



HAL
open science

Steps towards decisiveness in social choice

Justin Kruger

► **To cite this version:**

Justin Kruger. Steps towards decisiveness in social choice. Operations Research [math.OC]. Université Paris sciences et lettres, 2019. English. NNT : 2019PSLED050 . tel-03220645

HAL Id: tel-03220645

<https://theses.hal.science/tel-03220645>

Submitted on 7 May 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

THÈSE DE DOCTORAT

de l'Université de recherche Paris Sciences et Lettres
PSL Research University

Préparée à Université Paris-Dauphine

Steps towards decisiveness in social choice

École doctorale n°543

Spécialité **INFORMATIQUE**

Soutenue par
Justin Kruger
le **26/06/2019**

Dirigée par **Jérôme Lang**



COMPOSITION DU JURY :

M. Stéphane Airiau, Paris-Dauphine
Co-encadrant de thèse

M. Markus Brill, TU Berlin
Membre du jury

Mme. Edith Elkind, Oxford University
Membre du jury

M. Jean Lainé, CNAM
Rapporteur

M. Jérôme Lang, Paris-Dauphine
Directeur de thèse

M. Jean-François Laslier, CNRS
Rapporteur

M. Nicolas Maudet, LIP6,
Président du jury

M. Remzi Sanver, Paris-Dauphine
Membre du jury

Contents

Introduction	3
Overview	4
Common definitions	6
1 Discipline aggregation, manipulation and decisiveness	8
1.1 The combined format for sports climbing	11
1.2 Formal description of the setting	14
1.3 Manipulation in the discipline aggregation framework	17
1.4 The likelihood of undesirable profiles	21
1.5 Impossibilities that suggest a trade-off between decisiveness and manipulability	23
1.6 How easy is it to determine whether or not to manipulate	32
1.7 Final remarks on discipline aggregation	38
2 Relieving the tension between neutrality and resoluteness using parallel uni-	43
verses	
2.1 Binary trees	47
2.1.1 Families of binary trees	48
2.1.2 Using binary trees to define social choice functions	50
2.2 Identities and properties of tree based rules	51
2.2.1 Singleton trees	52
2.2.2 Parallel universe trees	53
2.2.3 Argmax applied to trees	57
2.3 Success of tree-based argmax rules as refinements	59
2.4 Assessment of the possible universes method	60
3 Maximally decisive social choice	62
3.1 The informational content of families of functions	64
3.1.1 Examples of informational bases	65
3.1.2 Informational functions that express positional scoring rules	68
3.2 Maximally decisive rules with respect to bases	74
3.2.1 Borda and maximal decisiveness	75
3.2.2 Finding maximally decisive rules	78
3.3 Final remarks on maximally decisive social choice	80
4 Domain restrictions and independence of irrelevant alternatives	83
4.1 Definitions	84
4.2 A counterexample on restricted domains	85

4.3	The Blau equivalence	87
4.4	Violating the Blau equivalence	90
4.5	Blau's impossibility	93
4.6	Final remarks on domain restrictions and independence	97
5	Combining ordinal and evaluative information in social choice	99
5.1	The preference-approval model	101
5.2	An impossibility in aggregating preference-approvals	103
5.3	Possibilities from relaxing decomposability	106
5.3.1	Relaxing the decomposability of the approval aggregation	107
5.3.1.1	Well-behaved approval aggregators	107
5.3.1.2	Maximally discriminating approval aggregators	108
5.3.1.3	Relaxing decomposability with three or more alternatives	109
5.3.2	Relaxing decomposability between rankings and approvals	110
5.3.2.1	Shortlist by elementary approval aggregators then rank	111
5.3.2.2	Shortlist by non-decomposable approval aggregator	111
5.3.2.3	Perform ranking aggregation then approve using ranking	112
5.3.2.4	Borda with a movable zero determined by the zero-line	112
5.4	Extending the impossibility to more evaluation levels	113
5.5	Final remarks	115
	Conclusion	119
	Glossary	121

Introduction

This thesis is constructed from a series of articles concerning **social choice theory**. Some may be further specified as using the ideas of **computational** social choice theory. An analysis of the words “social”, “choice”, “theory” and “computational” provides a general idea of the type of results contained within.

The word “theory” indicates that social choice theory is a formal field of study. In particular, this thesis is mostly concerned with *mathematically* formal issues—as opposed to empirical issues. However, although the work presented here is not dependent upon observations of the real world, it does involve modelling potentially real situations (this is explicitly the case for the simulations of Chapter 1).

The “social” part connects us to what is modelled, which is, in a rough sense, society. Of course, the concept of society is somewhat amorphous. In our models we attempt to capture one aspect of society that is relevant to collective choice: the preferences that each agent has. Such preferences can be expressed in many ways. The traditional approach of social choice theory, due to Arrow (1950), models society as a **profile of ordinal preferences**. This approach, and others that we consider, specifies precisely what are to be taken as the relevant possible states of society.

The “choice” part refers to some method for selecting or for ranking alternatives. For our purposes, such a method is a function whose domain is the possible states of society. The output of the function might be a chosen alternative. We call such functions social choice functions. Other functions may output a ranking over the alternatives. We call such functions social welfare functions. We are less concerned with descriptive accounts of how this is done in practice, and more with the normative properties of potential methods—we study what choices *should* be made by a social choice procedure.

We now come to “computational”. Problems of a social choice theoretic nature have been found to be applicable to problems traditionally associated with computer science. In the opposite direction, computational methods have been applied to the traditional problems of social choice theory. This cross-fertilisation has led to the birth of computational social choice. Interdisciplinary by nature, this field of study now treats many problems from disciplines such as economics, computer science, artificial intelligence, mathematics, political science and philosophy. Across these diverse settings, a number of specific techniques have been developed that have a particularly computational-social-choice-theoretic nature. Of particular note is the application of the theory of **computational complexity**, which may (for example, negatively) show that a supposedly attractive social choice procedure is infeasible in practice (it may also be used positively to mitigate a negative result). Results using this

and other such concepts contribute to an understanding of the issues of social choice theory, not merely through performing simulations or calculations *using* computers, but thanks to the formal *concepts* and *ideas* of *theoretical* computer science.

Within the general framework of computational social choice theory, this thesis focuses on two main issues: how **expressive** the preferences of the agents are, and how **decisive** the choice procedure is. The first theme may be thought of as concerning *how much information we have to deal with*. The second theme considers what is done with this information. Now, the whole point of a choice procedure is to make a choice: if it often fails to do so, it will not be much use in practice. Along these lines, the second theme concerns *how small can we make the choice set, or how refined the output ranking*.

All of the chapters touch upon, to a lesser or greater degree, one or both of the themes of expressive-ness and decisive-ness. An overview of the specific content of each of the chapters follows.

Chapter summaries

Chapter 1 is intended to be an introduction to computational social choice theory that at the same time provides novel results. In it we apply the ideas of computational social choice theory to the problem of determining the winners for sports climbing at the 2020 Olympics. The method that will be used has been designed by the International Federation of Sports Climbing; we call this method **inverse-Borda-Nash**. Our nomenclature is intended to be descriptive—the method may be thought of as an inverse version of Borda scoring where the winning alternatives are those with *minimal* product scores.

Inverse-Borda-Nash has some undesirable behaviour: in particular, an athlete can increase a teammate's ranking by performing worse; we say that it is **manipulable**. Simulations suggest that cases where manipulation can occur are not unlikely. Given that the proposed method is manipulable, one may wonder whether there is a different, better method for ranking the athletes. We formalise this problem by giving a mathematical description of the properties we want the method to satisfy. Unfortunately, we cannot satisfy all the desired properties at the same time: we have an **impossibility**. Impossibility results of this type abound in social choice theory to such an extent that there are recognisable families of methods that have been applied in attempt to lessen their negative implications. One such family is to use the theory of computational complexity from theoretical computer science to show that, although undesirable behaviour may be theoretically possible, it is not computationally feasible. In this vein, we prove that it is **computationally hard** to detect a manipulation for inverse-Borda-Nash.

The reason that our problem in Chapter 1 can be considered in terms of social choice theory is due to the information that is available to rank the athletes. The International Federation of Sports Climbing expressly wants the final ranking to be produced from an aggregation of three *ordinal* rankings. (Their reasons for this echo traditional concerns about the difficulty of making **interpersonal comparisons of utility**.) This allows us to make a recap of the traditional issues of computational social choice theory that is nonetheless novel. Further, we will also introduce what we believe to be a novel problem, at least in the way that it is formulated here: the issue of decisiveness. Intuitively, a single gold medallist is preferable to a tied situation with multiple gold medallists. This chapter presents the open problem of whether or not there is a tension between the decisiveness of the choice method and the other desirable properties. This problem seems to require a more formal definition of what it means for a choice method to be decisive or not.

Chapter 2 expressly attempts to make choice methods more decisive. The technique used is to count permutations—a technique inspired by the “parallel universe” tie-breaking version of **instant runoff voting**. In Chapter 2 we apply this technique to the class of **tournament solutions**. Tournament solutions take as input a complete directed graph over alternatives: this may be considered an informational restriction, as the majority preferences of society over the alternatives is less expressive than a full profile of preferences. Tournament solutions are apt to be made more decisive as they tend to output large sets. By counting permutations we certainly succeed in making them more decisive; however, depending on the precise manner in which this is done, we sometimes get non-intersecting choice sets, and it is not clear why we should prefer one choice method to another. Another drawback is the computationally demanding nature of our technique. Intuitively, tournament solutions provide large outcome sets in part because they do not use a lot of information. Using the techniques of this chapter we can create solution concepts that are much more decisive than those that are normally studied.

In Chapter 3 we look more generally at how decisive choice methods can be made given different informational assumptions. This combines the two major themes of expressiveness and decisiveness. The framework of this chapter is ordinal: we start with a full ordinal profile, and consider the effect of using *less* information than this. We develop a definition of **maximal decisiveness** with respect to the available information. This is intended to identify whether or not a choice method is as selective as possible, given some natural restrictions. (An example restriction would be to require that the choice to be based on tournaments, as was the case in Chapter 2.) In particular, our definition of maximal decisiveness is designed around **neutrality**, the idea that all alternatives should be treated equally.

Chapter 4 considers information in two ways that are classic to social choice theory: **independence of irrelevant alternatives** and **domain restrictions**. Independence of irrelevant alternatives is a *local* informational restriction: instead of a global look at what information

is used by a rule or a family of rules, it concerns how much information from an *individual* profile is required to calculate *part* of the output of a social welfare function. With respect to social welfare functions, the strictest restriction of this type is **binary independence**, which requires that we can determine which of two alternatives is preferred in the output only by considering the pairwise comparisons of these two alternatives in the profile. This can then be weakened to ternary independence, quaternary independence, etc. It is known that under the full domain these are not real weakenings as they all collapse into binary independence (except for independence over the whole set of alternatives which is trivially satisfied). In Chapter 4 we investigate whether this still happens under restricted domains. A restricted domain is an express limit on the expressiveness: some profiles are supposed to not exist. We show that for different restricted domains the different levels of independence described above may or may not be equivalent, and specify when and to what extent different versions of independence collapse into the same condition.

Chapter 5 returns more speculatively to the question of what information should be collected from society. We contrast the traditional ordinal approach—that we have been using throughout the preceding chapters—with what we call the **evaluative** approach. The evaluative approach involves aggregating evaluations of the alternatives made by the agents. For example, approval voting, where society is supposed only to approve or disapprove each alternative, is evaluative. We prove an impossibility in combining the ranking and evaluative approaches. Our results suggest that there is a deep incompatibility between the two approaches.

Common definitions

The following definitions are applied consistently throughout the thesis, except for some slight variations in the first chapter.

Throughout we consider a finite set of objects A . With the exception of Chapter 1, this is called the set of **alternatives**. It is from this set that “society” makes a choice or creates a ranking. Typically $|A| = m$ for some positive integer m . We refer to individual alternatives using $a_1, a_2, \dots, a_m, a, b, c, \dots, x, y, z, \dots$

The elements of A are variously ranked, approved, evaluated or chosen. We write $\mathcal{W}(A)$ for the set of complete and transitive binary relations, or complete preorders, over A . We use R to refer to elements of $\mathcal{W}(A)$. We typically use infix notation for binary relations, for example writing xRy instead of $(x, y) \in R$. A complete preorder R is a linear order if and only if R is also antisymmetric. We use $\mathcal{L}(A)$ for the set of linear orders over A . We use R^* for the asymmetric part of R , i.e. for all $x, y \in A$, xR^*y if and only if xRy and not yRx . The

asymmetric part of a complete preorder is also a strict weak order, and the asymmetric part of a linear order is a strict linear order; we use P to denote such irreflexive relations when they do not correspond to some already defined R . We write $\mathcal{W}^*(A)$ and $\mathcal{L}^*(A)$ for the set of strict weak orders and the set of strict linear orders respectively. We sometimes display a (strict) linear order by simply writing a list starting with most preferred elements and ending with least preferred. For example, $adbc$ corresponds to the linear order R such that aRd , dRb and bRc (or to the corresponding strict linear order P depending on context).

We also consider an index set N which, with the exception of Chapter 1, we refer to as the set of **agents**. Typically $|N| = n$. The set N may be thought of as the “members” of “society”. Different members of society may have different opinions concerning the different alternatives, which we typically model by assigning each agent a weak order preference over the alternatives. We call a full vector of preferences, denoted $\mathbf{R} \in \mathcal{W}(A)^N$, a profile. In general, we use bold typeface for vectors. For $i \in N$ and $\mathbf{R} \in \mathcal{W}(A)^N$, we use R_i to refer to the i th coordinate of \mathbf{R} . By (\mathbf{R}_{-i}, R') we mean the profile \mathbf{R} with the coordinate R_i replaced by R' .

A **social choice function** takes preferences over alternatives and produces a non-empty set of choices; $f: \mathcal{W}(A)^N \rightarrow 2^A \setminus \emptyset$. Sometimes the domain will be restricted or even modified entirely. A **social welfare function** takes preferences over alternatives and produces a ranking of alternatives; $F: \mathcal{W}(A)^N \rightarrow \mathcal{W}(A)$.

A **permutation** on some set X is an onto function $g: X \rightarrow X$. We use typically consider permutations of alternatives or agents with the notation $\sigma: A \rightarrow A$ and $\rho: N \rightarrow N$. We abuse notation and apply permutations to preferences and profiles as follows. For $R \in \mathcal{W}(A)$ we write $\sigma(R) = \{(\sigma(x), \sigma(y)) : (x, y) \in R\}$. For $(R_1, \dots, R_n) \in \mathcal{W}(A)^N$ we write $\sigma((R_1, \dots, R_n)) = (\sigma(R_1), \dots, \sigma(R_n))$ and $\rho((R_1, \dots, R_n)) = (R_{\rho(1)}, \dots, R_{\rho(n)})$.

Chapter 1

Discipline aggregation, manipulation and decisiveness

This chapter serves both as an introduction to the themes of this thesis, and as a recap of the basic ideas of social choice theory under a fresh interpretation. It is based upon joint work with Sebastian Schneckeburger that was presented at COMSOC 2018 and AAMAS 2019. In it we apply the traditional social choice framework to the problem of aggregating the competition results of multiple disciplines in a sports competition. This provides a recap of many of the theoretical results of social choice theory; in fact we will retread many of the historical advancements made within the field. Ultimately, our work raises the issue of decisiveness, one of the main themes of this thesis.

The 2020 Olympic Games in Tokyo will inaugurate ten new gold medals; one male and one female in each of five new events: karate, skateboarding, surfing, baseball and sports climbing. As it is the first time that these events will appear in the Olympics, for each event it is necessary to specify for the first time how the medal winners will be decided. This presents a particular problem for sports climbing, because sports climbing has not existed in a *unique* competition format before its introduction as an Olympic event. Instead, there are *three* prevalent types of competitive climbing: bouldering, lead-climbing and speed-climbing. It is these distinct types of sports climbing that we refer to as disciplines. Each discipline requires different skills and measures the performance of athletes using different methods. Thus sports climbing is to be a composite event, similar to the pentathlon—but the manner in which medal winners are determined for the pentathlon is by no means suitable here.

For sports climbing, it is desired that *only the rankings* of the athletes in each of the disciplines are used to determine the medal winners. In contrast, for the pentathlon quantities such as time, height and distance, measured independently for each athlete, form the basis for the final ranking. Given the “only rankings” constraint, the problem of designing a procedure for choosing the medal winners is similar to that of designing a procedure of social choice theory. Of course, the terminology is different: what is traditionally the set of voters has become the set of **disciplines**; the set of candidates becomes the set of **athletes**. But the underlying problem is the same: how should we aggregate multiple rankings into a single, definitive ranking?

For the sports climbing event at the 2020 Olympics, the International Federation of Sports Climbing (2017) has decided to rank the athletes according to the product of their ranks in the disciplines. They have various reasons for this choice, which we will discuss later. Also, the method will be tested in some trial events before the Olympics in order to confirm its suitability. On the other hand, there are theoretical situations where the method gives undesirable results: it can produce **ties** among a large number of athletes, which cannot always be broken in an obvious manner; and it is potentially open to a type of **manipulation**.

In social choice theory, undesirable possibilities are often called **paradoxes**. The most famous example is Condorcet's (1785) paradox, which shows that it is possible that a majority prefers a first alternative to a second, that a majority prefers the second to a third, and also that a majority prefers the third to the first.¹ We will see another such **majority cycle** later in this chapter (Figure 1.1). Of course, just stating individual examples of undesirable behaviour would be a somewhat haphazard method of analysis, a more general and principled method of analysis is preferable.

In order to assess *all* relevant possible procedures for determining the medal winners, one can try to encompass them in a single framework. In order to do so, we treat them as *mathematical functions* from lists of rankings to a final output ranking. We call such functions **discipline aggregators**.² The method of formalising social choice problems in this manner was pioneered by Arrow (1950); it allows us to assess the possible procedures for selecting medal winners based upon their formal properties.

There are a number of properties that a discipline aggregator may satisfy—following Arrow, in the academic literature on social choice theory such desirable properties are typically called **axioms**. We want our function to satisfy such desirable properties. For example, if a first athlete performs better than a second in *all* the disciplines, surely the first athlete should be ranked better than the second. On the other hand, there may be undesirable possibilities for certain functions, for example, a function may be manipulable in some way or another. The basic idea behind our definitions of manipulation is that, in certain situations, an athlete can aid a teammate by performing *worse* than they normally would have done.³ If such an situation is theoretically possible, an athlete may indeed perform such a manipulation in

¹Other examples of so-called paradoxes include the Ostrogorski paradox (Rae and Daudt, 1976), Anscombe's paradox (Anscombe, 1976), and the multiple election paradox (Brams et al., 1998).

²Although "discipline aggregator" may be thought of as a synonym for "social welfare function" mathematically speaking, our nomenclature in this chapter serves as a reminder that there is a difference in interpretation here.

³This idea of manipulation is *not* an alias of some previously studied concept from social choice theory. In particular, it is not simply Gibbard-Satterthwaite manipulation. Gibbard-Satterthwaite manipulation is "performed" by a discipline, and involves *any possible change* to the *single* ranking of this discipline. In contrast, the manipulations we consider in this chapter are "performed" by an athlete, and involve a *restricted change* to that athlete's ranking *throughout the whole profile*.

reality, which we take to be an undesirable event. We want our function to preclude such possibilities.

A formal description of the problem also makes it easier to determine *how often* any given paradoxical situation will occur. This question has a potentially empirical character; however there is sometimes a lack of real life data: this can be the case for real-life preferences of agents, but is especially so for our particular discipline aggregation problem, as events using the combined format for sports climbing have only started to be run last year (2018). We can, however, perform **simulations** given some assumptions on how likely each possible state is (Garman and Kamien, 1968; Gehrlein and Fishburn, 1976).

We will give a few different definitions describing situations where a discipline aggregation function can be manipulated. Under our most permissive definitions of manipulation, only unreasonable functions can completely prevent such manipulation from occurring. That is, we prove an **impossibility result**: any function that prevents manipulation must violate some other desirable condition. Such a negative result, if accepted, seems to imply that some sort of trade-off must be made.

Impossibility results often arise in social choice theory,⁴ and there are many different techniques that have been developed in attempt to mitigate their negative conclusions. Perhaps the most direct attack against such a result is to argue that the identified conditions are not really necessary for a rule to be reasonable. For the most part, the properties we define are obviously desirable. The point of weakness is in our definitions of manipulation, which indeed takes the place of the much discussed—and challenged—property of **independence of irrelevant alternatives**.⁵ Various weakened versions of this property have been proposed in the literature.

Indeed, with a slightly less permissive definition of when manipulation is possible, there are functions that satisfy our desirable properties. However, the function that exemplifies the possibility is, intuitively, not **decisive** enough. Although it satisfies our given properties, it is nonetheless an unreasonable rule. This can only occur because our given properties are quite weak—they are *necessary* for a rule to be reasonable, *but not sufficient*; precisely, they are missing a condition that ensures a minimal level of decisiveness.

This raises the issue of determining a condition that ensures a minimal level of decisiveness. Although the function exemplifying the possibility is *intuitively* indecisive, there is no *formal* specification of this property. It is not obvious how to define a *reasonable* condition of decisiveness that this function does not satisfy. In fact, there is an *unreasonable*, very strong decisiveness condition that we can impose: we can require that there are no ties at all in the

⁴Impossibility results are the subject of a Chapter 1 in the Handbook of Social Choice and Welfare (Campbell and Kelly, 2002).

⁵In this thesis, Chapter 4 contributes yet more discussion of this property.

output; that a linear order is returned—this is known as **resoluteness**. With this property added, the impossibility reoccurs—the question is whether or not there is a sensible, weaker decisiveness condition will have the same effect. So, more generally, this chapter motivates further investigation of the concept of decisiveness; a project that will be continued in later chapters. Decisiveness is especially important for sports competitions, but remains important for social choice more generally: after all, the whole point of social choice theory is to make a choice!

The rest of this chapter proceeds as follows. In Section 1.1 we describe the proposed combined format for sports climbing, and give an example that shows that this method may be manipulated and may produce ties. We then provide a formal description of the basic setting in Section 1.2, which contains most of the definitions necessary for our results. However, our description of the various types of manipulation we consider merits its own section: Section 1.3. Having completed our definitions, we then start our analysis in earnest. Our simulation results, in Section 1.4, show that the probability of ties under the proposed method is quite low, but that on the other hand there is a high potential for manipulation. Our theoretical results of Section 1.5 show that it is impossible to completely rule out manipulation, although we give a possibility for a plausible weakened version of our requirements concerning manipulation. This possibility, however, raises the issue of the decisiveness of discipline-aggregators. In Section 1.6 we prove that it is computationally hard to determine whether or not one can manipulate the proposed method if we allow the numbers of athletes and disciplines to increase. Section 1.7 concludes this chapter.

1.1 The combined format for sports climbing

For the 2020 Olympics the International Federation of Sports Climbing (2017) has devised the following combined format for sports climbing. Twenty athletes—of which each country can have up to two representatives—will be involved in the main event, which itself is split into two rounds. A qualification round first reduces the twenty athletes to six. These six finalists then compete in a separate final round to determine the medals. Although in the qualification round it is only necessary to *select* six athletes, in fact the procedure is identical to that of the final round and produces a full ordering over the contending athletes, the six best ranked are then selected.

So in each round a ranking over all the athletes is produced in the same manner: the contending athletes compete in all three disciplines, thereby producing three linear orders over the athletes. For each discipline any given specific athlete has a **rank**, which is the number of other athletes that defeat her plus one; a given athlete's overall **score** is the product of the athlete's ranks in each discipline. These product scores provide a final ordering over the

competitors, with lower scores being better. If two athletes receive the same score the **tie** can be broken in favour of the athlete that performs better than the other in more disciplines. That is, ties can be broken depending upon the so-called majority relation—we will see an example demonstrating how this works, but, in short: if more than two athletes receive the same score, ties are broken first by the transitive closure of the majority relation; where there are cycles, an external linear-order tiebreaker is used.

Unfortunately, the combined Olympic format may lead to tied situations if more than two athletes have the same product score and the pairwise comparisons form a cycle. Ties are problem for both rounds of the competition: in the final round, it is desirable to have a single gold winner; in the qualifying round, ties necessitate an extra method to determine which athletes progress. An example demonstrates the problem—Table 1.1 shows potential ranks of seven of the athletes and bounds on the ranks of the other thirteen athletes after the qualification round.

Table 1.1: Possible ranks of seven of the athletes after the first round of competition.

Athlete	Ranking in discipline			
	Speed	Bouldering	Lead	Product
Alice	3	3	6	54
Barbara	1	5	12	60
Chloe	4	15	1	60
Dorothy	15	1	4	60
Emma	2	10	3	60
Faye	5	6	2	60
Grace	6	2	5	60
others	≥ 7	≥ 4	≥ 7	≥ 196

Only six of the athletes can proceed to the final round. This means that one of the six athletes with tied product scores—that is, one of the named athletes excepting Alice—must be eliminated from the competition (along with all the other unnamed athletes). However, it is not clear which of the six should be eliminated; when we compare each pair by which athlete defeats the other in more disciplines, we produce a cycle. This is displayed graphically in Figure 1.1: the athletes are represented by the first letter of their name, and an arrow points from one athlete to another if the first performs better than the second in more disciplines.

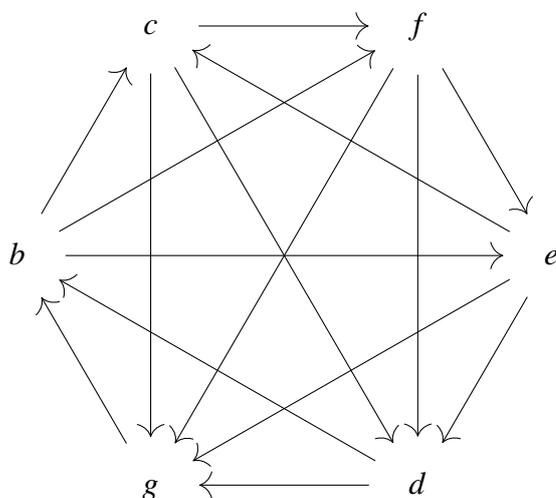


Figure 1.1: The majority relation between b , c , d , e , f and g —the athletes are labelled by the first letter of their name—given the results from Table 1.1.

The problem is that there is no **Condorcet loser**; that is, no athlete that is—in terms of majorities—defeated by all the other athletes. There is a parallel concept of a **Condorcet winner**: an athlete who defeats all other athletes according to the majority relation. The proposed method considers situations without Condorcet winners or losers as tied; thus the situation described in Table 1.1 requires the use of an **external tiebreaker** to decide which athlete is eliminated.⁶ Suppose that the external tiebreaker eliminates Barbara. If the remaining athletes perform the same as in the first round, the ranks and products will be as displayed in Table 1.2.

