losing candidate influences who wins the election. This is precisely ruled out by set-contraction, as the choice set is completely unaffected by unchosen alternatives. If f satisfies set-contraction, it is also idempotent in the sense that $f \circ f = f$ and repeated application of the SCF does not further narrow down the choice set.

set-expansion A choice function f satisfies *set-expansion* if for all $A, B \in \mathcal{F}$,

$$f(A) = f(B) \quad \Rightarrow \quad f(A \cup B) = f(A).$$

Set-contraction demands that any set that is chosen from some feasible set has to be chosen in every feasible subset in which it is contained. Set-expansion requires that any set that is chosen from two feasible sets has to be chosen from their union.

Similar to Theorem 2.1, there is a strong relationship between consistency and rationalizability. (Note, however, that, in contrast to Theorem 2.1, set-expansion is not required.)

Theorem 7.9 (Brandt and Harrenstein, 2011)

Let f be a choice function.

$$f \text{ is set-rationalizable} \quad \Leftrightarrow \quad f \text{ satisfies set-contraction.}$$

Proof.

$\Rightarrow$ Let f be a set-rationalizable SCF and $\succsim$ the witnessing binary relation on sets. Now consider $V, W \in \mathcal{F}$ such that $f(V) \subseteq W \subseteq V$. It needs to be shown that $f(W) = f(V)$. As $\succsim$ set-rationalizes f, there is no $Y \subseteq V$ such that $Y \succ X$. Moreover, as $W \subseteq V$, there is no $Y \subseteq W$ such that $Y \succ X$. Therefore, $X \in Max(\mathcal{P}^*(W), \succsim)$ and, by set-rationalizability, $X = f(W)$.

$\Leftarrow$ Let f be an SCF that satisfies set-contraction. We will show that the extended base relation $\succsim_f$ set-rationalizes f, i.e., for all $A \in \mathcal{F}$, $Max(\mathcal{P}^*(W), \succsim) = \{f(A)\}$. To this end, consider $Y \in \mathcal{P}^*(A) \setminus \{f(A)\}$. Then, $f(A) \subseteq f(A) \cup Y \subseteq A$ and set-contraction implies that $f(A) = f(f(A) \cup Y)$. By definition of $\succsim_f$, $f(A) \succsim_f Y$. Moreover, since $\succsim_f$ is antisymmetric, $f(A) \succ_f Y$. Hence, $Max(\mathcal{P}^*(A), \succsim) = \{f(A)\}$ and $\succsim_f$ set-rationalizes f.

□

stability We will say that a choice function is *stable* if it satisfies set-contraction and set-expansion. Such a function is stable in the sense that neither adding nor removing unchosen alternatives affects the choice set.

Proposition 7.7

Quasi-transitive rationalizability implies stability.

Proof. Exercise. □

Hence, we have the following implications:

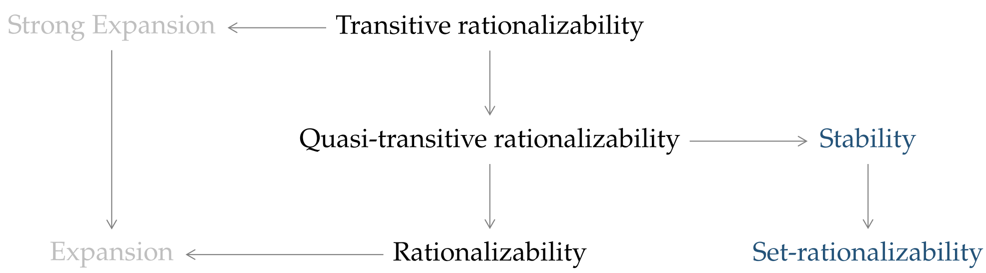

A natural concern at this point may be whether any interesting set-rationalizable or even stable SCFs, except for the trivial SCF that always returns all alternatives, exist. As a matter of fact, almost all of the SCFs we have studied so far violate set-rationalizability. The profile in the margin shows this for all non-trivial monotonic scoring rules. Every such rule will choose $\{a\}$ from $\{a, b, c\}$ but $\{a, b\}$ from $\{a, b\}$. The same is true for Baldwin's rule, Black's rule, Kemeny's rule, maximin, Young's rule, and Copeland's rule. Moreover, similar—albeit more complicated examples—can be constructed for UC and BA.

3	2	1
a	b	c
c	a	b
b	c	a

However, TC satisfies set-contraction and set-expansion and thus is stable (Exercise 7.13)! In the next section, we will introduce an appealing stable refinement of the uncovered set.

7.6 The Bipartisan Set

The last majoritarian SCF we consider strikes a pleasing balance between discriminative power, desirable axiomatic properties, and efficient computability.

In order to define this SCF, we need to introduce probability distributions over alternatives. Let $p\colon A \to [0,1]$ with $\sum_{x\in A} p(x) = 1$ be a *probability distribution* that assigns probabilities to alternatives in tournament $(A, \succ_M)$. By $\Delta(A)$, we denote the set of all probability distributions over A.

probability distribution

$\Delta(A)$

The *value* of distribution p against alternative $x \in A$ is the difference between the probability that p picks an alternative that dominates x and the probability that p picks an alternative that is dominated by x. Formally,

value of a distribution

$$v_p(x) = \sum_{y\in \bar{D}(x)} p(y) - \sum_{y\in D(x)} p(y).$$

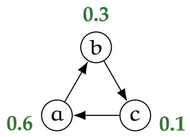

In general, values thus lie within the interval $[-1, 1]$. If $\nu_p(x) \geqslant 0$, then the alternative selected by p is at least as likely to dominate x as to be dominated by x. For the three-cycle tournament given in the margin and the distribution p denoted by the green numbers, we have that $\nu_p(a) = 0.1 - 0.3 = -0.2$, $\nu_p(b) = 0.6 - 0.1 = 0.5$, and $\nu_p(c) = 0.3 - 0.6 = -0.3$.

We first show that the sum of all values, weighted by p, is 0. This will be useful for the further analysis of distributions.

Lemma 7.3

Let $(A, \succ_M)$ be a tournament and $p \in \Delta(A)$. Then, $\sum_{x \in A} p(x) \cdot \nu_p(x) = 0$.

Proof.

$$\sum_{x \in A} p(x) \cdot \nu_p(x) \overset{\text{Def. of } \nu_p}{=} \sum_{x \in A} p(x) \cdot \left(\sum_{y \in \bar{D}(x)} p(y) - \sum_{y \in D(x)} p(y) \right)$$
$$= \sum_{x,y \in A \colon y \succ_M x} p(x) \cdot p(y) - \sum_{x,y \in A \colon x \succ_M y} p(x) \cdot p(y) = 0.$$

The second equality follows from the fact that every outgoing edge of an alternative is an incoming edge for another alternative. The argument is thus reminiscent of the handshaking lemma in graph theory. □

In the example given above, we have $-0.6 \cdot 0.2 + 0.3 \cdot 0.5 - 0.1 \cdot 0.3 = -0.12 + 0.15 - 0.03 = 0$. In this example, it also holds that $\sum_{x \in A} \nu_p(x) = 0$, but this is only true if $(A|_{supp(p)}, \succ_M)$ is regular, i.e., every alternative in $(A|_{supp(p)}, \succ_M)$ has the same Copeland score.

Lemma 7.3 implies that every distribution does well against some alternative.

Corollary 7.1

For every tournament $(A, \succ_M)$ and $p \in \Delta(A)$, there exists $x \in A$ such that $\nu_p(x) \geqslant 0$.

Proof. Since the convex combination $\sum_{x \in A} p(x) \cdot \nu_p(x)$ is non-negative (in fact, it is 0), the value $\nu_p(x)$ has to be non-negative for at least one $x \in A$. □

optimal distribution

We are interested in lotteries that perform well against *all* alternatives. To this end, we say that a distribution p is *optimal* if $\nu_p(x) \geqslant 0$ for all $x \in A$. In other words, p beats or ties every alternative in expectation, which means that p can be understood as a randomized weak Condorcet winner.

An optimal distribution p can be viewed as the solution of a homogeneous system of linear inequalities with integer coefficients. Since the vertices of a rational polyhedron are rational, this implies that if an optimal distribution exists, there is also a rational-valued optimal distribution. Moreover, by homogeneity, one can multiply all fractions by a common denominator and obtain a non-zero solution in non-negative integers.

Lemma 7.3 entails that an optimal distribution has value 0 against all alternatives in its support. In other words, p "ties" against every alternative in its support.

Corollary 7.2

For every tournament $(A, >_M)$ and optimal distribution $p \in \Delta(A)$,

$$p(x) > 0 \quad \Rightarrow \quad \nu_p(x) = 0.$$

Proof. The equality from Lemma 7.3 shows that the sum of the products of probabilities and values is zero. Probabilities are non-negative by definition. The value of p against any x is also non-negative when p is optimal. We thus have

$$\sum_{x \in A} \underbrace{p(x)}_{\geqslant 0} \cdot \underbrace{\nu_p(x)}_{\geqslant 0} = 0.$$

□

Remarkably, optimal distributions always exist, and they are unique in the absence of majority ties. While Corollary 7.1 has established that

$$\forall p \in \Delta(A)\, \exists x \in A \colon \nu_p(x) \geqslant 0,$$

we now show that

$$\exists p \in \Delta(A)\, \forall x \in A \colon \nu_p(x) \geqslant 0,$$

and, moreover, there is a unique such p.

Theorem 7.10 (Laffond et al., 1993; Fisher and Ryan, 1995a)

Every tournament admits a unique optimal probability distribution.

Proof. **Existence of optimal distributions** follows from the Minimax Theorem (von Neumann, 1928). In order to have a self-contained presentation, we prove the existence of optimal distributions by induction on $|A|$. In contrast to standard proofs of the Minimax Theorem, which are based on separating hyperplane arguments, fixed-point theorems, LP duality, or Farkas' lemma, this one only uses elementary mathematics. The *base case* is trivial: if $A = \{x\}$, then the only possible distribution is the one with $p(x) = 1$. This distribution is optimal since $\bar{D}(x) = D(x) = \emptyset$ and $\nu_p(x) = 0$.
Now let $|A| > 1$. By the *induction hypothesis*, we can assume that for all $B \subset A$ with $|B| = |A| - 1$, there is an optimal distribution for $(B, >_M|_B)$. Let

$$p \in \arg \max_{p \in \Delta(A)} \min_{x \in A} \nu_p(x). \qquad (*)$$

Such a p exists because $\nu_p(x)$ is linear in p. Observe that an optimal distribution for $(A, >_M)$ exists iff $\nu_{\min} = \min_{x \in A} \nu_p(x) = 0$. We assume for contradiction that $\nu_{\min} < 0$.

Corollary 7.1 has shown that there is some $z \in A$ such that $v_p(z) \geqslant 0$. Moreover, by the induction hypothesis, there is an optimal distribution q for $(A \setminus \{z\}, \succ_M|_{A\setminus\{z\}})$, i.e. there is

$$q \in \Delta(A) \text{ with } q(z) = 0 \text{ and } v_q(x) \geqslant 0 \text{ for all } x \in A \setminus \{z\}.$$

Now define

$$r = (1-\varepsilon)\cdot p + \varepsilon \cdot q \text{ for } 0 < \varepsilon \leqslant 1.$$

Then, for all $x \in A \setminus \{z\}$,

$$v_r(x) \geqslant (1-\varepsilon)\cdot v_{\min} + \varepsilon \cdot 0 = (1-\varepsilon)\cdot v_{\min} > v_{\min}.$$

Furthermore, for small enough ε ($\varepsilon < -v_{\min}/|v_q(z)|$),

$$v_r(z) = \underbrace{(1-\varepsilon)\cdot v_p(z)}_{\geqslant 0} + \varepsilon \cdot v_q(z) > v_{\min}.$$

As a consequence, $\min_{x\in A} v_r(x) > v_{\min}$, which contradicts (*): p was chosen as to maximize $v_{\min} = \min_{x\in A} v_p(x)$.
To show the **uniqueness of optimal distributions**, assume for contradiction that there is a tournament $(A, \succ_M)$ and two optimal distributions $p, q \in \Delta(A)$ with $p \neq q$. We may assume without loss of generality that $\mathrm{supp}(p) = \mathrm{supp}(q) =: B$. Otherwise, let $p' = \lambda_1 \cdot p + (1-\lambda_1)\cdot q$ and $q' = \lambda_2 \cdot p + (1-\lambda_2)\cdot q$ for $\lambda_1, \lambda_2 \in (0,1)$ with $\lambda_1 \neq \lambda_2$. Then, $\mathrm{supp}(p') = \mathrm{supp}(q')$ and both p' and q' are optimal. In general, the definition of optimality immediately implies that the convex combination of two optimal lotteries is optimal.
Now let $r(x) = p(x) - q(x)$ for all $x \in B$. Even though r is not a distribution, let us consider the value of r by considering the following straightforward extension of the definition of the value. For all $x \in B$,

$$\begin{aligned} v_r(x) &= \sum_{y\in\bar{D}(x)} r(y) - \sum_{y\in D(x)} r(y) \overset{\text{Def. of } r}{=} \sum_{y\in\bar{D}(x)} \big(p(y)-q(y)\big) - \sum_{y\in D(x)} \big(p(y)-q(y)\big) \qquad \text{(i)}\\ &= \left(\sum_{y\in\bar{D}(x)} p(y) - \sum_{y\in D(x)} p(y)\right) - \left(\sum_{y\in\bar{D}(x)} q(y) - \sum_{y\in D(x)} q(y)\right) = \underbrace{v_p(x)}_{\overset{\text{Cor. 7.2}}{=}0} - \underbrace{v_q(x)}_{\overset{\text{Cor. 7.2}}{=}0}\\ &= 0 \end{aligned}$$

Moreover, since both p and q are distributions,

$$\sum_{x\in B} r(x) \overset{\text{Def. of } r}{=} \sum_{x\in B} \big(p(x)-q(x)\big) = \sum_{x\in B} p(x) - \sum_{x\in B} q(x) = 1 - 1 = 0 \qquad \text{(ii)}$$

(i) and (ii) define a homogeneous system of linear equations. Since $p \neq q$, this system has a non-zero solution r. Moreover, since all coefficients are integers, the system has to have a solution in integers. We can assume without loss of generality that there is

an integer solution r^* such that $r^*(b)$ is odd for some $b \in B$. For any integer solution r for which this is not the case, we can repeatedly divide $r(x)$ by 2 for all $x \in B$ until the desired property is satisfied.
Now, let $x \in B$. Then,

$$0 \overset{(ii)}{=} \sum_{y \in B} r^*(y) \overset{A=\{x\}\cup\bar{D}(x)\cup D(x)}{=} r^*(x) + \sum_{y\in\bar{D}(x)} r^*(y) - \underbrace{\sum_{y\in D(x)} r^*(y)}_{\overset{(i)}{=}\sum_{y\in\bar{D}(x)} r^*(y)} = r^*(x) + \underbrace{2\cdot \sum_{y\in\bar{D}(x)} r^*(y)}_{\text{even}}.$$

As a consequence, $r^*(x)$ has to be even for all $x \in B$, which contradicts the fact that $r^*(b)$ is odd. Hence, our initial assumption that $p \neq q$ was incorrect. □

Optimal distributions are closely connected to optimal mixed strategies in symmetric zero-sum games. The optimal distribution of a tournament corresponds to the unique Nash equilibrium (or maximin strategy) of the zero-sum game given by the skew-adjacency matrix of the tournament.

The *bipartisan set* (BP) is defined as the support of the optimal distribution, i.e.

bipartisan set (BP)

$$\text{BP}(A, \succ_M) = \{x \in A\colon p(x) > 0, p \text{ is optimal for } (A, \succ_M)\}.$$

The name "bipartisan set" is derived from a specific interpretation of the zero-sum game given by the skew-adjacency matrix of the tournament (see Laslier, 2000). In this interpretation, the alternatives correspond to policies that parties can adopt, and the voters have preferences over these policies. There are two political parties that can take mixed positions on the given policies. Each party aims at maximizing the probability that a majority of voters prefer their *ex post* policy to that of the other party. The optimal distribution corresponds to the optimal mixed strategy of each party and guarantees the same winning probability to each party.[6]

The bipartisan set and optimal distributions satisfy a number of interesting properties.

Theorem 7.11 (Laffond et al., 1993)

Let $(A, \succ_M)$ be a tournament and $p \in \Delta(A)$ the optimal distribution of $(A, \succ_M)$. Then,

(i) for all $x \in \text{BP}(A, \succ_M)$, $p(x)$ is a quotient of odd numbers,

(ii) $|\text{BP}(A, \succ_M)|$ is odd, and

(iii) for all $x \in A$, $\quad p(x) > 0 \quad \Leftrightarrow \quad v_p(x) = 0$.

[6]The uniqueness shown in Theorem 7.10 is of independent interest in game theory. Every tournament game, such as rock-paper-scissors, has a unique optimal mixed strategy. This does not hold when ties are allowed.

Proof. As to *(i)*, we know from Corollary 7.2 that

$$v_p(x) = \sum_{y\in \bar{D}(x)} p(y) - \sum_{y\in D(x)} p(y) = 0 \quad \text{for all } x \in \mathsf{BP}(A, \succ_M). \tag{*}$$

Since $\sum_{x\in \mathsf{BP}(A,\succ_M)} p(x) = 1$, this homogeneous system of linear equations has a non-zero solution p. Moreover, since all coefficients are integer numbers, the system has to have a solution in integer numbers. We can assume without loss of generality that there is an integer solution p^* such that $p^*(z)$ is odd for some $z \in \mathsf{BP}(A, \succ_M)$. (Otherwise, repeatedly divide $p^*(x)$ by 2 for all $x \in A$ until the desired property is satisfied.)
For any $x \in \mathsf{BP}(A, \succ_M)$,

$$\sum_{y\in A} p^*(y) \overset{A=\{x\}\cup\bar{D}(x)\cup D(x)}{=} p^*(x) + \sum_{y\in\bar{D}(x)} p^*(y) + \underbrace{\sum_{y\in D(x)} p^*(y)}_{\overset{(*)}{=}\sum_{y\in\bar{D}(x)} p^*(y)} = p^*(x) + \underbrace{2\cdot \sum_{y\in\bar{D}(x)} p^*(y)}_{\text{even}}.$$

Hence, $\sum_{y\in A} p^*(y)$ and $p^*(x)$ have to have same parity for any $x \in \mathsf{BP}(A, \succ_M)$. Since we know that $p^*(z)$ is odd, $\sum_{y\in A} p^*(y)$ has to be odd as well, and, moreover, $p^*(x)$ is odd for every $x \in \mathsf{BP}(A, \succ_M)$. This proves statement *(i)*.
$\sum_{y\in \mathsf{BP}(A,\succ_M)} p^*(y)$ is the odd sum of odd integers. Hence, the number of summands, and consequently $|\mathsf{BP}(A, \succ_M)|$ has to be odd. This proves statement *(ii)*.
As to *(iii)*, recall that the direction from left to right was shown in Corollary 7.2. In order to show the direction from right to left, let $x \in A$ with $v_p(x) = 0$. This implies that $\sum_{y\in\bar{D}(x)} p^*(y) = \sum_{y\in D(x)} p^*(y)$. Hence,

$$\underbrace{\sum_{y\in A} p^*(y)}_{\text{odd}} = p^*(x) + \underbrace{2\cdot \sum_{y\in D(x)} p^*(y)}_{\text{even}}$$

and consequently, $p^*(x)$ has to be odd and $p^*(x) > 0$. □

In other words, the bipartisan set always consists of an odd number of alternatives, all non-zero probabilities of the optimal distribution are "odd", and the optimal distribution strictly "beats" every alternative x not contained in its support (in the sense that $v_p(x) > 0$).[7]

Let us consider two example tournaments to get a better understanding of optimal distributions and their supports. The green numbers are the probabilities associated with the alternatives in the optimal distributions.

[7] Statement *(iii)* of Theorem 7.11 also follows from the uniqueness of optimal distributions (Theorem 7.10) and results from game theory, such as the equalizer theorem (e.g., Raghavan, 1994, p. 740) or the existence of quasi-strict equilibrium in zero-sum games (Jansen, 1981).

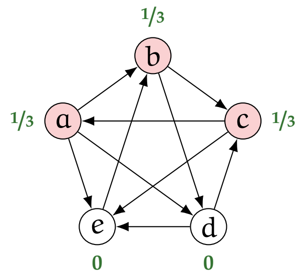

$\mathsf{BP}(A, \succ_M) = \{a, b, c\}$

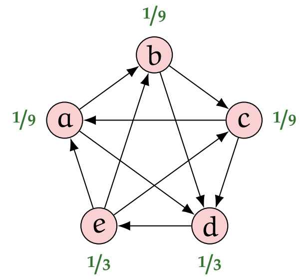

$\mathsf{BP}(A, \succ_M) = \{a, b, c, d, e\}$

In the first example, the value of p against a is $v_p(a) = 1/3 - (1/3 + 0 + 0) = 0$ while the value of p against d is $v_p(d) = (1/3 + 1/3) - (1/3 + 0) = 1/3$. In the second example, $v_p(a) = (1/9 + 1/3) - (1/9 + 1/3) = 0$ and $v_p(d) = (1/9 + 1/9 + 1/9) - 1/3 = 0$.