Table 1.2: Predicted results if the named athletes except Barbara progress to the second round and perform the same as in the first round as given in Table 1.1.

Athlete	Ranking in discipline			Product
	Speed	Bouldering	Lead	
Alice	2	3	6	36
Chloe	3	6	1	18
Dorothy	6	1	4	24
Emma	1	5	3	15
Faye	4	4	2	32
Grace	5	2	5	50

⁶The IFSC will use a “seeding list” to break ties that are not resolvable by pairwise comparisons. Such a seeding list is, in effect, an external linear order tiebreaker. For the final round the ranking of the qualification round will be used as a seeding list; for the qualification round a seeding list based on the prior qualification system will be used (Meyer, 2018).

There is another potential problem here: **manipulation**. Suppose now that Alice and Chloe have the same nationality. The predicted results suggest that Alice will not win a medal, and certainly not the gold, while her teammate Chloe is predicted to win the silver. However, if Alice deliberately performs worse than Chloe in the speed competition and all other ranks remain the same, Chloe will become the unique gold medal winner with a product score of 12.⁷ Thus, national loyalty may lead Alice to manipulate. In doing so she **spoils** the result for Emma, whose efforts would normally have been enough for a gold medal.

The above example demonstrates two *potential* problems with the proposed method: ties and manipulation. We have not yet assessed how often these problems are likely to occur, that is, whether they are a likely problem in practice. Nor do we know whether or not these problems are in some sense *necessary*: it may not be possible to design other discipline aggregation functions which avoid them. To address these issues we require a more formal description of the setting.

1.2 Formal description of the setting

Denote by $A = \{a, b, c, \dots\}$, $|A| = m$ the set of athletes and by $N = \{1, \dots, n\}$ the set of disciplines. We suppose that $m \geq 3$ and $n \geq 2$; note that this covers the specific cases of $m = 6$, $m = 20$ and $n = 3$. Denote the set of strict linear orders of A by $\mathcal{L}^*(A)$ and use P to denote a strict linear order over athletes. Denote the set of complete preorders of A by $\mathcal{W}(A)$. and use R to denote a complete preorder over the athletes. We use R^* to denote the asymmetric part of R —note that $P = R^*$ for some R .

All the athletes compete in each discipline $i \in N$, resulting in n strict linear orders P_i over A . A **profile** that summarises the results for each discipline is denoted by $(P_1, \dots, P_n) = \mathbf{P} \in \mathcal{L}^*(A)^N$. A **discipline aggregator** F uses these results to produce a complete preorder over the competitors: $F: \mathcal{L}^*(A)^N \rightarrow \mathcal{W}(A)$.

Athletes have a rank for each discipline and for the output complete preorder. Formally, for an ordering R over competitors, the rank of $a \in A$ is $r(R)(a) = r_R(a) = |\{x \in A : xPa, x \neq a\}| + 1$. Note that $r(R) = r(R^*)$. To simplify notation, for a discipline $i \in N$ we write $r_i = r(R_i) = r(P_i)$ when the context makes it clear which profile was intended. Because lower ranks are better, P_i and the natural ordering on ranks are inverted: for all $x, y \in A$ and $i \in N$, $xP_i y$ if and only if $r_i(x) < r_i(y)$. Note that the output of a discipline aggregator

⁷The manipulation described here would be feasible in practice. The method used to produce the speed-climbing ranking is a type of knockout tournament. Under the assumption that a defeats g and that c defeats f in the quarter finals, a will face c in the semifinals; if a deliberately loses then c will be guaranteed to be ranked at second or better at speed-climbing (International Federation of Sports Climbing, 2017).

can contain ties, though if two or more competitors are ranked first, no competitor is ranked second—a shared gold medal implies that no-one receives silver. We refer to athletes ranked first in the output as **winners**.

We can use a profile to make pairwise comparisons between athletes. For an arbitrary profile \mathbf{P} the **majority relation** $T_{\mathbf{P}} \subseteq A \times A$ is defined by, for all $x, y \in A$, $x \neq y$

$$xT_{\mathbf{P}}y \quad \text{if and only if} \quad |\{i \in N : xP_iy\}| \geq |\{i \in N : yP_ix\}|.$$

The majority relation is connected regardless of the parity of $|A|$. It may not be transitive: for a binary relation Q we write Q^+ for the transitive closure of Q , the smallest transitive relation that contains Q .

We now formally define the proposed discipline aggregator, insofar as it is determined by the profile of results. We call this function **inverse-Borda-Nash**,⁸ and denote it by $\text{ibn} : \mathcal{L}^*(A)^N \rightarrow \mathcal{W}(A)$. Define the binary relation $Q \subseteq A \times A$ by

$$xQy \quad \text{iff} \quad \begin{cases} \prod_{i \in N} r_i(x) > \prod_{i \in N} r_i(y) \\ \text{or} \\ \prod_{i \in N} r_i(x) = \prod_{i \in N} r_i(y) \text{ and } xT_{\mathbf{P}}y. \end{cases}$$

Define $\text{ibn}(\mathbf{P}) = Q^+ \cup \{(x, x) : x \in A\}$; because Q is connected this is a complete preorder.

Basic desirable properties

There are some basic desirable properties that a discipline aggregator F should satisfy. An athlete $x \in A$ clearly beats $y \in A$ in \mathbf{P} if for all $i \in N$, xP_iy . We say F satisfies the **clear winner condition** if whenever x clearly beats y in \mathbf{P} , then for $R = F(\mathbf{P})$ it is the case that

⁸This nomenclature is intended to be descriptive as it invokes Borda scores and the Nash product. The Nash product is sometimes described as providing a middle ground between the utility maximisation of additive methods and the maxi-min of egalitarian methods. However, this is not the case for the proposed method because the Borda scores are *inverted* i.e. smaller numbers are better. If the scores were added, this inversion would have no effect, but this is not the case for multiplication. For example, according to inverse-Borda-Nash, an athlete with rankings (1,1,4) beats an athlete with (2,2,2); whereas for traditional Borda scores the opposite is true: (19,19,19) would be considered better than (20,20,17). It is seen as an advantage of the method that it favours specialists—it is preferred that the winner of the combined format is a potential winner of world-cups in some individual discipline, rather than a generalist (Meyer, 2018).

We are not aware of any specific precedent for inverse Borda-Nash. This may be because it would become an “anti-fairness” approach when applied to social choice or social welfare. However, it can be subsumed under well-known concepts; it is equivalent to a **scoring rule** with weights $(\log(m), \log(m/2), \dots, \log(m/m-1), 0)$. The family of scoring rules are well studied within the literature, and the results of Proposition 1.1 below could be shown as corollaries to general results concerning this family.

xR^*y . The next two conditions impose symmetry restrictions, roughly speaking, they require that all athletes should be treated the same and that all disciplines should be treated the same. We say F satisfies **athlete-neutrality** if permuting the competitors in the profile similarly permutes the competitors in the output ranking: for any permutation $\sigma : A \rightarrow A$, given \mathbf{P} and \mathbf{P}' such that for all $a, b \in A, i \in N$ $aP_i b \Leftrightarrow \sigma(a)P'_i\sigma(b)$, then $aF(\mathbf{P})b \Leftrightarrow \sigma(a)F(\mathbf{P}')\sigma(b)$. We say F satisfies **discipline-neutrality** if permuting the disciplines in the profile has no effect on the output ranking: for any permutation $\rho : N \rightarrow A$, given \mathbf{P} and \mathbf{P}' such that for all $i \in N$, $P_i = P'_{\rho(i)}$, then $F(\mathbf{P}) = F(\mathbf{P}')$. Our last two properties limit how much a single discipline can determine the winner. For a discipline aggregator F , we say the gold is determined by $i \in N$ if, for any profile \mathbf{P} , $\max(F(\mathbf{P})) = \max(P_i)$. We say F is **not determined** if the gold is not determined by any $i \in N$. Similarly, for a discipline aggregator F , we say the gold is weakly determined by $i \in N$ if, for any profile \mathbf{P} and writing $R = F(\mathbf{P})$, $r_i(x) = 1$ implies $r_R(x) = 1$. We say F is **not weakly determined** if the gold is not weakly determined by any $i \in N$.

Each of these properties can be linked to axioms from social choice theory. The clear winner condition is called, for example, the ‘‘Pareto criterion’’ (Campbell and Kelly, 2002, p. 42).⁹ Athlete-neutrality becomes simply ‘‘neutrality’’, whereas discipline-neutrality is known as ‘‘anonymity’’. A discipline that determines the gold may be thought of as a ‘‘top-dictator’’: an agent whose top ranked alternative must be the top ranked alternative in the output. Note that it is easier for a discipline aggregator to be not determined than not weakly determined; in particular, if a discipline aggregator satisfies discipline-neutrality then it is not determined, but may still be weakly determined.

As a final exposition of our definitions, we show that Inverse-Borda-Nash satisfies all these basic desirable properties.

Proposition 1.1. *Inverse-Borda-Nash satisfies the clear winner condition, satisfies athlete-neutrality, satisfies discipline-neutrality, and is not weakly determined.*

Proof. Clear winner: if an athlete is ranked better than another in all disciplines, it must have a smaller product of ranks, therefore will be ranked better in the output.

Athlete-neutrality: if we permute athletes, we also permute their product scores and the relation $T_{\mathbf{P}}$.

Discipline-neutrality: permuting disciplines has no effect on product scores or on the relation $T_{\mathbf{P}}$.

Not weakly determined: for an arbitrary discipline $i \in N$, take a profile where some athlete $a \in A$ comes first in all other disciplines and second in this discipline: $r_i(a) = 2$ and $r_j(a) = 1$

⁹We will adopt this terminology in later chapters in this thesis.

for all $j \in N$, $j \neq i$. If $n > 2$, then a is the unique winner, thus the discipline does not determine the gold. For $n = 2$, the profile displayed in Table 1.3 shows that i does not weakly determine the gold (we use the assumption that $m \geq 3$).

Table 1.3: A profile that shows that inverse-Borda-Nash is not weakly determined for $n = 2$. Athlete a is the overall winner, though athlete b wins for (arbitrary) discipline i .

Athlete	Ranking in discipline		Product
	i	$j \neq i$	
a	2	1	2
b	1	3	3
c	3	2	6
others	≥ 4	≥ 4	≥ 16

□

1.3 Manipulation in the discipline aggregation framework

In this section we develop formal definitions of manipulation, one of which encapsulates the example given in Section 1.1. These definitions are general enough that they apply to any discipline aggregator, including the specific cases of $m = 6$, $m = 20$ and $n = 3$ that will be in place for sports climbing at the 2020 Olympics. Of course, we do not cover *all* possible types of manipulation here.

We only consider manipulations that apply to the discipline aggregator. We are not concerned with doping, score fixing, betting, or any other of the many documented examples of foul play in sports competitions, problems that any designer of a discipline aggregator would find difficult to directly combat. A characteristic of the manipulations that we *do* consider is that they involve an athlete *deliberately performing badly*. However, we are not concerned with deliberate bad performances that cause restarts (Jimenez, 2012), or that otherwise abuse the rules *internal* to the competition. An example of manipulation that a scrupulous designer of discipline aggregators *might* be able to prevent is the following: a competitor, in the group stages of a knockout tournament, deliberately loses a match in order to face a weaker opponent in the following round. Such manipulation seems to have occurred at the 2012 Olympics (TheGuardian, 2012).¹⁰

For our purposes, a manipulating athlete changes the **sincere** profile, P , to some other manipulated profile, P' , where the only changes possible are that the manipulator is ranked worse in one or more of the disciplines. The manipulator only has incentive to make the changes if

¹⁰Another example that fits our definitions of manipulation is the phenomenon of “team orders” in Formula One racing (Gpupdate, 2008).

the manipulated profile P' is more desirable to her than the sincere profile. So, for manipulation to be a problem, two conditions need to be fulfilled: the athlete needs to be able to force the manipulated profile P' to occur, and they need to prefer it to the real profile P . These two conditions involve a few assumptions.

We assume that the athlete knows what the results would be if everyone were to perform to their best abilities—that is, the athlete knows what the sincere profile is. This may be justified by the claim that the athlete knows roughly how good the other athletes are, especially in the specific case of the final round of the sports climbing event where the first round can be taken as a proxy. Relatedly, we assume that no other athletes will manipulate, so that the sincere profile would indeed be the result if the manipulator did not manipulate. This is a simplifying assumption: we have to start somewhere, and we don't want to start by considering levels of rationality.¹¹ We also assume that the athlete can precisely control *how much worse* they perform. In general, this is perhaps unrealistic, though there are specific cases where this is extremely reasonable: the speed competition runs head-to-head, so for certain configurations of profiles it would be easy to perform a specific manipulation that results in a specific ranking. Altogether, our assumptions are no stronger than those of what is probably the most important theorem concerning manipulation in social choice theory: the Gibbard-Satterthwaite Theorem.

We write the condition restricting which profiles an athlete can manipulate to—given a starting sincere profile—as follows. From a sincere profile P , a profile P' is a possible manipulation by athlete a if¹²

(1.i) for all $i \in N$, $x \in A \setminus \{a\}$ and $y \in A$, $xP_i y$ implies $xP'_i y$.

This **discipline manipulation condition** underlies all of our following definitions of manipulation. The distinct definitions of manipulation only differ in what is required for the profile P' to be preferred by a to P . The clearest reason that an athlete may prefer one profile to another is if it results in a better output ranking for the athlete herself. Supposing that a manipulates from P to P' , and writing $R = F(P)$ and $R' = F(P')$, this condition requires that

$$r_{R'}(a) < r_R(a).$$

We refer to this as an **egoistic** manipulation.¹³ Egoistic manipulation is impossible under inverse Borda-Nash: if an athlete worsens their ranking in one or more disciplines, she receives a strictly larger product score, whereas all other athletes receive at most the same product

¹¹Our analysis can be considered **decision theoretic** as opposed to **game theoretic**.

¹²Note that we do not consider the possibility that two or more agents manipulate together. In such a case, a manipulator may be able to rank higher than in the sincere profile.

¹³Egoistic manipulation can be compared to traditional definitions of **monotonicity**. However, note that the focus is on *rank* as opposed to pairwise comparisons between athletes.

score as before, thus any athletes that she was weakly defeated by according to the sincere profile will still weakly defeat her according to the non-sincere profile. Other rules, however, do permit egoistic manipulation.¹⁴

In contrast to egoistic manipulation, where an athlete improves their own output ranking, an athlete may manipulate in order to improve the output ranking of *some other athlete*. One may suppose that this other athlete is a teammate, friend, or comes from the same country. Given profiles and P and P' as in the discipline manipulation condition, an athlete a manipulates for athlete b if

$$(1.ii) \quad r_{R'}(b) < r_R(b).$$

We call this **altruistic manipulation**.¹⁵ Technically, this definition *includes* egoistic manipulation as the special case when $a = b$. Of course, if $a \neq b$, an altruistic manipulation may not result in a more desirable outcome for the manipulating athlete. We now give two different, but not necessarily incompatible, reasons why an altruistic manipulation may be desirable for the manipulating athlete, the first of which comes in a strong and weak version.

The first reason seems at first sight obvious: if the altruistic manipulation involves no sacrifice on the part of the manipulator a . This seemingly amounts to the condition

$$(1.iii) \quad r_{R'}(a) \leq r_R(a).$$

We will call such altruistic manipulations **without sacrifice**. However, there is a slight subtlety involved here; although we have said that the output ranks are the most important thing for discipline aggregation, pairwise comparisons may still have a secondary importance: it seems preferable to win the gold uniquely than to share it with other athletes. Even if an athlete can manipulate without sacrifice, they may nonetheless not want to do so because they end up sharing their rank with more athletes than before. The following condition rules out this possibility,

$$(1.iv) \quad |\{x \in A : xR'a\}| \geq |\{x \in A : xRa\}|,$$

if this is satisfied as well as (1.iii) above we say that a manipulation is **completely without sacrifice**. Because it occurs in fewer profiles, manipulation completely without sacrifice is easier to prevent than manipulation (just) without sacrifice.

The example of manipulation given in Section 1.1 is not a manipulation without sacrifice. The idea instead is that the manipulating athlete recognises that she probably won't get a

¹⁴Cf. non-monotonic binary trees, Chapter 2.

¹⁵One might also consider the possibility of **spiteful manipulation**, whereby an athlete manipulates in order to *worsen* the rank of a specific target athlete.

higher ranking than her teammate, but can nonetheless manipulate to aid her teammate. The condition that ensures this is the following

$$(1.v) \ r_{R'}(b) < r_R(a).$$

We give manipulation of this type the name **spoiler**, this refers to the fact that a poorly ranked athlete spoils the fair result concerning other, better ranked, athletes. The condition ensures that b performs better than a would have done.¹⁶

We now formally state the full definitions of the two main types of manipulation that we are interested in, the first of which splits into its strong and a weak versions.

Definition 1.2 (Non-sacrificial manipulation). Let $F(\mathbf{P}) = R$ and $F(\mathbf{P}') = R'$, and $a, b \in A$. Competitor a can **manipulate without sacrifice**, for competitor b , from the profile \mathbf{P} to the profile \mathbf{P}' if

$$(1.i) \text{ for all } i \in N, x \in A \setminus \{a\} \text{ and } y \in A, xP_i y \text{ implies } xP'_i y$$

$$(1.ii) \ r_{R'}(b) < r_R(b)$$

$$(1.iii) \ r_{R'}(a) \leq r_R(a).$$

Such a manipulation is **completely** without sacrifice if it also satisfies

$$(1.iv) \ |\{x \in A : xR'a\}| \geq |\{x \in A : xRa\}|.$$

Definition 1.3 (Spoiler manipulation). Let $F(\mathbf{P}) = R$ and $F(\mathbf{P}') = R'$, and $a, b \in A$. Athlete a can **spoil**, for athlete b , from the profile \mathbf{P} to the profile \mathbf{P}' if

$$(1.i) \text{ for all } i \in N, x \in A \setminus \{a\} \text{ and } y \in A, xP_i y \text{ implies } xP'_i y$$

$$(1.ii) \ r_{R'}(b) < r_R(b)$$

$$(1.v) \ r_{R'}(b) < r_R(a).$$

When both the athlete b and the profiles are implicit and need not be specified, we will say simply that athlete a **spoils**.

Two final remarks on our definitions of manipulation: first, for some pairs of profiles these definitions overlap; second, they do not cover all conceivable manipulations of the discipline aggregator. In particular, there is the example of altruistic manipulation in the *first round* of the sports climbing event: here, the desirable outcome is that the manipulator and their target are both ranked better than seventh (and not tied if ranked sixth).

¹⁶Manipulation without sacrifice and spoiler manipulation amount to the egoistic case if we set $a = b$, but manipulation completely without sacrifice does not; in fact there are cases of egoistic manipulation that are not cases of manipulation completely without sacrifice.

1.4 The likelihood of undesirable profiles: simulation results

At the time of typing, the combined format for sports climbing has not been used in many real-life events. We can, however, generate many possible outcomes for events and see what one might expect. Such simulation methods have a long history in social choice theory, where large scale data concerning individuals' preferences has not always been readily available. We consider the results of such simulations in this section. In short, our simulations suggest that for inverse-Borda-Nash, although ties are unlikely to be a problem, potential for manipulation occurs with a high probability.

We generated profiles with six athletes and three disciplines, the same numbers as in the final round of the Olympics sports climbing competition. The generated profiles form three groups: in the first group, for each discipline every possible strict linear order is equally likely—profiles are drawn from the **impartial culture**. For profiles in the second group there is a positive correlation in an athlete's results across the three disciplines. In the final group there is positive correlation between two disciplines and negative correlation with the third. This third culture conforms best to our actual expectations for the competition because the two disciplines of bouldering and lead climbing have an intersection of athletes at the top level, whereas top level speed climbers do not typically compete in the other disciplines.

For a profile based on the impartial culture we independently select three strict linear orders, each uniformly at random from the set of all possible linear orders. Our positively correlated culture uses the Plackett-Luce (1975; 1959) model¹⁷ with initial odds

$$2^1 : 2^2 : 2^3 : 2^4 : 2^5 : 2^6.$$

Label the athletes as a_j for $j \in \{1, 2, 3, 4, 5, 6\}$. Writing $t = \sum_{i=1}^6 i^2$, each athlete a_j has a $p_j = j^2/t$ probability of being ranked first. The idea is that the initial odds also represent the strengths of the athletes; in particular, we suppose that each athlete is two times stronger than her closest competitor. Given that the athletes in a set B defeat the athletes in $A \setminus B$, the probabilities that determine the winners within $A \setminus B$ should only depend upon the relative strengths of the athletes within $A \setminus B$. So if athlete a_k , $k \neq j$ ranks first, then a_j has a

$$\frac{j^2}{t - k^2} = \frac{p_j}{1 - p_k}$$

probability of being ranked second. If a_k ranks first and a_l ranks second, $l \neq j$, $l \neq k$, then a_j has a

$$\frac{j^2}{t - k^2 - l^2} = \frac{p_j}{1 - p_k - p_l}$$

¹⁷This model's ability to forecast Formula One results has been studied by Henderson et al. (2018).

Table 1.4: The number of randomly generated profiles that involved ties, were subject to spoiler manipulation, were subject to manipulation without sacrifice, were subject to manipulation completely without sacrifice, and that were subject to any of these three types of manipulation. 100,000 six agent profiles were generated for each culture. A profile was considered subject to manipulation if any of three disjoint random pairs of agents could manipulate.

Culture	Ties	Spoiler manipulation	Without sacrifice	Strict without sacrifice	Any manipulation
Impartial	632	37,730	47,807	47,326	59,660
Positive correlation	779	13,792	43,723	41,597	46,964
Negative correlation	526	44,826	48,741	48,350	63,151

probability of being ranked third. The positive correlation arises because we suppose that the athletes have the same strengths for each discipline; a profile consists of three independently generated strict linear orders using the same initial odds. A negatively correlated profile is created by taking a positively correlated profile and reversing the strict linear order of the last discipline.¹⁸

We randomly generated 100,000 profiles of each type. A profile counts as tied if at least one tie occurs at any ranking level—we do not count the number of distinct ties nor how many athletes are involved in each tie. To count manipulations, we first randomly pair the athletes into three disjoint pairs. A profile counts as manipulable if at least one of the pairs can manipulate. We perform the count separately for spoiler manipulation, manipulation without sacrifice, manipulation completely without sacrifice, and for any of these types of manipulation. The results are presented in Table 1.4.

According to our simulations, it is very unlikely that there will be a tie at any level of the output complete preorder in the final round of the competition. We also ran simulations for twenty athlete profiles obtaining similar results.¹⁹ This strongly suggests that a tie in the actual competition is very unlikely to occur: note that each of our models exhibits a

¹⁸The associated culture best represents our, somewhat naive, expectations for the competition—we expect lead climbing and bouldering to be positively correlated with each other and negatively correlated with speed climbing.

¹⁹For profiles with twenty athletes, of the 100,000 profiles we generated for each culture, 208 profiles had ties for the impartial culture, 1108 profiles had ties for the positively correlated culture, and only 92 profiles had ties for the negatively correlated culture that conforms best to our expectations for the actual competition.

high degree of symmetry; and such symmetries are intuitively more likely to lead to tied situations. Indeed, it has been shown that the impartial culture maximises the probability for majority cycles (Tsetlin et al., 2003), one of the necessary conditions for a tie. However, from our simulations we see more ties for the positively correlated culture; we conjecture that this is because two opposing criteria need to be fulfilled for there to be a tie: there need to be majority cycles, but these must occur among athletes with the same *scores*. Regardless, the incidence of ties is low even for the positive culture. Of the three cultures, we see fewest ties with the negatively correlated culture which best represents our expectations for the competition.

In contrast to the low incidence of ties, there does seem to be a high potential for manipulation, of both kinds. For each culture approximately half the profiles are manipulable.²⁰ We make two further observations: first, the incidence of manipulation without sacrifice that is not also completely without sacrifice is very small—the value obtained when subtracting the value of column five from the value of column four. A loose interpretation is that for inverse-Borda-Nash there is not much difference between the stronger and weaker versions of the without sacrifice property. Second, spoiler manipulation seems less likely under the positively correlated culture. An intuitive explanation for the lower incidence of spoiler manipulation for the positively correlated culture is the following: for this culture it is more likely that one athlete in a pair will *always* be ranked above their teammate, in which case the worse ranked athlete cannot spoil. Nevertheless, even for the positive culture there is a non-negligible potential for spoiler manipulation (more than 10% of the generated profiles).

1.5 Impossibilities that suggest a trade-off between decisiveness and manipulability: axiomatic results

In this section we investigate the general theoretical properties of discipline aggregators. Our investigation leads to the introduction of decisiveness as a desirable condition.

Inverse-Borda-Nash is manipulable. However, this doesn't mean that a better discipline aggregator exists. Our first result of this section is the following: one cannot prevent both kinds of manipulation without violating a desirable condition. This is an impossibility result in the classical sense. Apart from the conditions concerning manipulation, the result uses only the

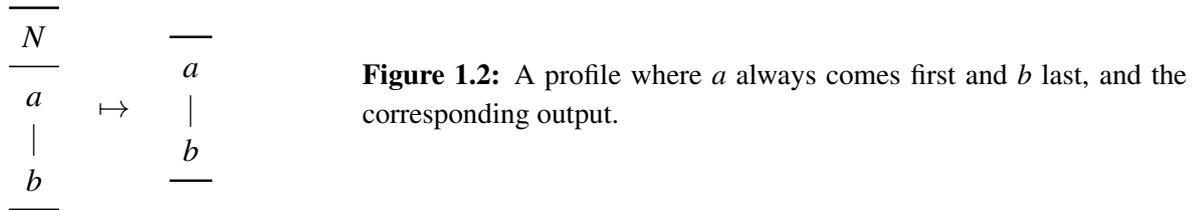
²⁰We also tested profiles with twenty athletes for manipulation. Each culture resulted in higher counts of potential manipulation than in the six athlete case. Of course, to fully address the issue of manipulation in the qualification round would require other modifications: here a manipulation is desirable if and only if it moves the target athlete below the sixth place threshold; it can also be noted that we no longer have an argument for why the non-manipulated profile is common knowledge.

weakest of our conditions; in particular, it does not rely upon the relatively strong conditions of athlete-neutrality and discipline-neutrality, and uses the weaker “not determined” condition as opposed to “not weakly determined”. It thus shows that almost any reasonable discipline aggregator must allow for some kind of manipulation, some of the time.