Like BA, BP is a refinement of UC (and thus of TC).

Proposition 7.8 $\mathcal{D} = \mathcal{D}_{\mathsf{TOUR}}$

BP $\subseteq$ UC.

Proof. Let $(A, \succ_M)$ be a tournament, p an optimal distribution, and $x, y \in A$ two alternatives such that x covers y. Recall that x covering y has immediate consequences on the dominions and dominators of x and y.

$$x \mathrel{C} y \overset{\text{Def. of C}}{\Leftrightarrow} D(y) \subseteq D(x) \wedge \bar{D}(x) \subseteq \bar{D}(y). \tag{*}$$

This entails that $v_p(y) \geqslant v_p(x)$, in particular

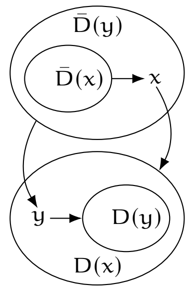

$$\begin{aligned} v_p(y) - v_p(x) &\overset{\text{Def. of } v_p}{=} \sum_{z \in \bar{D}(y)} p(z) - \sum_{z \in D(y)} p(z) - \sum_{z \in \bar{D}(x)} p(z) + \sum_{z \in D(x)} p(z) \\ &\overset{(*)}{=} \sum_{z \in \left(\bar{D}(y) \setminus \bar{D}(x)\right) \cup \left(D(x) \setminus D(y)\right)} p(z) \geqslant 0. \end{aligned}$$

Hence,

$$v_p(x) \leqslant v_p(y) - \sum_{z \in \left(\bar{D}(y) \setminus \bar{D}(x)\right) \cup \left(D(x) \setminus D(y)\right)} p(z).$$

Now assume for contradiction that $y \in \mathsf{BP}(A, \succ_M)$. Per definition of BP, we thus have $p(y) > 0$. Corollary 7.2 implies that $v_p(y) = 0$. Since $y \in D(x) \setminus D(y)$ and $p(y) > 0$, $v_p(x) \leqslant -p(y) < 0$, which contradicts the optimality of p. □

BP is much more discriminating than TC, UC, and BA. For these SCFs, the average number of chosen alternatives over all labeled tournaments goes to infinity as the number of alternatives increases (Fey, 2008). BP, on the other hand, chooses exactly half of the alternatives on average when averaging over all labeled tournaments of the same fixed size (Fisher and Reeves, 1995). Of course, for realistic distributions of preferences, tournaments are not uniformly distributed, and all of the mentioned majoritarian SCFs are much more discriminating. For reasonable parametrizations of standard preference distribution models (including impartial culture) and at least ten alternatives, all common majoritarian SCFs, except TC, discard at least 75% of the alternatives on average (Brandt and Seedig, 2016).

Like TC, UC, and BA, BP is strongly Condorcet-consistent.

Proposition 7.9 $\mathcal{D} = \mathcal{D}_{\mathsf{TOUR}}$

BP is a strong Condorcet extension.

Proof. BP is a Condorcet extension, because it refines the Condorcet extension UC. For strong Condorcet-consistency, let p be the optimal distribution of $(A, \succ_M)$ and $\mathrm{BP}(A, \succ_M) = \{x\}$. Then, $p(x) = 1$ and $p(y) = 0$ for all $y \in A \setminus \{x\}$. If there was some $y \in A$ with $y \succ_M x$, then $\nu_p(y) = -1 < 0$, contradicting the optimality of p. Hence, $x \succ_M y$ for all $y \in A \setminus \{x\}$, and x is a Condorcet winner. □

The bipartisan set and the Banks can be disjoint, as shown by Brandt and Grundbacher (2026) via a tournament with 36 alternatives. This tournament also has the remarkable property that all alternatives in the Banks set have a below-average Copeland score. By contrast, it follows from a result by Laffond et al. (1993) that at least one alternative in the bipartisan set has an average Copeland score or higher.

As mentioned at the beginning of the section, BP satisfies several desirable properties. Two of them are stability and monotonicity.

Proposition 7.10 $\mathcal{D} = \mathcal{D}_{\mathsf{TOUR}}$

BP is stable and monotonic.

Proof. Stability of BP follows relatively straightforwardly from the definition and uniqueness of the optimal distributions. When removing an alternative $x \in A$ with $p(x) = 0$ from tournament $(A, \succ_M)$, p remains optimal in $(A \setminus \{x\}, \succ_M)$ (set-contraction). Similarly, if p is optimal in both $(A, \succ_M)$ and $(B, \succ_M)$, then $\mathrm{supp}(p) \subseteq A \cap B$ and p remains optimal in $(A \cup B, \succ_M)$ (set-expansion).
Showing that BP is monotonic is a bit trickier. Let $P, P' \in \mathcal{D}_{\mathsf{TOUR}}$, $A \in \mathcal{F}$, $x \in \mathrm{BP}(A, \succ_M)$, $y \in A$, and $i \in N$ such that $P_{-i} = P'_{-i}$, $y \succ_i x$ and $x \succ'_i y$. Let p and p' be the optimal distributions in $(A, \succ_M)$ and $(A, \succ'_M)$, respectively. We will show that $p = p'$, which, of course, implies that $x \in \mathrm{BP}(A, \succ'_M)$. If $y \notin \mathrm{BP}(A, \succ_M)$, the optimality of p is unaffected by strengthening x against y (with $p(y) = 0$), and $p = p'$. If $y \in \mathrm{BP}(A, \succ_M)$, assume for contradiction that $x \notin \mathrm{BP}(A, \succ'_M)$ and consider two cases. If $y \in \mathrm{BP}(A, \succ'_M)$, similar to the previous argument, the optimality of p' in $(A, \succ'_M)$ is unaffected by strengthening y against x (with $p'(x) = 0$), and $p = p'$, contradicting $p(x) > 0$ and $p'(x) = 0$. If $y \notin \mathrm{BP}(A, \succ'_M)$, set-contraction implies that $\mathrm{BP}(A, \succ'_M) = \mathrm{BP}(A \setminus \{y\}, \succ'_M)$, which is at variance with $p(y) > 0$ and $p'(y) = 0$. □

It is tempting to define a refinement of BP that only returns the alternatives with *maximal* probability in the optimal distribution. However, such an SCF would violate several of the key properties of BP, such as set-contraction, set-expansion, and monotonicity.

BP can be computed in polynomial time by solving a linear feasibility program.

Interior-point methods achieve polynomial runtime for linear feasibility. More precisely, the runtime of state-of-the-art methods is $O(\sqrt{n_v} \cdot n_c \cdot n_v)$ where n_v is the number of variables and n_c the number of constraints. For computing BP where $n_v = m$ and $n_c = O(m^2)$ we thus get a runtime in $O(m^{3.5})$.

procedure $\mathsf{BP}(A, \succ_M)$

$$p \in \{p \in \mathbb{Q}^A : \begin{array}{ll} v_p(x) \geqslant 0 & \forall x \in A \\ \sum_{x \in A} p(x) = 1 & \\ p(x) \geqslant 0 & \forall x \in A\} \end{array}$$

return $\{x \in A : p(x) > 0\}$

As mentioned above, BP selects half of the alternatives on average, where the average is taken over all labeled tournaments of the same fixed size. For two majoritarian SCFS f and f', f is *more discriminating* than f' if there exists k such that f selects fewer alternatives than f' on average, where the average is taken over all labeled tournaments of size k. discriminability

Theorem 7.12 (Brandt et al., 2018b) $\mathcal{D} = \mathcal{D}_{\mathsf{TOUR}}$

There is no more discriminating stable majoritarian SCF than BP.

However, BP is not *the* finest majoritarian SCF satisfying stability. No such SCF exists.

BP can be characterized as the unique most discriminating majoritarian SCF satisfying stability, monotonicity, and two further axioms. For up to 7 alternatives, the two additional axioms are not required. It is unknown whether they are required in general.[8]

[8] Kiefer (2024) has constructed a stable, monotonic, and majoritarian SCF for up to 8 alternatives, which is different from BP but as discriminating as BP. Whether the definition of this SCF can be extended to larger numbers of alternatives while preserving all these properties is unclear.

7.7 Key Takeaways

Majoritarian SCFs

- The top cycle is characterized by strong expansion and satisfies stability.
- The uncovered set is characterized by expansion.
- The Banks set is characterized by dominator expansion.
- Computing the Banks set is NP-hard.
- The bipartisan set is stable.
- No stable majoritarian social choice function returns fewer alternatives on average than the bipartisan set.

The set-theoretic relationships between majoritarian SCFs are depicted in this Venn-like diagram. If the ellipses for two SCFs f and f' are disjoint, this signifies that $f(A, \succ_M) \cap f'(A, \succ_M) = \emptyset$ for some tournament $(A, \succ_M)$. All shown SCFs are monotonic. Stable SCFs are marked in red.

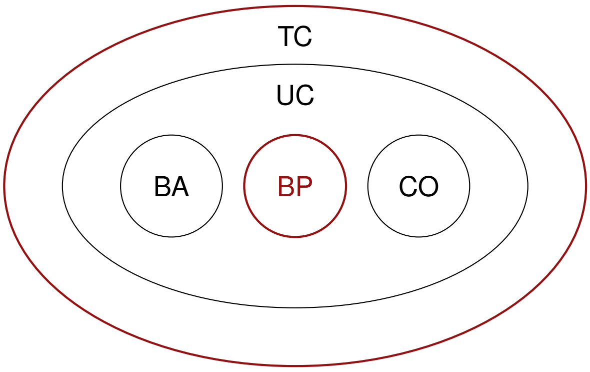

7.8 Further Reading

The formal systematic study of majoritarian SCFs was initiated by Moulin (1986). Comprehensive overviews of majoritarian SCFs were provided by Laslier (1997), Brandt (2009), Hudry (2009), and Brandt et al. (2018a).

The characterization given in Theorem 7.2 has been slightly modified from the original version by Bordes (1976), who used positive responsiveness and a condition called minimality, which rules out that the choice set strictly contains a dominant set, instead of inclusion-minimality.

A prominent variant of TC is defined by taking the transitive closure of the *strict* majority relation. It is known as the Schwartz set or GOCHA and refines TC (Schwartz, 1972, 1986). Duggan and Le Breton (2001) propose another variant of the top cycle based on the unique *mixed saddle* of the zero-sum game given by the skew-adjacency matrix of the majority graph. The mixed saddle is nested in between the top cycle and the Schwartz set. Both the Schwartz set and the mixed saddle coincide with the top cycle in the absence of majority ties.

When permitting majority ties, there are various covering relations, such as *McKelvey covering*, *Bordes covering*, and *Gillies covering*, leading to different variants of uncovered sets (see, e.g., Dutta and Laslier, 1999; Peris and Subiza, 1999; Brandt and Fischer, 2008; Duggan, 2013). Brandt et al. (2016b) have actually shown a stronger statement than Theorem 7.5. They proved that for every preference profile P, one can find another profile P' such that $\succsim_{\mathsf{M}} = \succsim'_{\mathsf{M}}$ and $\mathsf{UC}(\mathsf{P}) = \mathsf{PO}(\mathsf{P}')$, where UC is the McKelvey uncovered set. Computer-generated counterexamples demonstrate that this statement does not hold for a constant number of voters and cannot be extended to establish the majority-equivalence of the McKelvey covering relation and the Pareto dominance relation.

McKelvey covering
Bordes covering
Gillies covering

BA was originally proposed in the context of strategic voting in binary trees, where voters compare pairs of alternatives to each other. The Banks set consists precisely of those alternatives that may result as outcomes of the *amendment agenda*, an influential subset of voting trees studied in political science (Banks, 1985). The top cycle corresponds to the set of outcomes of general voting trees. For a more extensive introduction to the literature on agenda implementation, the reader is referred to Moulin (1988a, Chapter 9), Laslier (1997, Chapter 8), Austen-Smith and Banks (2005, Chapter 4), and Brandt et al. (2018a, Section 3.4.2). Banks and Bordes (1988) discusses four different generalizations of the *Banks set* BA that can deal with majority ties.

amendment agenda

The inductive proof showing the existence of optimal distributions (Theorem 7.10) is based on inductive proofs of the minimax theorem by Loomis (1946) and Owen (1967). In the graph-theoretic literature dealing with tournaments, optimal distributions are sometimes referred to as *winning densities* or *losing distributions* of the converse tournament (e.g., Fisher and Ryan, 1992, 1995a,b,c; Fisher and Reeves, 1995; Bang-Jensen and Gutin, 2018). They have been used to prove *Dean's conjecture*: every tournament contains an alternative that dominates at least as many alternatives in two steps as it does in a single step. Fisher (1996) has shown that some alternative in the bipartisan set of the converse tournament has this property. Whether such alternatives always exist in general directed graphs is an important open problem, known as the *second neighborhood problem*, proposed by the mathematician Paul Seymour in 1990 (see also Bang-Jensen and Gutin, 2018, pp. 45–51).

When permitting majority ties, optimal distributions are no longer unique. A natural generalization of the bipartisan set that can cope with majority ties is to return all alternatives that are contained in the support of *some* Nash equilibrium of the zero-sum game given by the skew-adjacency matrix (Dutta and Laslier, 1999). This generalization is called the *sign-essential set* to distinguish it from its margin-based variant, the *essential set*. The sign-essential set inherits most of the attractive properties of the bipartisan set, such as Pareto-optimality, monotonicity, and set-contraction. However, it does not satisfy set-expansion.

sign-essential set
essential set

7.9 Exercises

7.1 *Copeland's rule*

Prove or disprove the following statements.

(a) $\mathsf{CO} \subseteq \mathsf{UC}$.

(b) CO is a strong Condorcet extension.

(c) For each tournament $(A, \succ_M)$ and all $x, y \in A$, x covers y if and only if

$$x \succ_M y \text{ and } |\mathsf{CO}(\{x, y, z\}, \succ_M)| = 1 \text{ for all } z \in A.$$

7.2 *Top Cycle and Uncovered Set*

Let $\mathcal{T}$ be the class of tournaments $(\{a_1, \dots, a_k\}, \succ_M)$ for $k > 3$ such that $a_i \succ_M a_j$ if $i < j$ except that $a_k \succ_M a_1$. Identify the top cycle and the uncovered set for each tournament in $\mathcal{T}$.

☆ **7.3** *Hamiltonian cycle*

A *Hamiltonian cycle* is a cycle that visits each vertex of a directed graph exactly once. Show that the directed graph $G = (\mathsf{TC}(A, P), \succsim_M|_{\mathsf{TC}(A,P)})$ contains a Hamiltonian cycle.

Hint: Consider a cycle of maximal length in G and show that it contains all alternatives in $\mathsf{TC}(A, P)$. We also allow directed cycles of length 1 and 2.

☆ **7.4** *Iterating the uncovered set*

Prove the following statements.

(a) Let $(A, \succ_M)$ be a tournament. There is a tournament $(A', \succ'_M)$ such that $(B, \succ'_M|_B) = (A, \succ_M)$ for $B = \mathsf{UC}(A', \succ'_M)$ iff $|A| = 1$ or $(A, \succ_M)$ does not have a Condorcet winner.

(b) Let $\mathsf{UC}^1 = \mathsf{UC}$ and $\mathsf{UC}^{k+1} = \mathsf{UC}(\mathsf{UC}^k)$ for all $k > 1$. $\mathsf{UC}^k \neq \mathsf{UC}^\ell$ for all $k, \ell > 0$, $k \neq \ell$.

7.5 *Refining the uncovered set*

Construct a tournament $(A, \succ_M)$ for each of the following statements.

(a) $\mathsf{TC}(\mathsf{UC}(A, \succ_M)) \neq \mathsf{UC}(A, \succ_M)$.

☆ (b) $\mathsf{UC}^2(A, \succ_M) \cap \mathsf{CO}(A, \succ_M) = \emptyset$.

7.6 *Finding a Banks loser*

Identify an alternative in the tournament below that is not contained in the Banks set.

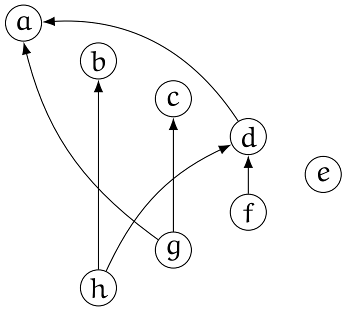

Fun fact: This is a minimal tournament in which the Banks set and the bipartisan set are not contained in each other (Brandt et al., 2015b).

7.7 *Properties of the Banks set*

Prove the following statements.

(a) BA is monotonic.

(b) In the tournament $(A, >_M)$ below, $\mathsf{BA}(A, >_M) \cap \mathsf{CO}(A, >_M) = \emptyset$.

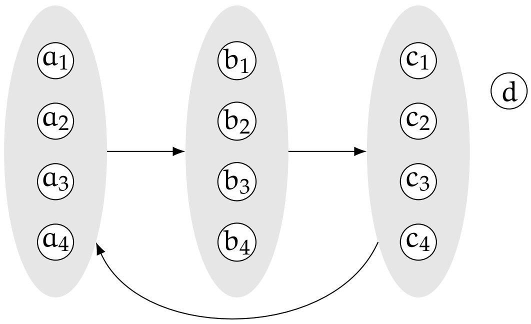

Fun fact: This is a minimal tournament in which the Banks set and the Copeland set are disjoint (Hudry, 1999).

☆ 7.8 *The structure of the Banks set* *(Miller et al., 1990)*

Prove that $\mathsf{BA}(A, >_M) = \mathsf{TC}(\mathsf{BA}(A, >_M))$ for every tournament $(A, >_M)$.

7.9 *Monotonicity*

Show that TC, UC, and BA are monotonic.

7.10 *Condorcet winners*

Prove the following statement for majoritarian SCFs f.

(a) If f satisfies expansion, then $\mathsf{Cond} \subseteq f$.

(b) If f is stable, then f is a strong Condorcet extension.

7.11 *Extended base relation*

Let $U = \{a, b, c\}$ and consider the choice function f given below. Determine the extended base relation $\succsim_f$ on *sets* of alternatives. Is f set-rationalizable?

X	$f(X)$
$\{a, b\}$	$\{a, b\}$
$\{b, c\}$	$\{c\}$
$\{a, c\}$	$\{a\}$
$\{a, b, c\}$	$\{a\}$

7.12 *Set-consistency*

Show that contraction, set-contraction, expansion, and set-expansion are logically independent.

Hint: For every pair of axioms, construct a choice function that satisfies one axiom but not the other and *vice versa*. For contraction and expansion, such choice functions were already constructed in Exercise 2.2. $|U| = 3$ suffices.

7.13 *Stability*

Prove the following statements.

(a) TC is stable.

(b) PO is stable.

(c) UC^∞ satisfies set-contraction but violates set-expansion.

7.14 *A stable non-majoritarian SCF*

Define the following set-based generalization of Pareto dominance: $y \in U$ is dominated in $A \in \mathcal{F}$ if there is $X \subseteq A$ such that for all $i \in N$, there is $x \in X$ with $x \succ_i y$. Let f be the SCF that returns all undominated alternatives in A. Prove the following statements.

(a) f satisfies contraction but violates expansion.

(b) f is stable.

Hint: First find an alternative, much simpler representation of f.

Part III

Strategic Manipulation

My scheme is intended only for honest men.

Chevalier de Borda, 1784

8 Resolute SCFs

Learning Outcomes

- Which resolute social choice functions are immune to strategic manipulation?
- Which resolute social choice functions do not suffer from strategic abstention?
- How about approval voting and median voting in their respective domains?

To discuss whether an SCF can be manipulated by submitting insincere preferences or by abstaining from the election, we need to check whether an agent prefers one outcome returned to another outcome. Since we only know the agents' preferences over single alternatives, our initial analysis will be restricted to resolute SCFs.

Obviously, assuming that the true preferences of all voters are known is unrealistic. We only know what the voters tell us and, as we have seen in Section 1.3, voters are sometimes better off by misrepresenting their preferences.

Some readers may wonder why manipulation is a problem at all. It is difficult—if not impossible—to detect and usually perfectly legal. So why not just accept the fact that voters will be lying about their preferences?