Theorem 1.4. *No discipline aggregator prevents spoiler manipulation, prevents manipulation without sacrifice, satisfies the clear winner condition and is not determined.*

Proof. We proceed by showing that if both types of manipulation are prevented and the clear winner condition is satisfied then there a determined discipline, that is, a discipline $i \in N$ such that, for $a \in A$ such that $r_i(a) = 1$, aPx for all $x \neq a$. The proof follows the structure of the proof of Arrow’s theorem by Reny (2001).

Take an arbitrary discipline aggregator F that prevents spoiler manipulation and manipulation without sacrifice and that satisfies the clear winner condition. Consider any profile where a comes first in all disciplines and b comes last. By the clear winner condition a must be ranked first and b last. We express this fact as in the Figure 1.2.



Now consider moving b up in the first discipline. So long as b does not cross above a , a must still be uniquely ranked first, as otherwise the agent c that b becomes ranked above can spoil for the new winner from P^1 to P^2 . Note that this argument does not rely on the fact that a wins in all disciplines.

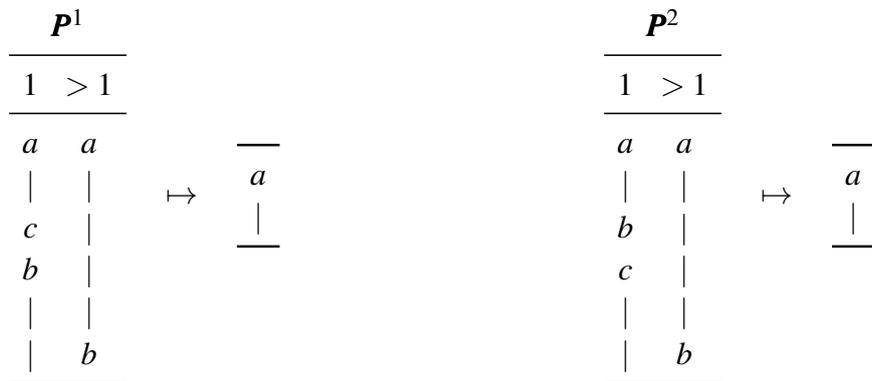


Figure 1.3: A profile P^1 for which the position of b is changed, but not with respect to the unique winner a , leading to another profile P^2 for which a is still the unique winner.

If we rank b above a , either a remains the unique winner or there is some other set of winners

X . In the second case there is some $x \in X, x \neq a$. If we continue to successively rank b first for the remaining disciplines, then at some point b must become the unique winner by the clear winner condition—in particular when b is ranked first in all disciplines: thus the second disjunct of the previous sentence must be fulfilled at some point; there must be some profile which outputs a top, but moving b above a gives some other set of winners X . Label the discipline for which this happens i^* , and label the respective profiles as \mathbf{P}^3 and \mathbf{P}^4 . Note that in \mathbf{P}^3 athlete b has been moved up to be directly below a , the same argument as above implies that a is still the winner in the output.

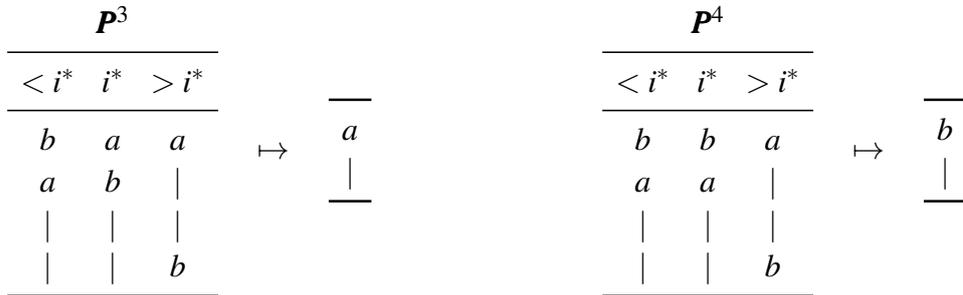


Figure 1.4: A critical profile \mathbf{P}^3 and discipline i^* such that if athlete b is ranked above athlete a to create \mathbf{P}^4 , then athlete b becomes the unique winner.

We know that $a \notin X$, otherwise a could spoil without sacrifice for x from \mathbf{P}^3 to \mathbf{P}^4 . This implies that $b \in X$, as otherwise b could spoil for a from \mathbf{P}^4 to \mathbf{P}^3 . This implies that $x \notin X$ for $x \neq a, b$, as otherwise b could manipulate without sacrifice for x from the profile where b is the unique winner.

In \mathbf{P}^4 we can move a down in the profile without changing the output winner b (otherwise a could spoil), we display this as \mathbf{P}^5 . Also consider the profile \mathbf{P}^6 created from \mathbf{P}^5 by moving a up one place in discipline i^* .

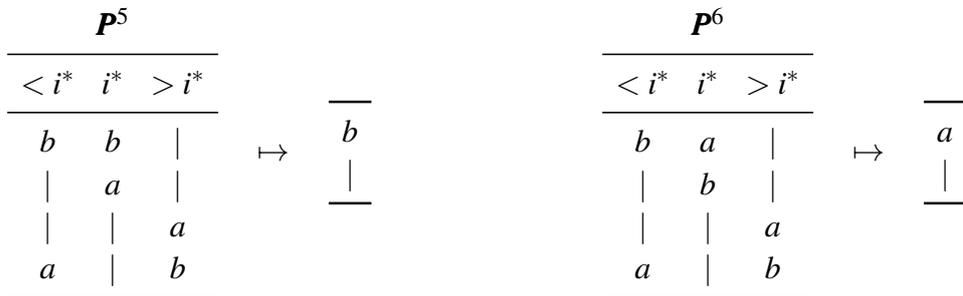


Figure 1.5: Profiles showing that i^* plays an important role for whether or not athlete a wins in the output.

We claim that a must be the unique winner in \mathbf{P}^6 . First, note that if neither a nor b were ranked first for \mathbf{P}^6 , then a can spoil for b from \mathbf{P}^6 to \mathbf{P}^5 . If b is ranked first but not uniquely ranked

first, then b can spoil without sacrifice from \mathbf{P}^5 to \mathbf{P}^6 . If b is uniquely ranked first, then at some point in stepwise changes from \mathbf{P}^6 to \mathbf{P}^3 some other athlete must perform a spoiler manipulation. Thus as b is not ranked first a is amongst the winners. If a were not unique, a could spoil without sacrifice from \mathbf{P}^3 to \mathbf{P}^6 .

Take some third alternative $c \neq a, b$. The profile \mathbf{P}^7 is obtained from \mathbf{P}^6 by moving b and c down in the profile. Here the unique winner is still a , as otherwise b or c could spoil. Create \mathbf{P}^8 by moving a to be ranked last in all disciplines except i^* .

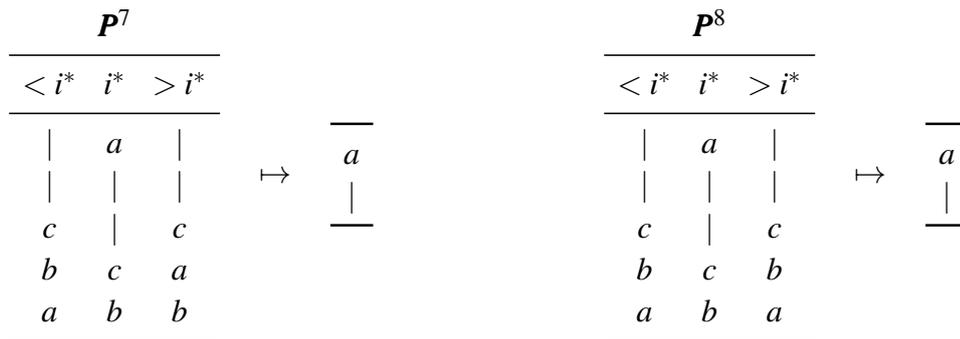


Figure 1.6: Profiles for which a is ranked first in i^* , for which a must also be the winner in the output.

In the profile \mathbf{P}^8 , alternative c is a clear winner over b , so b cannot be ranked first. If a were not ranked first then b could spoil for a from \mathbf{P}^8 to \mathbf{P}^7 . If any other athlete is ranked first, then a can manipulate without sacrifice from \mathbf{P}^7 to \mathbf{P}^8 . Thus a must be the unique winner in \mathbf{P}^8 .

In general, for any profile where a wins in discipline i^* , a must be uniquely ranked first in the output, as otherwise there would be some chain of changes from \mathbf{P}^8 to the profile in question, one of which would be a spoiler manipulation for the new winning athlete. As a was chosen arbitrarily, for each alternative x there is a discipline i_x such that whenever x wins in i_x , x must be uniquely ranked first. As two alternatives x and y cannot both be ranked first, $i_x = i_y$ for all $x, y \in A$, thus i^* determines the gold. □

The proof of Theorem 1.4 closely follows that of Reny (2001), who presents Arrow’s impossibility and the Gibbard-Satterthwaite theorem side by side. Although we consider manipulation, the result is, in terms of its formal shape, closer to the presentation of Arrow’s result than to that of the Gibbard-Satterthwaite result. Requiring the impossibility of both forms of manipulation takes the place of **binary independence**, though this requirement does not imply binary independence.²¹ Consider the discipline aggregator that returns the total pre-order where a is ranked uniquely first and all other alternatives jointly second if a is first in

²¹By binary independence we mean the specific version of Arrow’s independence of irrelevant alternatives that requires of a social welfare function F that, for all pairs $\{a, b\} \subseteq A$, for any two profiles \mathbf{R}, \mathbf{R}' such that for all $i \in N$ $xR_i y \Leftrightarrow xR'_i y$, $x F(\mathbf{R}) y \Leftrightarrow x F(\mathbf{R}') y$ (note that this definition works only because $R, R', F(\mathbf{R})$ and $F(\mathbf{R}')$ are complete, cf. the definition of 2-IND given in Chapter 4).

all disciplines, and a is ranked second and all other alternatives jointly first otherwise; this violates binary independence but prevents both kinds of manipulation. The no-manipulation requirement is closer to **Condorcet independence of irrelevant alternatives** (Yu, 2015). This means that our impossibility is not simply a corollary of Arrow’s theorem.

Theorem 1.4 is tight for the four conditions, in the sense if any one is removed there is a discipline-aggregator that satisfies the other three. **Dictatorships**, for which the ranking of a single discipline is copied, are determined but satisfy the other three conditions. **Constant functions** violate only the clear winner condition (except the function that always ranks every athlete first, which is also weakly determined). We define a function below that we call iterative first place elimination that only allows manipulation without sacrifice. Before this, we describe a function that only allows spoiler manipulation: this proceeds sequentially, at stage t , remove the athlete who is ranked last in discipline t modulo n , and rank this athlete below the other athletes remaining in the profile.²² Arguably the best exposition of this process is given by the example in Table 1.5. Formally, we define **iterative successive last removal**, $\text{isr}: \mathcal{L}^*(A)^N \rightarrow \mathcal{W}(A)$ as follows. For an arbitrary profile \mathbf{P} , let $\text{lose}_t(\mathbf{P}) = \{a \in A : r_t(a) \text{ is maximal}\}$. Let $\mathbf{P}^1 = \mathbf{P}$, and for $t \geq 1$ recursively define \mathbf{P}^{t+1} as the restriction of \mathbf{P}^t to $A \setminus \text{lose}_{t \bmod n}(\mathbf{P}^t)$. Writing $P = \text{isr}(\mathbf{P})$, for $x, y \in A$, define xRy if and only if there are integers $s, t \leq m$ such that $s \geq t$ and $x \in \text{lose}_{s \bmod n}(\mathbf{P}^s)$ and $y \in \text{lose}_{t \bmod n}(\mathbf{P}^t)$.

Proposition 1.5. *Iterative successive last removal prevents manipulation without sacrifice, satisfies the clear winner condition, is not weakly determined and satisfies athlete-neutrality.*

Proof. Prevents manipulation without sacrifice: suppose an athlete “manipulates” by performing worse in a profile but also that she does not get a worse output ranking. Thus she is removed at the same point t and has output rank $m - t + 1$. Clearly, all the partial profiles after this point will be the same as in the non-manipulated case. As she was not removed before t , this means that for all the partial profiles at stage $s < t$ she was not ranked last in discipline s modulo n , this means that she did not change the athlete who was ranked last in this discipline, thus the loser at this stage will be the same.

The clear winner condition: if a is better than b in all disciplines then a cannot be removed before b .

Not weakly determined: for disciplines $i \neq 1$, consider the profile where the athlete ranked first in i is ranked last in 1. For discipline 1 consider the profile where the athlete ranked first in 1 is ranked second last in 2. \square

We cannot satisfy all our properties simultaneously. However, if we strengthen manipulation without sacrifice to manipulation completely without sacrifice there are functions that work.

²²Cf. the Coombs rule described by, for example, Grofman and Feld (2004) or in Freeman et al. (2014).

Table 1.5: The process of iterative successive last removal applied to the final round given the results of Table 1.1, supposing that Faye was eliminated by the external tie-breaker and the remaining athletes perform as in the first round.

Ranking	Speed	Bouldering	Lead	Final ranking
1st	Barbara	Dorothy	Chloe	
2nd	Emma	Grace	Emma	
3rd	Alice	Alice	Dorothy	
4th	Chloe	Barbara	Grace	
5th	Grace	Emma	Barbara	
6th	Dorothy	Chloe	Alice	
1st	Barbara	Grace	Chloe	
2nd	Emma	Alice	Emma	
3rd	Alice	Barbara	Grace	
4th	Chloe	Emma	Barbara	
5th	Grace	Chloe	Alice	
6th				Dorothy
1st	Barbara	Grace	Emma	
2nd	Emma	Alice	Grace	
3rd	Alice	Barbara	Barbara	
4th	Grace	Emma	Alice	
5th				Chloe
6th				Dorothy
1st	Barbara	Grace	Emma	
2nd	Emma	Barbara	Grace	
3rd	Grace	Emma	Barbara	
4th				Alice
5th				Chloe
6th				Dorothy
1st	Barbara	Barbara	Emma	
2nd	Emma	Emma	Barbara	
3rd				Grace
4th				Alice
5th				Chloe
6th				Dorothy
1st	Barbara	Barbara	Barbara	Barbara
2nd				Emma
3rd				Grace
4th				Alice
5th				Chloe
6th				Dorothy

Table 1.6: The process of iterative first place elimination applied to the final round given the results of Table 1.1, supposing that Faye was eliminated by the external tie-breaker and the remaining athletes perform as in the first round.

Ranking	Speed	Bouldering	Lead	Final ranking
1st	Barbara	Dorothy	Chloe	
2nd	Emma	Grace	Emma	
3rd	Alice	Alice	Dorothy	
4th	Chloe	Barbara	Grace	
5th	Grace	Emma	Alice	
6th	Dorothy	Chloe	Barbara	
1st				Barbara, Chloe, Dorothy
2nd				(none)
3rd				(none)
4th	Emma	Grace	Emma	
5th	Alice	Alice	Grace	
6th	Grace	Emma	Alice	
1st				Barbara, Chloe, Dorothy
2nd				(none)
3rd				(none)
4th				Emma
5th	Alice	Grace	Grace	
6th	Grace	Alice	Alice	
1st				Barbara, Chloe, Dorothy
2nd				(none)
3rd				(none)
4th				Emma
5th				Grace
6th	Alice	Alice	Alice	Alice

The function we define proceeds in stages, determining the top ranked candidates then removing them from the profile. The function may be described as a back-to-front version of instant runoff voting applied using two **quota rules**. If an athlete is ranked first in strictly more than half the disciplines, then she is the unique winner with respect to the athletes in the profile—the quota is $(m + 1)/2$. Otherwise, any athlete that has at least one first place ranking in the profile is a joint winner—that is, the quota is 1. The winners are removed from the profile, and the procedure repeats. We name this **iterative first place elimination**, $\text{ifpe}: \mathcal{L}^*(A)^N \rightarrow \mathcal{W}(A)$. See Table 1.6 for an example. Formally, for an arbitrary profile \mathbf{P} , let

$$\text{localwin}(\mathbf{P}) = \begin{cases} \{a\} & \text{if } \exists a \in A, |\{i \in N : r_i(a) = 1\}| > n/2 \\ \{x \in A : \exists i \in N, r_i(x) = 1\} & \text{otherwise.} \end{cases}$$

Let $\mathbf{P}^1 = \mathbf{P}$. For $t \geq 1$, recursively define \mathbf{P}^{t+1} as the restriction of \mathbf{P}^t to $A \setminus \text{localwin}(\mathbf{P}^t)$. Writing $R = \text{ifpe}(\mathbf{P})$, for $x, y \in A$, define xRy if and only if there are integers s, t such that $s \leq t$ and $x \in \text{localwin}(\mathbf{P}^s)$ and $y \in \text{localwin}(\mathbf{P}^t)$.

Proposition 1.6. *Iterative first place elimination prevents spoiler manipulation, prevents manipulation completely without sacrifice, satisfies the clear winner condition, is not weakly determined if $n \geq 3$, satisfies athlete-neutrality and satisfies discipline-neutrality.*

Proof. Prevents spoiler manipulation: an athlete clearly cannot affect any of the partial profiles \mathbf{P}^t starting from $t = 1$ until she is ranked first in one of the disciplines. So suppose that the athlete is ranked first in the partial profile. There are two possibilities

- 1) The athlete becomes a winner for this partial profile, in which case she cannot spoil because she does as well as the other athlete for which she spoils.
- 2) The athlete is not a winner for this partial profile, in which case a different athlete is ranked first in more than half the disciplines; this second athlete will be the winner no matter how the putative manipulator changes her ranking.

Prevents manipulation completely without sacrifice: we know that an athlete a cannot affect athletes that get better output ranks. Suppose that \mathbf{P} is the partial profile for which $a \in \text{win}(\mathbf{P})$. First suppose a is ranked first in more than half the disciplines: if she performs worse in enough of these disciplines she will no longer be the unique winner, but this violates the strictness condition. Otherwise, she will be removed from the profile in the next step, thus any other changes to her ranking do not affect the output. Now suppose a is ranked first in less than half the disciplines. If a performs worse in a discipline i for which $r_i(a) > 1$, this will not affect the output ranking as a is removed from the profile in the next round. If a performs worse in a discipline i for which $r_i(a) = 1$, there are three possibilities:

- 1) A different athlete becomes the unique winner, thus a is ranked lower in the output.

- 2) A new athlete becomes a winner, violating the strictness condition.
- 3) The winners remain the same, thus the same athletes will be removed from this profile and the output will not change.

Satisfies the clear winner condition: if a is ranked better than b in all disciplines, it is not possible that b become ranked first in a partial profile if a is still contained in the profile.

Not weakly determined: here we require the condition that $n \geq 3$; for arbitrary $i \in N$ consider the profile in Table 1.7.

Table 1.7: Profile demonstrating that, if the number of disciplines is greater than or equal to three, iterative first place elimination is not weakly determined.

Athlete	Ranking in discipline		Product
	$\{i\}$	$N \setminus \{i\}$	
a	2	1	2
b	1	2	2^{n-1}
others	≥ 3	≥ 3	$\geq 3^n$

□

Even though it satisfies our state properties, iterative first place elimination is nonetheless unsatisfactory. The intuitive problem is that it is *indecisive*, where we use the (slightly imprecise) term **decisive** to describe how often ties are produced in the output. A strong version does provide a clear definition: a **resolute** function always produces a linear order. Requiring this level of decisiveness recreates the impossibility because it makes manipulation without sacrifice equivalent to the strong version. However, resoluteness is *too strong*—it implies that either athlete-neutrality or discipline-neutrality is violated. Given a completely symmetric profile, the results of the competition are not enough to tell the athletes apart. In such a case it seems reasonable that the discipline aggregator outputs ties, and to suppose that conditions *external* to the profile will be used to break them.

We can note some differences in focus for the multi-discipline interpretation, in comparison to traditional social choice: it stresses the importance of having a minimal *rank*; an athlete is mainly concerned with the number of athletes ranked *strictly higher* than her in the output complete preorder. The other difference is a particular importance on the “decisiveness” of the function; how often ties are output at any ranking level. Authors often sidestep the issue of ties in order to obtain their main results, by supposing that there is an exogenous linear order tiebreaker or by restricting the output to linear orders, but this is obviously unsatisfactory for our purposes as it is the question of ties themselves that we are interested in. We are not aware of a good reference for this subject, but it provides a natural further step: starting from resolute functions, to what extent must we relax decisiveness in order to prevent manipulation?

1.6 How easy is it to determine whether or not to manipulate: complexity results

The fact that manipulation is possible does not necessarily imply that it is *feasible*: there may be a variety of barriers to manipulation in practice. One such barrier, expressed using the theoretical tools of computer science, is that it may be *too computationally demanding* to determine how to manipulate. In fact, prior even to this, we show that it is too demanding to even determine whether or not a profitable manipulation can be carried out. Such ideas have been pursued in computational social choice theory following the seminal work of Bartholdi et al. (1989). In this section we apply these ideas taken from computational social choice to our definitions of manipulation.

What does it mean to say that a problem is too hard? Suppose that we have an algorithm—a specific set of instructions—that solves the problem in, for example, ten simple steps: in such a case the problem doesn't seem too hard. Of course, the length of time taken by an algorithm typically varies depending upon the input; indeed, just to read all the input data when there are more than ten disciplines will take more than ten steps. So an algorithm that fully reads the input takes *at least* linear time; in the other direction, we say that the algorithm *is* linear if this is in a sense the most time that it takes, i.e. if the number of steps required is bounded from above by a (fixed) multiple of the size of the input. More generally, a problem is considered tractable if, for some (fixed) *polynomial*, there is an algorithm that always solves the problem in less steps than the number given when the input size is inserted into the polynomial—the words “always” and “less” mean that we consider the amount of time that it would take to find a solution *in the worst case*. Some algorithms take more time than this. In particular, some algorithms take *exponentially* many steps with respect to the size of the input. This means that above a certain size of input the algorithm will certainly take too long. If the *only* algorithms that can solve a problem are exponential, the problem may be considered *too hard*, or in other words, *intractable*.

Proving that a problem can only be solved by an exponential algorithm is not the only method used to show that a problem is too hard. There is another class of problems, nondeterministic polynomial problems (NP), which are commonly considered to serve as a boundary for tractability. Although the exact difficulty of problems in this class is not known, it is sometimes possible to prove that a given problem is *at least as hard* as any other problem in this class. Roughly speaking, it is shown that if one has an algorithm for the problem in question, then this can be used to solve any of the problems in the class. A problem that can be so used is called NP-hard. Under the famous assumption that $P \neq NP$ —that some NP problems are not polynomial—this means that for any polynomial there are infinitely many inputs whose computation time will be larger than the polynomial; a proof that a problem

is NP-hard is typically taken as evidence that the problem is intractable. In this section we show that determining whether or not there is a manipulation without sacrifice is NP-hard.

Given that we have managed to perform simulations, it is obviously easy in practice to determine whether or not a manipulation is possible. Although this is not of direct importance to the case of sports climbing, which involves three disciplines and at most twenty candidates, it may well be important for other applications of our model. We precisely state our decision problem concerning manipulation as follows.

CAN-MANIPULATE-WITHOUT-SACRIFICE:

Input: A set of athletes A , a set of disciplines N , a profile $\mathbf{P} \in \mathcal{L}(A)^N$, two specific athletes a and b .

Question: Is there a manipulation without sacrifice by a for b from \mathbf{P} to some other profile?

We show that CAN-MANIPULATE-WITHOUT-SACRIFICE becomes a difficult question to answer as the number of athletes and disciplines increases.

Proposition 1.7. CAN-MANIPULATE-WITHOUT-SACRIFICE is NP-complete.

Proof. For membership, take the profile that is manipulated to as a certificate and verify that this actually is a manipulation, i.e. that b does actually end up ranked better, and that a does not end up ranked worse.

For hardness, we provide a polynomial reduction from EXACT-3-COVER, which is the following NP-complete problem:

EXACT-3-COVER:

Input: A set $X = \{x_1, \dots, x_{3t}\}$ of elements, a set $\mathcal{X} = \{X_1, \dots, X_s\}$ of subsets of X such that for each $Y \in \mathcal{X}$, $|Y| = 3$.

Question: Is there a subset $\mathcal{X}' \subseteq \mathcal{X}$ such that $\bigcup \mathcal{X}' = X$ and for any distinct pair $X_i, X_j \in \mathcal{X}'$, $X_i \cap X_j = \emptyset$?

Suppose we are given an instance of EXACT-3-COVER. We use this to construct—in time polynomial in the size of the original instance—a new instance of CAN-MANIPULATE-WITHOUT-SACRIFICE, such that an algorithm solving the manipulation decision problem returns “yes” for this new instance if and only if EXACT-3-COVER returns “yes” for the original instance.

The athletes in the instance are $\{a, b, c, d\} \cup X \cup X' \cup X''$. The sets X' and X'' are of fixed cardinality; there are $|X| + 14$ athletes in total. The purpose of these athletes in the reduction is as follows: a is the potential manipulator; b is the athlete she attempts to help; c is

an athlete who defeats b according to the starting profile, but who will be tied with b if a manipulates in at least t disciplines; d is an athlete defeated by b according to the starting profile but who will defeat b if a manipulates in strictly more than t disciplines; each $x_i \in X$ is defeated by a according to the starting profile, but threatens to defeat a in the course of the manipulation—in particular a will only be able to improve the ranking of x_i at most once; and $X' = \{x'_1, \dots, x'_8\}$ and $X'' = \{x''_1, x''_2\}$ contain dummy athletes who pad rankings and ensure that cycles are of the right length. In the starting profile, each discipline that a potentially changes her ranking in corresponds to some $Y \in \mathcal{X}$. For each of these disciplines, if a changes her ranking then she improves the ranking of each $x_i \in Y$ (though by different amounts: see Figure 1.8). We design the starting profile in such a manner that there is a successful manipulation without sacrifice if and only if EXACT-3-COVER returns “yes” for the original problem.

We construct the starting profile from two parts: a base subprofile P^{base} and a padding subprofile P^{pad} . In the base subprofile the athletes $\{b, c, d\}$ all receive the same product score, and the athletes $\{a\} \cup X$ all receive the same product score; it also contains all those disciplines for which a might want to worsen her ranking in order to manipulate for b . The padding profile creates the required differences in the product scores of the athletes.