One problem with manipulation is that better-informed agents will be more successful at pulling the outcome in their direction. Manipulable SCFs reward voters who spend resources on finding out the others' preferences and computing beneficial manipulations. These resources may not be spread evenly across the population, and thus, manipulation introduces fairness issues. Another, more fundamental problem, is that manipulable SCFs may return socially undesirable outcomes. All of the statements we presented so far were made under the assumption that we have access to the true preferences of the voters. It is very hard to predict or give any formal guarantees about what is going to happen when voters can vote insincerely. The voters' problem of finding a beneficial—or even an optimal—manipulation is not an optimization problem. Whether a strategy is optimal depends on the strategies of the other voters, which introduces game-theoretic considerations into social choice. As an extreme example, recall the silly rule for two alternatives we defined in Section 3.3 and let us make it resolute by assuming that $\mathcal{D} = \mathcal{S}^N$ and n is odd. The outcome of this SCF depends on the parity of the number of voters who share the same preference. As pointed out in Section 3.3, it can *always* be manipulated by any unhappy voter! The underlying game does not admit a pure Nash equilibrium, and the only mixed equilibrium is for everyone to randomize uniformly between both statements about one's preferences. The silly rule counters the common argument that strategic voting "regulates" itself.

Let us now define manipulability and strategyproofness. A resolute SCF f is manipulable by voter i if there exist $P, P' \in \mathcal{D}$ and a feasible set A such that $P_{-i} = P'_{-i}$ and $f(A, P') \succ_i f(A, P)$. f is *strategyproof* if it is not manipulable by any voter.

strategyproofness

This definition essentially assumes that voters are aware of each other's preferences. Three arguments to support this assumption are that *(i)* it avoids the complexity of formal models that rely on questionable premises about voters' beliefs, *(ii)* statements about *non-manipulability* become stronger because they hold even when all preferences were known, and *(iii)* in some settings (e.g., voting in a committee after a debate, in which participants express their personal stances), it is not unreasonable to assume that voters are indeed well-informed about each other's preferences.

Our investigation of strategyproof SCFs will follow a similar approach to our study of Arrovian SCFs. Strategyproofness can be easily achieved for two alternatives, most naturally by using majority rule. The restricted domains we studied in Chapter 5 also admit attractive strategyproof SCFs. In the unrestricted domain of weak or strict preferences involving more than two alternatives, however, a sweeping impossibility shows that strategyproofness can only be attained by dictatorial rules. This result and its proof are closely linked to Arrow's impossibility. We conclude the chapter by examining strategic abstention, which draws a sharp line between scoring rules and Condorcet extensions when insisting on resoluteness.

8.1 Two Alternatives

For resolute SCFs on two alternatives, strategyproofness and monotonicity are equivalent.

Proposition 8.1 $m = 2$

A resolute SCF is strategyproof iff it is monotonic.

Proof. Since strategyproofness is defined as *non*-manipulability, it is more convenient to work with its negation. Similarly, monotonicity of a resolute SCF f is violated if the outcome of f changes from a to b, even though a is reinforced. Formally, f violates monotonicity iff there are $P, P' \in \mathcal{D}$ with $P_{-i} = P'_{-i}$, $\succsim_i \neq \succsim'_i$, $f(P) = \{a\}$, and $f(P') = \{b\}$ such that

$$(b \succ_i a \ \wedge \ a \succsim'_i b) \ \vee \ (b \succsim_i a \ \wedge \ a \succ'_i b).$$

This is also precisely the condition for the manipulability of f: in the case of the first disjunct, voter i prefers b to a and can manipulate from P to P', in the case of the second disjunct, voter i prefers a to b and can manipulate from P' to P.

We have thus shown that for resolute SCFs, non-monotonicity is equivalent to non-strategyproofness. □

It follows that majority rule on two alternatives with lexicographic tie-breaking is strategyproof. There are further strategyproof SCFs, in particular, the threshold rules discussed on page 34 with lexicographic tie-breaking. These include the Pareto rule with lexicographic tie-breaking, which becomes the unanimity or veto rule: return alternative a unless everyone prefers b. We can leverage May's theorem (Theorem 3.3) to characterize majority rule via strategyproofness. Restricting the domain to strict preferences with an odd number of voters ensures that majority rule is resolute. The statement then immediately follows from Theorem 3.2 and the observation that positive responsiveness and monotonicity coincide for resolute SCFs.

Theorem 8.1 $\mathcal{D} = \mathcal{S}^N, m = 2, n$ odd

Majority rule is the only resolute SCF that satisfies anonymity, neutrality, and strategyproofness.

8.2 Restricted Domains

Cond with appropriate tie-breaking is strategyproof in subdomains of $\mathcal{D}_{\text{QTRANS}}$. For example, the following SCFs discussed in Chapter 5 are strategyproof in their respective domains.

- Approval voting with lexicographic tie-breaking is strategyproof when $\mathcal{D} \subseteq \mathcal{W}^N_{\text{DICH}}$.
- Median voting is strategyproof when $\mathcal{D} \subseteq \mathcal{S}^N_{\text{SP}(\geqslant)}$ and n is odd.

Also, median voting with ties broken according to the linear order $\geqslant$, is strategyproof. Note, however, that breaking ties lexicographically or taking the average of the two medians (i.e., the alternative in the middle of the two medians) will result in violations of strategyproofness.

The following remarkable theorem characterizes all anonymous strategyproof SCFs for domain $\mathcal{S}^N_{\text{SP}(\geqslant)}$ in terms of median voting with "phantom voters." Phantom voters are virtual voters added to the preference profile before computing the median. As with regular voters, it suffices to consider their peaks.

Theorem 8.2 (Moulin, 1980) $\mathcal{D} = \mathcal{S}^N_{\text{SP}(\geqslant)}$

The only anonymous, strategyproof, and resolute SCFs are median voting rules (Cond) with $n + 1$ constant phantom voters.

To get more insight into this characterization, let us consider some special cases. When placing all phantoms at the same position $x \in U$, the resulting SCF will invariably return x, because a majority of real and virtual voters share the same peak. Such constant SCFs are trivially strategyproof. Placing pairs of phantoms at each of the two ends of the linear spectrum has no effect on the median. When placing one phantom at each end, the resulting SCFs are guaranteed to be Pareto-optimal because the remaining $n - 1$ phantoms only represent a minority of all real and virtual voters. Putting $\lceil n/2 \rceil$

phantoms at the left of the spectrum and $\lfloor n/2 \rfloor$ phantoms at the right end results in standard median voting with ties broken according to the linear order $\geqslant$. Uniformly distributing the phantoms along the linear spectrum results in an SCF that tends to return alternatives that are closer to the middle of the spectrum. This can be useful to increase the minimum satisfaction of all voters.

When $\mathcal{D} = \mathcal{D}_{\text{TRANS}} \cap \mathcal{S}^N$ and n is odd, $\succsim_M$ is antisymmetric and transitive, and Cond is a resolute SCF that returns the Condorcet winner.

Theorem 8.3 (Campbell and Kelly, 2003) $\mathcal{D} = \mathcal{D}_{\text{TRANS}} \cap \mathcal{S}^N, n$ odd

Cond is the only non-dictatorial, Pareto-optimal, strategyproof, and resolute SCF.

This result is reminiscent of Theorem 5.1 on page 64, where Cond was characterized within domain $\mathcal{D}_{\text{TRANS}}$ using the conditions from Arrow's impossibility.

8.3 The Gibbard-Satterthwaite Theorem

In this section, we consider the unrestricted preference domain. Which of the non-dictatorial SCFs we studied so far allow for strategyproof resolute refinements? As we will see, none of them do. Even reducing the domain to strict preferences offers no relief.

Let us first consider the case where $\mathcal{D} = \mathcal{S}^N$, i.e., there is a fixed set of voters with strict preferences. It turns out that strategyproofness is closely connected to a strengthening of monotonicity.

strong monotonicity

An SCF f is *strongly monotonic* if $f(P') = f(P)$ for all $P, P' \in \mathcal{D}$, and i with $P = P'$ except $b \succ_i a$ and $a \succ'_i b$ for some a, b with $b \notin f(P)$. In other words, f is invariant under the weakening of unchosen alternatives.

Strong monotonicity would not be called strong monotonicity if it did not imply monotonicity. Monotonicity demands that $a \in f(P)$ implies $a \in f(P')$. If $b \notin f(P)$, this follows immediately from strong monotonicity. When $b \in f(P)$, assume for contradiction that $a \notin f(P')$. Then, strong monotonicity can be applied from P' to P, as the unchosen alternative a is weakened. We thus get that $f(P) = f(P')$, contradicting $a \in f(P)$ and $a \notin f(P')$.

For resolute SCFs, strong monotonicity is equivalent to the following condition: Let $P, P' \in \mathcal{D}$, $i \in N$ such that $P_{-i} = P'_{-i}$, and $x \in A$.

$$f(P) = \{x\} \text{ and } (\forall z\colon x \succ_i z \Rightarrow x \succ'_i z) \quad \Rightarrow \quad f(P') = \{x\}. \tag{*}$$

In other words, x remains the winner if everything that lies below x remains below x. This equivalence holds because repeated application of strong monotonicity allows for three kinds of changes in the preferences of voter i:

(i) alternatives above x may be rearranged arbitrarily,

(ii) alternatives above x may be moved below x, and

(iii) alternatives below x can be rearranged arbitrarily.

We will often use strong monotonicity to move x or alternatives above x to the top of a preference ranking without affecting the choice set. The following theorem shows that strong monotonicity and strategyproofness are equivalent for resolute SCFs with strict preferences.

Theorem 8.4 (Muller and Satterthwaite, 1977) $\mathcal{F} = \{\mathcal{U}\}, \mathcal{D} = \mathcal{S}^N$

Let f be a resolute SCF. Then,

$$f \text{ is strategyproof} \quad \Leftrightarrow \quad f \text{ is strongly monotonic.}$$

Proof. Let f be a resolute SCF. Strategyproofness and strong monotonicity are both conditions based on two profiles $P, P' \in \mathcal{D}$ such that there is some voter $i \in N$ and $P_{-i} = P'_{-i}$. Strategyproofness then requires that $\neg(f(P') \succ_i f(P))$ while the implication required by strong monotonicity is given above (*).

$\Rightarrow$ Assume for contradiction that f is strategyproof but violates strong monotonicity. Then, there must be $x, y \in A$ such that $f(P) = \{x\}$, $\forall z \in A\colon x \succ_i z \Rightarrow x \succ'_i z$, and $f(P') = \{y\} \neq \{x\}$. Strategyproofness implies that $x \succ_i y$, as otherwise, there was a manipulation from P to P'. By assumption $x \succ_i y$ implies $x \succ'_i y$. However, this means that voter i can manipulate from P' to P, a contradiction.

$\Leftarrow$ Assume for contradiction that f is strongly monotonic but manipulable. Then there must be $x, y \in A$ such that $f(P) = \{x\}$, $f(P') = \{y\}$, and $y \succ_i x$. Now define a new relation that is identical to $\succsim_i$, except that alternative y is moved to the top of the ranking, i.e.

$$\succsim''_i = \succsim_i \setminus \{(z,y)\colon z \neq y\} \cup \{(y,z)\colon z \neq y\}.$$

$\succsim_i$	$\succsim''_i$	$\succsim'_i$
$\vdots$	y	$\vdots$
y	$\vdots$	$\vdots$
$\vdots$	$\vdots$	y
x	x	$\vdots$
$\vdots$	$\vdots$	$\vdots$

Recall that $f(P) = \{x\}$ and $y \succ_i x$. Hence, one can transform $\succsim_i$ to $\succsim''_i$ by only weakening unchosen alternatives, namely by moving all alternatives above y down. Hence, strong monotonicity implies that $f(P'') = \{x\}$. Now recall that $f(P') = \{y\}$. Again, one can transform $\succsim'_i$ to $\succsim''_i$ by only weakening unchosen alternatives. This time, all alternatives except y are moved below y and then rearranged exactly as in $\succsim''_i$. Strong monotonicity then implies that $f(P'') = \{y\}$, which contradicts our previous conclusion that $f(P'') = \{x\}$.

□

The applications of strong monotonicity in the previous proof show how strong this condition is, but so is strategyproofness.

For two alternatives, monotonicity and strong monotonicity coincide. Hence, Proposition 8.1 becomes a corollary of Theorem 8.4.

Only very few resolute SCFs are strategyproof. As a warm-up for the sweeping Gibbard-

Satterthwaite impossibility, let us first show that strategyproofness is incompatible with Condorcet-consistency.

Proposition 8.2 $m \geqslant 3, n \geqslant 3$

No resolute Condorcet extension is strategyproof.

Proof. Let $m, n = 3$ and f be a strategyproof (and hence strongly monotonic) resolute SCF. Assume for contradiction that f is Condorcet-consistent and consider the Condorcet cycle preference profile P on the left.

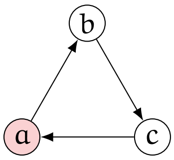

1	1	1
a	b	c
b	c	a
c	a	b

P

1	1	1
a	b	c
c	c	a
b	a	b

P'

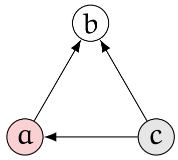

We can assume without loss of generality that $f(P) = \{a\}$ (highlighted in red). Since the profile is completely symmetric, the subsequent arguments would work just as well for $f(P) = \{b\}$ and $f(P) = \{c\}$. Now let the first voter swap alternatives b and c. This results in profile P' on the right. Since we obtain P' by weakening the unchosen alternative b in the first voter's preference relation, strong monotonicity implies that $f(P') = \{a\}$. This contradicts Condorcet-consistency as c is a Condorcet winner in P' (highlighted in gray).
The proof can be generalized to $m > 3$ by ranking all additional alternatives at the bottom of the individual rankings. For odd $n > 3$, we can add pairs of voters with completely opposed preferences. For even $n \geqslant 6$, we can clone the three voters in the original profiles and again add pairs of voters with opposing preferences. Proving the statement for $n = 4$ is tedious and left to the avid reader. □

Proposition 8.2 already entails that all resolute refinements of the SCFs we studied in Chapter 7, such as TC, UC, BA, and BP, are manipulable. Unfortunately, this simple negative result can be strengthened significantly. The relationship between Proposition 8.2 and the upcoming impossibility is similar to that between the Condorcet-May impossibility (Theorem 3.4) and Arrow's impossibility (Theorem 4.1).

non-imposition

Recall that an SCF is *non-imposing* if it is "onto" (surjective), i.e., for every $x \in U$, there is P such that $f(P) = \{x\}$. Non-imposition is weaker than each of Condorcet-consistency, Pareto-optimality, and neutrality.

Theorem 8.5 (Gibbard, 1973; Satterthwaite, 1975) $\mathcal{F} = \{U\}, \mathcal{D} = \mathcal{S}^N, m \geqslant 3$

Every non-imposing, strategyproof, and resolute SCF is dictatorial.

Proof. We prove the theorem by reducing it to Arrow's impossibility theorem. Given a non-imposing, strategyproof, and resolute SCF f, we construct an SWF g that satisfies Pareto-optimality and IIA. Hence, g is dictatorial, which, as we will see, implies that f is dictatorial, too.

SCF **f**	non-imposition	strategyproofness		Dictatorship
$\downarrow$				$\uparrow$
SWF **g**	Pareto-optimality	IIA	$\Rightarrow$	Dictatorship

For nonempty $S \subseteq U$ and $\succsim_i \in \mathcal{S}$, let $\succsim_i^S$ be the ranking obtained by moving all alternatives in S to the top of $\succsim_i$ and not changing the internal ordering.
For profiles, we write $P^S = (\succsim_1^S, \dots, \succsim_n^S)$.
We first show that the SCF f satisfies Pareto-optimality. Assume for contradiction that $x \succ_i y$ for all i and $f(P) = \{y\}$. Strong monotonicity implies that $f(P^{\{x\}}) = \{y\}$. There exists P' such that $f(P') = \{x\}$ because of non-imposition. Strong monotonicity implies that $f(P'^{\{x\}}) = \{x\} = f(P^{\{x\}})$, a contradiction.
Hence, $f(P^S) \subseteq S$ for all $P \in \mathcal{D}$.
We construct the collective preference relation $\succsim$ returned by SWF g for a given profile $P \in \mathcal{D}$ by moving pairs of alternatives $x, y \in U$ to the top of all individual rankings. In particular, for all $x, y \in U$, we let

$$x \succsim y \quad \Leftrightarrow \quad x \in f(P^{\{x,y\}}).$$

Since $f(P^{\{x,y\}}) \subseteq \{x, y\}$, $\succsim$ is complete. Moreover, since f is resolute, $\succsim$ is antisymmetric. For g to be a proper SWF, $\succsim$ also needs to be transitive. This can be shown as follows. Let $x, y, z \in U$. Without loss of generality, we may assume that $f(P^{\{x,y,z\}}) = \{x\}$ and $f(P^{\{y,z\}}) = \{y\}$. Then, by strong monotonicity, $f(P^{\{x,y\}}) = \{x\}$ and $f(P^{\{x,z\}}) = \{x\}$. Hence, $x \succ y$, $y \succ z$, and $x \succ z$.
We will now show that g satisfies the axioms of Arrow's theorem (Theorem 4.2).

- Pareto-optimality
  If for two alternatives $x, y \in A$, $x \succ_i y$ for all $i \in N$, then $f(P^{\{x,y\}}) = \{x\}$ since f is Pareto-optimal and, hence, $x \succ y$.

- Independence of irrelevant alternatives (IIA)
  Let $x, y \in U$ and $P, P' \in \mathcal{D}$ such that $P|_{\{x,y\}} = P'|_{\{x,y\}}$. Then,

$$x \succsim y \quad \Leftrightarrow \quad x \in f(P^{\{x,y\}}) \quad \Leftrightarrow \quad x \in f(P'^{\{x,y\}}) \quad \Leftrightarrow \quad x \succsim' y.$$

Arrow's theorem thus implies that g is dictatorial.
Let i be the dictator for SWF g and consider $P \in \mathcal{D}$ with $Max(U, \succsim_i) = \{x\}$. Since g is dictatorial, $x \succ y$ for all $y \neq x$. Let $f(P) = \{z\}$. By strong monotonicity, $f(P^{\{x,z\}}) = \{z\}$. This implies that $z \succsim x$, which is only possible if $x = z$. Hence, i is a dictator for f. □

Theorem 8.5 also holds for weak preferences. Note that, just like for Arrow's theorem (Theorem 4.1), SCFs, for which voter i is a dictator, are no longer unique when preferences are weak. Dictatorship merely demands that a subset of voter i's most-preferred alternatives has to be returned. For resolute SCFs, this subset must be a singleton. However, not all dictatorial SCFs are strategyproof. If the dictator is completely indifferent, dictatorship has no consequences whatsoever for the social choice. In contrast to Theorem 8.5, the following corollary, which is identical to Theorem 8.5 except that $\mathcal{D} = \mathcal{W}^N$, does not completely characterize all dictatorial SCFs.

Corollary 8.1 $\mathcal{F} = \{U\}, \mathcal{D} = \mathcal{W}^N, m \geqslant 3$

Every non-imposing, strategyproof, and resolute SCF is dictatorial.

Proof. Let f be a non-imposing, strategyproof, and resolute SCF. Then, f is also strategyproof and resolute within $\mathcal{S}^N$. In order to apply Theorem 8.5, we need to show that f also satisfies non-imposition within $\mathcal{S}^N$. Let $x \in U$. Non-imposition implies that there is $P \in \mathcal{W}^N$ such that $f(P) = \{x\}$.
Define $\hat{P} = (\hat{\succsim}_1, \dots, \hat{\succsim}_n) \in \mathcal{S}^N$, which is identical to P except that x is moved to the top of each preference relation and all ties are broken arbitrarily. Further let $\hat{P}_i = (\hat{\succsim}_1, \dots, \hat{\succsim}_i, \succsim_{i+1}, \dots, \succsim_n)$ for $0 \leqslant i \leqslant n$ and assume for contradiction that $1 \leqslant k \leqslant n$ is the smallest index such that $f(\hat{P}_k) \neq \{x\}$. Then, f would be manipulable by Voter k, going from $\hat{P}_k$ to $\hat{P}_{k-1}$. Hence, $f(\hat{P}) = \{x\}$.
By Theorem 8.5, some voter, say Voter 1, is a dictator for f within $\mathcal{S}^N$. Now let $P \in \mathcal{W}^N$ be an arbitrary profile of weak preferences and $B = \max(U, \succsim_1)$. It remains to show that $f(P) \subseteq B$. To this end, let $\hat{P} = (\hat{\succsim}_1, \dots, \hat{\succsim}_n)$ be some profile in $\mathcal{S}^N$ such that $x \succ_1 y$ and $y \succ_i x$ for all $x \in B, y \in U \setminus B$, and $i \in N \setminus \{1\}$. Define $\hat{P}_i = (\hat{\succsim}_1, \dots, \hat{\succsim}_i, \succsim_{i+1}, \dots, \succsim_n)$ for $0 \leqslant i \leqslant n$ and let $0 \leqslant k \leqslant n$ be the smallest index such that $f(\hat{P}_k) \in B$. Such a k must exist because $f(\hat{P}_n) = f(\hat{P}) \in B$. If $k = 1$, Voter 1 can manipulate f by going from $\hat{P}_0$ to $\hat{P}_1$. If $k > 1$, Voter k can manipulate f by going from $\hat{P}_k$ to $\hat{P}_{k-1}$. Hence, $k = 0$, which implies that $f(\hat{P}_0) = f(P) \subseteq B$. □

serial dictatorships

Common examples of dictatorial and strategyproof SCFs for weak preferences are *serial dictatorships*. A serial dictatorship is based on a fixed ordering of the voters such that one voter after another narrows down the set of admissible social choices to his most-preferred ones among the ones that have not been excluded so far. If more than one alternative remains after going over all voters, ties can be broken lexicographically.

8.4 Abstention

Strategic abstention is a phenomenon similar to strategic manipulation. As we have seen in Section 1.3, sometimes a voter can be better off by abstaining from an election. To study strategic abstention, we need to allow for variable sets of voters by letting $\mathcal{D} \subseteq \mathcal{W}^{\subseteq N}$.

participation

An SCF satisfies *participation* if there is no $P = (\succsim_1, \dots, \succsim_n)$ and $i \in N_P$ such that $f(P_{-i}) \succ_i f(P)$.

The following SCFs discussed in Chapter 5 satisfy participation in their respective domains.

- Approval voting with lexicographic tie-breaking satisfies participation when $\mathcal{D} \subseteq \mathcal{W}^{\subseteq N}_{\text{DICH}}$.

- Median voting with ties broken according to the linear order $\geqslant$ satisfies participation when $\mathcal{D} \subseteq \mathcal{S}^{\subseteq N}_{\text{SP}(\geqslant)}$.

For unrestricted preferences, participation is far less prohibitive than strategyproofness. As mentioned in Section 1.3, plurality, Borda's rule, and monotonic scoring rules in general satisfy participation.

Theorem 8.6 $\mathcal{D} = \mathcal{S}^N$

Every monotonic scoring rule with lexicographic tie-breaking satisfies participation.

Proof. Assume for contradiction that there is a monotonic scoring rule f (defined by score vector s with $s_1 \geqslant s_2 \geqslant \ldots \geqslant s_m$) that fails participation when using lexicographic tie-breaking. Hence, there is a profile $P \in \mathcal{D}$ and voter $i \in N$ such that $f(P_{-i}) \succ_i f(P)$. Now, let $f(P_{-i}) = \{x\}$ and $f(P) = \{y\}$. The score of x in P_{-i} is maximal and, since $x \succ_i y$, voter i assigns at least as many points to x as to y. Hence, in P, the score of x is at least as large as the score of y. If it is strictly greater, $f(P) \neq \{y\}$. If both scores are equal, they are also equal in P_{-i}, which means that x is lexicographically smaller than y. Again, $f(P) \neq \{y\}$. As a consequence, the assumption that $f(P_{-i}) \succ_i f(P)$ was wrong. □

Are there interesting SCFs other than scoring rules that satisfy participation? A large class of SCFs is ruled out by the following theorem.

Theorem 8.7: No-Show Paradox (Moulin, 1988b) $\mathcal{D} = \mathcal{S}^N, m \geqslant 4, n \geqslant 12$

No resolute Condorcet extension satisfies participation.

Proof. The original proof by Moulin (1988b) requires at least $n \geqslant 25$ voters. We here give a slightly simplified version of the minimal proof by Brandt et al. (2017), which was found with the help of a computer. The complete proof for $m = 4$ and $n = 12$ is visualized in the diagram below. Assume for contradiction that f is a Condorcet extension that satisfies participation.

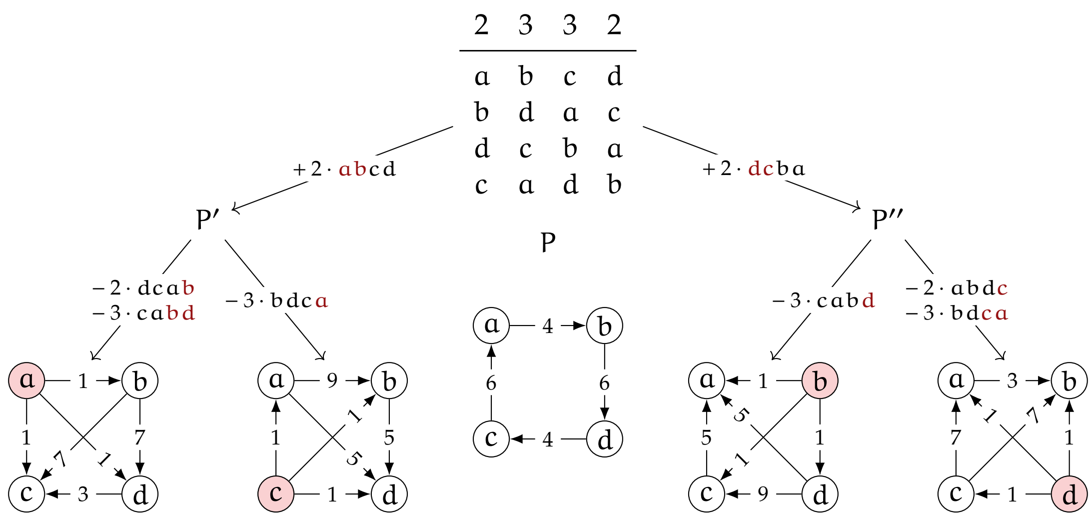

Consider the 10-voter preference profile P in the middle. First, observe that P remains fixed when renaming alternatives according to $\sigma = \left(\begin{smallmatrix} a & b & c & d \\ d & c & b & a \end{smallmatrix}\right)$ and then reordering voters. We can, therefore, assume without loss of generality that $f(P) \in \{a, b\}$. (An explicit analysis in case $f(P) \in \{c, d\}$ is shown in the right-hand side of the diagram.) We proceed to the left-hand side by adding two voters with preferences $a \succ b \succ c \succ d$, one after the other, resulting in profile P'. Since f satisfies participation and both voters prefer a and b to c and d, $f(P') \in \{a, b\}$. We now distinguish between two cases.
If $f(P') = \{a\}$, we consecutively remove three voters with preferences $b \succ d \succ c \succ a$ (note that these voters are indeed contained in the profile). Participation implies that the outcome remains a, their least preferred alternative. However, as shown in the weighted majority tournament, alternative c is a Condorcet winner in the resulting profile, yielding a contradiction.
If $f(P') = \{b\}$, we consecutively remove five voters with the preferences listed in the diagram. Participation implies that either b or d are returned in the resulting profile. However, as shown in the weighted majority tournament, alternative a is a Condorcet winner, again yielding a contradiction.
The proof can be easily extended to more than four alternatives by placing the additional alternatives arbitrarily below a, b, c, and d in all considered profiles. In these profiles, f cannot return any of the additional “dummy” alternatives. (If that were the case, we could remove one voter after another, and, due to participation, in each of the reduced profiles, a dummy would be returned. Eventually, a dummy would be returned in a single-voter sub-profile, which clashes with Condorcet-consistency.)

□

The bounds given in Theorem 8.7 are tight. When $m \leqslant 3$, maximin and Nanson’s rule (with lexicographic tie-breaking) satisfy participation and have been characterized using participation (Brandt et al., 2025). When $n \leqslant 11$, an artificial Condorcet extension that satisfies participation has been provided by Brandt et al. (2017). This SCF is a refinement of TC that is identical to maximin in 99.8% of all profiles. A soccer team of 11 players could use this Condorcet extension without worrying about strategic abstention.

Note that Theorem 8.7 strongly relies on resoluteness. Selecting a single winner for profile P of its proof is already at variance with anonymity or neutrality. When considering irresolute SCFs and making careful assumptions about what the voters know about the mechanism that eventually breaks ties between selected alternatives, Condorcet extensions can be immune to strategic abstention (see Chapter 9).

8.5 Hardness of Manipulation

An interesting stream of research in the early days of computational social choice has investigated the computational complexity of finding a beneficial manipulation. The idea is to use computational hardness (in particular, NP-hardness) as a shield against manipulation. As already mentioned in Section 1.3, for some SCFs, finding a beneficial manipulation can be a difficult problem. In fact, this problem was shown to be NP-hard for Nanson's rule (Davies et al., 2014), instant-runoff (Bartholdi, III and Orlin, 1991), and a refinement of Copeland's rule (Bartholdi, III et al., 1989). There are many more hardness results when considering weighted voting and coalitional manipulation (see, e.g., Conitzer et al., 2007; Faliszewski et al., 2009, 2010; Faliszewski and Procaccia, 2010; Conitzer and Walsh, 2016). The idea is that, as with established cryptographic methods, manipulation is possible but computationally infeasible. A major problem with this approach is that NP-hardness is a worst-case measure and, in contrast to cryptography, where we can, for example, pick large numbers that are particularly hard to factorize, we do not have control over the instances. It has been shown that many SCFs can be manipulated efficiently in restricted domains or under reasonable distributional assumptions about preferences (Procaccia and Rosenschein, 2007; Zuckerman et al., 2009; Faliszewski et al., 2011, 2014; Brandt et al., 2015a). As a matter of fact, drawing on earlier work by Friedgut et al. (2011) for $m = 3$ and Isaksson et al. (2012) for neutral SCFs, Mossel and Rácz (2015) have shown the following result for $m \geqslant 3$ and *all* resolute SCFs that are sufficiently different from strategyproof SCFs such as dictatorships or constant SCFs:

The probability that a uniformly sampled preference profile and a uniformly sampled preference relation for one voter constitute a manipulation instance is at most polynomially small in m *and* n*.*

This suggests that trying random manipulations until a beneficial one has been found works well under fairly general conditions. These results cast doubt on using computational hardness as a shield against manipulation.

8.6 Key Takeaways

Resolute SCFs

- For two alternatives, majority rule is strategyproof and satisfies participation.
- For dichotomous preferences, approval voting is strategyproof and satisfies participation.
- For single-peaked preferences, median voting is strategyproof and satisfies participation.
- For unrestricted preferences, only dictatorships are non-imposing, strategyproof, and resolute.
- For unrestricted preferences, no resolute Condorcet extension satisfies participation.

8.7 Further Reading

Interestingly, the two original proofs of the Gibbard-Satterthwaite theorem (Theorem 8.5) differ significantly. While Satterthwaite gives a direct proof, Gibbard reduces the statement to Arrow's impossibility (Theorem 4.2) as we did, too. Many other proofs have been given since then (e.g., Schmeidler and Sonnenschein, 1978; Beja, 1993; Benoît, 2000; Sen, 2001; Reny, 2001; Cato, 2009; Svensson and Reffgen, 2014). The proof of Corollary 8.1 is due to Schmeidler and Sonnenschein (1978).

Many extensions of Theorem 8.5 explore what happens when the assumption of resoluteness is dropped. Some of these results are discussed in the next chapter.

8.8 Exercises

8.1 *Top- and bottom-ranked alternatives*

Let f be a strategyproof, non-imposing, and resolute SCF, $m \geqslant 3$, and $n \geqslant 2$. Moreover, consider two distinct alternatives $x, y \in A$ and a voter $i \in N$. Prove the following statements without using the Gibbard-Satterthwaite theorem:

☆ (a) If there is a profile P such that $f(P) = \{x\}$ and all voters $j \in N \setminus \{i\}$ top-rank y, there is a profile P' such that $f(P') = \{x\}$ and all voters $j \in N \setminus \{i\}$ bottom-rank x.

(b) If there is a profile P' such that $f(P') = \{x\}$ and all voters $j \in N \setminus \{i\}$ bottom-rank x, then $f(P'') = \{x\}$ for all profiles P'' in which voter i top-ranks x.

☆ **8.2** *A direct proof of the Gibbard-Satterthwaite theorem*

An SCF f is k-*unanimous* if $f(P) = \{x\}$ for all $P \in \mathcal{D}$ and $x \in U$ such that x is top-ranked by at least $n - k$ voters.

Give a direct proof of Theorem 8.5 by showing the following implications for a strategyproof, non-imposing, and resolute SCF f when $m \geqslant 3$ and $n \geqslant 2$.

(a) If f is non-dictatorial, it is 1-unanimous.

(b) If f is k-unanimous for $0 < k < \frac{n}{2}$, it is $k + 1$-unanimous.

(c) No SCF is k-unanimous for $k \geqslant \frac{n}{2}$.

Doubt is not a pleasant condition, but certainty is absurd.

Voltaire, 1770

9 Irresolute SCFs

To study strategyproofness and participation for irresolute SCFs, we introduce preference extensions that extend voters' preferences over alternatives to preferences over sets. Formally, a *preference extension* maps every $\succsim \in \mathcal{W}$ to an *incomplete* relation over sets $\hat{\succsim} \subseteq \mathcal{P}^*(U) \times \mathcal{P}^*(U)$ such that for all $x, y \in U$,

preference extension

$$\{x\} \mathrel{\hat{\succsim}} \{y\} \quad \Leftrightarrow \quad x \succsim y.$$

Typical preference extensions we consider will return incomplete relations because preference relations over alternatives do not contain enough information to decide whether one set should be preferred to another. For example, it is unclear whether a voter with preferences $a \succ b \succ c$ will prefer $\{a, c\}$ to $\{b\}$ or not. Nevertheless, many common SCFs can be manipulated for *any* such extension because there are manipulations from profiles with a single winner to profiles with another single winner. The SCFs suffering from these kinds of "single-winner manipulations" include Borda's rule, instant-runoff, plurality with runoff, Baldwin's rule, Black's rule, Kemeny's rule, maximin, Young's rule, and Copeland's rule. Plurality is notably absent from this list because a single manipulator cannot change the outcome from one single winner to another. Define the trivial extension as the extension that only allows the comparison of singletons, as explained above.

Given some preference extension $\hat{\succsim}$, an SCF f is $\hat{\succsim}$*-strategyproof* if there are no $P, P' \in \mathcal{D}$ and $i \in N$ such that $P_{-i} = P'_{-i}$ and $f(P') \mathrel{\hat{\succ}_i} f(P)$.

$\hat{\succsim}$-strategyproofness

The following proposition shows that strong Condorcet extensions (such as TC, UC, BA, and BP) are strategyproof with respect to the trivial extension.

Proposition 9.1 $\mathcal{D} = \mathcal{W}^N$

Every strong Condorcet extension is $\hat{\succsim}$-strategyproof when $\hat{\succsim}$ is the trivial preference extension.

Proof. The trivial extension $\hat{\succsim}$ only allows the comparison of singletons. Recall that strong Condorcet-consistency of f means that $f(P) = \{x\}$ iff x is a Condorcet winner. Let $P, P' \in \mathcal{D}$ and $i \in N$ such that $P_{-i} = P'_{-i}$, $f(P) = \{a\}$, and $f(P') = \{b\}$ with $a \neq b$. Since a is a Condorcet winner in P and b is a Condorcet winner in P', we have $a \succ_M b$ and $b \succ'_M a$. Since only voter i changed his preferences, this means that $a \succsim_i b$. Voter i is not better off in P' than in P. □

It is not difficult to define further SCFs that are strategyproof with respect to the trivial extension. For example, this is the case for every SCF that always returns at least two alternatives.

$\hat{\succsim}$-participation

Participation can also be extended to irresolute SCFs using preference extensions. Given some preference extension $\hat{\succsim}$, an SCF f satisfies $\hat{\succsim}$-*participation* if there is no $P \in \mathcal{W}^{\subseteq N}$ and $i \in N_P$ such that $f(P \setminus \{(i, \succsim_i)\}) \hat{\succ}_i f(P)$.

Proposition 9.2 $\mathcal{D} \in \{\mathcal{S}^{\subseteq N}, \mathcal{W}^{\subseteq N}\}$

Every $\hat{\succsim}$-strategyproof and majoritarian SCF satisfies $\hat{\succsim}$-participation.

Proof. Let f be an SCF that violates participation. We will show that f can also be strategically manipulated. Let $P \in \mathcal{D}$ and $i \in N$ such that agent i is better off by abstaining, i.e., $f(P_{-i}) \hat{\succ}_i f(P)$. Let us first consider the easy case when $\mathcal{D} = \mathcal{W}^{\subseteq N}$. We can exploit that majoritarian SCFs cannot distinguish between indifferent and absent voters. Voter i can manipulate f by pretending to be completely indifferent, as

$$f(P_{-i}) \overset{\text{maj.}}{=} f(P_{-i} \cup \{(i, U \times U)\}).$$

The case when $\mathcal{D} = \mathcal{S}^{\subseteq N}$ can be addressed by defining a copy of profile P, $P^* = \{(n+1, \succsim_1), (n+2, \succsim_2), \dots, (2n, \succsim_n)\}$, and letting $\succsim_i^{-1}$ denote the inverted preference relation of voter i. We then have that

$$f(P_{-i} \cup \{(i, \succsim_i^{-1})\} \cup P^*) \overset{\text{maj.}}{=} f(P_{-i}) \quad \hat{\succ}_i \quad f(P) \overset{\text{maj.}}{=} f(P \cup P^*).$$

Hence, f can be strategically manipulated. □

In the following two sections, we will study two concrete preference extensions. The idea is to exploit the voters' uncertainty how ties are broken to establish strategyproofness.

9.1 Kelly's Extension

The first extension we consider is based on the assumption that a single alternative is eventually chosen, but the voters know no more about the tie-breaking mechanism than that any alternative may be selected. Under this assumption, the preferences over choice sets are given by *Kelly's preference extension* $\succsim^K \subseteq \mathcal{P}^*(U) \times \mathcal{P}^*(U)$. For all $X, Y \in \mathcal{P}^*(U)$,

Kelly's preference extension

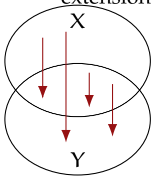

$$X \succsim^K Y \quad \Leftrightarrow \quad \forall x \in X, y \in Y \colon x \succsim y.$$

It follows that the asymmetric part $\succ^K$ of $\succsim^K$ is defined such that

$$X \succ^K Y \quad \Leftrightarrow \quad \forall x \in X, y \in Y \colon x \succsim y \land \exists x \in X, y \in Y \colon x \succ y.$$

$\succsim^K$ is an incomplete relation whenever there are at least three alternatives. For example, when $a \succ b \succ c$, we have that

$$\{a\} \succ^K \{a, b\} \succ^K \{b\} \succ^K \{b, c\} \succ^K \{c\}.$$

However, $\{a, c\}$ and $\{b\}$ are incomparable. Perhaps surprisingly, also $\{a, b\}$ and $\{a, b, c\}$ are incomparable. This follows from our absence of assumptions about how ties are broken. It may be possible that b is selected from $\{a, b\}$ and that a is selected from $\{a, b, c\}$.

As pointed out at the beginning of this chapter, many common SCFs are manipulable for *any* preference extension, including Kelly's extension. However, the following proposition shows that strong monotonicity is sufficient for $\succsim^K$-strategyproofness.

Proposition 9.3 (Brandt, 2015) $\mathcal{D} = \mathcal{S}^N$

Every strongly monotonic SCF is $\succsim^K$-strategyproof.

Proof. Let f be a strongly monotonic SCF and assume, for contradiction, that f can be $\succsim^K$-manipulated. Then there are profiles $P, P' \in \mathcal{D}$ and $i \in N$ such that $P_{-i} = P'_{-i}$ and $f(P') \succ^K_i f(P)$. In particular, $f(P) \neq f(P')$.
Let $x \in f(P')$ such that $y \succsim_i x$ for all $y \in f(P')$. Note that, if $f(P) \cap f(P') \neq \emptyset$, then $f(P) \cap f(P') = \{x\}$. Next, we partition U into the so-called strict lower contour set $L = \{y \in U : x \succ_i y\}$ and the weak upper contour set $H = U \setminus L$ of x with respect to voter i. Define a new preference relation by letting

$$\succsim^*_i \;=\; \succsim_i|_L \,\cup\, \succsim'_i|_H \,\cup\, \{(y, z) : y \in H, z \in L\}.$$

The upper part of $\succsim^*_i$ is ranked as in $\succsim'_i$ and the lower part as in $\succsim_i$. Further, let $P^* = P \setminus (i, \succsim_i) \cup (i, \succsim^*_i)$.
We now show $f(P) = f(P^*) = f(P')$. First observe that $f(P) \subseteq L \cup \{x\}$ because $f(P') \succ^K_i f(P)$ and $y \succsim_i x$ for all $y \in f(P')$. We can now reorder all alternatives in H by weakening unchosen alternatives in $\succsim_i$ to reach $\succsim^*_i$. Therefore, by strong monotonicity, $f(P) = f(P^*)$.
To show that $f(P') = f(P^*)$, we observe that $f(P') \subseteq H$ because $y \succsim_i x$ for all $y \in f(P')$. We can now weaken unchosen alternatives in $\succsim'_i$ to reach $\succsim^*_i$. Strong monotonicity then implies that $f(P') = f(P^*) = f(P)$, contradicting our initial assumption. □

While the Gibbard-Satterthwaite Theorem (Theorem 8.5) established that only imposing or dictatorial *resolute* SCFs can satisfy strong monotonicity, there *are* reasonable strongly monotonic *irresolute* SCFs. A simple example is PO, which satisfies strong monotonicity because weakening Pareto-dominated alternatives does not affect the set of Pareto-dominated alternatives. The following proposition can serve as a useful tool to show that SCFs are strongly monotonic.