For the base profile, for each triple $Y \in \mathcal{X}$ we create $3t + v + 1$ disciplines and the associated rankings, where v is the smallest natural number such that $3t + v + 1$ is divisible by nine. Thus $P^{\text{base}} = \{P_Y^i : Y \in \mathcal{X}, i \in \{1, \dots, 3t + v + 1\}\}$. For each triple Y only the ranking P_Y^1 can be profitably changed by a during a manipulation for b , we thus will refer to P_Y^1 as corresponding to Y . The remaining rankings make athletes within two specific sets have the same product scores. For P_Y^1 , the two dummy athletes x''_1 and x''_2 are placed at the top of the ranking, then c , a , and d , then (in arbitrary order) the elements in the triple Y , then b , then (in arbitrary order) the remaining elements of $X \setminus Y$, and finally the elements in X' . To ensure that athletes in $\{b, c, d\}$ have the same product scores, and that athletes in $\{a\} \cup X$ have the same product scores, we also add rankings P_Y^i for $i = \{2, \dots, 3t + v + 1\}$ that cycles the positions of the athletes in these two sets, also requiring that for each of these rankings b performs better than a . The dummy athletes X' ensure that these cycles can be realised—Figure 1.7 demonstrates the cycles, and shows how the dummy athletes in X' allow for b to perform better than a in these cases. In particular, elements from X' are added to the set $\{a\} \cup X$ that is cycled through so that the cardinality of the created set is divisible by nine. This means that we only need to cycle through the set containing a exactly once while we cycle through the set containing b three times slower. We can then place a sixth in P_Y^2 , seventh in P_Y^3 and eighth in P_Y^4 , and place b third in these three rankings: this means that for a given $Y \in \mathcal{X}$, the only ranking P_Y^i for which a is ranked better than b is that where $i = 1$. For a given $Y \in \mathcal{X}$, we create less than $|X| + 9$ rankings: in total the number of disciplines in the base profile is bounded from above by $(|X| + 9)|\mathcal{X}|$.

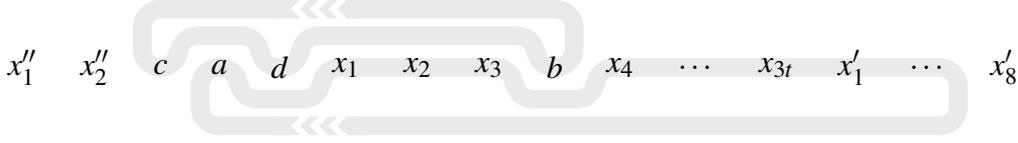


Figure 1.7: The ranking $P_{\{x_1, x_2, x_3\}}$ corresponding to the set $\{x_1, x_2, x_3\} \in \mathcal{X}$. Leftwards athletes perform better than rightwards athletes. The required cycling through the ranking is indicated by the thick arrows. It can be ensured that a is ranked below b for all other rankings in the cycling if we cycle according to the higher arrow once every three times we cycle according to the lower arrow. In the cycle indicated by the lower arrow, of the dummy athletes we cycle through only the $X''' \subseteq X'$ required such that $1 + |X| + |X'''|$ is divisible by nine; this ensures that both cycles end at the same point.

For the padding part of the profile P^{pad} , we want to ensure, firstly, that the following ratios between the product scores of the athletes hold:

- (A) $\prod r_i(x'') \leq \prod r_i(c)$ for all $x'' \in X''$,
- (B) $\prod r_i(c) = (8/9)^t \prod r_i(b)$,
- (C) $\prod r_i(b) = (9/10)^t \prod r_i(d)$,
- (D) $\prod r_i(d) \leq \prod r_i(a)$,
- (E) $\prod r_i(a) = (4/9)^t (4/5) \prod r_i(x)$ for all $x \in X$,
- (F) $\prod r_i(a) < (4/9)^t \prod r_i(x')$ for all $x' \in X'$;

secondly, that b is not ranked worse than a in any of the rankings; and thirdly, that the number of disciplines required is bounded by a polynomial in the size of the original input. Let us assume for now that we can create the required padding profile. Because the elements in $\{b, c, d\}$ have the same product scores as each other in the base profile, and ditto for the elements in $\{a\} \cup X$, and because elements in X' have larger product scores and elements in X'' smaller, (A) to (F) also apply in the full profile consisting of the combination of the base and padding profiles. We now show that this full profile is manipulable if and only if the original problem has an exact covering by three sets.

Suppose that we have a “yes” instance of EXACT-3-COVER, so there exists an exact covering \mathcal{X}' of X . Let a worsen her position, so that she is directly below b , in exactly the t disciplines corresponding to \mathcal{X}' —an example of how this affects each one of these t rankings is displayed in Figure 1.8—we claim that this is a successful manipulation without sacrifice by a for b . In the resulting profile the product score of c remains the same, but the product score of b is multiplied by $(8/9)^t$. Because, by (B), in the original profile the product score of c is $(8/9)^t$ times that of b , in the resulting profile b and c have the same product score. Similarly, the product score of d is multiplied by $(4/5)^t$; so (B) and (C) imply that the product scores of a , b and c are all the same in the resulting profile. We now need to show that none of the athletes who were ranked worse than b become ranked better than b , and, to ensure that the

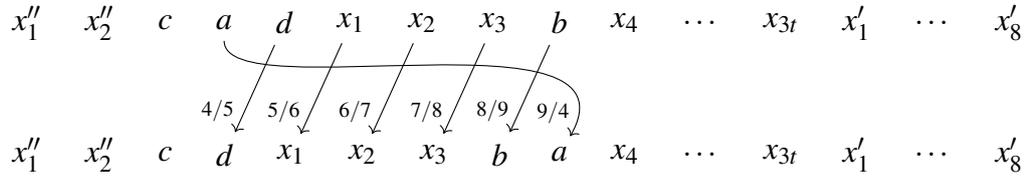


Figure 1.8: The ranking of a single discipline before and after a manipulates. The numbers display the ratio, if not equal to one, of an athlete's product scores before and after this particular manipulation.

manipulation is without sacrifice, to show that none of the athletes who were ranked worse than a become ranked better than a . In fact, because those athletes ranked worse than b were either a or themselves ranked worse than a , and because the product score of a has increased, it suffices to show the without sacrifice part. To see that this holds, note that for each $x \in X$ the product score is decreased by at most $5/6$, because they only have their ranking changed in exactly one discipline and $5/6 < 6/7 < 7/8$. Because for a the product score is increased by $(9/4)^t$ and because $4/5 < 5/6$, condition (E) implies that these athletes remain worse ranked than a . To be precise, the product score of a in the manipulated profile is $24/25$ times that of the product score of each $x \in X$ in the manipulated profile. Similarly, for each $x' \in X'$ the product score doesn't change, thus (F) implies that a still defeats x' as required.

In the other direction, suppose that there is a successful manipulation without sacrifice. As the only athletes that b is ranked worse than in the non-manipulated profile are x''_1 , x''_2 and c , and because it is impossible that b becomes ranked at least as high as x''_1 or x''_2 through a manipulation of a (by (A) and the fact that x''_1 and x''_2 perform best in each ranking of the base profile), this means that b must be ranked at least as high as c in the manipulated profile, which means that a must change her ranking in at least t disciplines where she is ranked better than b . However, a cannot change her ranking in more of these disciplines, otherwise d would become ranked strictly above b and the manipulation would not be successful. Also, of the disciplines that she changes her ranking in, she can only improve the ranking of any given x_i at most once, otherwise this x_i would multiply her product score by a value less than $(7/8)^2$ —itself less than $(4/5)$ —so x_i would be ranked better than a and the manipulation would not be without sacrifice. Thus the disciplines in which a does manipulate correspond to an exact covering of the set X as required.

It remains to show that the desired padding profile can be constructed. This construction is facilitated by the following technical lemma.

Lemma 1.8 (Assign a specific ratio of product scores). *Given disjoint sets X^I, X^{II}, X^{III} such that $X^I \cup X^{II} \cup X^{III} = A$, distinct athletes $x, y \in X^{II}$, and natural numbers*

$$p, q \in \{|X^I| + 1, |X^I| + 2, \dots, |X^I \cup X^{II}|\},$$

there is a profile, over an index set I of cardinality the lowest common multiple of $|X^I|$, $|X^{II}|$ and $|X^{III}|$, such that

$$\begin{aligned} \prod_{i \in I} r_i(z) &= \prod_{i \in I} r_i(w) && \text{for } j \in \{I, III\}, z, w \in X^j, \\ \prod_{i \in I} r_i(z) &< \prod_{i \in I} r_i(w) && \text{for } j < k, j, k \in \{I, II, III\}, z \in X^j, w \in X^k \\ \prod_{i \in I} r_i(x) &= (p/q) \prod_{i \in I} r_i(z) && \text{for } z \in X^{II} \setminus \{x, y\} . \end{aligned}$$

Proof of Lemma 1.8. For $p = q$, we create a profile that cycles through the three sets of alternatives separately: choose arbitrary orderings of X^I , X^{II} , and X^{III} ; concatenate the list representations of these to create the first ranking in the profile. Now, cycle the positions of the elements within the three sets, each time adding the new ranking to the profile: after as many steps as the lowest common multiple of the cardinalities of these three sets we will reach a ranking such that the next ranking is again the first ranking. For each set, each element in this set is in each position the same number of times, thus elements within a set have the same product scores.

In fact, the $p \neq q$ case is dealt with in almost the same manner, we only modify one of the rankings in the cyclical profile. Create a base ranking starting with elements of X^I in arbitrary order, then elements of X^{II} such that x is at position q and y is at position p and with the other elements in arbitrary order, and finally elements of X^{III} in arbitrary order. The profile is built up of cyclings through this ranking, with the base ranking replaced once by the ranking where the positions of x and y are swapped. The cycles go individually through X^I , X^{II} and X^{III} . ■

Example 1.9. Suppose that $X^I = \{a, b\}$, $X^{II} = \{x, y, c\}$, $X^{III} = \{d\}$, $p = 3$, and $q = 5$. The base ranking is $abxcd$. The profile constructed according to the proof of Lemma 1.8 contains the rankings $(abxcd, baxycd, abcxyd, bayxcd, abxycd, bacxyd)$: for this profile a and b have the same product score of 2^3 —which is also the smallest, c has a product score of $(3.4.5)^2$, x has a product score of $3^3 \cdot 4^2 \cdot 5 = (3/5)(3.4.5)^2$, y has a product score of $3 \cdot 4^2 \cdot 5^3$, and d has the largest product score of 6^6 .

We repeatedly apply Lemma 1.8 in order to get the required ratios between the product scores of the athletes. For some of these applications we let $X^I \supseteq X''$: this ensures that for these subprofiles (A) holds; when this is not the case we will occasionally require some extra compensatory padding.

In order to achieve the ratio given by (B), we use the above lemma $3t$ times. First we apply it t times with the values $X^I = \emptyset$, $X^{II} = X'' \cup \{b, c, d\}$, $X^{III} = X \cup X' \cup \{a\}$, $x = x''$, $y = c$, $p = 3$, and $q = 4$. Note here that x''_1 actually has the smallest product score in this

partial profile, and x_2'' has the joint second smallest, alongside b and d —no compensatory padding is required. Note also that the number of disciplines required for each application of the lemma is bounded (conservatively) from above by $(|X| + 14)^3$, because the lowest common multiple of three numbers that sum to a value is certainly less than the cube of that value. We now again apply the lemma t times, with almost identical values: $X^I = \emptyset$, $X^{II} = X'' \cup \{b, c, d\}$, $X^{III} = X \cup X' \cup \{a\}$, $x = x_1''$, $y = c$, $p = 3$, and $q = 2$. Combined with the previous t applications, this gets the desired ratio of $(4/3)^t (2/3)^t = (8/9)^t$ for (B). However, for these applications x_1'' has a larger product score than b , c and d ; and x_2'' has a larger product score than c . To redress this we apply the lemma t times with $X^I = X''$, $X^{II} = \{b, c, d\}$, $X^{III} = X \cup X' \cup \{a\}$, and $p = q$ (thus x and y can be chosen arbitrarily). Because $3t = |X|$, the number of disciplines used up to now is certainly bounded from above by $(|X| + 14)^4$.

The process for (C) involves an almost identical application of the lemma $3t$ times, only for the first $2t$ applications here $x = x_1''$, $y = d$ and (p, q) takes the values $(2, 3)$ and $(5, 3)$ for the first and second t applications respectively. Again, the final t applications compensate for the larger product score of x_1'' (and to a certain extent for the product score of x_2'') in the t preceding applications: we again apply the lemma t times with $X^I = X''$, $X^{II} = \{b, c, d\}$, $X^{III} = X \cup X' \cup \{a\}$, and $p = q$.

In total, the number of disciplines used up to now is bounded from above by $2(|X| + 14)^4$.

The inequality of part (D) holds because for all the applications of the lemma, we place d in some set X^i and a in some set X^j , for $i, j \in \{I, II, III\}$, such that $i < j$.

For part (E), we apply the lemma $t + 1$ times. For the first t applications, let $X^I = \{b, c, d\} \cup X''$, $X^{II} = \{a\} \cup X \cup X'$ and $X^{III} = \emptyset$, and let $x = a$, $y = x_1'$ and $p = 8$, $q = 18$. Note that for this subprofile x_1' has the largest product score of all the athletes. We also apply the lemma once with the same values except $p = 8$ and $q = 10$. To again give a conservative upper bound, we require less than $3(|X| + 14)^4$ disciplines.

And in fact this suffices; it is not hard to see that (F) is also satisfied. Because for all of these rankings b is ranked better than a , and their amount is polynomially bounded from above, we have the required padding profile P^{pad} . \square

1.7 Final remarks on discipline aggregation

In this chapter we have seen a novel interpretation of Arrow's traditional social choice framework involving the aggregation of linear orders. Under this interpretation what are traditionally thought of as candidates become the agents of the model. These agents can strategize in a specific manner: they can worsen their own position within one or more of the in-

put strict linear orders. This interpretation captures the problem of aggregating multiple ranked competitions. In particular we consider the proposed method for sports climbing at the 2020 Olympics. Simulations suggest that, although ties are unlikely to occur, this proposed method is potentially open to manipulation. Although it is impossible to completely prevent manipulation under its least restrictive definitions, a small further assumption about the sacrifices that athletes are unwilling to make means that we can find other methods that prevent manipulation. This possibility is demonstrated through a novel method which is, however, too susceptible to ties to be practical.

Our interpretation is novel to the best of our knowledge. Other work concerning manipulation in sports competitions includes work concerning manipulating seedings (Russell and Van Beek, 2011), and tends to be of a more operations-research nature than social-choice-theoretic. Wright (2014) provides a survey. We are not aware of other work that explicitly considers candidates as agents in the way that we do—our work is distinct from the strand of literature which considers manipulation by strategic candidacy (Dutta et al., 2001; Lang et al., 2013; Brill and Conitzer, 2015). Of course, there are strong similarities between our results and more traditional work in social choice theory, and there may be implicit connections that we have not determined, for example with Yu’s (2015) definition of Condorcet independence of irrelevant alternatives or Sanver and Zwicker’s (2009) definition of one-way monotonicity. Our types of manipulation are also novel in terms of their formal definitions, although there are some similarities with work that looks at restricted manipulations, for example where voters aim to be in some sense truthful (Obraztsova and Elkind, 2012). In particular, it should be noted that our version of manipulation is not just a restricted version of coalitional manipulation; rather, the ranking in *every* discipline can be changed in a restricted manner. This means that we have an orthogonal result to Proposition 1 of Conitzer and Sandholm (2002), which states that if calculating the output of a rule is easy, then fixing the number of athletes leads to easy decision problems concerning manipulations; for our definitions of manipulation it is instead a fixed number of disciplines that leads to easy decision problems.

Our interpretation fits well into Arrow’s framework. Table 1.8 contains a comparison of the terminology used in this chapter with traditional terminology. Arguably, the problem of aggregating multiple disciplines is better served by this framework than the typical problems of social choice theory. The linear order profile is the input in practice.²³ There are no questions, as there are for social choice theory, about whether eliciting full linear or-

²³Rather than aggregating linear order profiles, a method could aggregate points determined by each individual’s performance, as in the modern pentathlon. The final score for an athlete would be independent of the performance of other athletes, sidestepping our issue of manipulation. This approach was discarded by the International Federation of Sports Climbing because (1) it is considered too complex for the audience and (2) it is difficult, perhaps impossible, to assign points in a balanced way across the disciplines (Meyer, 2018). We thus take it as given that the input is ordinal.

Table 1.8: A comparison of the terminology used for the discipline aggregation framework with that used for social choice theory.

Discipline aggregation	Social choice theory
Athletes	Alternatives
Disciplines	Agents
Discipline aggregator	Social welfare function
Clear winner condition	Pareto criterion
Athlete-neutrality	Neutrality
Discipline-neutrality	Anonymity
Not determined	No top-dictator
No egoistic manipulation	Monotonicity
No altruistic manipulation	No specific identical concept, similarities with binary independence and Gibbard-Satterthwaite manipulation
No non-sacrificial manipulation	
No spoiler manipulation	

ders is problematic, let alone if linear order preferences are suitable or even sensible—cf. competing approaches like approval voting (Brams and Fishburn (1978), for a recent extensive treatment see Laslier and Sanver (2010)) and majority judgment (Balinski and Laraki, 2011).²⁴ The required output is also obviously a complete preorder, whereas in social choice theory often what is desired is a *choice*, requiring a social choice function as opposed to a social welfare function. Sometimes it is not obvious that manipulation is actually undesirable for social choice theory, especially when one considers iterative manipulation (Meir et al., 2010; Lev and Rosenschein, 2012; Obratzsova et al., 2015). Indeed, “manipulation” is a misnomer, a better term is *strategic behaviour*. In contrast, for sports competitions manipulation is aptly named and clearly undesirable in and of itself, whether because it goes against the spirit of the competition or because it cheapens the spectacle. Concerning the information requirements for manipulation, we have argued that the qualification round can be used as a proxy for the results in the final round. This is to some degree unrealistic—the athletes will not perform exactly the same—however it is certainly not less realistic than the traditional Gibbard-Satterthwaite assumption of common knowledge of all preferences of all agents. We also observe that cheating *does* occur in the Olympics, including in the form of deliberate bad performances (Jimenez, 2012; TheGuardian, 2012). In another arena, Formula 1,

²⁴Interestingly, Balinski and Laraki (2011) use the example of Olympic figure skating as part of their argument against the ordinal approach. In the past the ranking of skaters was produced by aggregating multiple *ordinal* rankings given by multiple judges. The particular method has since been replaced, and it is argued that reason underlying the replacement is that the method violates independence. Balinski and Laraki thus go in the opposite direction to us: they use experience from the Olympics and apply it to social choice theory.

altruistic manipulation of the type we consider—more precisely, spoiler manipulation—has been observed to occur (Gpupdate, 2008).²⁵

Although it fits well into the traditional framework, there are two ways in which the focus of the discipline aggregation problem differs from that the traditional view: it stresses the importance of having a minimal *rank*; and it places particular importance on the decisiveness of the ranking method, that is, how often ties are output at any ranking level. Authors often sidestep the issue of ties in order to obtain their main results, perhaps by supposing that there is an external linear order tiebreaker, or by restricting the output to linear orders (Zwicker, 2016, p. 33). Alternative approaches, such as dealing directly with set-valued outcomes (Barberà et al., 2004; Brandt, 2011; Brandt and Brill, 2011), or using a randomised mechanism to break ties (Obraztsova et al., 2011), are also unsatisfactory for our purposes because it is the precisely the problem of *minimising the number of ties* itself that we are interested in. Although decisiveness is also required for elections, the sidestepping of the issue is perhaps justified when there are large electorates. However, for our problem the small number of disciplines makes the issue of decisiveness particularly important. So although even for a small election involving a hundred voters it seems likely that for a given profile a rule will distinguish the candidates, this is no longer the case for a “population” of three disciplines.

An interesting extension of our model would be to apply a *sequential protocol approach*, where one considers partial revelation of the profile in a sequential manner. This is precisely how the Olympic sports climbing event will unfold, though in practice there will be measures put in place to isolate the athletes from the partial results during the progress of competition for a single discipline. Such a sequential extension can also be applied to other competition formats. There is already a literature of related results concerning necessary and possible winners (stemming from Konczak and Lang (2005)) upon which such an extension may be based. Another avenue for further study would be to introduce the issue

²⁵The work of this chapter is in part due to the extremely pertinent and useful remarks made by reviewers before its presentation at COMSOC 2018. In particular we should address a remark that was made concerning Formula 1: that the type of manipulation that we consider is less “important” than the phenomenon where a racer blocks the passage of cars of rival teams in order to preserve an advantage for a teammate. In our opinion this exemplifies the distinction between *manipulation* and *strategic behaviour*. In Formula 1 a single “blocking move” is sanctioned Benson (2012). Allowing blocking may be considered desirable because it adds a strategic level for the competitors and improves the spectacle for viewers; one might consider blocking to be strategic behaviour. On the other hand, the type of behaviour we consider—which although rare, also seems to occur in Formula 1 (Gpupdate, 2008)—is *clearly undesirable*—the reviewer used the term “scandalous”—and merits its designation as manipulation. Now, strategic behaviour between teammates plays an important role in many seemingly individualistic competitions, e.g. in various cycling events teammates draft behind each other. However, such strategic behaviours are typically domain specific and thus difficult to consider in the general manner that we treat manipulation; they would arguably be better investigated under the heading of operations research.

of bribery (Faliszewski et al., 2006); in particular the notion of swap bribery (Elkind et al., 2009) seems applicable. This would extend the relevance of any results beyond the case of altruistic teammates.

The question that arises out of this work concerns the decisiveness of social choice rules that prevent manipulation. It seems obvious that extremely indecisive rules are unsatisfactory. This is implicit, for example, in the impossibility theorem of Wilson (1972), which shows that the conditions of Arrow's theorem *without* the Pareto criterion results in either a dictatorship, and anti-dictatorship or the null rule which always returns complete indifference. In this case it is clear that complete indifference is an undesirable outcome, so the impossibility is effectively recreated with a different set of properties: any level of decisiveness at all leads to dictatorship. However, for rules that are *only sometimes* indecisive, it is reasonable to have recourse to an external tiebreaker if the candidates have performed symmetrically. We have a possibility result that requires an intuitively indecisive method, and, on the other hand, we know that requiring *resoluteness* leads to an impossibility. But it is not clear how to describe the situation in between these two options. Indeed, at this point, we are not even sure how to go about measuring the decisiveness of a rule.

Chapter 2

Relieving the tension between neutrality and resoluteness using parallel universes

This chapter introduces a novel family of rules whose purpose is to be more decisive than the rules on which they are based, while retaining the property of neutrality.¹ The aim is to push the limits of the trade-off between decisiveness on the one hand and neutrality (and implicitly anonymity) on the other. Our work is largely inspired by Conitzer et al. (2009), who consider this trade-off for instant runoff voting, and Brill and Fischer (2012), who consider this trade-off for ranked pairs. It is based upon the notion of **parallel universes**, which was introduced in Conitzer et al. (2009).

The idea behind parallel universes is the following: multiple possible universes arise when there are multiple possible implementations of a given method, each of which is undertaken in a slightly different manner; these implementations are run in parallel and their outputs combined. Let us take instant runoff voting, applied to the profile with three alternatives and six voters displayed in Table 2.1, as an example.

1	2	3	4	5	6
<i>a</i>	<i>a</i>	<i>b</i>	<i>b</i>	<i>c</i>	<i>c</i>
<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>b</i>	<i>a</i>
<i>b</i>	<i>b</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>b</i>

Table 2.1: An example profile. Each column corresponds to the ranking of an agent, with higher alternatives preferred to lower alternatives.

Instant runoff voting successively eliminates the alternative that receives the least first place votes. However, in the above profile, there is a three-way tie: *a*, *b* and *c* are all ranked first two times. In one universe, *a* may be eliminated, whereas in others *b* or *c* may be eliminated. If we suppose that alternative *c* is eliminated, then we again end up with a tie between *a* and *b*: again, there are different possible universes, in one of which *a* is selected, in another *b* is selected.

What, exactly, does a universe correspond to? The question is open to interpretation and will depend upon the particular method in question. For instant runoff voting, one possible answer is that a universe corresponds to a linear order tiebreaker. Different linear order

¹It is based upon work by Kruger and Airiau (2017).

tiebreakers can sometimes produce different outcomes, and thus correspond to different universes. The tree of possibilities for each linear order tiebreaker is given in Figure 2.1.

Considered individually, each possible universe produces a single outcome, thus instant runoff voting within a given universe is **resolute**. However, we can see that the profile is symmetric between a and b : **neutrality**, if imposed, would require that if a is chosen, then b must be chosen as well. Given a specification of what universes are possible, rather than supposing that we live in one fixed universe, we can return all alternatives that win in any possible universe. If we have a suitable definition of which universes are possible, this will even make the rule neutral. Taking the *union* of all possible universes in this way is a general method for making many resolute, non-neutral rules neutral, at the price of losing resoluteness. We call this the parallel universe concept.

Another possibility is to consider possible universes in a probabilistic manner. Applied to the tree in Figure 2.1, the idea would be that we choose which branch to take depending on some random process, flipping a coin for example. For instant runoff voting, if each possible branch is chosen with equal probability, this amounts to determining the probability based upon how many linear orders tiebreakers a given alternative is selected. We thus have a method that is based upon *counting* linear orders. In such a manner we get both neutrality and (ex post) resoluteness, but randomness is perhaps too high a price for many practical situations, in particular where *legitimacy* is supposed to be established through the choice procedure. Agents may find it difficult to accept the result of a random decision, no matter how fair it is ex ante.

A way of derandomising, while still counting possible universes, is to select those alternatives that win in *most* universes. To return to our above example, we saw that a was only selected in one universe, whereas it can be seen (Figure 2.1) that c wins in four universes. Indeed, using the three alternative profile as a template, examples can be constructed such that some alternative wins for $m! - (m - 1)!$ universes, whereas another alternative only wins for one universe, under instant runoff voting. It seems a little unreasonable to treat the alternative that wins in only $1/m!$ of the universes the same as the alternative that wins in $(m - 1)/m$ of the universes.

We have seen that the general principles of taking parallel universes and, further, argmax alternatives can be applied to rules that are normally considered in a non-neutral and resolute form, specifically instant runoff voting and ranked pairs. This general principle is, however, also applicable to slightly different types of choice methods. In particular, we focus in this chapter upon **tournament solutions**. There are two reasons for our choice. First, tournament solutions can be described, in a slightly imprecise manner, as **indecisive**, that is, they tend to output large sets of alternatives. Second, a large class of tournament solutions are implementable by binary trees (Banks, 1985; Laslier, 1997; Horan, 2013). That is, there are

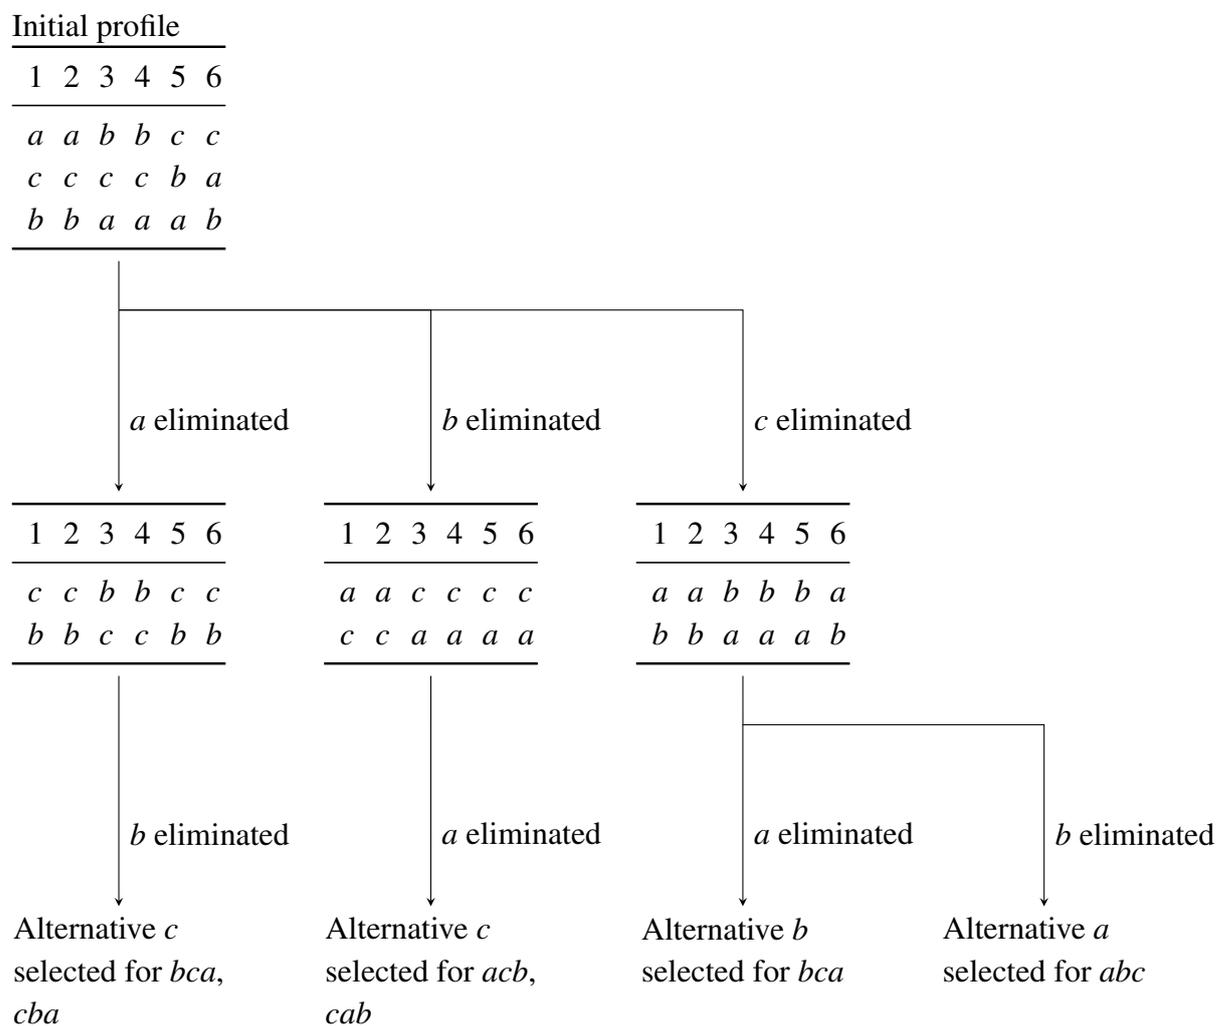


Figure 2.1: The possible instant runoff outcomes for a profile under different tiebreakers, where the rightmost alternative in the tiebreaker amongst those alternatives with fewest first place votes is eliminated in the first round. In the first round there is a three-way tie between *a*, *b* and *c*, leading to three different possibilities depending upon which of these alternatives is eliminated. Two of these possibilities lead to *c* being selected, whereas the third is subject to a further tie between *a* and *b*. This means that *c* is the winner for four linear order tiebreakers, whereas *a* and *b* each win for only one linear order tiebreaker.

natural ways of seeing tournament solutions as parallel extensions of a resolute non-neutral rule. Such binary trees provide a natural object on which to apply permutations, which can be identified with the linear order tiebreakers of different universes.

For the rest of this chapter, the input to the social choice function is not a profile of preferences, but rather a **tournament**. Agents are not explicitly present in a tournament, but their existence can be taken as implicit: it is well known that for any tournament there is a corre-

sponding preference profile such that the **majority relation** coincides with this tournament (McGarvey, 1953). In this chapter we consider functions of type $f: \mathcal{T}(A) \rightarrow 2^A \setminus \emptyset$ where $\mathcal{T}(A)$ is the set of tournaments over A . We use T to refer to arbitrary tournaments. For some properties it will be necessary to have variable electorates, in such a case f will in fact refer to a family of functions, each applicable to tournaments over different sets of alternatives; we distinguish such families from individual functions by referring to them as tournament solutions.

We draw tournaments by placing an arrow between two alternatives. We impose the convention that if there is no arrow, we assume it goes to the right. Thus the two diagrams in Figure 2.2 display the same tournament. In fact, this tournament is the smallest tournament that is non-trivial in the sense that it is neither completely symmetric nor transitive. We will refer to this particular tournament with the name T^4 .

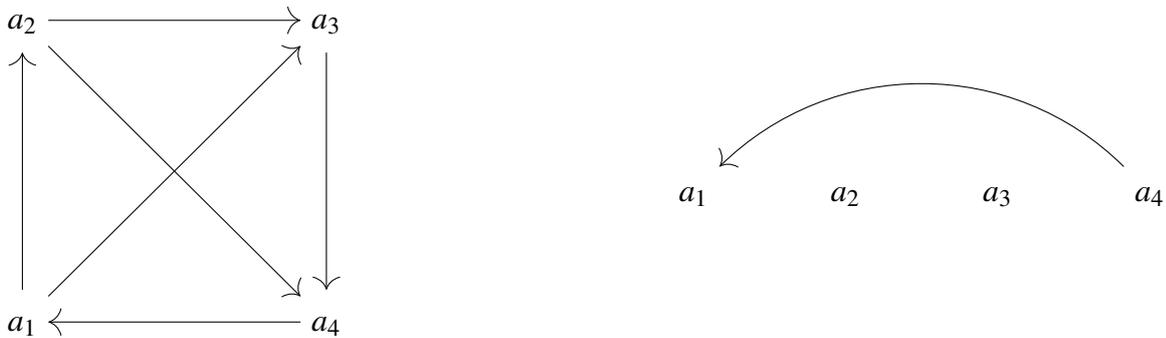


Figure 2.2: Two representations of the tournament T^4 . An arrow from x to y represents xTy . When no arrow is drawn between two alternatives x and y , and x is to the left of y , then xTy .

The methods that we use to select alternatives from a tournament are based upon **binary trees**. A binary tree is a structure containing nodes, each with two children who are either themselves nodes or leaves of the tree. An example of a graphical representation of a tree is drawn in Figure 2.3: note that we only write the subscript label of each alternative.

One manner of reducing a tree to produce a smaller tree is to compare two leaf siblings according to the tournament relation, and to replace their parent with a leaf corresponding to whichever of the two alternatives “defeats” the other. This process can be repeated until only one alternative remains. Some of the steps in the reduction process of the above tree using the tournament T^4 are displayed in Figure 2.4.

The structure of the rest of this chapter is as follows. In Section 2.1 we provide definitions for the rules we consider in this chapter. All these rules are based on binary trees, but more general rules will require families of binary trees, of which we will define some particular

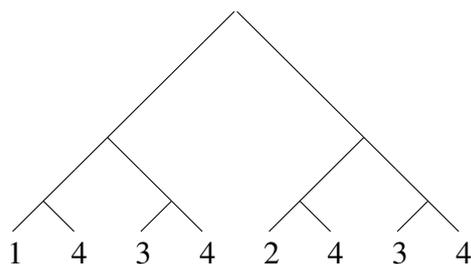


Figure 2.3: An example binary tree. (Specifically, the Banks tree over four alternatives, see Section 2.1.1).

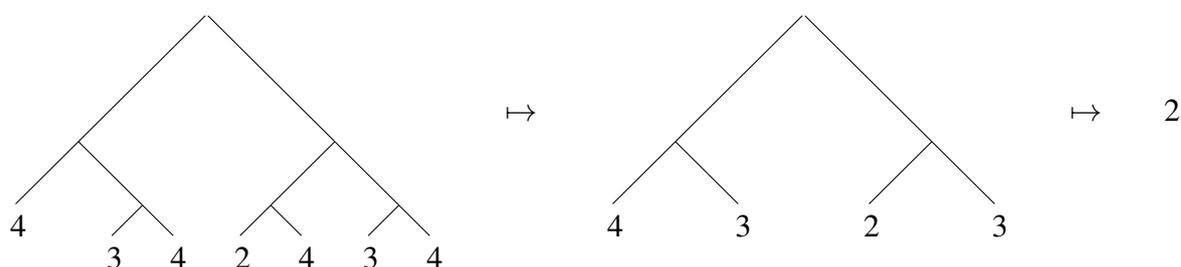


Figure 2.4: Some steps in the reduction of the tree in Figure 2.3 by the tournament T^4 . If a leaf dominates its sibling, this leaf replaces the siblings' parent.

instances by recursion. We then show how to use single trees or families of trees to define tournament solutions. In Section 2.2 we provide theoretical results for the defined tournament solutions: we identify if these are identical to other rules and determine their axiomatic properties. In Section 2.3 we assess how successful our novel rules are at providing more decisive choice than the rules upon which they are based. We assess our results on trees in Section 2.4.

2.1 Binary trees

The specific choice functions that we consider in this chapter are based on binary trees whose leaves are alternatives in A . We represent the set of such trees by the context-free grammar generated by the rules below (in Backus-Naur form).

$$\tau ::= a \in A \quad | \quad (\tau\tau)$$

Parentheses delimit subtrees; directly adjacent elements within parentheses are sibling leaves. We use τ_A for the set of all such trees with leaves from A . We often omit outermost and leftmost parentheses, i.e. we assume left associativity. We use $\tau \in \tau_A$ to refer to an arbitrary

tree. If every element in A occurs as a leaf in τ , we say that τ is complete. If some element in A occurs as two or more leaves in τ , we say that τ is repetitive.

2.1.1 Families of binary trees

As well as individual trees, we consider *families* of binary trees for any set of alternatives A . To define these families, we define a function g that takes a list of alternatives as input and returns a tree whose leaves are alternatives in the list. Technically, for all A with $|A| \geq 2$ we define functions $g_A: \mathcal{L}(A) \rightarrow \tau_A$, where $\mathcal{L}(A)$ is the set of linear orders, each interpreted as a list of alternatives. The function g is defined as, for $R \in \mathcal{L}(A)$, $g(R) = g_A(R)$.

We define five families of trees using the following five functions. These five functions apply to linear orders over alternatives: for trees we only write the subscript label of the alternative; we write i instead of a_i . These functions are recursive in that they form subtrees from functions applied to shorter lists. To terminate the recursion, for a list x we set $g(x)$ to be the tree with the single leaf x .

Simple trees: $\text{st}(12\dots m) = 1(\text{st}(2\dots m))$.

We follow Laslier's (1997) nomenclature. Such trees are sometimes called voting caterpillars (Fischer et al., 2009).

Banks trees: $\text{bn}(12\dots m) = \text{bn}(13\dots m)(\text{bn}(23\dots m))$.

These trees are named after Banks (1985). Note we do not place parentheses around the left subtree because we assume left associativity.

Two-leaf trees: $\text{tt}(12\dots m) = 12$.

These trees are incomplete for $|A| > 2$. This is the only family we define that contains incomplete trees.

Balanced tree: $\text{ft}(12\dots m) = \text{ft}(1\dots \lfloor \frac{m}{2} \rfloor) (\text{ft}(\lceil \frac{m}{2} \rceil \dots m))$.

These trees have minimal height over the space of complete trees, where height is given by the maximal nesting of parentheses. Formally, the height of a tree with a single leaf is zero; the height of tree $(\tau\tau')$ is one plus the maximum of the heights of the subtrees τ and τ' . Lang et al. (2007) define balanced voting trees with a slightly less restrictive definition.

Iterative Condorcet trees: $\text{ic}(12\dots m) = 12\dots m\text{ic}(2\dots m)$.

The above definition is inspired by the work of Altman and Kleinberg (2010). Recall that we assume left associativity; an iterative Condorcet tree contains no parentheses when represented as strings.

Figure 2.5 contains graphical representations of these trees for four alternatives.

2.1.2 Using binary trees to define social choice functions

We define choice functions from a binary tree in three ways. The first involves competing siblings against each other until one alternative remains, and is thus a resolute tournament function. The second considers each possible permutations of the alternatives as a different parallel universe, and returns all possible winners in any parallel universe. The third involves counting universes, and returns those alternatives which win in most universes. Formal definitions for each follow.

First, alternatives compete against their siblings according to the majority relation, the winners moving up to the parent node until we have a final winner at the root. Formally, for a tree $\tau \in \tau_A$, the function $\tau: \mathcal{T}(A)^N \rightarrow A$ has the following recursive definition:

$$xy(T) = \begin{cases} x & \text{if } xTy \\ y & \text{otherwise} \end{cases} \quad \tau\tau'(T) = \tau(T)\tau'(T)(T).$$

Effectively, the tournament T defines a binary algebraic operation, a tree describes an algebraic term, and the result of applying the binary operation throughout the term is the reduced form of the tree.

Second, we define a tournament function by take the union of our first function for all possible permutations. Following Conitzer et al. (2009), we use the terminology “universe” to refer to each possible permutation. Thus this method forms the parallel universe solution concept. We indicate this with a superscript PU; formally, for a family of trees g with $g_A: \mathcal{L}(A) \rightarrow \tau_A$, we create a function of type $g_A^{\text{PU}}: \mathcal{T}(A)^N \rightarrow 2^A$. Define $g_A^{\text{PU}}(T) = \{g_A(R)(T) : R \in \mathcal{L}(A)\}$. The outcomes of many of these parallel universe versions are well studied—the question of what tournament rules are implementable by parallel universes in this manner has been studied by Horan (2013), who gives necessary and sufficient conditions. Not that a single family of trees actually defines a tournament function for any set A : this means that we will be able to consider properties which require a variable agenda.

Third, we define a tournament function as follows: count the number of permutations that result in each winner; instead of taking the union over all universes, only return those alternatives that win in *most* universes. We refer to this as the *argmax* solution concept, and use AM as the corresponding superscript. For $g_A: \mathcal{L}(A) \rightarrow \tau_A$, we define the function of type $g_A^{\text{AM}}: \mathcal{T}(A)^N \rightarrow 2^A$ by

$$g_A^{\text{AM}}(T) = \operatorname{argmax}_{x \in A} | \{R \in \mathcal{L}(A) : g_A(R)(T) = x\} |.$$

This family of rules appears to be novel. As in the parallel universe case, a given family of trees actually defines a rule over variable agendas.

2.2 Identities and properties of tree based rules

This section is mostly composed of two kinds of results: the first show that the rules we have defined output certain known sets; the second describe axiomatic properties that these rules satisfy. Some of the rules coincide with one or other of the following five tournament solution concepts.

The set of Condorcet non-losers.

This contains all $a \in A$ such that there exists a $b \in A$ such that aTb .

The top cycle.

This is the smallest set $B \subseteq A$ such that bTa for all $b \in B$ and $a \in A \setminus B$.

Given a tournament T , we say $a \in A$ covers $b \in A$ if aTb and for all $x \in A$ such that bTx , aTx .

The uncovered set.

This contains all and only alternatives that are not covered by any other alternative.

The Banks set.

This contains maximal alternatives (by T) of maximal (by inclusion) subsets of A that are transitive according to T .

The Copeland set.

This contains alternatives that dominate the most other alternatives, i.e. it is the set $\operatorname{argmax}_{a \in A} |\{x \in A : aTx\}|$.

The first three concepts are increasing in terms of decisiveness: for any tournament, each always outputs a subset of the former, and sometimes a proper subset. So the last two concepts are the most decisive of the five, however it is not obvious how to compare them to each other, because they may produce completely disjoint solution sets. So although it appears, for example, that the Copeland set tends to be smaller, by cardinality, than the Banks set, even a formal statement to this effect could not necessarily be taken as meaning that Copeland is more decisive. We will discuss more why this is the case in the following chapter.

We now list five axiomatic properties, two of which are close variants. The first three apply to an arbitrary tournament function $f: \mathcal{T}(A) \rightarrow 2^A \setminus \{\emptyset\}$. The last two require the possibility of a variable agenda, they thus apply to families of tournament functions for all possible sets of alternatives, or tournament solutions. To express these we require some additional notation: $T^{(a,b)}$ is T with the relation between a and b inverted. For the last two properties we will require the notation of a summary of a tournament. Let $T' \in \mathcal{T}(\{a_1, \dots, a_m\})$, and, for each $i \in \{1, \dots, m\}$, $T^i \in \mathcal{T}(A_i)$ for pairwise disjoint sets of alternatives A_i . We write $T = (T'; T^1, \dots, T^m)$ to refer to the tournament over $\bigcup A_i$ defined as follows: for each $a \in A_i$ and $b \in A_j$, if $i = j$ then (aTb iff $aT^i b$); whereas if $i \neq j$ then (aTb iff $aT' b$). The final

property also requires the idea of the restriction of a binary relation to a subset: for $B \subseteq A$ we write $T \upharpoonright_B = \{(a, b) \in T : a, b \in B\}$. The idea behind our first condition is that reinforcing a winning alternative should not make it lose.

Monotonicity: if $a \in f(T)$ and bTa , then $a \in f(T^{(a,b)})$.

The idea behind the second condition is that a clearly best alternative should be selected.

Condorcet consistency: if aTx for all $x \neq a$, then $f(T) = \{a\}$.

We now present two conditions: the first of these can be applied to tournament functions (and thus also to tournament solutions); the second is a strengthening of the first that can only be applied to tournament solutions. The idea behind both these conditions is that making restricted changes to tournaments should not completely change the output solution (Laffond et al., 1996).

Weak composition consistency: take any $X \subseteq A$ such that for any $x, x' \in X$, $y \in A \setminus X$ we have xTa if and only if $x'Ta$. Then for any $\{a, b\} \subseteq X$ it is the case that

- 1) $f(T) \cap (A \setminus X) = f(T^{(a,b)}) \cap (A \setminus X)$ and
- 2) $f(T) \cap X \neq \emptyset$ if and only if $f(T^{(a,b)}) \cap X \neq \emptyset$.

Composition consistency: for any $T = (T'; T^1, \dots, T^m)$, index i and alternative $a \in A_i$, it is the case that $a \in f(T)$ iff $a_i \in f(T')$ and $a \in f(T^i)$.

The final condition, like composition consistency, can only be applied to tournament solutions. Indeed, it explicitly concerns variable agendas; it explicitly places restrictions on the output when the set of alternatives expands or contracts. Our definition follows Brandt et al. (2016), p61.

Stability: for all tournaments T and all nonempty $X, Y, Z \subseteq A$ with $Z \subseteq X \cap Y$,

$$Z = f(T \upharpoonright_X) = f(T \upharpoonright_Y) \text{ if and only if } Z = f(T \upharpoonright_{X \cup Y}).$$

All of the tournament solution concepts and properties defined above are treated in more detail by Laslier (1997) and Brandt et al. (2016). We now consider, in turn, the three different types of tree rules with respect to these rules and properties, starting with singleton trees.

2.2.1 Singleton trees

Singleton trees are technically not identifiable with tournament functions, because they are in fact of a different type, outputting a single alternative instead of a set. Of course, it is natural to consider the output of a singleton tree as a singleton, thereby allowing for a comparison with tournament functions. However, other tournament functions are typically neutral and thus irresolute. Thus, tournament functions that are not defined in terms of trees typically

do not coincide with singleton tree functions. Indeed, the most explanatory description of a singleton tree function is just the tree itself.

With respect to properties, a singleton tree can only be applied to tournaments defined on a superset of its leaves, and indeed are best interpreted as applying to a fixed agenda. Thus it makes little sense to apply properties that require a variable candidature to a function defined by a singleton tree. The following general properties are easily verified.

Observation 2.1. Rules defined by individual trees are resolute and non-neutral. Such a rule satisfies Condorcet consistency iff it is complete. If a tree is non-repetitive, its rule is monotonic.

The main property of interest for this subsection is monotonicity. Non-repetition is not a necessary condition for monotonicity; the Banks tree provides an example of a monotonic but also repetitive tree.

Proposition 2.2. *For any number of alternatives, bn is monotonic.*

Proof sketch. This result is well known (Banks, 1985), so we only provide a high level description of the proof. Consider the following intuitive definition of the Banks tree. First, we suppose that the rightmost alternative in the Banks tree is the preliminary winner. We then successively examine the other alternatives (as arguments in the recursive definition from right to left), potentially setting them as new preliminary winners. For an alternative to become the new preliminary winner, it must defeat every member of the set of previous preliminary winners. After all the alternatives have been tested, we select the current preliminary winner. Clearly, if an alternative was selected then it defeated all previous preliminary winners: changing only this alternative so that it defeats more alternatives will not change the fact that it is selected. \square

2.2.2 Parallel universe trees

The first results of this subsection concern the identities of parallel universe rules based on families of trees.

Proposition 2.3. *We have the following equivalences between parallel universe rules and known rules:*

- 1) st^{PU} returns the top cycle.
- 2) bn^{PU} returns the Banks set (Banks, 1985).
- 3) tt^{PU} returns the set of Condorcet non-losers.
- 4) The tournament solution ic^{PU} is identical to the top cycle.

Proof. Only (4) deserves comment, the other results are well known or obvious. The binary tree ic implements the iterative Condorcet rule, described by Altman and Kleinberg (2010), into a binary tree rule. This rule successively removes alternatives from the tournament, according to some given linear order, until the contracted tournament has a Condorcet winner; this Condorcet winner is then selected. If we parallelize over all possible linear orders, this produces every element in the top cycle: for an arbitrary element in the top cycle, consider successively removing the elements following the cycle starting with the alternative dominated by the arbitrary element. It is similarly easy to see that no element *not* in the top cycle can be produced by this procedure.

To see that the iterative Condorcet tree $ic(1 \dots m)$ implements this, note that after the first $m - 1$ comparisons if m is still a possible winner it will have defeated every other alternative, if not it has been eliminated as it does not reappear higher in the tree; in the latter case the process then continues by comparing $m - 1$ against all alternatives $< m - 1$, etc. Figure 2.6 contains an example. □

Most of the above results can be found in the literature on social choice theory. Only the parallel universe version of the iterative Condorcet tree is novel, although it does not define a novel rule, as it coincides with the parallel universe version of the simple tree, i.e. the top

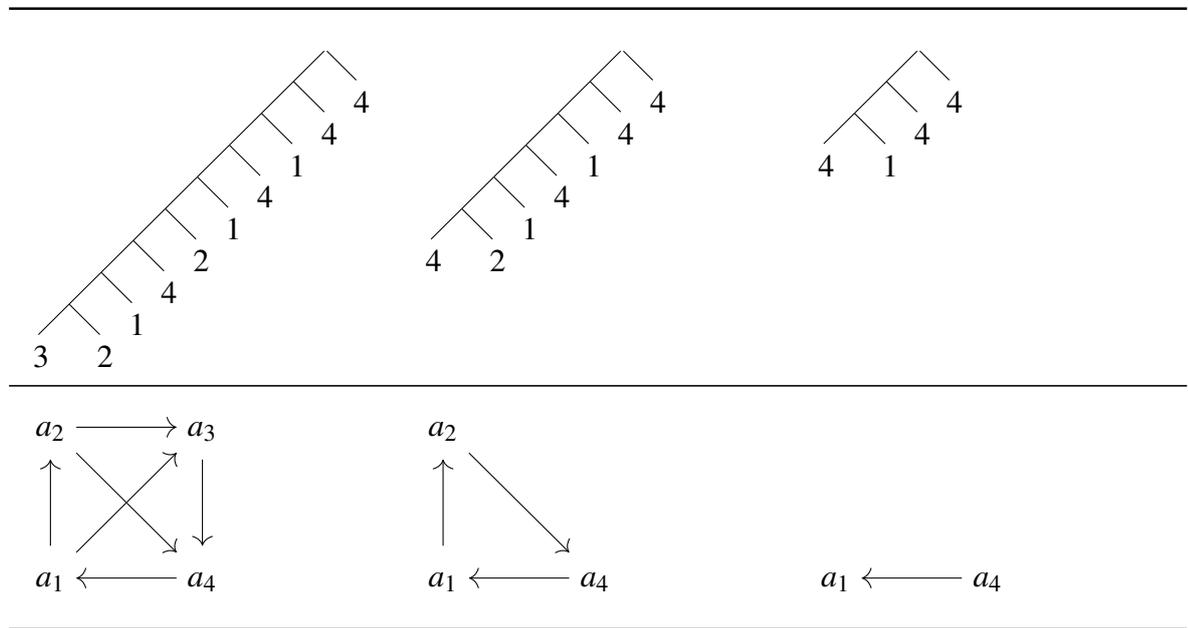


Figure 2.6: At top: partial reductions of the tree $ic(3214) = 32142141444$. Below that: the tournament T^4 and the tournaments created by the removal of alternative a_3 and then alternative a_2 , at which point there is a Condorcet winner. This illustrates the equivalence of the iterative removal of alternatives until there is a Condorcet winner with the tree ic .

cycle.

We have not, however, considered balanced trees yet. The set of parallel universe winners for the balanced tree can be a proper subset of the Top Cycle. In particular, to win in a universe an alternative must dominate at least $\lfloor \log_2 m \rfloor$ alternatives. Thus this solution concept does not always produce a superset of the uncovered set, as for arbitrarily many alternatives we can construct a tournament with an uncovered alternative that only defeats one other alternative.

Nor need the parallel universe version of the balanced tree be a subset of the uncovered set: consider the tournament in Figure 2.7 to the right. Here a_2 is covered by a_1 , but a_2 is also the output of a balanced tree, specifically $154(32)(T) = 2$.

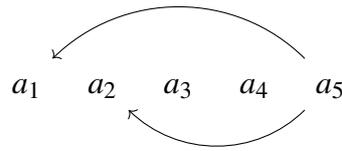


Figure 2.7: A tournament for which a covered alternative, a_2 , is the output of a balanced tree.