Proposition 9.4 (Brandt, 2015)

Every SCF that satisfies monotonicity, set-contraction, and IIA satisfies strong monotonicity.

Proof. Let f be an SCF that satisfies monotonicity, set-contraction, and IIA. Moreover, let $P, P' \in \mathcal{D}$, $i \in N$, $A \in \mathcal{F}$, and $x, y \in A$ such that $P = P'$ except that $x \succ_i y$ and $y \succ'_i x$.

Moreover, assume that $x \notin f(A, P)$. We can show that f satisfies strong monotonicity by proving that $f(A, P) = f(A, P')$. Monotonicity of f implies that $x \notin f(A, P')$. For, otherwise, $x \in f(A, P')$ and, by monotonicity, $x \in f(A, P)$, a contradiction. We then have that

$$f(A, P) \overset{\text{set-contraction}}{=} f(A \setminus \{x\}, P) \overset{\text{IIA}}{=} f(A \setminus \{x\}, P') \overset{\text{set-contraction}}{=} f(A, P').$$

□

TC and BP satisfy monotonicity and set-contraction (Proposition 7.10 and Exercise 7.9) and hence also strong monotonicity and $\succsim^K$-strategyproofness. It can also be shown that UC is $\succsim^K$-strategyproof, even though it violates strong monotonicity. It is open whether BA is $\succsim^K$-strategyproof.

As discussed in Section 7.8, there are extensions of UC and BP that can handle majority ties. $\succsim^K$-strategyproofness thus extends to the general domain of strict preferences $\mathcal{S}^{\subseteq N}$. By Proposition 9.2, these SCFs also satisfy $\succsim^K$-participation.

The following theorem shows that Kelly-strategyproofness requires a certain degree of indecisiveness.[1]

Theorem 9.1 (Barberà, 1977) $\mathcal{F} = \{\mathcal{U}\}, \mathcal{D} = \mathcal{S}^N, m \geqslant 4$

Every Pareto-optimal, $\succsim^K$-strategyproof, and positively responsive SCF is dictatorial.

Positive responsiveness requires a very high degree of decisiveness. Of the SCFs studied in this book, only Borda and Black's rule satisfy positive responsiveness. Some SCFs can be made positively responsive by breaking ties using Borda scores, e.g., TC and MM.

For weak preferences, $\succsim^K$-strategyproofness is much more restrictive.

Theorem 9.2 (Brandt et al., 2022a) $m \geqslant 3, n \geqslant 4$

Let f be $\succsim^K$-strategyproof. Then, f violates Condorcet-consistency. Moreover, if f is non-imposing, it cannot be majoritarian and returns a Condorcet loser of at least one profile.

9.2 Fishburn's Extension

Kelly's extension allows for positive results under very restrictive assumptions about the agent's knowledge of the tie-breaking mechanism. A very natural refinement of Kelly's extension is based on the assumption that there is a linear tie-breaking order unknown to the voters (e.g., a chairman who picks his most preferred alternative from the choice set). Under this assumption, the preferences over choice sets are given by Fishburn's

[1] Theorem 9.1 actually uses a much weaker extension than Kelly's, which merely requires that $x \succ y$ implies that $\{x\} \succ \{x, y\} \succ \{y\}$.

preference extension $\succsim^F \subseteq \mathcal{P}^*(U) \times \mathcal{P}^*(U)$. For all $X, Y \in \mathcal{P}^*(U)$,

$$X \succsim^F Y \quad \Leftrightarrow \quad (\forall x \in X \setminus Y, y \in Y\colon x \succsim y) \wedge (\forall x \in X, y \in Y \setminus X\colon x \succsim y).$$

In contrast to Kelly's extension, tie-breaking has to be consistent across different sets. While it is possible that the chairman picks some alternative in $X \setminus Y$ from X and another one from Y, or some alternative from X and another one in $Y \setminus X$ from Y, it is impossible that he picks two different alternatives in $X \cap Y$ from X and Y. The chairman's choices have to be rationalizable. By contrast, Kelly's extension does not impose such a constraint on the chairman's choices.[2]

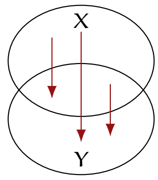

Just like $\succsim^K$, $\succsim^F$ is an incomplete relation. For any $X, Y \in \mathcal{P}^*(U)$, $X \succsim^K Y \Rightarrow X \succsim^F Y$ and hence $\succsim^K \subseteq \succsim^F$.

$a \succ b \succ c$, for example, implies that $\{a, b\} \succ^F \{a, b, c\} \succ^F \{b, c\}$.

TC can be characterized using $\succsim^F$-strategyproofness and a strengthening of non-imposition. An SCF f satisfies *set non-imposition* if for every $X \in \mathcal{P}^*(A)$, there is $P \in \mathcal{D}$ such that $f(A, P) = X$.

set non-imposition

Theorem 9.3 (Brandt and Lederer, 2023) $\mathcal{D} = \mathcal{S}^{\subseteq N}$

TC is the only margin-based SCF satisfying $\succsim^F$-strategyproofness and set non-imposition.

TC is also the finest margin-based $\succsim^F$-strategyproof SCF. Since TC violates Pareto-optimality, this can be interpreted as an impossibility theorem, showing the incompatibility of $\succsim^F$-strategyproof and Pareto-optimality for margin-based SCFs.

Few other SCFs are known to be $\succsim^F$-strategyproof, e.g., PO.

The following impossibilities were shown with the help of computers.

Theorem 9.4 (Brandl et al., 2019) $\mathcal{D} = \mathcal{D}_{\text{TOUR}}, m \geqslant 5$

There is no Pareto-optimal majoritarian SCF that satisfies $\succsim^F$-participation.

For weak preferences, $\succsim^F$-strategyproofness entails severe restrictions.

Theorem 9.5 (Brandt et al., 2022c) $\mathcal{D} = \mathcal{W}^N, m \geqslant 3, n \geqslant 3$

There is no strongly Pareto-optimal and $\succsim^F$-strategyproof anonymous SCF.

An interesting open question is whether anonymity can be weakened to non-dictatorship in Theorem 9.5.

[2]Note that none of the voters can serve as the chairman. From the perspective of this voter, the SCF would be resolute, and Theorem 8.5 applies.

9.3 Even-Chance Extension

9.4 Optimist and Pessimist Extension

Duggan and Schwartz (1992, 2000)

9.5 Complete Preferences over Sets

Barberà et al. (2001)

9.6 Further Reading

9.7 Exercises

9.1 *Plurality*

Show that plurality fails to be $\succsim^K$-strategyproof.

9.2 *Transitivity of Kelly's and Fishburn's extensions*

Show that $\succsim^K$ and $\succsim^F$ are transitive for every transitive relation $\succsim$.

9.3 $\succsim^F$*-strategyproofness,* TC*, and* UC

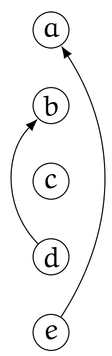

(a) Use the tournament on the right to prove that UC is $\succsim^F$-manipulable.
Hint: The goal is to manipulate to a profile where UC returns $\{a, b, d\}$.

(b) Show that TC satisfies $\succsim^F$-strategyproofness.
Hint: Let P, P' differ in the preference relation of a single voter $i \in N$. Then, consider the cases *(i)* $\mathrm{TC}(P') \subseteq \mathrm{TC}(P)$, *(ii)* $\mathrm{TC}(P) \subseteq \mathrm{TC}(P')$, and *(iii)* $\mathrm{TC}(P) \not\subseteq \mathrm{TC}(P')$ and $\mathrm{TC}(P') \not\subseteq \mathrm{TC}(P)$. Show for each case that voter i cannot manipulate by deviating from P to P'.

9.4 $\succsim^K$*-strategyproofness of* BP

Show that BP satisfies strong monotonicity and thus is $\succsim^K$-strategyproof.

9.5 *Strategyproofness for majoritarian SCFs*

Let $\mathcal{D} = \mathcal{S}^N$. Show that every majoritarian SCF f that satisfies $\succsim^K$-strategyproofness and non-imposition is Condorcet-consistent.

Hint: First show that f is 0-unanimous.

Part IV

Convex Sets of Alternatives

Casting the lot puts an end to disputes and decides between powerful contenders.

Solomon, c. 900 BC

10 Probabilistic Social Choice

Sonnenschein (1971)
Fishburn (1982)
von Neumann and Morgenstern (1947)
Harsanyi (1955)
Pattanaik and Peleg (1986)
Zeckhauser (1969)
Kalai and Schmeidler (1977); Hylland (1980b)
Hylland (1980a)
Gibbard (1977)
Kreweras (1965); Fishburn (1984a)
Brandl et al. (2016)
Brandl and Brandt (2020)

11
Budget Aggregation

Appendices

A

Notation and Terminology

Choice Theory	
$U = \{a, b, \dots\}$	Universe of m alternatives
$\mathcal{P}^*(U)$	Set of nonempty subsets of U ($2^U \setminus \{\emptyset\}$)
$\mathcal{F} = \{A, B, \dots\}$	Set of feasible subsets (aka agendas) of U
$f: \mathcal{F} \to \mathcal{P}^*(U)$	Choice function ($\forall A \in \mathcal{F}\colon f(A) \subseteq A$)
$\succsim$	Binary relation (typically on U)
$\succ$	Strict part of $\succsim$ ($x \succ y \Leftrightarrow x \succsim y \wedge \neg(y \succsim x)$)
$\sim$	Indifference part of $\succsim$ ($x \sim y \Leftrightarrow x \succsim y \wedge y \succsim x$)
$\succsim\|_A$	Restriction of $\succsim$ to A ($\succsim \cap (A \times A)$)
$Max(A, \succsim)$	Set of maximal elements ($\{x \in A\colon \nexists y \in A,\ y \succ x\}$)
$\succsim_f$	Base relation of f ($x \succsim_f y \Leftrightarrow x \in f(\{x, y\})$)
$X \succsim_f Y$	Extended base relation ($X \succsim_f Y \Leftrightarrow X = f(X \cup Y)$)
$\succsim$ is **transitive**	$\forall x, y, z \in U\colon x \succsim y \wedge y \succsim z \Rightarrow x \succsim z$
$\succsim$ is **quasi-transitive**	$\forall x, y, z \in U\colon x \succ y \wedge y \succ z \Rightarrow x \succ z$
$\succsim$ is **acyclic**	$\forall x_1, \dots, x_n \in U\colon x_1 \succ x_2 \wedge \dots \wedge x_{n-1} \succ x_n \Rightarrow \neg(x_n \succ x_1)$
f is **rationalizable**	$\exists \succsim \subseteq U \times U\colon \forall A \in \mathcal{F}\colon f(A) = Max(A, \succsim)$
f satisfies **contraction**	$\forall A, B \in \mathcal{F}\colon B \subseteq A \Rightarrow f(A) \cap B \subseteq f(B)$
f satisfies **expansion**	$\forall A, B \in \mathcal{F}\colon f(A) \cap f(B) \subseteq f(A \cup B)$
f satisfies **strong expansion**	$\forall A, B \in \mathcal{F}\colon B \subseteq A \wedge f(A) \cap B \neq \emptyset \Rightarrow f(B) \subseteq f(A)$
f satisfies **WARP**	f satisfies contraction and strong expansion
f is **set-rationalizable**	$\exists \succsim \subseteq \mathcal{P}^*(U) \times \mathcal{P}^*(U)\colon \forall A, X \in \mathcal{P}^*(U)\colon X = f(A) \Leftrightarrow X \in Max(\mathcal{P}^*(A), \succsim)$
f satisfies **set-contraction**	$\forall A, B \in \mathcal{F}\colon f(A) \subseteq B \subseteq A \Rightarrow f(A) = f(B)$
f satisfies **set-expansion**	$\forall A, B \in \mathcal{F}\colon f(A) = f(B) \Rightarrow f(A \cup B) = f(A)$
f is **stable**	f satisfies set-contraction and set-expansion

Basics of Social Choice	
$f: \mathcal{F} \times \mathcal{D} \to \mathcal{P}^*(U)$	Social choice function (SCF, $f(A, P) \subseteq A$)
$N = \{1, \dots, n\}$	Finite set of voters (aka electorate)
$\mathcal{W}$	Weak preference relations over U ($\{\succsim \subseteq U \times U\colon \succsim$ is transitive and complete$\}$)

$\mathcal{S}$	Strict preference relations over U ($\{\succsim \in \mathcal{W}\colon \forall x,y \in U\colon x \sim y \Rightarrow x = y\}$)
$\mathcal{D}$	Domain (e.g., $\mathcal{W}^N$)
$P = (\succsim_1, \dots, \succsim_n)$	Preference profile
N_P	Voters present in P (e.g., $N_P = N$ if $P \in \mathcal{W}^N$)
$P\|_A$	Restriction of P to A ($(\succsim_{i_1}\|_A, \dots, \succsim_{i_n}\|_A)$)
P_{-i}	P without voter i ($P \setminus \{(i, \succsim_i)\}$)
N_{xy}	Set of voters who weakly prefer x to y ($\{i \in N\colon x \succsim_i y\}$)
n_{xy}	Number of voters who weakly prefer x to y ($\|N_{xy}\|$)
$\succsim_M$	Pairwise majority relation ($x \succsim_M y \Leftrightarrow n_{xy} \geqslant n_{yx}$)
$\mathrm{PO}(A, P)$	Set of Pareto-optimal alternatives in A
$\mathrm{Cond}(A, P)$	Set of weak Condorcet winners ($\mathrm{Max}(A, \succsim_M)$)
f is **trivial**	$\forall A \in \mathcal{F}, P \in \mathcal{D}\colon f(A,P) = A$
f is a **refinement** of f'	$\forall A \in \mathcal{F}, P \in \mathcal{D}\colon f(A,P) \subseteq f'(A,P)$
f is **anonymous**	$\forall A \in \mathcal{F}, P, P' \in \mathcal{D}\colon (\exists$ bijection $\pi\colon N_P \to N_{P'}$ s.t. $\forall i \in N_P\colon \succsim_i = \succsim'_{\pi(i)}) \Rightarrow f(A,P) = f(A,P')$
f is **neutral**	$\forall A, B \in \mathcal{F}, P, P' \in \mathcal{D}$ with $N_P = N_{P'}\colon (\exists$ bijection $\pi\colon A \to B$ s.t. $\forall i \in N_P, \forall x,y \in A\colon x \succsim_i y \Leftrightarrow \pi(x) \succsim'_i \pi(y)) \Rightarrow f(B,P') = \{\pi(x)\colon x \in f(A,P)\}$
x **Pareto-dominates** y	$\forall i \in N_P\colon x \succ_i y$
x is **Pareto-optimal**	$\forall y \in A\colon \neg(y$ Pareto-dominates $x)$
f is **non-imposing**	$\forall A \in \mathcal{F}, x \in A\colon \exists P \in \mathcal{D}\colon f(A,P) = \{x\}$
f is **resolute**	$\forall A \in \mathcal{F}, P \in \mathcal{D}\colon \|f(A,P)\| = 1$
f is **monotonic**	$\forall P, P' \in \mathcal{D}, i \in N, a \in U\colon (P_{-i} = P'_{-i} \wedge \forall x,y \in U \setminus \{a\}\colon (x \succsim_i y \Leftrightarrow x \succsim'_i y) \wedge (a \succsim_i y \Rightarrow a \succsim'_i y) \wedge (a \succ_i y \Rightarrow a \succ'_i y)) \Rightarrow (\forall A \in \mathcal{F}\colon a \in f(A,P) \Rightarrow a \in f(A,P'))$
f is **positively responsive**	(as above) $\Rightarrow (\forall A \in \mathcal{F}\colon a \in f(A,P) \wedge P\|_A \neq P'\|_A \Rightarrow \{a\} = f(A,P'))$
x is a **Condorcet winner**	$\forall y \in A \setminus \{x\}\colon x \succ_M y$
x is a **weak Condorcet winner**	$\forall y \in A\colon x \succsim_M y$

Monotonicity$_2$, positive responsiveness$_2$, Pareto-optimality$_2$, etc. denote weakenings of axioms that only require the given property for feasible sets of size 2.

Arrow's Theorem

$g\colon \mathcal{D} \to \mathcal{W}$	Social welfare function (SWF)
$\succsim = g(P)$	Collective preference relation
$G \subseteq N$	Subset of voters
f satisfies **IIA**	$\forall A \in \mathcal{F}, P, P' \in \mathcal{D}\colon P\|_A = P'\|_A \Rightarrow f(A,P) = f(A,P')$
f is **Pareto-optimal**	$\forall A \in \mathcal{F}, P \in \mathcal{D}, x, y \in A\colon (\forall i \in N\colon x \succ_i y) \Rightarrow y \notin f(A,P)$
f is **dictatorial**	$\exists i \in N\colon \forall A \in \mathcal{F}, P \in \mathcal{D}\colon f(A,P) \subseteq \mathrm{Max}(A, \succsim_i)$
f is **reversely dictatorial**	$\exists i \in N\colon \forall A \in \mathcal{F}, P \in \mathcal{D}\colon f(A,P) \subseteq \mathrm{Max}(A, \precsim_i)$
g satisfies **IIA**	$\forall P, P' \in \mathcal{D}, x, y \in U\colon P\|_{\{x,y\}} = P'\|_{\{x,y\}} \Rightarrow \succsim\|_{\{x,y\}} = \succsim'\|_{\{x,y\}}$

g is **Pareto-optimal**	$\forall P \in \mathcal{D}, x, y \in U\colon (\forall i \in N\colon x \succ_i y) \Rightarrow x \succ y$
g is **dictatorial**	$\exists i \in N\colon \forall P \in \mathcal{D}, x, y \in U\colon x \succ_i y \Rightarrow x \succ y$
$a\ D_G\ b$	$\forall P \in \mathcal{D}\colon (\forall i \in G\colon a \succ_i b) \Rightarrow a \succ b$
$a\ \tilde{D}_G\ b$	$\forall P \in \mathcal{D}\colon (\forall i \in G\colon a \succ_i b \wedge \forall j \notin G\colon b \succ_j a) \Rightarrow a \succ b$
G is **decisive**	$\forall x, y \in U\colon x\ D_G\ y$
i is a **weak dictator**	$\forall P \in \mathcal{D}, x, y\colon x \succ_i y \Rightarrow x \succsim y$
i is a **quasi-dictator**	i is a weak dictator and $\forall G \subseteq N\colon i \in G \wedge \lvert G\rvert \geqslant 2 \Rightarrow G$ is decisive
G is an **oligarchy**	G is decisive and $\forall i \in G\colon i$ is a weak dictator
G is a **collegium**	$G = \bigcap_{G' \subseteq N \text{ is decisive}} G' \neq \emptyset$

The definitions for SWFs can be transferred to equivalent definitions for SCFs by using that $x \succsim y \Leftrightarrow x \succsim_f y$.

Restricted Domains

$\mathcal{S}^N$	Domain of strict preference profiles
$\mathcal{S}^{\subseteq \mathbb{N}}$	$\bigcup_{N' \in \mathcal{P}^*(\mathbb{N}),\ N' \text{ is finite}} \mathcal{S}^{N'}$ where $N \subseteq \mathbb{N}$
$\mathcal{D}_{\text{TOUR}}$	$\{P \in \mathcal{S}^{\subseteq \mathbb{N}}\colon \forall x, y \in U\colon n_{xy} \neq n_{yx}\}$
$\mathcal{D}_{\text{TRANS}}$	$\{P \in \mathcal{W}^N\colon\ \succsim_M$ is transitive$\}$
$\mathcal{D}_{\text{QTRANS}}$	$\{P \in \mathcal{W}^N\colon\ \succsim_M$ is quasi-transitive$\}$
$\mathcal{W}^N_{\text{DICH}}$	Dichotomous profiles ($\{\succsim \in \mathcal{W}\colon \forall x, y, z \in U\colon x \succ y \Rightarrow z \sim x \vee z \sim y\}^N$)
$\mathcal{S}^N_{\text{SP}(\geqslant)}$	Single-peaked profiles w.r.t. $\geqslant$ ($\{\succsim \in \mathcal{S}\colon \forall x, y, z \in U\colon (x > y > z) \vee (z > y > x) \Rightarrow (x \succ y \Rightarrow y \succ z)\}^N$)
$\mathcal{S}^N_{\text{SC}(\geqslant)}$	Single-caved profiles w.r.t. $\geqslant$ ($\{\succsim \in \mathcal{S}\colon \forall x, y, z \in U\colon (x > y > z) \vee (z > y > x) \Rightarrow (y \succ x \Rightarrow z \succ y)\}^N$)
$\mathcal{D}$ is **Cartesian**	$\exists \mathcal{R} \subseteq \mathcal{W}\colon \mathcal{D} = \mathcal{R}^N$
$\mathcal{D}$ is **value-restricted**	$\mathcal{D} = \mathcal{R}^N$ with $\mathcal{R} \subseteq \mathcal{S}$ s.t. $\forall x, y, z \in U$, there is some alternative, say x, such that for all $\succsim \in \mathcal{R}$, $(x \succ y) \vee (x \succ z)$ or $(y \succ x) \vee (z \succ x)$ or $((x \succ y) \wedge (x \succ z)) \vee ((y \succ x) \wedge (z \succ x))$.