This means that solution concepts based on balanced trees are somewhat atypical: most other tournament solutions that have been studied *are* refinements of the uncovered set.

We could have given a different definition for the balanced trees. If $\log_2 |X|$ is not an integer there are multiple non-repetitive tree structures that have minimum height. Our choice of definition not only minimises the height but also minimises the difference between the number of leaves in any node's left and right subtrees, other definitions produce different functions. For example, consider the two tournaments and the two trees in Figure 2.8.

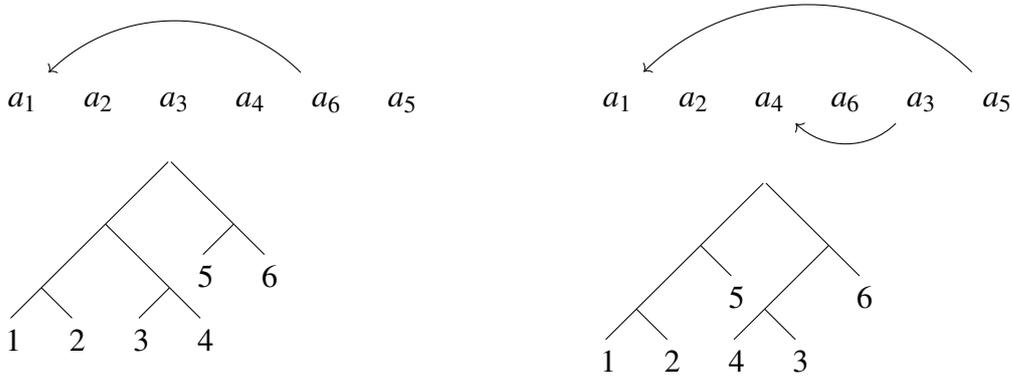


Figure 2.8: Two tournaments and two different balanced trees. The trees output different solution sets that are not contained in each other.

Both the trees are of minimal height for complete trees over six alternatives. Here a_6 is the winner for the left tree and left tournament, and also for the right tree and right tournament. However, there are no permutations of the labels of the trees which make 6 the winner for a tree and its opposite tournament. Thus the different balanced trees produce distinct sets of parallel universe winners that are not contained in each other. Ultimately, as was the case

for singleton trees, it appears that the best description of the identity of the parallel universe version of the fair tree is in terms of the tree itself.

We now consider the properties of parallel universe rules. In general, parallel universe versions of binary trees may not be monotonic. Laslier (1997) provides a large counterexample.² Figure 2.9 provides a smaller example with only 5 alternatives involving the tournament T defined as $a_5 T a_4 T a_3 T a_2 T a_1$ and $i T j$ for all other $i < j$; the tournament $T^{\langle a_4, a_5 \rangle}$; and the binary tree $\tau = 412(453)4$.

T :

$$a_1 \leftarrow a_2 \leftarrow a_3 \leftarrow a_4 \leftarrow a_5$$

$T^{\langle 3,4 \rangle}$:

$$a_1 \leftarrow a_2 \leftarrow a_3 \leftarrow a_4 \quad a_5$$

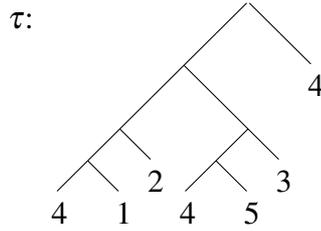


Figure 2.9: A tournament and a second version of this tournament with a_4 reinforced, and a tree for which a_4 is output according to the first tournament but for which no permutation outputs a_4 as the winner after a_4 is reinforced. This demonstrates a violation of monotonicity for parallel universe versions of binary trees.

It can be verified that $\tau(T) = 4$. In fact there are four permutations of the tree that select this alternative. Now reinforce alternative a_4 , obtaining $T^{\langle a_4, a_5 \rangle}$. We have verified by computer that there are no permutations of τ that select a_4 for the tournament $T^{\langle a_4, a_5 \rangle}$. That is to say, $a_4 \in \tau^{\text{PU}}(T)$ but $a_4 \notin \tau^{\text{PU}}(T^{\langle 3,4 \rangle})$.

Although monotonicity does not hold in general for parallel universe binary tree rules, the particular parallel universe rules we define inherit monotonicity from the singleton tree of which they are composed.

Observation 2.4. For a family of trees g_A , if for every $R \in \mathcal{L}(A)$ the singleton tree $g_A(R)$ is monotonic, then so to is g_A^{PU} .

It is also known that parallel universe binary tree rules are weakly composition consistent, consult Moulin (1986) or Laslier (1997) for proof.

The final property we consider is stability. Brandt et al. (2016) show that the top cycle is stable but that Banks is not. It is easy to find a counterexample to stability for the set of Condorcet non-losers: for example, consider any transitive tournament with at least three alternatives. Only the balanced tree remains: the parallel universe version of this rule can be seen to be non-stable using tournament T^4 of Figure 2.2.

²More precisely, a family of counterexamples.

2.2.3 Argmax applied to trees

Of the argmax versions of our five families of trees, only tt^{AM} is a known rule, namely Copeland. Note that Copeland is not implementable by parallel universe trees (Horan, 2013). The other argmax rules are novel. We have seen that $\text{st}^{\text{PU}} = \text{ic}^{\text{PU}}$; this identity does not carry across to the argmax versions.

Proposition 2.5. *The tournament solutions st^{AM} and ic^{AM} are distinct.*

Proof. The tournament T in Figure 2.10 provides the required example: $\text{st}^{\text{AM}}(T) = \{d\}$ and $\text{ic}^{\text{AM}}(T) = \{a\}$. This tournament is one of the tournaments for which we display the results in Table 2.2. The solutions were found with the aid of a computer.

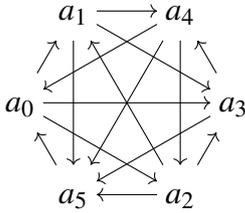


Figure 2.10: A tournament whose output shows that two families of binary trees, the simple tree st and the iterative Condorcet tree ic , although equivalent for parallel universes, are distinct for argmax. □

This suggests that we have to choose our refinement with care: in fact, there doesn't seem to be any principled reason to choose one of these versions over the other in order to obtain a refinement of the top cycle.

The big question here is whether or not the argmax versions of our rules also satisfy monotonicity. Unfortunately, it seems difficult to prove the monotonicity property here, which is a basic desirable property. The only rule for which we are sure that this holds is the argmax version of tt , which corresponds to Copeland. It would be somewhat surprising if monotonicity did not also apply to the other argmax versions of the rules.

We can show that no argmax rule is composition consistent, for $m \geq 5$.

Proposition 2.6. *There is no family of trees τ such that τ^{AM} satisfies weak composition consistency.*

Proof. Consider the tournaments $T = (T'', T^1, 2, 3)$ and $T' = (T'', T^2, 2, 3)$ such that T'' and T^1 are cyclic tournaments with 3 alternatives, T^2 is a transitive tournament, and 2 and 3 are tournaments with single alternatives. Suppose for a contradiction that there is a tree τ such that τ_{AM} is composition consistent. This must make all five alternatives winners in T , by composition consistency and neutrality. However, in the second, only one alternative in T^2 can win (by Condorcet consistency), and 2 and 3 must also win. Now, consider the number of permutations for which 2 and 3 win for the tree τ under tournaments T and T' : these must be the same. Similarly, the number of permutations for which some alternative in T^1 wins

must be the same as the number of permutations for which some alternative in T^2 wins. As in T all alternatives are winning, 2 and 3 must win 20 times. But in T' the Condorcet winner of T^2 must win 60 times, contradiction. \square

None of our argmax rules are stable either. Consider the tournament T^4 of Figure 2.2: each argmax solution concept returns at most two alternatives, but removing one of the others results in a symmetric tournament that must return all three alternatives for any neutral rule.

2.3 Success of tree-based argmax rules as refinements

Assuming that our argmax rules satisfy the basic property of monotonicity, the other issue at hand is how effective they actually are at refining the set of winners. Let us start by considering tt : for tt^{PU} , the set of Condorcet non-losers is an extremely indecisive solution concept. However, tt^{AM} corresponds to the *Copeland set*, which is much more decisive even than many other solution concepts we consider. Here moving to the argmax universe version provides a large gain in decisiveness.

The same seems to be the case for the other solution concepts. We have tested this on some example tournaments. The outcome of all of our rules is only affected by alternatives in the top cycle. Thus we restrict attention to **non-reducible tournaments**, where there is a cycle throughout the whole tournament. Moon (1968) provides a list of all non-isomorphic tournaments involving six alternatives, from which we see that there are only thirty-four such non-reducible tournaments. We applied our rules to all of these, and compared them with the (intuitively decisive) Markov solution concept. The specific results are found in Table 2.2.

From Table 2.2 it can be verified that all these solution concepts are distinct. We can also see that the Banks set contains three alternatives 14 times, four alternatives 8 times, five alternatives 9 times and six alternatives 3 times. In contrast bn^{AM} outputs a single winner 32 times, two winners 1 time and three winners 2 times. Both st^{AM} and ic^{AM} get similar (though distinct) results. Copeland, equivalent to tt^{AM} , outputs a single winner 18 times, two winners 7 times, three winners 5 times and four winners 4 times. Thus it appears that the argmax rules are significantly more decisive than the full parallel universe versions.

2.4 Assessment of the possible universes method applied to tree based tournament functions

In this chapter we have described a method for making tournament solutions more resolute using a technique of counting permutations. The particular attraction of the argmax rules is that they are more decisive than their parallel version counterparts. We summarise the status of our rules in Table 2.3, and their properties in Table 2.4. However, the real test of their attractiveness hinges upon whether or not they satisfy monotonicity, which we have not been able to prove or disprove.

There is also the issue of the computational complexity of these rules. Naive implementations of parallel or argmax rules with as many universes as linear orders run in factorial time, which is prohibitively slow. However, there is a counterpoint here: approximating argmax winners will be easy to do using Monte-Carlo methods. This echoes the fact that, although determining the set of Banks winners is NP-hard (Woeginger, 2003), a simple greedy algorithm can find *some* Banks winner in polynomial time (Hudry, 2004). It would also be interesting to consider reducing the amount of possible universes. For example, instead of considering all linear order tiebreakers, consider only “cycles” of tiebreakers, e.g. the set $\{12\dots m, 23\dots m1, \dots, m1\dots(m-1)\}$. Although we are no longer certain to retain neutrality (although there may be special cases where this occurs), considering such universes corresponds to a natural weakening of neutrality. The selection according to both parallel universes and argmax over *restricted* universes will be a subset, and likely a proper subsets, of the selection according to the parallel universe rule over all universes. This provides a potentially tractable method for making rules more decisive.

Although decisiveness seems to be an attractive property, it is not the sole aim of a choice function. If there is little reason to choose one alternative over another, then perhaps they should be tied. In particular, we have seen methods above that are designed to make particular rules more decisive. However, they do not do so in a consistent manner. Indeed, it is hard to fully justify the foundations behind the selection process that refines the particular solutions. Ultimately, this comes back into the question of information, and how it is used to distinguish alternatives.

Table 2.3: Summary of identities of rules based on trees.

Family of trees	Parallel universe	Argmax
Simple tree	Top cycle (Laslier, 1997)	new
Banks tree	Banks (Banks, 1985; Laslier, 1997)	new
Fair tree	Cup rule	new
Two-leaf tree	Condorcet non-losers	Copeland
Iterative Condorcet tree	Top cycle (Altman and Kleinberg, 2010)	new

Table 2.4: Known properties for rules based on trees.

Family of trees	Monotonic		Condorcet consistent	
	PU	AM	PU	AM
Simple tree	✓	?	✓	✓
Iterative Condorcet tree	✓	?	✓	✓
Banks tree	✓	?	✓	✓
Balanced tree	✓	?	✓	✓
Two-leaf tree	✓	✓	✗	✓
Family of trees	Composition consistent		Stable	
	PU	AM	PU	AM
Simple tree	weakly	✗	✓	✗
Iterative Condorcet tree	weakly	✗	✓	✗
Banks tree	✓	✗	✗	✗
Balanced tree	weakly	✗	✗	✗
Two-leaf tree	weakly	✗	✗	✗

Chapter 3

Maximally decisive social choice

In this chapter we investigate the notion of **decisiveness**¹ formally. What do we mean by the concept of decisiveness, and why are we interested in it? Decisiveness provides the “choice” in social choice—after all, the ultimate purpose of social choice is to actually pick something from some options. The more decisive a social choice function, the more effective it is at guiding our actions and decisions. Thus we want our rules to select small sets; a study of decisiveness may tell us just how small we can make our selections.

It may be wondered why we should introduce this concept, given that **resoluteness** has been well studied in social choice theory, going back at least to Moulin (1983). Our response is that resoluteness is *too strong* a property. Moulin’s work shows that resolute rules that satisfy anonymity and neutrality only exist if the number of alternatives, m , cannot be written as the sum of non-trivial divisors of the number of individuals, n . This means that for many sets of alternatives and agents resoluteness is an unattainable goal. Even in cases where it is possible, it may not be desirable: consider the general example of a profile with $n = 2^k$ and $m = 3$ in Table 3.1.

Table 3.1: A profile for which the cardinality of alternatives and agents makes it possible to have a resolute social choice function, but for which the resolute outcome is undesirable.

1 to $2^{k-1} - 1$	2^{k-1} to $2^k - 2$	$2^k - 1$	2^k
a	b	c	c
b	a	a	b
c	c	b	a

We know that at least one resolute anonymous and neutral social welfare choice function exists for these numbers, however such a rule must select c in the above profile, which is not

¹Our choice of terminology deserves some justification, especially given that the term “decisiveness” is often used with a different meaning in proofs of Arrow’s theorem. Fishburn (1977) uses the term “discriminability”, but this rather seems to suggest that a social choice function tells a lot of alternatives apart: selecting two out of four alternatives actually discriminates between four distinct pairs of alternatives, while selecting one out of four alternatives only discriminates between three distinct pairs. We consider this point again in Chapter 5. Another choice would have been “selectiveness”, as used by Sato (2009).

a desirable outcome.² In this profile, if we select a , then anonymity and neutrality force us to also select b . The profile is *intrinsically tied*.³ Nonetheless, we can still apply an effective notion of decisiveness here: given anonymity and neutrality, there are only three possible outcomes: $\{a, b, c\}, \{a, b\}, \{c\}$. The idea is to rule out $\{a, b, c\}$ because it is possible to select a proper subset of this outcome.

Anonymity forces an equivalence class on the space of profiles: two profiles for which we permute the agents must output the same outcome. Although neutrality can be considered as forming another equivalence class, as is done for the impartial, anonymous and neutral model described by Egecioğlu and Giritligil (2013), we rather consider neutrality in a slightly different manner, as a group action on the space of profiles. Alternatives in the same orbit are considered indistinguishable, as per Bubboloni and Gori (2014).

However, the requirement of maximal decisiveness given anonymous information is again too strong. Suppose we add another alternative to the profile in Table 3.1 and create the profile of Table 3.2.

Table 3.2: Profile with two alternatives that it is difficult to choose between.

1 to $2^{k-1} - 1$	2^{k-1} to $2^k - 2$	$2^k - 1$	2^k
a	b	c	d
b	a	a	b
c	c	b	a
d	d	d	c

Now a and b are distinguishable. However, it does not seem that the addition of d has made any real change to what the output should be.⁴

In what follows neutrality is a central property for the definition of what it means to be maximally decisive. Anonymity is less central: we construe anonymity as only one of many

²Bubboloni and Gori (2014) address this issue by also requiring that the rules satisfy a majoritarian property. This will further reduce the values of n and m for which there are suitable rules.

³In practice, such a tied profile will probably ultimately use an external tiebreaker to choose a single alternative in the case of ties. For example, from a tied set, one may (i) select the alternative ranked highest by some fixed strict linear order, or (ii) select the alternative ranked highest by some specific agent. Randomisation is also possible—one may (iii) randomly select an alternative in the output set, or (iv) randomly select an agent and use their ordering to break the tie. However, such tiebreakers are here considered as secondary processes, *external* to the social choice rule itself. While such tiebreakers are probably necessary for practical implementations of an irresolute choice functions, they are not part of the definition of a choice function proper, and do not affect decisiveness.

⁴This also connects with ideas of local independence conditions. We consider such properties in Chapter 4.

types of **informational restrictions**. It is the combination of neutrality with an informational restriction that provides our notion of maximal decisiveness. There are a number of other informational bases that we will consider. Most of these will imply anonymity, in such cases we can consider functions that apply to profiles with varying amounts of agents, i.e. variable electorate social choice functions. We write \mathcal{N} for the set containing each finite and non-empty initial segment of \mathbb{N} , i.e. $\mathcal{N} = \{\{0\}, \{0, 1\}, \{0, 1, 2\}, \{0, 1, 2, 3\}, \dots\}$. In the following we use \mathcal{D} to refer to the domain in a general fashion, i.e. $\mathcal{D} \in \{\mathcal{L}^*(A)^N, \mathcal{L}^*(A)^{\mathcal{N}}\}$. In this chapter, our aim is to formalise decisiveness within the simple ordinal framework of social choice theory. We thus here consider functions of the form and consider functions of type $f: \mathcal{D}^{\mathcal{N}} \rightarrow A$. Note if $\mathbf{P} \in \mathcal{L}^*(A)^{\mathcal{N}}$ then $\mathbf{P} \in \mathcal{L}^*(A)^N$ for some N . This framework is in the Arrovian tradition: there are no interpersonal comparisons, nor cardinal values of utility.⁵

The rest of this chapter is divided into three sections: in the first of these, we formalise the idea of information bases, and give various examples. This will allow us, in the section after, to give a better definition of maximal decisiveness, which takes into account what information is considered relevant. The final section gives some closing remarks.

3.1 The informational content of families of functions

By an **(informational) basis** we mean a partition of the space of profiles into equivalence classes.

Definition 3.1. A **basis** is the equivalence relation determined by some symmetric and reflexive relation $\sim \subset \mathcal{D} \times \mathcal{D}$.

Definition 3.2. We say that a social choice function f is **expressible** by a basis \sim if for all $\mathbf{P}, \mathbf{P}' \in \mathcal{D}$, $\mathbf{P} \sim \mathbf{P}'$ implies that $f(\mathbf{P}) = f(\mathbf{P}')$.

A basis implicitly describes an environment of information. In order to compute the output of a rule that is expressible by this basis, only the relevant information concerning this environment is required. This is less information than that given by the environment involving complete profiles—unless the equivalence class of basis only contains singletons, a somewhat trivial case. Given that *any* equivalence relation defines a basis, there may not always be a nice interpretation of the environment of information.

The definition of informational bases may be considered quite a dry, abstract notion that it is hard to get an intuitive grasp upon. Sometimes, however, there *is* a nice interpretation

⁵It is in fact a slightly restricted version of the full ordinal framework: we allow ties in the output but not in the input. Allowing for ties only in the output provides the intricacy necessary to investigate decisiveness, which is already enough to raise non-trivial problems.

of what is going on behind a basis. This can typically be given in terms of a function that reduces profiles into a *compact representation*, e.g. a tournament. Under this interpretation, two profiles contain the same information if they have the same compact representation. Such functions in effect produce bases: profiles that map to the same object belong to the same equivalence class.

3.1.1 Examples of informational bases

Two well-known environments of restricted information go under the names of C1 information and C2 information. These concepts were first defined by Fishburn (1977) and are captured by the following two functions. Let $\text{pairs} = \{(x, y) : x, y \in A\}$. Define the majority relation of a profile

$$T : \mathcal{L}^*(A)^{\mathcal{N}} \rightarrow 2^{\text{pairs}} \quad \text{by} \quad (x, y) \in T(\mathbf{P}) \text{ iff } |\{i \in N : xP_i y\}| \geq |\{i \in N : yP_i x\}|.$$

Define the weighted tournament of a profile

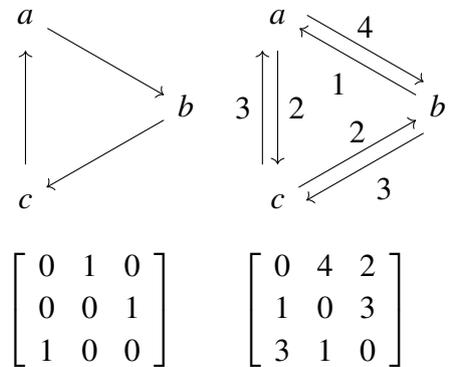
$$W : \mathcal{L}^*(A)^{\mathcal{N}} \rightarrow \mathbb{N}^{\text{pairs}} \quad \text{by} \quad W(\mathbf{P})(x, y) = |\{i \in N : xP_i y\}|.$$

The equivalence relation \sim_{C1} is defined by $\mathbf{P} \sim_{C1} \mathbf{P}'$ iff $T(\mathbf{P}) = T(\mathbf{P}')$. Similarly, the equivalence relation \sim_{C2} is defined by $\mathbf{P} \sim_{C2} \mathbf{P}'$ iff $W(\mathbf{P}) = W(\mathbf{P}')$. In general, we will be slightly loose about the distinction between an information environment, a basis, and a function that defines a basis; the context should make it clear what is meant. For instance, occasionally we will use C1 to refer to the basis \sim_{C1} .

Example 3.3. Consider the profile \mathbf{P}^1 involving three alternatives and five agents, two with preferences abc , two with preferences cab , and one with preferences bca . The tabular presentation, tournament, and weighted tournament of this profile are displayed in Figure 3.1 below.

1	2	3	4	5
a	a	c	c	b
b	b	a	a	c
c	c	b	b	a

Figure 3.1: Clockwise from above: the tabular representation of a profile, its tournament, its weighted tournament, the matrix corresponding to its weighted tournament, and the matrix corresponding to its tournament.



C2 is more expressive than C1; any rule that can be calculated using C1 information can also be calculated by C2 information. The formal property that corresponds to this is the fact that (the basis) C2 is a **refinement** of C1: for any pair of profiles $\mathbf{P}, \mathbf{P}' \in \mathcal{L}^*(A)^{\mathcal{N}}$ such that $\mathbf{P} \sim_{C2} \mathbf{P}'$ we also have $\mathbf{P} \sim_{C1} \mathbf{P}'$. In the other direction, there are rules that can be calculated using C2 information that cannot be calculated using C1 information. In general, the relative expressivity of two bases is simply a question of refinement.

Observation 3.4. A basis \sim can express all rules expressed by a basis \sim' if and only if \sim refines \sim' .

Functions expressible by C1 include the top cycle, the uncovered set, and the Banks set. Functions expressible by C2 include all those expressible by C1, Kemeny (defined as a union choice rule), ranked pairs, and Borda.

We know the range of the function T . A theorem by McGarvey (1953) shows that for any possible tournament, there is a profile that maps to it.

Theorem 3.5 (McGarvey). *Take arbitrary A . For every complete and asymmetric relation $S \subseteq \text{pairs}$, there is some profile $\mathbf{P} \in \mathcal{L}^*(A)^{\mathcal{N}}$ such that $T(\mathbf{P}) = S$.*

The range of W is difficult to determine. However, a slightly modified function can be defined that *almost* contains the same amount of information: define the symmetric weighted tournament of a profile

$$\text{sym}W: \mathcal{L}^*(A)^{\mathcal{N}} \rightarrow \mathbb{Z}^{\text{pairs}} \quad \text{by} \quad \text{sym}W(\mathbf{P})(x,y) = W(\mathbf{P})(x,y) - W(\mathbf{P})(y,x).$$

Example 3.6. We consider again the profile from Example 3.3. This produces the following symmetric weighted tournament and representative matrix. Note that in the graph we only need to label positive edges, the reverse edges are labelled with the negative of these values.

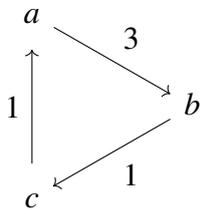


Figure 3.2: Left: a symmetric weighted tournament. Right: the matrix corresponding to this symmetric weighted tournament.

$$\begin{bmatrix} 0 & 3 & -1 \\ -3 & 0 & 1 \\ 1 & -1 & 0 \end{bmatrix}$$

Debord (1987) shows that all (and only) symmetric weighted tournaments whose weights all have the same parity correspond to some profile. Note that if we know how many agents were involved in the original profile, then we can determine the weighted tournament from

a weighted symmetric tournament.⁶ It is only if we do not know the number of agents that information is lost when a weighted tournament is turned into a symmetric weighted tournament.

Theorem 3.7 (Debord). *Take arbitrary A . For every function $g: \text{pairs} \rightarrow \mathbb{Z}$ such that (i) for all $(x, y) \in \text{pairs}$, $g(x, y) = -g(y, x)$, and (ii) the range of g only contains numbers of the same parity, there is some profile $\mathbf{P} \in \mathcal{L}^*(A)^{\mathcal{N}}$ such that $\text{sym}W(\mathbf{P}) = g$.*

Let us call the informational basis determined by symmetric weighted tournaments symC2 . The relation between symC2 and C2 is subtle: they are equivalent if we consider them bases over $\mathcal{L}^*(A)^N$ for some fixed electorate N . This corresponds to the fact that if we know the number of agents, then we can reconstruct the weighted tournament from the symmetric weighted tournament. However, this is not the case for the variable electorate domain, where C2 is more expressive than symC2 , as the following (somewhat arbitrary) example shows.

Example 3.8. Define a social choice function that returns the Borda winners (see Definition 3.11) if there is some pair of alternatives (x, y) for which $W(\mathbf{P})(x, y) \geq 5$, and that otherwise returns the top cycle. As the condition and both functions can be determined using C2 information, this function as a whole can be determined with C2 information. The winners for the profile \mathbf{P}^1 of Example 3.3 will be $\{a, b, c\}$, however if we add three symmetric agents with preferences abc , bca and cab , to create the profile \mathbf{P}^2 , displayed in Table 3.3, then there is a unique winner $\{a\}$. However, the symmetric weighted tournament is the same for both profiles.

Table 3.3: The profile of Figure 3.2 with three agents with cyclic preferences added.

1	2	3	4	5	6	7	8
a	a	c	c	b	a	b	c
b	b	a	a	c	b	c	a
c	c	b	b	a	c	a	b

It is hard, however, to think of a *reasonable* social choice function that makes use of the extra information afforded by C2 in comparison to symC2 . Intuitively, the extra expressivity of C2 gained through being able to determine the size of the electorates is not relevant to fair social choice.

Of course, there are informational environments that are even more expressive than C2 . For **voting situations** the information giving the names of agents is removed, leaving only the counts of each strict linear order; this is captured by the function

$$\text{vs}: \mathcal{L}^*(A)^{\mathcal{N}} \rightarrow \mathbb{N}^{\mathcal{L}^*(A)} \quad \text{defined by} \quad \text{vs}(\mathbf{P})(P) = |\{i \in N : P_i = P\}|.$$

⁶It is also important that we have assumed strict linear order profiles.

Observation 3.9. The range of vs is $\mathbb{N}^{\mathcal{L}^*(A)}$.

The basis associated with voting situations can express any anonymous rule, including those not expressible by C2 information, such as instant runoff voting (see the introduction to Chapter 2 for a description) or Dodgson (reference).

In general, we want to impose anonymity, thus most bases that we consider are coarsenings of voting situations. The following are important exceptions—though not because they produce reasonable rules. We remove all information from the profile with the function

$$\text{null}: \mathcal{L}^*(A)^{\mathcal{N}} \rightarrow \{\emptyset\}.$$

Obviously, the only choice functions expressible by the associated basis are constant. Now, suppose we have a fixed electorate N . For each agent $i \in N$ we remove information involving the other agents with the projection

$$\text{proj}_i: \mathcal{L}^*(A)^N \rightarrow \mathcal{L}^*(A) \quad \text{defined by} \quad \text{proj}_i(\mathbf{P}) = P_i.$$

The choice functions expressible by the associated basis include not only the dictatorship of i , but also reverse dictatorships and various other such functions. Note that although a dictatorship and reverse dictatorship for i are expressible by the same basis, their outputs never coincide (except in the trivial case where $|A| = 1$).

We have defined the functions T , W , $\text{sym}W$, vs , null , and the projections proj_i for each agent i . Each of these functions maps from the set of profiles into some other set of objects, and can be used to define a basis or informational environment, respectively tournaments, weighted tournaments, symmetric weighted tournaments, voting situations, trivial information, and single agent based information. A basis in turn is associated with the family of social choice functions that it expresses. We now go in the opposite direction; we start with a family of rules, and try to determine a basis which can be used to express this family and an associated function.

3.1.2 Informational functions that express positional scoring rules

The family of positional scoring rules is quite distinct from the families given by C1 and C2 information.

The traditional definition of a positional scoring rule makes reference to a **scoring vector** $\mathbf{x} \in \mathbb{R}^A$. Each individual assigns each alternative the score in the position of the vector that corresponds to the position of the alternative in the individual's ranking. We then sum up each alternative's scores to produce a final summed score. Alternatives with maximal

summed scores are selected. Formally, the scoring function with vector \mathbf{x} is the function $f_{\mathbf{x}}: \mathcal{D} \rightarrow 2^A \setminus \{\emptyset\}$ defined by

$$f_{\mathbf{x}}(\mathbf{P}) = \operatorname{argmax}_{b \in A} \sum_{i \in N} x_{r_i(a)} \quad \text{where} \quad r_i(a) = |\{b \in A : bP_i a\}| + 1.$$

Observation 3.10. For each scoring vector \mathbf{x} , the associated positional scoring rule on the variable electorate domain, $f_{\mathbf{x}}: \mathcal{L}^*(A)^{\mathcal{N}} \rightarrow 2^A \setminus \{\emptyset\}$, is a unique social choice function up to affine transformations of the vector; i.e. a scoring vector \mathbf{x}' such that there are no $\alpha, \beta \in \mathbb{R}$, $\alpha > 0$ such that $\alpha\mathbf{x} + \beta = \mathbf{x}'$ defines a different social choice function.

Note that the above observation does not apply when we have a fixed electorate. But, given a variable electorate, it means that we can define a canonical normalised vector for each scoring rule: $\mathbf{x} \in [0, 1]^A$ such that 0 and 1 are each values of some coordinate in \mathbf{x} .⁷ We will refer to such rules as positional scoring rules.

Positional scoring rules are not expressible by any of the bases that we have seen thus far, with one notable exception, the Borda social choice function. In canonical form the Borda vector is

$$\text{BV} = \left(1, \frac{|A| - 2}{|A| - 1}, \dots, \frac{1}{|A| - 1}, 0 \right).$$

Definition 3.11. The Borda social choice function, or just Borda for short, is the scoring function f_{BV} .

The obvious way to create a reduced informational environment that can express all scoring rules is to transform the profile into a matrix that records how many times each alternative is ranked in each position. We call this the positional matrix of a profile: it associates each column with an agent, and for each row counts how many times the agent is ranked in that position. We will assume that the agents are $a_1, a_2, \dots, a_{|M|}$ and that agent a_i is associated with column i . We formally define

$$\text{posM}: \mathcal{D} \rightarrow \mathbb{N}^{A \times A} \quad \text{by} \quad \text{posM}(\mathbf{P})_{ji} = |\{k \in N : r_k(a_i) = j\}|.$$

We define the basis \sim_{posM} in the obvious manner, and will sometimes simply refer to it as posM .

⁷Except in the trivial case where all the coordinates have the same value. Note that our definition is a slight generalisation of some traditional definitions, which tend to require that the vector is weakly decreasing throughout and strictly decreasing at least once.

Example 3.12. Let $A = \{a, b, c\}$, $N = \{1, 2, 3, 4, 5, 6\}$, and \mathbf{P}^3 be as given in Table 3.4.

Table 3.4: A profile in which all alternatives are ranked third twice.

1	2	3	4	5	6
a	a	a	a	b	c
b	b	c	c	c	b
c	c	b	b	a	a

$$posM(\mathbf{P}^3) = \begin{bmatrix} 4 & 1 & 1 \\ 0 & 3 & 3 \\ 2 & 2 & 2 \end{bmatrix}$$

Consider an arbitrary scoring vector $(\alpha_1, \alpha_2, \alpha_3)$. The total summed score for a is $4\alpha_1 + 2\alpha_3$, and for both b and c is $\alpha_1 + 3\alpha_2 + 2\alpha_3$. Under positional scoring the selection will be determined by which of these two sums is larger; in effect, then, which of $4\alpha_1$ and $\alpha_1 + 3\alpha_2$ is larger.

The information given by $posM$ can be used to express any scoring rule. In fact, it can express many more functions than these. In general, any partial ordering over the *set of vectors with natural number coordinates* corresponds to a generalised scoring rule which returns maximal alternatives according to this ordering.⁸ The following theorem, due to Fine and Fine (1974) determines the range of $posM$.

Theorem 3.13 (Fine). *Let M be a matrix of natural numbers and p be a natural number such that each row in M sums to p and each column in M sums to p . Then there is some profile \mathbf{P} such that $posM(\mathbf{P}) = M$.*⁹

We now look for a coarser basis which still expresses scoring rules. There are two obvious methods for reducing the information of $posM$. The first is scaling—matrices which are scalar multiples of each other can be considered equivalent. To obtain a canonical form of scaled matrices, we normalise such that all the columns sum to 1. As a result we will end up with matrices of type $\mathbb{Q}^{A \times A}$ rather than $\mathbb{N}^{A \times A}$: for the rest of the subsection we will suppose that all our matrices are of the former type, because this is a superset of the latter type. The normalisation itself is

$$\text{norm}: \mathbb{Q}^{A \times A} \rightarrow [0, 1]^{A \times A} \quad \text{defined by} \quad \text{norm}(M) = \left(\frac{1}{\sum_{i \in N} M_{1i}} \right) M.$$

We use npM to refer to the function $\text{norm} \circ posM$ and to this function's associated basis and informational environment.

The second method for reducing information concerns cases where all the alternatives are ranked in a given position at least once. This means that some of the scores assigned to the alternatives will cancel out. This leads to the following type of reduction.

⁸Note that multiple partial orderings will correspond to the same choice function: we do not have an equivalent result to Proposition 3.10 here.

⁹Note that the definition of the positional matrix given here transposes that of Fine and Fine (1974).

We write Z_i for a matrix with 1s in the row i and zeros elsewhere. The reduction of a (positive) matrix involves taking away as many Z_i rows as necessary until there is a zero in each row. Formally, define

$$\text{red}: \mathbb{Q}^{A \times A} \rightarrow \mathbb{Q}^{A \times A} \quad \text{by} \quad \text{red}(M) = M - \sum_{i \in N} \alpha_i C_i$$

$$Z_i = \begin{bmatrix} 0 & \cdots & 0 \\ \vdots & & \vdots \\ 0 & \cdots & 0 \\ 1 & \cdots & 1 \\ 0 & \cdots & 0 \\ \vdots & & \vdots \\ 0 & \cdots & 0 \end{bmatrix} \quad (\text{ith row})$$

where each $\alpha_i = \min\{M_{ij} : j \in N\}$, using the traditional definitions of matrix subtraction and scalar multiplication. We use rpM to refer to the composed function $\text{red} \circ \text{pos}M$ and the associated basis and informational environment.

Speaking somewhat loosely, these two reductions *commute*; the order of reduction is unimportant: this is due to the distribution of scalar multiplication over matrices. More precisely, $\text{norm} \circ \text{red} = \text{norm} \circ \text{red} \circ \text{norm}$. For the complete reduction we write $nrpM = \text{norm} \circ \text{red} \circ \text{pos}M$, which comes with its associated basis and informational environment.

Example 3.14. We again consider $A = \{a, b, c\}$ and $N = \{1, 2, 3, 4, 5, 6\}$ and the profile \mathbf{P}^3 from Example 3.12.

$$\text{pos}M(\mathbf{P}^3) = \begin{bmatrix} 4 & 1 & 1 \\ 0 & 3 & 3 \\ 2 & 2 & 2 \end{bmatrix} \quad \text{red} \circ \text{pos}M(\mathbf{P}^3) = \begin{bmatrix} 3 & 0 & 0 \\ 0 & 3 & 3 \\ 0 & 0 & 0 \end{bmatrix} \quad \text{rpos}(\mathbf{P}^3) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 0 \end{bmatrix}$$

It is obviously possible to express any scoring rule using $nrpM$ information. What about the range of this function? The following result is corollary to Theorem 3.13.

Corollary 3.15. *Take arbitrary A . For every matrix $M \in 0, 1^A \times [0, 1]^A$ with rational values such that (i) the values in each column sum to 1 and (ii) each row contains a 0 there is a profile $\mathbf{P} \in \mathcal{L}^*(A)^{\mathcal{N}}$ such that $\text{rpos}(\mathbf{P}) = M$.*

Let us compare the move from full positional information (pM) to scaled and reduced positional information ($nrpM$) with the move from weighted tournaments (W) to symmetric weighted tournaments ($\text{sym}W$). The move to $nrpM$ involves two reductions (normalisation and subtraction of rows) whereas $\text{sym}W$ only involves one. With $\text{sym}W$, if the size of the electorate is known, then we can reconstruct the full weighted tournament. This is not the case for $nrpM$; we can only reconstruct the full positional matrix if *only one* of the reductions is applied—i.e. from npM or rpM , if we know the size of the electorate, we can reconstruct pM , but we cannot reconstruct pM from $nrpM$. This suggests that these two reductions are distinct in an important way. Indeed, this is the case, as there are functions that can be expressed by npM and not by rpM , and vice-versa.

An example of a rule that is expressible by npM but not by rpM is the social choice function that returns those alternatives with the best median *position* (or, equivalently, the best median score, for any strictly decreasing score vector).¹⁰ We directly define this in terms of the npM matrix: the median score of an alternative is the first number such that sum of the values in the rows less than or equal to this number is at least a half. Formally, this is equivalent to: $med(a, M) = \min\{n \in \{1, \dots, |A|\} : \sum_{i=1}^n M_{ia} \geq 1/2\}$. The function

$$f_{\text{med}} : \mathcal{D} \rightarrow 2^A \setminus \{\emptyset\} \quad \text{is defined by} \quad f_{\text{med}}(\mathbf{P}) = \underset{a}{\operatorname{argmin}} (med(a, npM(\mathbf{P}))).$$

Example 3.16. Recall the profile \mathbf{P}^3 from Example 3.12. Its full positional and scaled positional matrices are displayed below.

$$posM(\mathbf{P}^3) = \begin{bmatrix} 4 & 1 & 1 \\ 0 & 3 & 3 \\ 2 & 2 & 2 \end{bmatrix} \quad npM(\mathbf{P}^3) = \begin{bmatrix} 2/3 & 1/6 & 1/6 \\ 0 & 1/2 & 1/2 \\ 1/3 & 1/3 & 1/3 \end{bmatrix}$$

The median winner is a .

Define a new profile \mathbf{P}^4 by adding three groups of three symmetric agents, for a total of nine agents. This produces the following matrices.

$$posM(\mathbf{P}^4) = \begin{bmatrix} 7 & 4 & 4 \\ 3 & 6 & 6 \\ 5 & 5 & 5 \end{bmatrix} \quad npM(\mathbf{P}^4) = \begin{bmatrix} 7/15 & 4/15 & 4/15 \\ 3/15 & 5/15 & 5/15 \\ 5/15 & 6/15 & 6/15 \end{bmatrix}$$

The median winners are a , b and c . Although the median winners for \mathbf{P}^3 and \mathbf{P}^4 are distinct, $rpM(\mathbf{P}^3) = rpM(\mathbf{P}^4)$ (which is easily verified), thus the median rule cannot be expressed by rpM .

An example of a rule expressible by rpM but not by npM is one that we call *near plurality*. We only provide a sketch: this returns both the plurality winners and any other alternative that is *one changed preference away* from being a plurality winner.

It is trivially the case that $symW$ can express all *and only* weighted tournament rules—it is how these rules are defined. This means, furthermore, that $symW$ is the *coarsest possible* basis for this class of rules. However, scoring rules were not defined with respect to $nrpM$, and, except in trivial cases, $nrpM$ is not *coarsest possible* basis for this family.

Example 3.17. Let $A = \{a, b, c\}$. Take \mathbf{P}^5 and \mathbf{P}'^5 such that

¹⁰The first use of the median for social choice seems to be Bassett and Persky (1999). General versions of such rules (ordered weighted averages) are studied by Goldsmith et al. (2014) and García-Lapresta and Martínez-Panero (2017).

$$\text{spos}(\mathbf{P}^5) = \begin{bmatrix} 1 & 1/3 & 0 \\ 0 & 2/3 & 1 \\ 0 & 0 & 0 \end{bmatrix} \quad \text{and} \quad \text{spos}(\mathbf{P}'^5) = \begin{bmatrix} 1 & 2/3 & 0 \\ 0 & 1/3 & 1 \\ 0 & 0 & 0 \end{bmatrix}.$$

Such profiles exist by Corollary 3.15. We observe that any scoring rule gives the same output to both these profiles. Suppose the scoring vector is (s_1, s_2, s_3) . If $s_1 > s_2$, then $\{a\}$ is returned; if $s_2 > s_1$ then $\{b\}$ is returned; and if $s_1 = s_2$ then $\{a, b, c\}$ is returned.

We can use the above example to show that (for $|A| = 3$) there is a coarser basis than $nrpM$ that expresses all positional scoring rules. Let B be the basis such that $\mathbf{P}^5 \sim_B \mathbf{P}'^5$, and all other profiles are in individual equivalence classes. Neither of $nrpM$ or B is a coarsening of the other, however there is a unique finest basis that coarsens both: the **infimum** or meet according to the partial order of refinement.

Over bases, refinement forms not only a partial order, but a complete lattice, thus both infimums and **supremums** are well-defined. With respect to refinement, for two bases B_1 and B_2 , we write $B_1 \wedge B_2$ for the infimum and $B_1 \vee B_2$ for the supremum. More generally, we write $\bigwedge X$ for the infimum (meet) of the set X and $\bigvee X$ for the supremum (join) of X .

Observation 3.18. If two bases B_1 and B_2 can both express a function, then so too can their meet $B_1 \wedge B_2$. More generally, if a family of bases B_i for $i \in I$ can express a function then so too can $\bigwedge_{i \in I} B_i$.

We can continue this idea in order to show that there is a unique coarsest basis that expresses all scoring vector rules. First, note that each specific rule f directly defines a basis, \sim_f , defined by $\mathbf{P} \sim_f \mathbf{P}'$ iff $f(\mathbf{P}) = f(\mathbf{P}')$. Such a basis \sim_f is *almost* a complete description of the rule—although it does not record to which alternatives a specific profile maps, it does associate this profile with all other profiles that map to the same alternative. The coarsest basis for positional scoring rules is $\bigwedge \{\sim_f : f \text{ is a positional scoring rule}\}$. We will refer to this basis as \sim_{scor} .

The existence of a coarser basis that expresses all positional scoring rules immediately implies that there are rules expressible by $nrpM$ that are not positional scoring rules. In fact, further, it can be the case that the coarsest basis for a family can express rules not in the family. For variable electorates this is the case for \sim_{scor} .

We define **lexicographic plurality**

$$f_{\text{lex}}: \mathcal{D} \rightarrow 2^A \setminus \{\emptyset\} \quad \text{by} \quad f_{\text{lex}}(\mathbf{P}) = \text{argleximax}_a [nrM(\mathbf{P})]_a$$

where by $[M]_a$ we mean the vector that is the column indexed by a , and by leximax we mean the ordering $>$ on vectors \mathbf{x}, \mathbf{y} such that $\mathbf{x} > \mathbf{y}$ iff there is $i \in \mathbb{N}$ such that $x_i > y_i$ and for all $j < i$, $x_j = y_j$.

On variable electorates this is not equivalent to any individual positional scoring rule. On the other hand, for any fixed electorate there is obviously an equivalent scoring rule with vector $(1/|N|, 1/|N|^2, 1/|N|^3, \dots)$. To see that lexicographic plurality, f_{lex} , is expressible by scoring rule information, we precisely need to show that for any two profiles that are distinguished by f_{lex} there is some scoring rule that distinguishes them. These profiles may have different sized electorates N and N' , without loss of generality suppose $N \geq N'$. Then the scoring rule with vector $(1/|N|, 1/|N|^2, 1/|N|^3, \dots)$ distinguishes the two profiles as required.

Here is a summary of the bases that we have seen in this section.

Basis	Notes
C1	Tournament information, majorities.
C2	Pairwise counts, how often an alternative is preferred to its pair.
symC2	Symmetric pairwise counts, how often an alternative is preferred minus how often the paired alternative is preferred.
vs	Voting situations, anonymous information.
null	No information, can only express constant functions.
proj	Projections, can express (rank positional) dictatorships.
posM	Information contained in a rank matrix, how often each alternative has each rank.
npM	Normalised rank matrix, invariant to scaling.
rpM	Reduced rank matrix, invariant to summing across any row, effectively changes count for a given rank for every alternative.
nrpM	Normalised and reduced rank matrix.
scor	Precisely the information need to express all scoring rules and no more.

In the following section we go on to define maximal decisiveness with respect to these.

3.2 Maximally decisive rules with respect to bases

An informational basis determines sets of profiles that must return the same outcomes. Whether or not a rule returns the smallest sets of alternatives possible, i.e. whether or not a rule is maximally decisiveness, is a question that concerns profiles that remain in the same set when alternatives are permuted.

Definition 3.19. A social choice function f is **maximally decisive** for a basis \sim if for any profile \mathbf{R} such that $a, b \in f(\mathbf{R})$, there is a permutation σ such that $\sigma(a) = b$ and $\mathbf{R} \sim \sigma(\mathbf{R})$.

Note that technically we don't require that a maximally decisive rule is expressible by the basis in question. If we need to make it precise that this is the case, we use the following term.

Definition 3.20. A social choice function is **properly** maximally decisive on \sim if it is maximally decisive on \sim and it is expressible by \sim .

Example 3.21. Constant rules are maximally decisive for and expressible by all bases. A dictatorship is similarly maximally decisive, but is not expressible by (for example) voting situations. We can however turn dictatorships into neutral, properly decisive social choice functions for voting situations.

When dealing with maximal decisiveness, it is often easier to think in terms of the specific representations of the informational bases in questions, rather than in terms of profiles. This and various other techniques make it slightly easier to deal with the abstract concept of maximal decisiveness, we now apply some of these to a specific social choice function: Borda.

3.2.1 Borda and maximal decisiveness

It is perhaps instructive to see some examples of *failures* of maximal decisiveness. Our first example shows that Borda is not maximally decisive for symmetric weighted tournaments by considering permutations applied to a weighted tournament, rather than to the underlying profiles.

Example 3.22. Borda is not maximally decisive for symC2 . The families of profiles determined by the weighted graph in Figure 3.3 provides a counterexample.

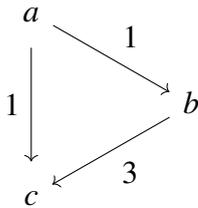


Figure 3.3: A weighted tournament that shows that Borda is not maximally decisive for symC2 . The output according to Borda is $\{a, b\}$. However, note that if we move a to b 's position in the graph, then a ends up with different edges. This implies that for any profile P that maps to this weighted tournament, if we permute a to b in the profile, then the resulting profile will map to a different weighted tournament.

Our second example shows that Borda is not maximally decisive for the positional scoring rule basis because there is a neutral rule that is more decisive than Borda on some specific profile.

Example 3.23. Borda is not maximally decisive for scoring vector information (i.e. \sim_{scor}). Consider, for example, the equivalence class of profiles that contain the following profile.

1	2
a	c
b	b
c	a

Table 3.5: A profile that shows that Borda is not maximally decisive for scoring vector information. Here the Borda output is $\{a, b, c\}$. However there are clearly other scoring rules that select only b : as these rules are neutral, this suffices to show that Borda is not maximally decisive.

We tried to choose the coarsest possible bases in the above examples because this gives the strongest results. This can be seen by the contraposition of the following.

Observation 3.24. If a function f is maximally decisive for a basis \sim then it is also maximally decisive for all coarsenings \sim' of \sim .

However, it is the case that Borda is maximally decisive for some bases: in particular the basis that is trivially defined in terms of the rule. As well as this rather trivial case, Borda is also maximally decisive for other, more interesting, bases.

Proposition 3.25. *Borda is properly maximally decisive for the meet of symmetric weighted tournament information and scaled positional information: that is, Borda is properly maximally decisive for $\text{symC2} \wedge \text{npM}$.*

Proof. The “properly” part is due to the fact that Borda is expressible by symC and by npM , thus, in line with Observation 3.18, Borda is expressible by $\text{symC2} \wedge \text{npM}$.

For maximal decisiveness, take an arbitrary profile \mathbf{P} for which Borda outputs both a and b . Suppose also that $|A| \geq 4$; for fewer alternatives the proof is simpler. As a general overview of the strategy, we add agents to the profile in order to create a profile in the same equivalence class for which a and b have the same weighted majorities, i.e. $W(\mathbf{P})(a, b) = W(\mathbf{P})(b, a)$ and for all $x \in A \setminus \{a, b\}$, $W(\mathbf{P})(a, x) = W(\mathbf{P})(b, x)$. The permutation that switches a and b and leaves the other alternatives fixed fulfils the condition of Definition 3.19. In particular, we will add four agents of \mathbf{P}^7 displayed in Table 3.7, which is constructed from \mathbf{P}^6 displayed in Table 3.6.

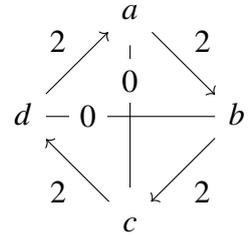
\mathbf{P}^6			
1	2	3	4
a	d	b	c
b	c	a	d
\mathbf{x}	\mathbf{x}'	\mathbf{x}	\mathbf{x}'
c	b	d	a
d	a	c	b

Table 3.6: A profile whose symmetric weighted tournament only contains zeros. Here \mathbf{x} refers to an arbitrarily ordered vector containing all the remaining alternatives, and \mathbf{x}' is the inverse of \mathbf{x} . The symmetric weighted tournament of this profile contains only zeros for weights because agents 1 and 2, and also agents 3 and 4, have inverse rankings.

We now permute a and c for agents 2 and 3. This results in the following profile.

P^7			
1	2	3	4
a	d	b	c
b	a	c	d
x	x'	x	x'
c	b	d	a
d	c	a	b

Table 3.7: A profile whose scaled positional matrix is the same as the profile in Table 3.6. Only candidates a and b have changed ranking in this profile. The weighted tournament of this profile does not contain only zeros: see the partial weighted tournament to the right.



We can assume that the weighted majorities in P are of even parity, if not, double the profile so that this is the case (by npM the doubled profile is in the same equivalence class). Assume, without loss of generality, that in P alternative b has a weak pairwise majority over a . If this majority is strict, add as many copies of P^7 as required so that neither a nor b has a strict pairwise majority over each other. This profile is in the same equivalence class as P , via the profile where copies of P^6 are added instead. Note furthermore that a and b have the same Borda scores as each other in this profile. Let us call this resultant profile P' .

The end of the proof involves a similar addition of agents to ensure that a and b have the same pairwise majority over every other alternative $x \neq a, b$.

Write $X = A \setminus \{a, b\}$, and consider the vector $\mathbf{v} \in \mathbb{Z}^X$ defined by $v_x = W_{P'}(a, x) - W_{P'}(b, x)$. If $\mathbf{v} = (0, \dots, 0)$, we are done. Otherwise, consider x such that the absolute value $|v_x|$ is maximal. Suppose also that v_x is positive—the negative case is symmetric. There must be some $y \in X$ such that v_y is negative, because otherwise a and b would have different Borda scores. We now add a permuted version of P^7 that has the (partial) weighted tournament displayed in Figure 3.4 to the profile.

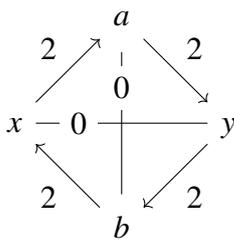


Figure 3.4: A (partial) weighted tournament. As for the profile in Table 3.7, a profile that induces this weighted tournament be constructed so that it has the same positional matrix as the profile displayed in Table 3.6, thus this constructed profile belongs to the same informational class.

This creates a profile P'' —in which neither a nor b dominates the each other—and a corresponding vector \mathbf{v}' defined by $v'_x = W_{P''}(a, x) - W_{P''}(b, x)$. Note that if we order \mathbf{v} and \mathbf{v}' decreasingly, the second ordered vector is lexicographically smaller than the first, because we are dealing with finite numbers, repeating this procedure will eventually end up with a vector that contains only zeros as required. \square

This result is perhaps not too surprising given known characterisations of Borda in terms of conditions that are quite similar to the informational restrictions imposed by the basis. Our

proposition itself, however, is not quite a characterisation: the rule that outputs the Borda loser, for example, is also maximally decisive for the basis in question. Although there are various other violations of the Pareto criterion possible here, at least maximal decisiveness rules out the constant rule that always outputs all alternatives.

Note that in the above proposition we used *symmetric* weighted tournaments and *scaled* positional matrices. Both of these informational classes form bridges between electorates of variable sizes, and it may be wondered if we get the same results with a refinement that only uses one or the other of these bridges. We cannot, for instance, do without the “symmetric” part: see Table 3.8.