Scoring Rules and Condorcet Extensions

$s = (s_1, \ldots, s_m)$	Score vector
f_s	Scoring rule with score vector s
BO	Borda's rule
$(n_{xy})_{x,y \in U}$	Support matrix
$(n_{xy} - n_{yx})_{x,y \in U}$	Margin matrix
CO	Copeland's rule
MM	Maximin rule
YO	Young's rule
$h\colon \mathcal{D} \to \mathcal{P}^*(\mathcal{S})$	Social preference function (SPF)
$\succsim \in h(P)$	Collective preference relation

KE	Kemeny's rule
f satisfies **reinforcement**	$\forall A \in \mathcal{F}, P, P' \in \mathcal{D}$ with $N_P \cap N_{P'} = \emptyset\colon f(A,P) \cap f(A,P') \neq \emptyset \Rightarrow f(A,P) \cap f(A,P') = f(A, P \cup P')$
f satisfies **cancellation**	$\forall A \in \mathcal{F}, P \in \mathcal{D}\colon (\forall x,y \in A\colon n_{xy} = n_{yx}) \Rightarrow f(A,P) = A$
f is **Condorcet-consistent**	$x \in A$ is a Condorcet winner $\Rightarrow f(A,P) = \{x\}$
f is **strongly Cond.-cons.**	$x \in A$ is a Condorcet winner $\Leftrightarrow f(A,P) = \{x\}$
f is **majoritarian**	f is neutral and $\forall P, P' \in \mathcal{D}, A \in \mathcal{F}\colon (\succsim_M = \succsim_f) \wedge (\succ_M = \succ'_M \Rightarrow f(A,P) = f(A,P'))$
f is **margin-based**	f is neutral and $\forall P, P' \in \mathcal{D}, A \in \mathcal{F}\colon (\forall x,y \in A\colon n_{xy} - n_{yx} = n'_{xy} - n'_{yx}) \Rightarrow f(A,P) = f(A,P')$
x is a **Condorcet loser** in A	$\forall y \in A \setminus \{x\}\colon y \succ_M x$
h is **anonymous**	$\forall P, P' \in \mathcal{D}\colon (\exists$ bijection $\pi\colon N_P \to N_{P'}$ s.t. $\forall i \in N_P\colon \succsim_i = \succsim'_{\pi(i)}) \Rightarrow h(P) = h(P')$
h is **neutral**	$\forall P, P' \in \mathcal{D}$ with $N_P = N_{P'}\colon (\exists$ permutation $\pi\colon U \to U$ s.t. $P' = \pi(P)) \Rightarrow \pi(h(P)) = h(P')$
h satisfies **reinforcement**	$\forall P, P' \in \mathcal{D}$ with $N_P \cap N_{P'} = \emptyset\colon h(P) \cap h(P') \neq \emptyset \Rightarrow h(P) \cap h(P') = h(P \cup P')$
h is **majority-consistent**	$\forall P \in \mathcal{D}, \succsim \in h(P), x, y \in U$ with $x \succ y$ adjacent in $\succsim\colon x \succsim_M y$; moreover, $x \sim_M y$ implies $\succsim \setminus \{(x,y)\} \cup \{(y,x)\} \in h(P)$
h satisfies **LIIA**	$\forall P, P' \in \mathcal{D}, \succsim \in h(P), \succsim' \in h(P'), x, y \in U$ adjacent in both $\succsim$ and $\succsim'\colon P\vert_{\{x,y\}} = P'\vert_{\{x,y\}} \Rightarrow \succsim \setminus \{(x,y),(y,x)\} \cup \succsim'\vert_{\{x,y\}} \in h(P)$
h satisfies **Pareto-optimality**	$\forall P \in \mathcal{D}, x, y \in U, \succsim \in h(P)\colon (\forall i \in N\colon x \succ_i y) \Rightarrow x \succ y$

Majoritarian SCFs

$D(x)$	Dominion of x ($\{y \in A\colon x \succ_M y\}$)
$D^k(x)$	$\{y \in A\colon \exists x_0 = x, x_1, \dots, x_k = y\colon x_0 \succ_M x_1 \succ_M \dots \succ_M x_k\}$
$D^*(x)$	$\bigcup_{k \geq 0} D^k(x)$
$\bar{D}(x)$	Dominators of x ($\{y \in A\colon y \succ_M x\}$)
$\bar{D}^k(x)$	$\{y \in A\colon \exists x_0 = y, x_1, \dots, x_k = x\colon x_0 \succ_M x_1 \succ_M \dots \succ_M x_k\}$
$\bar{D}^*(x)$	$\bigcup_{k \geq 0} \bar{D}^k(x)$
$Dom(A, \succ_M)$	Set of dominant sets
$\succsim^*_M$	Transitive closure of $\succsim_M$ ($x \succsim^*_M y \Leftrightarrow x \in \bar{D}^*(y)$)
C	Covering relation ($x \mathrel{C} y \Leftrightarrow D(y) \subset D(x)$)
$Trans(A, \succ_M)$	Set of transitive sets
p	Probability distribution $p : A \to [0,1]$ with $\sum_{x \in A} p(x) = 1$
$\Delta(A)$	Set of all probability distributions over A ($\{p\colon A \to [0,1]\colon \sum_{x \in A} p(x) = 1\}$)
$v_p(x)$	Value of p against x ($\sum_{y \in \bar{D}(x)} p(y) - \sum_{y \in D(x)} p(y)$)
TC	Top cycle
UC	Uncovered set
BA	Banks set

BP	Bipartisan set
B is **dominant**	$\forall x \in B, y \in A \setminus B\colon x \succ_M y$
B is **transitive**	$\succ_M\|_B$ is transitive
f satisfies **dominator exp.**	$\forall A \in \mathcal{F}, x \in A$ with $\bar{D}(x) \neq \emptyset\colon f(\bar{D}(x)) \subseteq f(A)$
p is **optimal**	$\forall x \in A\colon \nu_p(x) \geqslant 0$

Strategic Manipulation

$\succsim_i^S$	Relation in which all $x \in S$ are moved to the top of $\succsim_i$
P^S	$(\succsim_1^S, \dots, \succsim_n^S)$
$\hat{\succsim}$	Preference extension over sets ($\hat{\succsim} \subseteq \mathcal{P}^*(U) \times \mathcal{P}^*(U)$)
$\succsim^K$	Kelly's extension ($X \succsim^K Y \Leftrightarrow \forall x \in X, y \in Y\colon x \succsim y$)
$\succsim^F$	Fishburn's extension ($X \succsim^F Y \Leftrightarrow (\forall x \in X \setminus Y, y \in Y\colon x \succsim y) \wedge (\forall x \in X, y \in Y \setminus X\colon x \succsim y)$)
f is **strategyproof**	$\nexists P, P' \in \mathcal{D}, i \in N\colon P_{-i} = P'_{-i} \wedge f(P') \succ_i f(P)$
f is **strongly monotonic**	$\forall P, P' \in \mathcal{D}, i \in N, x \in U \setminus f(P), y \in U\colon\ P = P'$ except $x \succ_i y$ and $y \succ'_i x \Rightarrow f(P') = f(P)$
f satisfies **participation**	$\nexists P \in \mathcal{D}, i \in N\colon f(P \setminus \{(i, \succsim_i)\}) \succ_i f(P)$
f is **$\hat{\succsim}$-strategyproof**	$\nexists P, P' \in \mathcal{D}, i \in N\colon P_{-i} = P'_{-i} \wedge f(P') \hat{\succ}_i f(P)$
f satisfies **$\hat{\succsim}$-participation**	$\nexists P = (\succsim_1, \dots, \succsim_n) \in \mathcal{D}, i \in N\colon f(P \setminus \{(i, \succsim_i)\}) \hat{\succ}_i f(P)$
f satisfies **set non-imposition**	$\forall A \in \mathcal{F}, X \in \mathcal{P}^*(A), \exists P \in \mathcal{D}\colon f(A, P) = X$

B

Algorithms and Computational Complexity

This appendix provides a brief review of some key concepts and terminology from theoretical computer science.

Complexity theory deals with *complexity classes* of problems that are computationally equivalent in a certain well-defined way. Typically, problems that can be solved by an algorithm whose running time is polynomial in the size of the problem instance are considered *tractable*, whereas problems that do not admit such an algorithm are deemed *intractable*. Formally, an algorithm is *polynomial-time* if there exists a $k \in \mathbb{N}$ such that its running time is in $O(n^k)$, where n is the size of the input. Here, $O(n^k)$ denotes the class of all functions that, for large values of n, grow no faster than $c \cdot n^k$ for some constant number c (this is the so-called "*Big-O notation*"). For instance, when $k = 1$, the running time is *linear*, and when $k = 2$, the running time is *quadratic* in n. The class of decision problems that can be solved in polynomial time is denoted by P, whereas NP (for "nondeterministic polynomial time") refers to the class of decision problems whose solutions can be *verified* in polynomial time. For instance, the problem of deciding whether a directed graph admits a path from vertex s to vertex t is in P, whereas deciding whether there exists a path from s to t which contains at least k edges is only known to be in NP. The famous *P≠NP* conjecture states that the hardest problems in NP do not admit polynomial-time algorithms and are thus not contained in P. Although this statement remains unproven, it is widely believed to be true.

complexity classes

polynomial time

Big-O notation

P≠NP

Hardness of a problem for a particular class intuitively means that the problem is no easier than any other problem in that class. Both membership and hardness are established in terms of *reductions* that transform instances of one problem into instances of another problem using computational means appropriate for the complexity class under consideration. In the context of this book, we will mostly be interested in reductions that can be computed in time polynomial in the size of the problem instances. Finally, a problem is said to be *complete* for a complexity class if it is both contained in and hard for that class. The above-mentioned problem of deciding whether two vertices of a directed graph are connected via a path that exceeds a certain length was shown to be NP-complete.

hardness

reductions

NP-completeness

Given the current state of complexity theory, we cannot prove the *actual* intractability of most algorithmic problems, but merely give *evidence* for their intractability. Showing NP-hardness of a problem is commonly regarded as very strong evidence for computational intractability because it relates the problem to a large class of problems for which no efficient, i.e., polynomial-time, algorithm is known, despite enormous efforts to find such algorithms. As a consequence, finding a shortest path is easy while finding a longest path is considered hard.

If computing an aggregation function is intractable, its applicability is seriously undermined. Essentially, it means that any known algorithm for computing this function is asymptotically as bad as one that exhaustively enumerates all possible solutions.

For more detailed introductions and more extensive overviews of complexity classes, the reader is referred to Garey and Johnson (1979), Papadimitriou (1994), Wegener (2005), Arora and Barak (2009), and Sipser (2013).

Bibliography

Ron Adelsman and Andrew Whinston. The equivalenece of three social decision functions. *RAIRO. Recherche opérationnelle*, 11(3):257–265, 1977.

Fuad Aleskerov, Daniel Karabekyan, M. Remzi Sanver, and Vyacheslav Yakuba. On the manipulability of voting rules: The case of 4 and 5 alternatives. *Mathematical Social Sciences*, 64:67–73, 2012.

Alnur Ali and Marina Meila. Experiments with Kemeny ranking: What works when? *Mathematical Social Sciences*, 2012. Forthcoming.

Michael Allingham. *Choice Theory. A Very Short Introduction*. Very Short Introductions. Oxford University Press, 2002.

Josh Alman, Ran Ruan, Virginia Vassilevska Williams, Yinzhan Xu, Zixuan Xu, and Renfei Zhou. More asymmetry yields faster matrix multiplication. In *Proceedings of the 36th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2025.

Noga Alon. Ranking tournaments. *SIAM Journal on Discrete Mathematics*, 20(1):137–142, 2006.

Sanjeev Arora and Boaz Barak. *Computational Complexity: A Modern Approach*. Cambridge University Press, 2009.

Kenneth J. Arrow. A difficulty in the concept of social welfare. *The Journal of Political Economy*, 58(4):328–346, 1950.

Kenneth J. Arrow. *Social Choice and Individual Values*. New Haven: Cowles Foundation, 1st edition, 1951. 2nd edition 1963.

Kenneth J. Arrow. *Social Choice and Individual Values*. New Haven: Cowles Foundation, 2nd edition, 1963. 1st edition 1951.

Kenneth J. Arrow, Amartya Sen, and Kotaro Suzumura, editors. *Handbook of Social Choice and Welfare*, volume 1. North-Holland, 2002.

Kenneth J. Arrow, Amartya Sen, and Kotaro Suzumura, editors. *Handbook of Social Choice and Welfare*, volume 2. North-Holland, 2011.

David Austen-Smith and Jeffrey S. Banks. *Positive Political Theory I: Collective Preference*. University of Michigan Press, 1999.

David Austen-Smith and Jeffrey S. Banks. *Positive Political Theory II: Strategy and Structure*. University of Michigan Press, 2005.

Haris Aziz and Nisarg Shah. Participatory budgeting: Models and approaches. In Tamás Rudas and Gábor Péli, editors, *Pathways Between Social Science and Computational Social Science: Theories, Methods, and Interpretations*, pages 215–236. Springer International Publishing, 2021.

Georg Bachmeier, Felix Brandt, Christian Geist, Paul Harrenstein, Keyvan Kardel, Dominik Peters, and Hans Georg Seedig. k-majority digraphs and the hardness of voting with a constant number of voters. *Journal of Computer and System Sciences*, 105: 130–157, 2019.

David Balduzzi, Karl Tuyls, Julien Perolat, and Thore Graepel. Re-evaluting evaluation. In *Proceedings of the 32nd Conference on Neural Information Processing Systems (NeurIPS)*, pages 3272–3283, 2018.

Joseph M. Baldwin. The technique of the Nanson preferential majority system of election. *Transactions and Proceedings of the Royal Society of Victoria*, 39:42–52, 1926.

Michel Balinski. Le scrutin. *Pour la Science*, 294:46–51, 2002.

Michel Balinski and Rida Laraki. *Majority judgment: measuring, ranking, and electing*. MIT press, 2011.

Miguel A. Ballester and Guillaume Haeringer. A characterization of the single-peaked domain. *Social Choice and Welfare*, 36:305–322, 2011.

Jørgen Bang-Jensen and G. Gutin, editors. *Classes of Directed Graphs*. Springer Monographs in Mathematics. Springer, 2018.

Jeffrey S. Banks. Sophisticated voting outcomes and agenda control. *Social Choice and Welfare*, 1(4):295–306, 1985.

Jeffrey S. Banks. Acyclic social choice from finite sets. *Social Choice and Welfare*, 12:293–310, 1995a.

Jeffrey S. Banks. Singularity theory and core existence in the spatial model. *Journal of Mathematical Economics*, 24(6):523–536, 1995b.

Jeffrey S. Banks and Georges Bordes. Voting games, indifference, and consistent sequential choice rules. *Social Choice and Welfare*, 5:31–44, 1988.

Salvador Barberà. The manipulation of social choice mechanisms that do not leave "too much" to chance. *Econometrica*, 45(7):1573–1588, 1977.

Salvador Barberà. Indifferences and domain restrictions. *Analyse & Kritik*, 29:146–162, 2007.

Salvador Barberà, Bhaskar Dutta, and Arunava Sen. Strategy-proof social choice correspondences. *Journal of Economic Theory*, 101(2):374–394, 2001.

Salvador Barberà, Dolors Berga, and Bernardo Moreno. Arrow on domain conditions: A fruitful road to travel. *Social Choice and Welfare*, 54:237–258, 2020.

Jean Pierre Barthélemy and Bernard Monjardet. The median procedure in cluster analysis and social choice theory. *Mathematical Social Sciences*, 1:235–267, 1981.

John Bartholdi, III and James B. Orlin. Single transferable vote resists strategic voting. *Social Choice and Welfare*, 8(4):341–354, 1991.

John Bartholdi, III and Michael A. Trick. Stable matching with preferences derived from a psychological model. *Operations Research Letters*, 5(4):165–169, 1986.

John 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.

Antoinette Baujard, Herrade Igersheim, Isabelle Lebon, Frédéric Gavrel, and Jean-François Laslier. Who's favored by evaluative voting? An experiment conducted during the 2012 French presidential election. *Electoral Studies*, 34:131–145, 2014.

Antoinette Baujard, Frédéric Gavrel, Herrade Igersheim, Jean-François Laslier, and Isabelle Lebon. How voters use grade scales in evaluative voting. *European Journal of Political Economy*, 55(C):14–28, 2018.

Antoinette Baujard, Roberto Brunetti, Isabelle Lebon, and Simone Marsilio. How people understande voting rules. *Public Choice*, 2025. Forthcoming.

Avraham Beja. Arrow and gibbard-satterthwaite revisited: Extended domains and shorter proofs. *Mathematical Social Sciences*, 25(3):281–286, 1993.

Jean-Pierre Benoît. The Gibbard-Satterthwaite theorem: a simple proof. *Economics Letters*, 69(3):319–322, 2000.

Jeremy Bentham. *An Introduction to the Principles of Morals and Legislation*. 1789.

Therese Biedl, Franz-Josef Brandenburg, and Xiaotie Deng. On the complexity of crossings in permutations. *Discrete Mathematics*, 309(7):1813–1823, 2009.

Duncan Black. On the rationale of group decision-making. *Journal of Political Economy*, 56(1):23–34, 1948.

Duncan Black. *The Theory of Committees and Elections*. Cambridge University Press, 1958.

Douglas H. Blair and Robert A. Pollak. Acyclic collective choice rules. *Econometrica*, 50 (4):931–943, 1982.

Douglas H. Blair and Robert A. Pollak. Polychromatic acyclic tours in colored multigraphs. *Mathematics of Operations Research*, 8(3):471–476, 1983.

Douglas H. Blair, Georges Bordes, Jerry S. Kelly, and Kotaro Suzumura. Impossibility theorems without collective rationality. *Journal of Economic Theory*, 13:361–379, 1976.

Julian H. Blau. The existence of social welfare functions. *Econometrica*, 25:302–313, 1957.

Julian H. Blau and Rajat Deb. Social decision functions and the veto. *Econometrica*, 45(4): 871–879, 1977.

J. M. Blin and Andres B. Whinston. Discriminant functions and majority voting. *Management Science*, 21(5):557–566, 1975.

Georges Bordes. Consistency, rationality and collective choice. *Review of Economic Studies*, 43(3):451–457, 1976.

Tilman Börgers and Yan-Min Choo. Revealed relative utilitarianism. CESifo Working Paper No. 6613, 2017.

Walter Bossert and Kotaro Suzumura. Thre greatest unhappiness of the least number. *Social Choice and Welfare*, 47(1):187–205, 2016.

V. Joseph Bowman and Claude S. Colantoni. Majority rule under transitivity constraints. *Management Science*, 19(9):1029–1041, 1973.

Steven J. Brams and Peter C. Fishburn. Approval voting. *The American Political Science Review*, 72(3):831–847, 1978.

Steven J. Brams and Peter C. Fishburn. Voting procedures. In Kenneth J. Arrow, Amartya Sen, and Kotaro Suzumura, editors, *Handbook of Social Choice and Welfare*, volume 1, chapter 4. Elsevier, 2002.

Steven J. Brams and Peter C. Fishburn. *Approval Voting*. Springer-Verlag, 2nd edition, 2007.

Florian Brandl. Belief-averaging and relative utilitarianism. *Journal of Economic Theory*, 198, 2021.

Florian Brandl and Felix Brandt. Arrovian aggregation of convex preferences. *Econometrica*, 88(2):799–844, 2020.

Florian Brandl and Dominik Peters. Approval voting under dichotomous preferences: A catalogue of characterizations. *Journal of Economic Theory*, 205, 2022.

Florian Brandl, Felix Brandt, and Hans Georg Seedig. Consistent probabilistic social choice. *Econometrica*, 84(5):1839–1880, 2016.

Florian Brandl, Felix Brandt, Christian Geist, and Johannes Hofbauer. Strategic abstention based on preference extensions: Positive results and computer-generated impossibilities. *Journal of Artificial Intelligence Research*, 66:1031–1056, 2019.

Felix Brandt. *Tournament Solutions – Extensions of Maximality and Their Applications to Decision-Making*. Habilitation Thesis, Faculty for Mathematics, Computer Science, and Statistics, University of Munich, 2009.

Felix Brandt. Minimal stable sets in tournaments. *Journal of Economic Theory*, 146(4): 1481–1499, 2011.

Felix Brandt. Set-monotonicity implies Kelly-strategyproofness. *Social Choice and Welfare*, 45(4):793–804, 2015.

Felix Brandt. Stochastic choice and dynamics based on pairwise comparisons. In Mark Voorneveld, Jörgen W. Weibull, Tommy Andersson, Roger B. Myerson, Jean-François Laslier, Rida Laraki, and Yukio Koriyama, editors, *One Hundred Years of Game Theory: A Nobel Symposium*, Econometric Society Monographs, pages 158–166. Cambridge University Press, 2026.

Felix Brandt and Chris Dong. On locally rationalizable social choice functions. *Theory and Decision*, 2025. Forthcoming.

Felix Brandt and Felix Fischer. Computing the minimal covering set. *Mathematical Social Sciences*, 56(2):254–268, 2008.

Felix Brandt and Florian Grundbacher. The Banks set and the bipartisan set may be disjoint. *Social Choice and Welfare*, 66(2):349–355, 2026.

Felix Brandt and Paul Harrenstein. Set-rationalizable choice and self-stability. *Journal of Economic Theory*, 146(4):1721–1731, 2011.

Felix Brandt and Patrick Lederer. Characterizing the top cycle via strategyproofness. *Theoretical Economics*, 18(2):837–883, 2023.

Felix Brandt and Hans Georg Seedig. On the discriminative power of tournament solutions. In *Selected Papers of the International Conference on Operations Research, OR2014*, Operations Research Proceedings, pages 53–58. Springer-Verlag, 2016.

Felix Brandt, Felix Fischer, Paul Harrenstein, and Maximilian Mair. A computational analysis of the tournament equilibrium set. *Social Choice and Welfare*, 34(4):597–609, 2010.

Felix Brandt, Markus Brill, Edith Hemaspaandra, and Lane A. Hemaspaandra. Bypassing combinatorial protections: Polynomial-time algorithms for single-peaked electorates. *Journal of Artificial Intelligence Research*, 53:439–496, 2015a.

Felix Brandt, Andre Dau, and Hans Georg Seedig. Bounds on the disparity and separation of tournament solutions. *Discrete Applied Mathematics*, 187:41–49, 2015b.

Felix Brandt, Markus Brill, and Paul Harrenstein. Tournament solutions. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 3. Cambridge University Press, 2016a.

Felix Brandt, Christian Geist, and Paul Harrenstein. A note on the McKelvey uncovered set and Pareto optimality. *Social Choice and Welfare*, 46(1):81–91, 2016b.

Felix Brandt, Christian Geist, and Dominik Peters. Optimal bounds for the no-show paradox via SAT solving. *Mathematical Social Sciences*, 90:18–27, 2017. Special Issue in Honor of Hervé Moulin.

Felix Brandt, Markus Brill, and Paul Harrenstein. Extending tournament solutions. *Social Choice and Welfare*, 51(2):193–222, 2018a.

Felix Brandt, Markus Brill, Hans Georg Seedig, and Warut Suksompong. On the structure of stable tournament solutions. *Economic Theory*, 65(2):483–507, 2018b.

Felix Brandt, Martin Bullinger, and Patrick Lederer. On the indecisiveness of Kelly-strategyproof social choice functions. *Journal of Artificial Intelligence Research*, 73: 1093–1130, 2022a.

Felix Brandt, Marie Matthäus, and Christian Saile. Minimal voting paradoxes. *Journal of Theoretical Politics*, 34(4):527–551, 2022b.

Felix Brandt, Christian Saile, and Christian Stricker. Strategyproof social choice when preferences and outcomes may contain ties. *Journal of Economic Theory*, 202:105447, 2022c.

Felix Brandt, Chris Dong, and Dominik Peters. Condorcet-consistent choice among three candidates. *Games and Economic Behavior*, 153:113–130, 2025.

Donald J. Brown. Aggregation of preferences. *Quarterly Journal of Economics*, 89(3): 456–469, 1975.

Donald E. Campbell and Jerry S. Kelly. Impossibility theorems in the Arrovian framework. In Kenneth J. Arrow, Amartya Sen, and Kotaro Suzumura, editors, *Handbook of Social Choice and Welfare*, volume 1, chapter 1. Elsevier, 2002.

Donald E. Campbell and Jerry S. Kelly. A strategy-proofness characterization of majority rule. *Economic Theory*, 22(3):557–568, 2003.

Burak Can and Ton Storcken. Update monotone preference rules. *Mathematical Social Sciences*, 65(2):136–149, 2013.

Ioannis Caragiannis, Edith Hemaspaandra, and Lane A. Hemaspaandra. Dodgson's Rule and Young's Rule. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 5. Cambridge University Press, 2016.

Susumu Cato. Another induction proof of the Gibbard–Satterthwaite theorem. *Economics Letters*, 105(3):239–241, 2009.

Pierre Charbit, Stephan Thomassé, and Anders Yeo. The minimum feedback arc set problem is NP-hard for tournaments. *Combinatorics, Probability and Computing*, 16(1): 1–4, 2007.

Irène Charon and Olivier Hudry. A survey on the linear ordering problem for weighted or unweighted tournaments. *4OR*, 5(1):5–60, 2007.

Irène Charon and Olivier Hudry. An updated survey on the linear ordering problem for weighted or unweighted tournaments. *Annals of Operations Research*, 175(1):107–158, 2010.

Herman Chernoff. Rational selection of decision functions. *Econometrica*, 22(4):422–443, 1954.

Marquis de Condorcet. *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. Imprimerie Royale, 1785. Facsimile published in 1972 by Chelsea Publishing Company, New York.

Vincent Conitzer. Computing Slater rankings using similarities among candidates. In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI)*, pages 613–619, 2006.

Vincent Conitzer and Toby Walsh. Barriers to manipulation in voting. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 6. Cambridge University Press, 2016.

Vincent Conitzer, Tuomas Sandholm, and Jérôme Lang. When are elections with few candidates hard to manipulate? *Journal of the ACM*, 54(3), 2007.

Vincent Conitzer, Rachel Freedman, Jobst Heitzig, Wesley H. Holliday, Bob M. Jacobs, Nathan Lambert, Milan Mossé, Eric Pacuit, Stuart Russell, Hailey Schoelkopf, Emanuel Tewolde, and William S. Zwicker. Social choice should guide AI alignment in dealing with diverse human feedback. In *Proceedings of the 38th Conference on Neural Informatino Processing Systems (NeurIPS)*, 2024.

Arthur H. Copeland. A 'reasonable' social welfare function. Mimeo, University of Michigan Seminar on Applications of Mathematics to the Social Sciences, 1951.

Don Coppersmith and Shmuel Winograd. Matrix multiplication via arithmetic progressions. *Journal of Symbolic Computation*, 9(3):251–280, 1990.

Claude d'Aspremont and Louis Gevers. Social welfare functionals and interpersonal comparability. In Kenneth J. Arrow, Amartya Sen, and Kotaro Suzumura, editors, *Handbook of Social Choice and Welfare*, volume 1, chapter 10, pages 459–541. Elsevier, 2002.

Jessica Davies, George Katsirelos, Nina Narodytska, Toby Walsh, and Lirong Xia. Complexity of and algorithms for the manipulation of Borda, Nanson's and Baldwin's voting rules. *Artificial Intelligence*, 217:20–42, 2014.

Michiel S. De Vries, Juraj Nemec, and David Špaček. *International Trends in Participatory Budgeting: Between Trivial Pursuits and Best Practices*. Palgrave Macmillan, 2022.

Rajat Deb. On Schwartz's rule. *Journal of Economic Theory*, 16:103–110, 1977.

Bernard Debord. Caractérisation des matrices des préférences nettes et méthodes d'agrégation associées. *Mathématiques et sciences humaines*, 97:5–17, 1987.

Amrita Dhillon and Jean-François Mertens. Relative utilitarianism. *Econometrica*, 67(3): 471–498, 1999.

Mostapha Diss and Vincent R. Merlin, editors. *Evaluating Voting Systems with Probability Models: Essays by and in Honor of William Gehrlein and Dominique Lepelley*. Studies in Choice and Welfare. Springer Berlin Heidelberg, 2021.

Jean-Paul Doignon and Jean-Claude Falmagne. A polynomial time algorithm for unidimensional unfolding representations. *Journal of Algorithms*, 16(2):218–233, 1994.

Michelle Döring, Markus Brill, and Jobst Heitzig. The river voting method. In *Proceedings of the 40th AAAI Conference on Artificial Intelligence (AAAI)*, 2026.

John Duggan. Uncovered sets. *Social Choice and Welfare*, 41(3):489–535, 2013.

John Duggan. Limits of acyclic voting. *Journal of Economic Theory*, 163:658–683, 2016.

John Duggan and Michel Le Breton. Mixed refinements of Shapley's saddles and weak tournaments. *Social Choice and Welfare*, 18(1):65–78, 2001.

John Duggan and Thomas Schwartz. Strategic manipulability is inescapable: Gibbard-Satterthwaite without resoluteness. Caltech Social Science Working Paper 817, 1992.

John Duggan and Thomas Schwartz. Strategic manipulability without resoluteness or shared beliefs: Gibbard-Satterthwaite generalized. *Social Choice and Welfare*, 17(1): 85–93, 2000.

François Durand. Coalitional manipulation of voting rules: Simulations on empirical data. *Constitutional Political Economy*, 34:390–409, 2023.

François Durand. Why instant-runoff voting is so resilient to coalitional manipulation: Phase transitions in the perturbed culture. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 658–666, 2025.

Bhaskar Dutta and Jean-François Laslier. Comparison functions and choice correspondences. *Social Choice and Welfare*, 16(4):513–532, 1999.

Cynthia Dwork, Ravi Kumar, Moni Naor, and D. Sivakumar. Rank aggregation methods for the web. In *Proceedings of the 10th International Conference on the World Wide Web (WWW)*, pages 613–622, 2001.

Federico Echenique, Nicole Immorlica, and Vijay V. Vazirani, editors. *Online and Matching-Based Market Design*. Cambridge University Press, 2023.

Abdelhalim El Ouafdi, Dominique Lepelley, and Hatem Smaoui. On the Condorcet efficiency of evaluative voting (and other voting rules) with trichotomous preferences. *Annals of Operations Research*, 289:227–241, 2020.

Edith Elkind, Martin Lackner, and Dominik Peters. Structured preferences. In Ulle Endriss, editor, *Trends in Computational Social Choice*, chapter 10. 2017.

Edith Elkind, Martin Lackner, and Dominik Peters. Preference restrictions in computational social choice: A survey. Technical report, https://arxiv.org/pdf/2205.09092.pdf, 2025.

Peter Emerson. The original borda count and partial voting. *Social Choice and Welfare*, 40: 353–358, 2013.

P. Erdős and Leo Moser. On the representation of directed graphs as unions of orderings. *Publications of the Mathematical Institute of the Hungarian Academy of Science*, 9:125–132, 1964.

Bruno Escoffier, Jérôme Lang, and M. Öztürk. Single-peaked consistency and its complexity. In *Proceedings of the 18th European Conference on Artificial Intelligence (ECAI)*, pages 366–370. IOS Press, 2008.

Piotr Faliszewski and Ariel D. Procaccia. AI's war on manipulation: Are we winning? *AI Magazine*, 31(4):53–64, 2010.

Piotr Faliszewski, Edith Hemaspaandra, Lane A. Hemaspaandra, and Jörg Rothe. A richer understanding of the complexity of election systems. In S. Ravi and S. Shukla, editors, *Fundamental Problems in Computing: Essays in Honor of Professor Daniel J. Rosenkrantz*. Springer-Verlag, 2009.

Piotr Faliszewski, Edith Hemaspaandra, and Lane A. Hemaspaandra. Using complexity to protect elections. *Communications of the ACM*, 53(11):74–82, 2010.

Piotr Faliszewski, Edith Hemaspaandra, Lane A. Hemaspaandra, and Jörg Rothe. The shield that never was: Societies with single-peaked preferences are more open to manipulation and control. *Information and Computation*, 209(2):89–107, 2011.

Piotr Faliszewski, Edith Hemaspaandra, and Lane A. Hemaspaandra. The complexity of manipulative attacks in nearly single-peaked electorates. *Artificial Intelligence*, 207: 69–99, 2014.

George R. Feiwel, editor. *Arrow and the Foundations of the Theory of Economic Policy*. MacMillan, 1987a.

George R. Feiwel, editor. *Arrow and the Ascent of Modern Economic Theory*. MacMillan, 1987b.

Dan S. Felsenthal. Review of paradoxes afflicting various voting procedures where one out of m candidates ($m \geqslant 2$) must be elected. In *Electoral Systems*, pages 19–91. Springer-Verlag, 2012.

Dan S. Felsenthal and Hannu Nurmi. *Voting Procedures for Electing a Single Candidate*. Springer-Verlag, 2018.

Dan S. Felsenthal and Hannu Nurmi. *Voting Procedures Under a Restricted Domain*. Springer-Verlag, 2019.

Mark Fey. Choosing from a large tournament. *Social Choice and Welfare*, 31(2):301–309, 2008.

Alexander Fidora and Carles Sierra, editors. *Ramon Llull: From the Ars Magna to Artificial Intelligence*. Artificial Intelligence Research Institute, Barcelona, 2011.

Felix Fischer, Olivier Hudry, and Rolf Niedermeier. Weighted tournament solutions. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 4. Cambridge University Press, 2016.

Peter C. Fishburn. *The Theory of Social Choice*. Princeton University Press, 1973.

Peter C. Fishburn. Condorcet social choice functions. *SIAM Journal on Applied Mathematics*, 33(3):469–489, 1977.

Peter C. Fishburn. Axioms for approval voting: Direct proof. *Journal of Economic Theory*, 19(1):180–185, 1978.

Peter C. Fishburn. Nontransitive measurable utility. *Journal of Mathematical Psychology*, 26(1):31–67, 1982.

Peter C. Fishburn. Probabilistic social choice based on simple voting comparisons. *Review of Economic Studies*, 51(4):683–692, 1984a.

Peter C. Fishburn. Discrete mathematics in voting and group coherence. *SIAM Journal of Algebraic and Discrete Methods*, 5(2):263–275, 1984b.

Peter C. Fishburn. *Interprofile Conditions and Impossibility*, volume 18 of *Fundamentals of pure and applied economics. Social choice theory section*. Harwood Academic, 1987.

Peter C. Fishburn. Retrospective on the utility theory of von Neumann and Morgenstern. *Journal of Risk and Uncertainty*, 2(2):127–157, 1989.

Peter C. Fishburn. Acyclic sets of linear orders. *Social Choice and Welfare*, 14(1):113–124, 1997.

Peter C. Fishburn. Acyclic sets of linear orders: A progress report. *Social Choice and Welfare*, 19(2):431–447, 2002.

Peter C. Fishburn and Steven J. Brams. Paradoxes of preferential voting. *Mathematics Magazine*, 56(4):207–214, 1983.

Peter C. Fishburn, William V. Gehrlein, and Eric Maskin. Condorcet proportions and Kelly's conjectures. *Discrete Applied Mathematics*, 1(4):229–252, 1979.

David C. Fisher. Squaring a tournament: A proof of Dean's conjecture. *Journal of Graph Theory*, 23(1):43–48, 1996.

David C. Fisher and Richard B. Reeves. Optimal strategies for random tournament games. *Linear Algebra and its Applications*, 217:83–85, 1995.

David C. Fisher and Jennifer Ryan. Optimal strategies for a generalized "scissors, paper, and stone" game. *American Mathematical Monthly*, 99(10):935–942, 1992.

David C. Fisher and Jennifer Ryan. Tournament games and positive tournaments. *Journal of Graph Theory*, 19(2):217–236, 1995a.

David C. Fisher and Jennifer Ryan. Tournament games and Condorcet voting. *Linear Algebra and its Applications*, 217:87–100, 1995b.

David C. Fisher and Jennifer Ryan. Probabilities within optimal strategies for tournament games. *Discrete Applied Mathematics*, 56(1):87–91, 1995c.

Zack Fitzsimmons and Edith Hemaspaandra. Kemeny consensus complexity. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 196–202, 2021.

Marc Fleurbaey, Maurice Salles, and John A. Weymark, editors. *Justice, Political Liberalism, and Utilitarianism: Themes from Harsanyi and Rawls*. Cambridge University Press, 2008.

Ehud Friedgut, Gil Kalai, Nathan Keller, and Noam Nisan. A quantitative version of the Gibbard-Satterthwaite theorem for three alternatives. *SIAM Journal on Computing*, 40 (3):934–952, 2011.

Wulf Gaertner. *Domain Conditions in Social Choice Theory*. Cambridge University Press, 2001.

Wulf Gaertner. *A Primer in Social Choice Theory: Revised Edition*. LSE Perspectives in Economic Analysis. Oxford University Press, 2009.

Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.

John Geanakoplos. Three brief proofs of arrow's impossibility theorem. *Economic Theory*, 26:211–215, 2005.

William V. Gehrlein. Condorcet's paradox. *Theory and Decision*, 15(2):161–197, 1983.

William V. Gehrlein. *Condorcet's Paradox*. Springer-Verlag, 2006.

William V. Gehrlein and Peter C. Fishburn. Coincidence probabilities for simple majority and positional voting rules. *Social Science Research*, 7(3):272–283, 1978.

William V. Gehrlein and Peter C. Fishburn. Proportions of profiles with a majority candidate. *Computers & Mathematics with Applications*, 5(2):117–124, 1979.

William V. Gehrlein and Dominique Lepelley. *Elections, Voting Rules and Paradoxical Outcomes*. Studies in Choice and Welfare. Springer-Verlag, 2017.

Allan Gibbard. Social choice and the Arrow conditions. Seminar report, Harvard University, 1969.

Allan Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41(4): 587–601, 1973.

Allan Gibbard. Manipulation of schemes that mix voting with chance. *Econometrica*, 45 (3):665–681, 1977.

Itzhak Gilboa. *Rational Choice*. MIT Press, 2010.

Donald B. Gillies. Solutions to general non-zero-sum games. In Albert W. Tucker and R. Duncan Luce, editors, *Contributions to the Theory of Games IV*, volume 40 of *Annals of Mathematics Studies*, pages 47–85. Princeton University Press, 1959.

Paul Gölz, Nika Haghtalab, and Kunhe Yang. Distortion of ai alignment: Does preference optimization optimize for preferences? In *Proceedings of the 39th Conference on Neural Informatino Processing Systems (NeurIPS)*, 2025.

Irving John Good. A note on Condorcet sets. *Public Choice*, 10(1):97–101, 1971.

James Green-Armytage, T. Nicolaus Tideman, and Rafael Cosman. Statistical evaluation of voting rules. *Social Choice and Welfare*, 46(1):183–212, 2015.

A. S. Guha. Neutrality, monotonicity, and the right of veto. *Econometrica*, 40:812–826, 1972.

G. Hägele and Friedrich Pukelsheim. The electoral writings of Ramon Llull. *Studia Lulliana*, 41(97):3–38, 2001.

Bengt Hansson and Henrik Sahlquist. A proof technique for social choice with variable electorate. *Journal of Economic Theory*, 13(2):193–200, 1976.

John C. Harsanyi. Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. *Journal of Political Economy*, 63(4):309–321, 1955.

Edith Hemaspaandra, Holger Spakowski, and Jörg Vogel. The complexity of Kemeny elections. *Theoretical Computer Science*, 349(3):382–391, 2005.

Hans G. Herzberger. Ordinal preference and rational choice. *Econometrica*, 41(2):187–237, 1973.

John R. Hicks and Roy G. D. Allen. A reconsideration of the theory of value. Part II. *Economica*, 1(2):196–219, 1934.

Wesley H. Holliday and Eric Pacuit. Split cycle: A new Condorcet-consistent voting method independent of clones and immune to spoilers. *Public Choice*, 197(1–2):1–62, 2023.

Olivier Hudry. A smallest tournament for which the Banks set and the Copeland set are disjoint. *Social Choice and Welfare*, 16(1):137–143, 1999.

Olivier Hudry. A note on "Banks winners in tournaments are difficult to recognize" by G. J. Woeginger. *Social Choice and Welfare*, 23(1):113–114, 2004.

Olivier Hudry. NP-hardness results for the aggregation of linear orders into median orders. *Annals of Operations Research*, 163(1):63–88, 2008.

Olivier Hudry. A survey on the complexity of tournament solutions. *Mathematical Social Sciences*, 57(3):292–303, 2009.

Olivier Hudry. On the computation of median linear orders, of median complete preorders and of median weak orders. *Mathematical Social Sciences*, 64:2–10, 2012.

Aanund Hylland. Strategyproofness of voting procedures with lotteries as outcomes and infinite sets of strategies. Mimeo, 1980a.

Aanund Hylland. Aggregation procedure for cardinal preferences: A comment. *Econometrica*, 48(2):539–542, 1980b.

Herrade Igersheim, François Durand, Aaron Hamlin, and Jean-François Laslier. Comparing voting methods: 2016 us presidential election. *European Journal of Political Economy*, 71:102057, 2022.

Ken-Ichi Inada. A note on the simple majority decision rule. *Econometrica*, 32(4):525–531, 1964.

Marcus Isaksson, Guy Kindler, and Elchanan Mossel. The geometry of manipulation: A quantitative proof of the Gibbard-Satterthwaite theorem. *Combinatorica*, 32:221–250, 2012.

M. J. M. Jansen. Regularity and stability of equilibrium points of bimatrix games. *Mathematics of Operations Research*, 6(4):530–550, 1981.

Ehud Kalai and David Schmeidler. Aggregation procedure for cardinal preferences: A formulation and proof of Samuelson's impossibility conjecture. *Econometrica*, 45(6): 1431–1438, 1977.

Viggo Kann. *On the Approximability of NP-complete Optimization Problems*. PhD thesis, Royal Institute of Technology, Stockholm, 1992.

Richard M. Karp. Reducibility among combinatorial problems. In Raymond E. Miller and James W. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103. Plenum Press, 1972.

Jerry S. Kelly. Voting anomalies, the number of voters, and the number of alternatives. *Econometrica*, 42(2):239–251, 1974.

Jerry S. Kelly. Social choice and computational complexity. *Journal of Mathematical Economics*, 17:1–8, 1988.

Claire Kenyon-Mathieu and Warren Schudy. How to rank with few errors: A PTAS for weighted feedback arc set on tournaments. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing (STOC)*, pages 95–103, 2007.

Hadi Khalaf, Flavio P. Calmon, Daniel Halpern, Ariel D. Procaccia, Itai Shapira, and Serena Lutong Wang. Robust AI evaluation through maximal lotteries. 2026. Working paper.

Levin Damian Kiefer. Local reversal symmetry and the bipartisan set. Master's thesis, Technical University of Munich, 2024.

Bettina Klaus, David F. Manlove, and Francesca Rossi. Matching under preferences. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 14, pages 333–355. Cambridge University Press, 2016.

Gerald H. Kramer. A dynamical model of political equilibrium. *Journal of Economic Theory*, 16(2):310–334, 1977.

Germain Kreweras. Aggregation of preference orderings. In *Mathematics and Social Sciences I: Proceedings of the seminars of Menthon-Saint-Bernard, France (1–27 July 1960) and of Gösing, Austria (3–27 July 1962)*, pages 73–79, 1965.

Takashi Kurihara. A simple characterization of the anti-plurality rule. *Economics Letters*, 168:110–111, 2018.

Martin Lackner and Piotr Skowron. *Multi-Winner Voting with Approval Preferences*. Springer-Verlag, 2023.

Gilbert Laffond, Jean-François Laslier, and Michel Le Breton. The bipartisan set of a tournament game. *Games and Economic Behavior*, 5(1):182–201, 1993.

Steven P. Lalley and E. Glen Weyl. Quadratic voting: How mechanism design can radicalize domcracy. *American Economic Assoication Papers and Proceedings*, 108:33–37, 2018.

Michael Lampis. Determining a slater winner is complete for parallel access to np. In *Proceedings of the 39th International Symposium on Theoretical Aspects of Computer Science (STACS)*, Leibniz International Proceedings in Informatics (LIPIcs), pages 45:1–45:14. LZI, 2022.

Marc Lanctot, Kate Larson, Yoram Bachrach, Luke Marris, Zun Li, Avishkar Bhoopchand, Thomas Anthony, Brian Tanner, and Anna Koop. Evaluating agents using social choice theory. Technical report, https://arxiv.org/pdf/2312.03121, 2025.

Jean-François Laslier. *Tournament Solutions and Majority Voting*. Springer-Verlag, 1997.

Jean-François Laslier. Interpretation of electoral mixed strategies. *Social Choice and Welfare*, 17:283–292, 2000.

Jean-François Laslier. And the loser is... plurality voting. In D. Felsenthal and Moshé Machover, editors, *Electoral Systems*, Studies in Choice and Welfare, pages 327–351. Springer, 2011.

Jean-François Laslier. On choosing the alternative with the best median evaluation. *Public Choice*, 153(3/4):269–277, 2012.

Jean-François Laslier. The strange "Majority Judgment". *Revue Économique*, 70(4):569–588, 2019.

Jean-François Laslier and M. Remzi Sanver, editors. *Handbook on Approval Voting*. Studies in Choice and Welfare. Springer-Verlag, 2010.

Michel Le Breton and John A. Weymark. Arrovian social choice theory on economic domains. In Kenneth J. Arrow, Amartya Sen, and Kotaro Suzumura, editors, *Handbook of Social Choice and Welfare*, volume 2, chapter 17. North-Holland, 2011.

Charles Leedham-Green, Klas Markström, and Søren Riis. The largest Condorcet domain on 8 alternatives. *Social Choice and Welfare*, 62(1):109–116, 2024.

Dominique Lepelley. Une caractérisation du vote à la majorité simple. *RAIRO. Recherche opérationnelle*, 26(4):361–365, 1992.

Dominique Lepelley and Hatem Smaoui. Comparing two ways for eliminating candidates in three-alternative elections using sequential scoring rules. *Group Decision and Negotiation*, 28:787–804, 2019.

Lynn H. Loomis. On a theorem of von Neumann. *Proceedings of the National Academy of Sciences (PNAS)*, 32(8):213–215, 1946.

R. Duncan Luce and Howard Raiffa. *Games and Decisions: Introduction and Critical Survey*. John Wiley & Sons Inc., 1957.

Antonin Macé. Voting with evaluations: Characterizations of evaluative voting and range voting. *Journal of Mathematical Economics*, 79:10–17, 2018.

Michael Mandler. *Dilemmas in Economic Theory*. Oxford University Press, 1999.

François Maniquet and Philippe Mongin. Approval voting and Arrow's impossibility theorem. *Social Choice and Welfare*, 44(3):519–532, 2015.

David F. Manlove. *Algorithmics of Matching Under Preferences*. World Scientific Publishing Company, 2013.

Andreu Mas-Colell and Hugo Sonnenschein. General possibility theorems for group decisions. *Review of Economic Studies*, 39(2):185–192, 1972.

Eric Maskin and Amartya Sen. *The Arrow Impossiblity Theorem*. Columbia University Press, 2014.

Eric Maskin and Amartya Sen. How to let the majority rule. *New York Times*, page SR7, May 1st 2016.

Roberto-Rafael Maura-Rivero, Marc Lanctot, Francesco Visin, and Kate Larson. Jackpot! Alignment as a maximal lottery. Technical report, https://arxiv.org/abs/2501.19266, 2025. Working paper.

Kenneth May. A set of independent, necessary and sufficient conditions for simple majority decisions. *Econometrica*, 20(4):680–684, 1952.

Robert M. May. Some mathematical remarks on the paradox of voting. *Behavioral Science*, 16(2):143–151, 1971.

David C. McGarvey. A theorem on the construction of voting paradoxes. *Econometrica*, 21(4):608–610, 1953.

Richard D. McKelvey. Intransitivities in multidimensional voting models and some implications for agenda control. *Journal of Economic Theory*, 12(3):472–482, 1976.

Iain McLean. E. J. Nanson, social choice and electoral reform. *Australian Journal of Political Science*, 31(3):369–386, 1996.

Iain McLean and Fiona Hewitt. *Condorcet: Foundations of Social Choice and Political Theory*. Edward Elgar Publishing, 1994.

Iain McLean and Arnold Urken, editors. *Classics of Social Choice*. University of Michigan Press, 1995.

Deepak K. Merchant and M. R. Rao. Majority decisions and transitivity: Some special cases. *Management Science*, 23(2):109–222, 1976.

Nicholas R. Miller. A new solution set for tournaments and majority voting: Further graph-theoretical approaches to the theory of voting. *American Journal of Political Science*, 24(1):68–96, 1980.

Nicholas R. Miller, Bernard Grofman, and Scott L. Feld. The structure of the Banks set. *Public Choice*, 66(3):243–251, 1990.

Elchanan Mossel and M. Z. Rácz. A quantitative gibbard-satterthwaite theorem without neutrality. *Combinatorica*, 35(3):317–387, 2015. ISSN 1439-6912.

Hervé Moulin. On strategy-proofness and single peakedness. *Public Choice*, 35(4):437–455, 1980.

Hervé Moulin. *The Strategy of Social Choice*. North-Holland, 1983.

Hervé Moulin. Choice functions over a finite set: A summary. *Social Choice and Welfare*, 2:147–160, 1985a.

Hervé Moulin. From social welfare ordering to acyclic aggregation of preferences. *Mathematical Social Sciences*, 9(1):1–17, 1985b.

Hervé Moulin. Choosing from a tournament. *Social Choice and Welfare*, 3(4):271–291, 1986.

Hervé Moulin. *Axioms of Cooperative Decision Making*. Cambridge University Press, 1988a.

Hervé Moulin. Condorcet's principle implies the no show paradox. *Journal of Economic Theory*, 45(1):53–64, 1988b.

Eitan Muller and Mark A. Satterthwaite. The equivalence of strong positive association and strategy-proofness. *Journal of Economic Theory*, 14(2):412–418, 1977.

Rémi Munos, Michal Valko, Daniele Calandriello, Mohammad Gheshlaghi Azar, Mark Rowland, Zhaohan Daniel Guo, Yunhao Tang, Matthieu Geist, Thomas Mesnard, Andrea Michi, Marco Selvi, Sertan Girgin, Nikola Momchev, Olivier Bachem, Daniel J. Mankowitz, Doina Precup, and Bilal Piot. Nash learning from human feedback. In *Proceedings of the 38th Conference on Neural Informatino Processing Systems (NeurIPS)*, 2024.

Kenjiro Nakamura. The core of a simple game with ordinal preferences. *International Journal of Game Theory*, 4(2):95–104, 1975.

Edward J. Nanson. Methods of election. *Transactions and Proceedings of the Royal Society of Victoria*, 19:197–240, 1883.

Shmuel I. Nitzan. *Collective Preference and Choice*. Cambridge University Press, 2010.

Hannu Nurmi. *Comparing Voting Systems*. Theory and Decision Library. Springer-Verlag, 1987.

Hannu Nurmi. *Voting Paradoxes and How to Deal with Them*. Springer-Verlag, 1999.

Guillermo Owen. An elementary proof of the minimax theorem. *Management Science*, 13 (9):765, 1967.

Z. Emel Öztürk. Consistency of scoring rules: A reinvestigation of composition-consistency. *International Journal of Game Theory*, 49:801–831, 2020.

Christos H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.

Vilfredo Pareto. *Manuale di Economia Politica*. Società Editrice Libraria, Milano, 1906.

Prasanta K. Pattanaik and Bezalel Peleg. Distribution of power under stochastic social choice rules. *Econometrica*, 54(4):909–921, 1986.

Josep E. Peris and Begoña Subiza. Condorcet choice correspondences for weak tournaments. *Social Choice and Welfare*, 16(2):217–231, 1999.

Marcus Pivato. Formal utilitarianism and range voting. *Mathematical Social Sciences*, 67: 50–56, 2014.

Charles R. Plott. Path independence, rationality, and social choice. *Econometrica*, 41(6): 1075–1091, 1973.

Charles R. Plott. Axiomatic social choice theory: An overview and interpretation. *American Journal of Political Science*, 20(3):511–596, 1976.

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.

Clemens Puppe. The single-peaked domain revisited: A simple global characterization. *Journal of Economic Theory*, 176:55–80, 2018.

Clemens Puppe and Arkadii Slinko. *Condorcet Domains: The Mathematics of Coherent Collective Decision-Making*. Studies in Choice and Welfare. Springer-Verlag, 2026.

Tirukkannamangai Echambadi Srinivasa Raghavan. Zero-sum two person games. In Robert J. Aumann and Sergiu Hart, editors, *Handbook of Game Theory with Economic Applications*, chapter 20. North-Holland, 1994.

Philip J. Reny. Arrow's theorem and the Gibbard-Satterthwaite theorem: a unified approach. *Economics Letters*, 70(1):99–105, 2001.

Simon Rey, Felicia Schmidt, and Jan Maly. The (computational) social choice take on indivisible participatory budgeting. Technical report, https://arxiv.org/pdf/2303.00621.pdf, 2025.

F. Ricci, Lior Rokach, and Bracha Shapira, editors. *Recommender Systems Handbook*. Springer, 3rd edition, 2022.

Lionel Robbins. *An Essay on the Nature and Significance of Economic Science*. Macmillan, 1932.

Ariel Rubinstein. *Lecture Notes in Microeconomic Theory*. Princeton University Press, 2006.

Paul A. Samuelson. A note on the pure theory of consumers' behaviour. *Econometrica*, 5: 61–71, 1938.

Mark A. Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217, 1975.

Lisa Sauermann. On the probability of a Condorcet winner among a large number of alternatives. Technical report, arXiv.org, https://arxiv.org/abs/2203.13713v1, 2022.

David Schmeidler and Hugo Sonnenschein. Two proofs of the Gibbard-satterthwaite theorem on the possibility of a strategy-proof social choice function. In Hans W. Gottinger and Werner Leinfellner, editors, *Decision Theory and Social Ethics*, volume 17 of *Theory and Decision Library*, pages 227–234. Springer-Verlag, 1978.

Paul J. H. Schoemaker. The expected utility model: Its variants, purposes, evidence and limitations. *Journal of Economic Literature*, 20(2):529–563, 1982.

Norman Schofield. Instability of simple dynamic games. *Review of Economic Studies*, 45 (3):575–594, 1978.

Markus Schulze. A new monotonic and clone-independent single-winner election method. *Voting Matters*, 17:9–19, 2003.

Thomas Schwartz. Rationality and the myth of the maximum. *Noûs*, 6(2):97–117, 1972.

Thomas Schwartz. Choice functions, "rationality" conditions, and variations of the weak axiom of revealed preference. *Journal of Economic Theory*, 14:414–427, 1976.

Thomas Schwartz. *The Logic of Collective Choice*. Columbia University Press, 1986.

Yohei Sekiguchi. A characterization of the plurality rule. *Economics Letters*, 116(3): 330–332, 2012.

Amartya Sen. A possibility theorem on majority decisions. *Econometrica*, 34(2):491–499, 1966.

Amartya Sen. Quasi-transitivity, rational choice and collective decision. *Review of Economic Studies*, 36(3):381–393, 1969.

Amartya Sen. *Collective Choice and Social Welfare*. North-Holland, 1970. Expanded edition 2017.

Amartya Sen. Choice functions and revealed preference. *Review of Economic Studies*, 38 (3):307–317, 1971.

Amartya Sen. Social choice theory: A re-examination. *Econometrica*, 45(1):53–89, 1977.

Amartya Sen. Social choice theory. In Kenneth J. Arrow and Michael D. Intriligator, editors, *Handbook of Mathematical Economics*, volume 3, chapter 22, pages 1073–1181. Elsevier, 1986.

Amartya Sen. Rationality and social choice. *American Economic Review*, 85(1):1–24, 1995.

Amartya Sen. *Collective Choice and Social Welfare*. Penguin, expanded edition, 2017.

Amartya Sen and Prasanta K. Pattanaik. Necessary and sufficient conditions for rational choice under majority decision. *Journal of Economic Theory*, 1:178–202, 1969.

Arunava Sen. Another direct proof of the Gibbard-Satterthwaite Theorem. *Economics Letters*, 70(3):381–385, 2001.

Kenneth A. Shepsle and Barry R. Weingast. Uncovered sets and sophisticated outcomes with implications for agenda institutions. *American Journal of Political Science*, 28(1): 49–74, 1984.

Paul B. Simpson. On defining areas of voter choice: Professor Tullock on stable voting. *Quarterly Journal of Economics*, 83(3):478–490, 1969.

Michael Sipser. *Introduction to the Theory of Computation*. Cengage Learning, 3 edition, 2013.

Patrick Slater. Inconsistencies in a schedule of paired comparisons. *Biometrika*, 48(3–4): 303–312, 1961.

John H. Smith. Aggregation of preferences with variable electorate. *Econometrica*, 41(6): 1027–1041, 1973.

Hugo Sonnenschein. Demand theory without transitive preference with applications to the theory of competitive equilibrium. In J. Chipman, Leonid Hurwicz, M. Richter, and Hugo Sonnenschein, editors, *Preferences, Utility and Demand*. Houghton Mifflin Harcourt, 1971.

Volker Strassen. Gaussian elimination is not optimal. *Numerische Mathematik*, 13:354–356, 1969.

Warut Suksompong and Nicholas Teh. Voting in divisible settings: A survey. In *Proceedings of the 40th AAAI Conference on Artificial Intelligence (AAAI)*, 2026. Forthcoming.

Lars-Gunnar Svensson and Alexander Reffgen. The proof of the Gibbard-Satterthwaite theorem revisited. *Journal of Mathematical Economics*, 55:11–14, 2014.

Pingzhong Tang and Fangzhen Lin. Computer-aided proofs of Arrow's and other impossibility theorems. *Artificial Intelligence*, 173(11):1041–1053, 2009.

T. Nicolaus Tideman. Independence of clones as a criterion for voting rules. *Social Choice and Welfare*, 4(3):185–206, 1987.

Jill van Newenhizen. The Borda method is most likely to respect the Condorcet principle. *Economic Theory*, 2:69–83, 1992.

Oriol Vinyals, Igor Babuschkin, Wojciech M. Czarnecki, Michaël Mathieu, Andrew Dudzik, Jun-Young Chung, David H. Choi, Richard Powell, Timo Ewalds, Petko Georgiev, Junhyuk Oh, Dan Horgan, Manuel Kroiss, Ivo Danihelka, Aja Huang, Laurent Sifre, Trevor Cai, John P. Agapiou, Max Jaderberg, Alexander S. Vezhnevets, Rémi Leblond, Tobias Pohlen, Valentin Dalibard, David Budden, Yury Sulsky, James Molloy, Tom L. Paine, Caglar Gulcehre, Z. Wang, Tobias Pfaff, Yuhuai Wu, Roman

Ring, Dani Yogatama, D. Wünsch, Katrina McKinney, Oliver Smith, Tom Schaul, Timothy Lillicrap, Koray Kavukcuoglu, Demis Hassabis, Chris Apps, and David Silver. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, 575:350–354, 2019.

John von Neumann. Zur Theorie der Gesellschaftspiele. *Mathematische Annalen*, 100(1): 295–320, 1928.

John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.

John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 2nd edition, 1947.

John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 3rd edition, 1953.

Siyi Wang, Qi Deng, Shiwei Feng, Hong Zhang, and Chao Ling. A survey on rank aggregation. In *Proceedings of the 33rd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 8281–8289, 2024.

Benjamin Ward. Majority rule and allocation. *Journal of Conflict Resolution*, 5(4):379–389, 1961.

Ingo Wegener. *Complexity Theory: Exploring the Limits of Efficient Algorithms*. Springer-Verlag, 2005.

Robert B. Wilson. Social choice theory without the Pareto principle. *Journal of Economic Theory*, 5:478–486, 1972.

Gerhard J. Woeginger. Banks winners in tournaments are difficult to recognize. *Social Choice and Welfare*, 20(3):523–528, 2003.

H. Peyton Young. An axiomatization of Borda's rule. *Journal of Economic Theory*, 9(1): 43–52, 1974.

H. Peyton Young. Social choice scoring functions. *SIAM Journal on Applied Mathematics*, 28(4):824–838, 1975.

H. Peyton Young. Condorcet's theory of voting. *The American Political Science Review*, 82 (4):1231–1244, 1988.

H. Peyton Young. Optimal voting rules. *Journal of Economic Perspectives*, 9(1):51–64, 1995.

H. Peyton Young and Arthur B. Levenglick. A consistent extension of Condorcet's election principle. *SIAM Journal on Applied Mathematics*, 35(2):285–300, 1978.

Richard Zeckhauser. Majority rule with lotteries on alternatives. *Quarterly Journal of Economics*, 83(4):696–703, 1969.

Michael Zuckerman, Ariel D. Procaccia, and Jeffrey S. Rosenschein. Algorithms for the coalitional manipulation problem. *Artificial Intelligence*, 173(2):392–412, 2009.

William S. Zwicker. Introduction to the theory of voting. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 2. Cambridge University Press, 2016.