1	2	3
<i>a</i>	<i>a</i>	<i>b</i>
<i>b</i>	<i>b</i>	<i>c</i>
<i>c</i>	<i>c</i>	<i>a</i>

Table 3.8: A profile that demonstrates that Borda is not maximally decisive for $C2 \wedge npM$. For this profile, Borda outputs $\{a, b\}$. However, the only profiles in the same equivalence class are scalings. For these, there is no permutation between *a* and *b*.

Borda is not maximally decisive for scoring information. However, there is an obvious way to make it more decisive, using a ranking version of lexicographic plurality.

Observation 3.26. Borda refined by a social welfare version of lexicographic plurality is properly maximally decisive for full positional information.

One may wonder if we can create a similar maximally decisive refinement of Borda for weighted tournament information. We do not know of any natural maximally decisive rules for symC2 information, let only ones that can furthermore be generalised to social welfare functions.

3.2.2 Finding maximally decisive rules

We now turn attention from a specific rule to specific bases, and ask if we can say anything about maximally decisive rules with respect to a basis. For many bases, it is difficult to even describe proper maximally decisive rules.¹¹

With respect to C1 information, it is not only the case that none of the common tournament solutions are not maximally decisive, but furthermore it is difficult to even come up with a description of a maximally decisive rule. The problem can be demonstrated by the tournament in Figure 3.5. The details of the construction do not interest us too much here, suffice it to say that almost all tournament solutions described in the literature can be shown to be non-maximally decisive using this tournament.

¹¹Of course, proper maximally decisive rules certainly exist, in an abstract sense, for every basis.

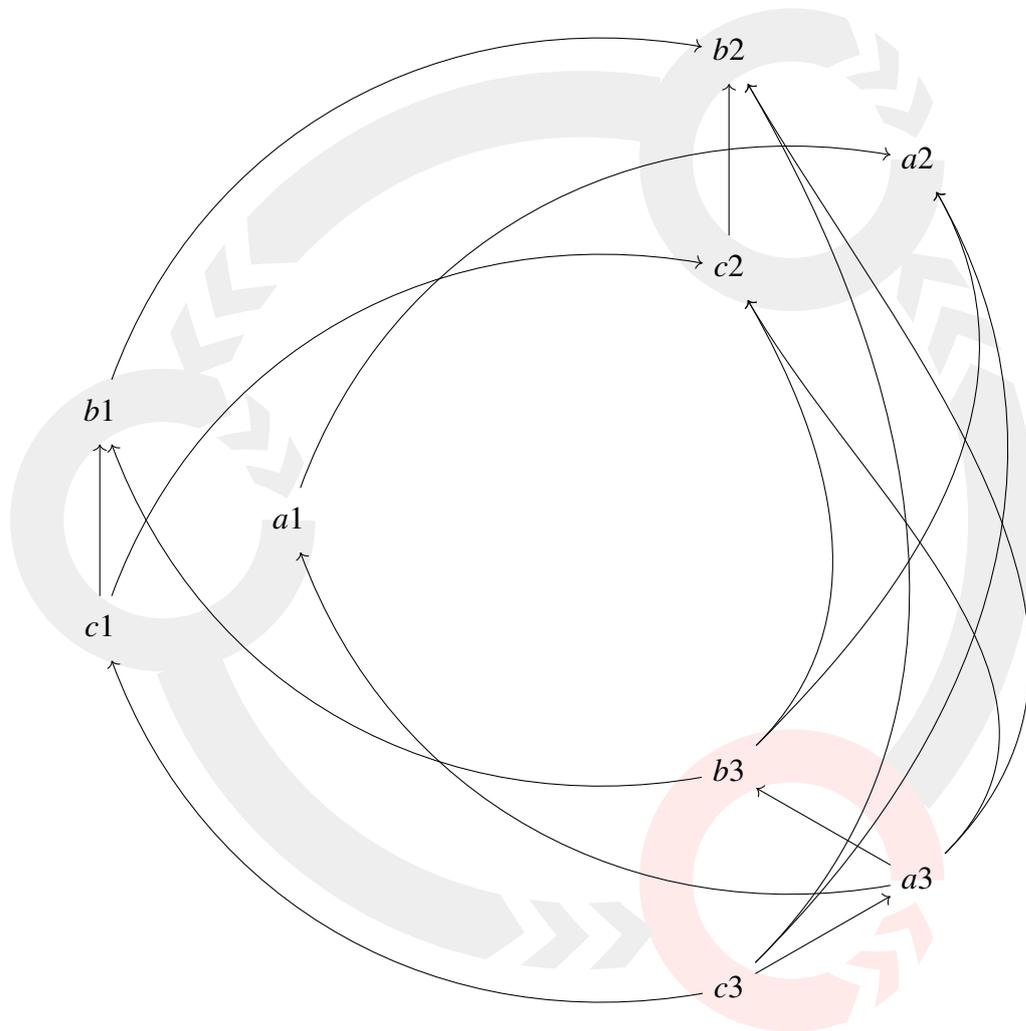


Figure 3.5: An almost symmetric tournament—undrawn arrows go down. The faded picture in the background displays the general shape of the tournament, with the asymmetry in red. This tournament contains nine alternatives x_i for $x \in \{a, b, c\}$ and $i \in \{1, 2, 3\}$. For each $i \in \{1, 2, 3\}$, there is a cycle between the triple of alternatives a_i, b_i and c_i . For $i \in \{1, 2\}$ the cycles are $a_i \rightarrow b_i \rightarrow c_i \rightarrow a_i$, whereas the cycle for $i = 3$ is inverted, $c_3 \rightarrow b_3 \rightarrow a_3 \rightarrow c_3$. The other relations are symmetric, and are defined as follows: for the pairs $(i, j) \in \{(1, 2), (2, 3), (3, 1)\}$, we have $x_i \rightarrow x_j$ for $x \in \{a, b, c\}$ and $x_j \rightarrow y_i$ for $x \neq y$. Due to the fact that the cycle through a_3, b_3 and c_3 goes in the opposite direction to the other triples, these three elements can be distinguished from the other triples, and it follows that the other two triples can also be distinguished from each other. However, it is difficult to describe a rule that selects only one of a_1, a_2 and a_3 , let alone to imagine a principled reason for selecting one and not the others.

An exception to this are the argmax rules of Chapter 2, which do tend to select a proper subset from this tournament. However, these do not form a consensus: the different rules disagree on which alternatives should be chosen.¹² Other profiles also show that even these rules are not maximally decisive, with the possible exception of argmax applied to the simple tree.

For the simpler bases, it is possible to characterise the proper maximally decisive rules in a more informative manner than this. For trivial information, the maximally decisive rules are precisely those constant rules that return singletons. Thus these rules are resolute.

Observation 3.27. Under the trivial basis, the only maximally decisive rules are constant singleton rules.

Similarly, any maximally decisive rule for single agent information must be resolute.

Observation 3.28. Under proj_i information, there is one maximally decisive rule for each function of the form $g: \mathcal{L}^*(A) \rightarrow A$, and this is defined by $g \circ \text{proj}_i$. Conversely, every maximally decisive rule under proj_i information corresponds to some function $g \circ \text{proj}_i$.

3.3 Final remarks on maximally decisive social choice

In this chapter we have introduced a new way of analysing whether or not a social choice function selects as small sets as possible. This notion, that we call maximal decisiveness, is intertwined the notion of what information is taken as relevant to the choice. With more information, one can expect a social choice function to select smaller sets. This allows us to define our concept of maximally decisiveness in such a way that does not conflict with anonymity and neutrality.

Anonymity is taken as only one type of informational constraint on social choice functions. In Section 3.1 we defined a general framework for the information used by functions, what we call the informational basis of a function. Within this framework we looked at known informational restrictions, C1 and C2 information, and defined our own versions of various positional information environments. Such informational restrictions really amount to equivalence classes; each equivalence class on the space of profiles determines a family of rules that can be expressed using the information expressed by that equivalence class. In the other direction, an informational basis can be determine in terms of a family of social choice functions.

¹²Argmax Banks returns six of the nine. With the simple tree and for iterative Condorcet, only three are returned, each a different subset of the six returned by argmax Banks. Argmax applied to the fair tree returns the three alternatives not returned by argmax Banks.

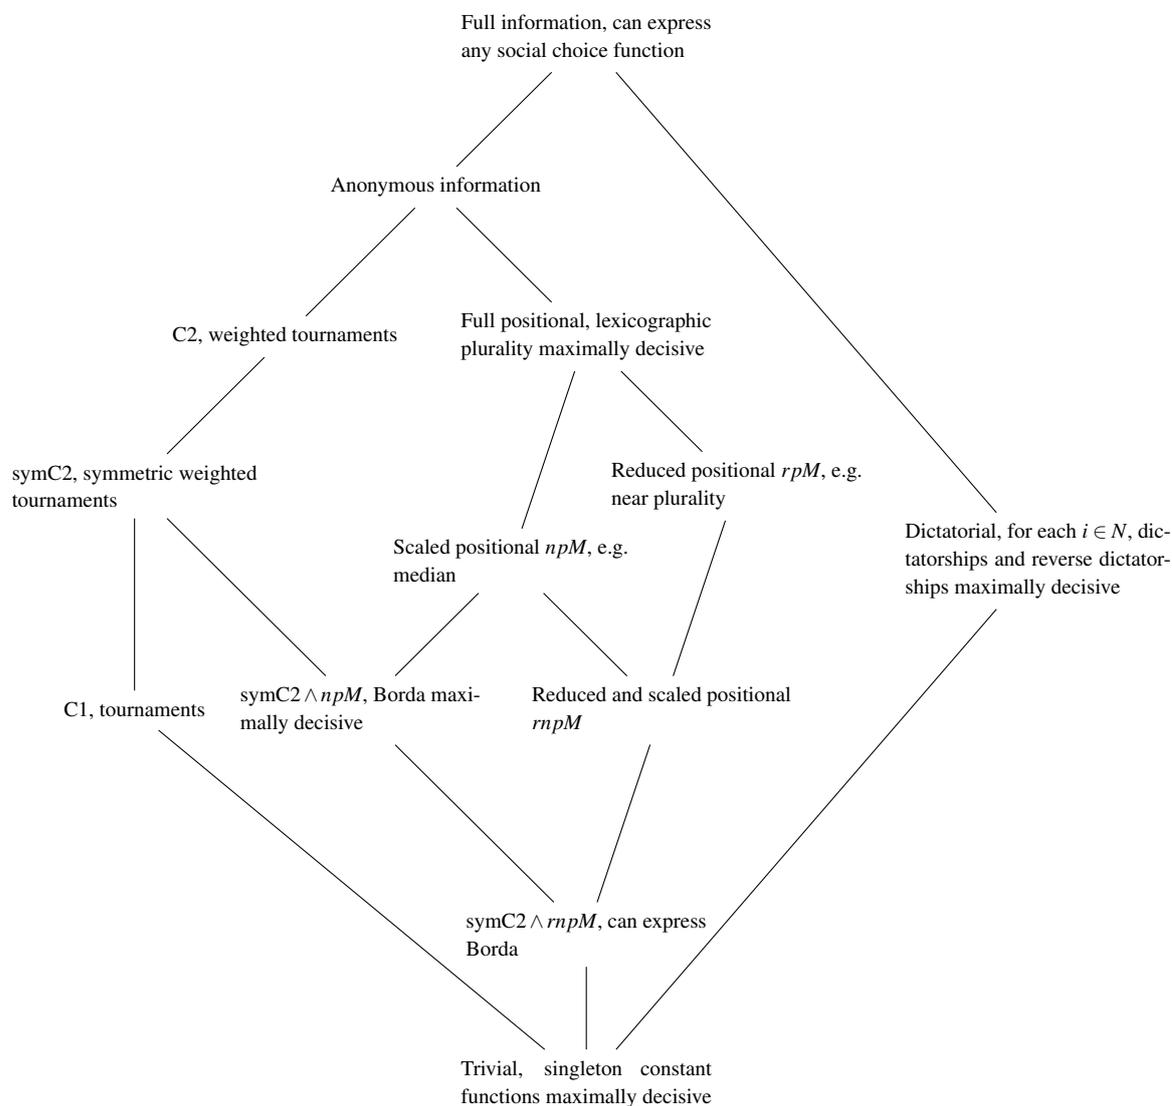


Figure 3.6: Informational bases ordered in terms of refinement. A line connects a higher basis with a lower basis if the higher basis refines the lower. For some bases, the social choice functions that it can express, or those functions which are maximally decisive for it are written.

In Section 3.2 we define the formal notion of maximally decisiveness. We investigated the Borda choice rule with respect to this concept, but in general we do not know too much about this concept. It seems to be easiest to consider positional-type methods using this concept. For other natural informational bases, such as those described by C1 and C2 information, maximal decisiveness seems to be an unreasonably strong condition to impose. Further investigation seems to be merited.

Figure 3.6 displays a number of different bases, ordered by refinement. Some examples of

maximally decisive rules are also displayed.

We can place the work in this chapter firmly in the Arrovian tradition of supposing that strict linear order preferences are the underlying objects of social choice. We do not admit other types of basic information, such as approval ballots. Some of these other possible options are considered in Chapter 5. Nor do we consider adding extra information, for example in the form of interpersonal comparisons—there is a long tradition of such a move in social choice theory. We go in the opposite direction, with a focus on neutrality as the basis of a rule: all alternatives should be treated equally. This in effect takes the place of independence conditions.

The work in this chapter can be contrasted with work considering how much information needs to be elicited from agents in order to compute the outcome of a social choice procedure. At a purely practical level, Conitzer and Sandholm (2005) determines the communication complexity of a number of social choice procedures: the (worst-case) number of bits required to be elicited from the agents in order to calculate the output of the procedure. On a more theoretical level Sato (2009; 2016) considers a measure of information based upon the message spaces for individual agents. Our measure of information is more holistic; it takes whole profiles into account, as opposed to looking at information from an agent-wise perspective.

Our notion has more similarity to another practical notion from the computational social choice literature: the compilation complexity of a social choice procedure. Defined by Chevaleyre et al. (2009), this determines how much space is required to store a partial profile, given that we will want to be able to calculate the outcome of a specific procedure when given the remainder of the profile. Our notion of the coarsest informational basis for a given family is similar: how much space is required to store a complete profile, given that we will want to be able to calculate the outcome when given any one of a family of procedures. Xia and Conitzer (2010) give the compilation complexity for many common social choice procedures.

However, our notion is more theoretical than practical. Let us look at things from an opposite point of view: *indecisiveness* may be informative. If for a particular profile there seems to be no reason to select one alternative over another, we may be able to infer something about the structure of this profile that makes this the case. This may help us to identify other profiles where this is less obviously the case, but where nonetheless two alternatives should only be selected together. Even more generally, this may help us to determine the underlying reasons why one alternative is chosen instead of another.

Chapter 4

Domain restrictions and independence of irrelevant alternatives

The work in this chapter concerns two types of informational restrictions in the traditional social choice theoretic setting where a profile of complete preorders is aggregated into a single complete preorder.¹ The first of these supposes that some profiles are not possible: it thus involves a **domain restriction** of a social welfare function. Such domain restrictions are often studied in social choice theory; in particular they are often observed to mitigate the negative conclusions of impossibility results. Perhaps the most celebrated example of this is that of the **single-peaked domain**, but many other examples exist. Gaertner (2001; 2002) provides surveys of the social choice theoretic literature on domain restrictions.

Instead of a global look at what information is used by a rule or a family of rules, we are concerned with how much information from a single profile is required to calculate *part* of the output of a social welfare function. It may be seen as desirable that only local information is required to calculate parts of the output. Independence properties ensure this by requiring that whenever individuals' preferences over a given set of alternatives remain the same, then so should the social preference.

Independence of irrelevant alternatives (IIA) is such an independence property. It is the central property of Arrow's (1950) celebrated impossibility. Arrow's original formulation requires independence over every subset of alternatives. However, there is an alternative formulation which requires independence only for subsets of cardinality two. Although this formulation is superficially weaker, it is straightforward to observe that independence over pairs implies independence over any set, and thus that the two formulations are equivalent.

Blau (1971) showed that independence over *larger* subsets implies independence over *smaller* subsets. Thus we observe what we call the **Blau equivalence**: roughly expressed, this states that all versions of independence have the same strength. The Blau equivalence implies that Arrow's result cannot be escaped by weakening independence to subsets of some fixed cardinality greater than two. In turn we call this stronger version of Arrow's result the **Blau impossibility**.

The Blau equivalence rests upon the assumption that the domain is complete. In this chap-

¹ The work in this chapter is based upon Kruger and Sanver (2018).

ter we consider the effects of imposing domain restrictions, and delineate when the Blau equivalence does and doesn't hold.

One set of motivations behind independence properties concerns the impracticality of determining or using complete profiles of preferences over anything more than a small set of alternatives. Grouped here are concerns about the cognitive burden on agents required to rank too many alternatives, the communication complexity and perhaps computation complexity involved with large profiles, and, even more fundamentally, the fact that in some situations one may not even know the full range of possible alternatives. Another motivation is more ideological: one may claim that the ranking over a given set of alternatives should be independent of some other alternatives because these other alternatives are irrelevant, and attempting to use them, as Borda scoring does, to give a pseudo-utility is wrong-headed. Our results in this chapter are probably of more interest to someone who is in favour of independence for practical reasons: we show that, for some domain restrictions, one may be able to escape impossibility results by using weaker forms of independence, which may still overcome the practical concerns.

Section 4.1 presents the setting formally, including a formal definition of the condition requiring that independence over sets of cardinality k is satisfied. Section 4.2 shows that under restricted domains, unlike under the full domain, increasing k may effectively weaken independence—hence the Blau equivalence may fail. Section 4.3 characterises the domains for which nominally different versions of independence have effectively different strengths, thus also effectively providing a characterisation for the Blau equivalence. In Section 4.4 we develop the idea of what we call the **Blau partition** of a domain. We show that there are many possible such partitions; thus that the Blau equivalence can fail in many different ways. We construct a particular domain in Section 4.5, making sure during its construction that it is subject to Arrow's impossibility. We define a social welfare function on this domain that satisfies a weaker version of independence alongside Arrow's other properties. This verifies that on some domains it is possible to escape Arrow's result through weakening independence. Section 4.6 makes some final remarks.

4.1 Definitions

In this chapter we consider social welfare functions with restricted domains: functions $F: D^N \rightarrow \mathcal{W}(A)$ for some $D \subseteq \mathcal{W}(A)$. We use $F^*(\mathbf{R})$ to denote the strict component of $F(\mathbf{R})$. Given a set $X \subseteq A$, the restriction of a preference R to X is $R \upharpoonright_X = \{(x, y) \in R : x, y \in X\}$. We use similar notation to apply restrictions to domains (sets of preferences) and preference profiles. Formally, for $X \subseteq A$, $D \subseteq \mathcal{W}(A)$ and $\mathbf{R} \in D^N$ we write $D \upharpoonright_X = \{R \upharpoonright_X : R \in D\}$ and $\mathbf{R} \upharpoonright_X = (R_i \upharpoonright_X)_{i \in N}$.

We say two preferences R and R' **coincide** X if $R \upharpoonright_X = R' \upharpoonright_X$. Similarly, preference profiles, and domains, coincide on X iff their restrictions to X are identical. A preference R over X **extend** a preference R' over $Y \subseteq X$ if these coincide on Y .

Most of this chapter is concerned with general properties of social welfare functions. The following are well-known examples of such properties. A social welfare function F is **dictatorial** if there is some $i \in N$ such that xR^*_iy implies $x[F^*(\mathbf{R})]y$. A social welfare function F satisfies the **Pareto criterion** if $xR^*_iy, \forall i \in N$ implies $x[F^*(\mathbf{R})]y$. These two properties, alongside independence of irrelevant alternatives which we formally define below, are the traditional conditions for the impossibility result of Arrow (1950). Along the lines of Blau (1971), we consider a generalised notion of what it means for a social welfare function to be “independent”. A social welfare function is k -IND, for an integer $2 \leq k \leq m$, if for every $X \subseteq A$ with $|X| = k$, for any two profiles $\mathbf{R}, \mathbf{R}' \in D^N$ that coincide on X , we have $F(\mathbf{R}) \upharpoonright_X = F(\mathbf{R}') \upharpoonright_X$. As discussed in the introduction to this chapter, independence of irrelevant alternatives is sometimes expressed as 2-IND and sometimes as the conjunction of all k -IND for $k = 2, \dots, m$; these formulations are equivalent.

4.2 A counterexample on restricted domains

One may question the whether there is any purpose, even at a formal level, in defining k -IND separately from 2-IND. The fact that independence over smaller sets implies independence over larger sets seems to have become common knowledge soon after Arrow’s original presentation of independence of irrelevant alternatives. Further, this “upwards” implication holds over any domain.

Proposition 4.1. *For any social welfare function F over any domain $D \subseteq \mathcal{W}(A)$, if F satisfies k -IND then it also satisfies l -IND, for any $2 \leq k < l \leq m$.*

Proof. Take two profiles \mathbf{R} and \mathbf{R}' such that $\mathbf{R} \upharpoonright_Y = \mathbf{R}' \upharpoonright_Y$ for some set Y of cardinality l . We want to show that $F(\mathbf{R}) \upharpoonright_Y = F(\mathbf{R}') \upharpoonright_Y$. For each distinct pair $\{x, y\} \subset Y$, there is a set X_{xy} such that $\{x, y\} \subseteq X_{xy} \subset Y$ and $|X_{xy}| = k$. Clearly $\mathbf{R} \upharpoonright_{X_{xy}} = \mathbf{R}' \upharpoonright_{X_{xy}}$, thus by k -IND $F(\mathbf{R}) \upharpoonright_{X_{xy}} = F(\mathbf{R}') \upharpoonright_{X_{xy}}$. This implies $F(\mathbf{R}) \upharpoonright_{\{x, y\}} = F(\mathbf{R}') \upharpoonright_{\{x, y\}}$. As this is the case for every pair $x, y \in Y$, $F(\mathbf{R}) \upharpoonright_Y = F(\mathbf{R}') \upharpoonright_Y$ as required. \square

This observation dates back to May (1954), though the fact that it applies in all domains is not explicitly noted.

In 1971 Blau proved the inverse “downwards” implication, thus showing that nominally weaker versions of independence imply nominally stronger versions. The Blau equivalence amounts to the combination of the upwards and downwards implications.

Theorem 4.2 (Paraphrase of Blau’s (1971) “Theorem 4”). *For any social welfare function F over the full domain $D = \mathcal{W}(A)$, if F satisfies l -IND then it also satisfies k -IND, for any $2 \leq k < l < m$.*

Proof. It suffices to show that when F satisfies l -IND it satisfies $(l - 1)$ -IND. Take F that satisfies l -IND. Take two arbitrary profiles, and suppose there is a set X of size $l - 1$ such that the two profiles coincide on this set. We suppose that the two profiles are identical except for one voter, they can thus be written (\mathbf{R}_{-i}, R_i) and (\mathbf{R}_{-i}, R'_i) . Extending to the case where multiple voters have different preferences only involves iteration of this case. Consider two alternatives $x, y \in A$ such that $x \neq y$, $x \notin X$ and $y \notin X$ (these exist as $|X| \leq m - 2$). Take a preference $R'' \in \mathcal{W}(A)$ such that this extend both $R_i \upharpoonright_{X \cup \{x\}}$ and $R'_i \upharpoonright_{X \cup \{y\}}$. Equivalently, $R'' \upharpoonright_{X \cup \{x\}} = R_i \upharpoonright_{X \cup \{x\}}$ and $R'' \upharpoonright_{X \cup \{y\}} = R'_i \upharpoonright_{X \cup \{y\}}$. There may be multiple possible R'' 's, but there is always at least one. Now by l -IND we have $F(\mathbf{R}, R_i) \upharpoonright_{X \cup \{x\}} = F(\mathbf{R}, R'') \upharpoonright_{X \cup \{x\}}$ and $F(\mathbf{R}, R'') \upharpoonright_{X \cup \{y\}} = F(\mathbf{R}, R'_i) \upharpoonright_{X \cup \{y\}}$. Thus $F(\mathbf{R}, R_i) \upharpoonright_X = F(\mathbf{R}, R'') \upharpoonright_X = F(\mathbf{R}, R'_i) \upharpoonright_X$. \square

However, Blau’s proof for the downward direction uses two conditions not present for the upward direction: first, it supposes the full domain; and second it requires a strict inequality $l < m$. The second of these two is necessary because m -IND is trivially satisfied by any social welfare function. We will be more interested in what happens when we relax the assumption of a full domain, which is used by Blau in the following manner: Blau’s proof takes two preferences and “connects” them through a third. This third preference is guaranteed to exist on the full domain. When we consider restricted domains, this is no longer always the case. In fact in the following domain this third “connecting” preference *never* exists, for $l = 3$.

Example 4.3. Fix $m = 4$. Let D be the domain containing the six linear orders displayed in Table 4.1. Every preference in this domain is a linear order. Further, every triple has six different possible orderings:

$$\forall x, y, z \in A \text{ such that } x, y, \text{ and } z \text{ are distinct, } |D|_{\{x, y, z\}} = 6.$$

Thus in a setting where preferences are restricted to linear orders, all triples are free.² A result established by Blau (1957) states that if all triples are free Arrow’s impossibility holds. Thus any social welfare function on this domain that satisfies 2-IND and the Pareto criterion is dictatorial.

Note that given any linear order over three alternatives, only one preference in D extends this linear order. No preference in the domain extends a non-linear complete preorder of three alternatives. That is to say, for $X \subset A$ with $|X| = 3$ and $R, R' \in D$, if $R \upharpoonright_X = R' \upharpoonright_X$ then $R = R'$. This implies that any social welfare function over this domain satisfies 3-IND

² Here “triple” means a subset of alternatives of size three. A triple is free if every possible ordering over this triple exists within the domain.

Table 4.1: A domain for which every triple is free. This means the Arrow’s result applies to this domain; however, the result no longer holds on this domain if we weaken the independence axiom.

R^I	R^{II}	R^{III}	R^{IV}	R^V	R^{VI}
a_2	a_3	a_3	a_2	a_1	a_4
a_1	a_1	a_4	a_4	a_4	a_1
a_3	a_2	a_2	a_3	a_2	a_3
a_4	a_4	a_1	a_1	a_3	a_2

trivially. For example, ranking by Borda scores provides a non-dictatorial social welfare function that satisfies the Pareto criterion. By the above paragraph this cannot satisfy 2-IND; more generally this shows that 3-IND does not imply 2-IND.

The above example implies that k -IND and l -IND for $k \neq l$ need not be equivalent. In the next section we first consider k -IND for some fixed value of k , before returning to the issue of comparing k -IND to l -IND for $k \neq l$.

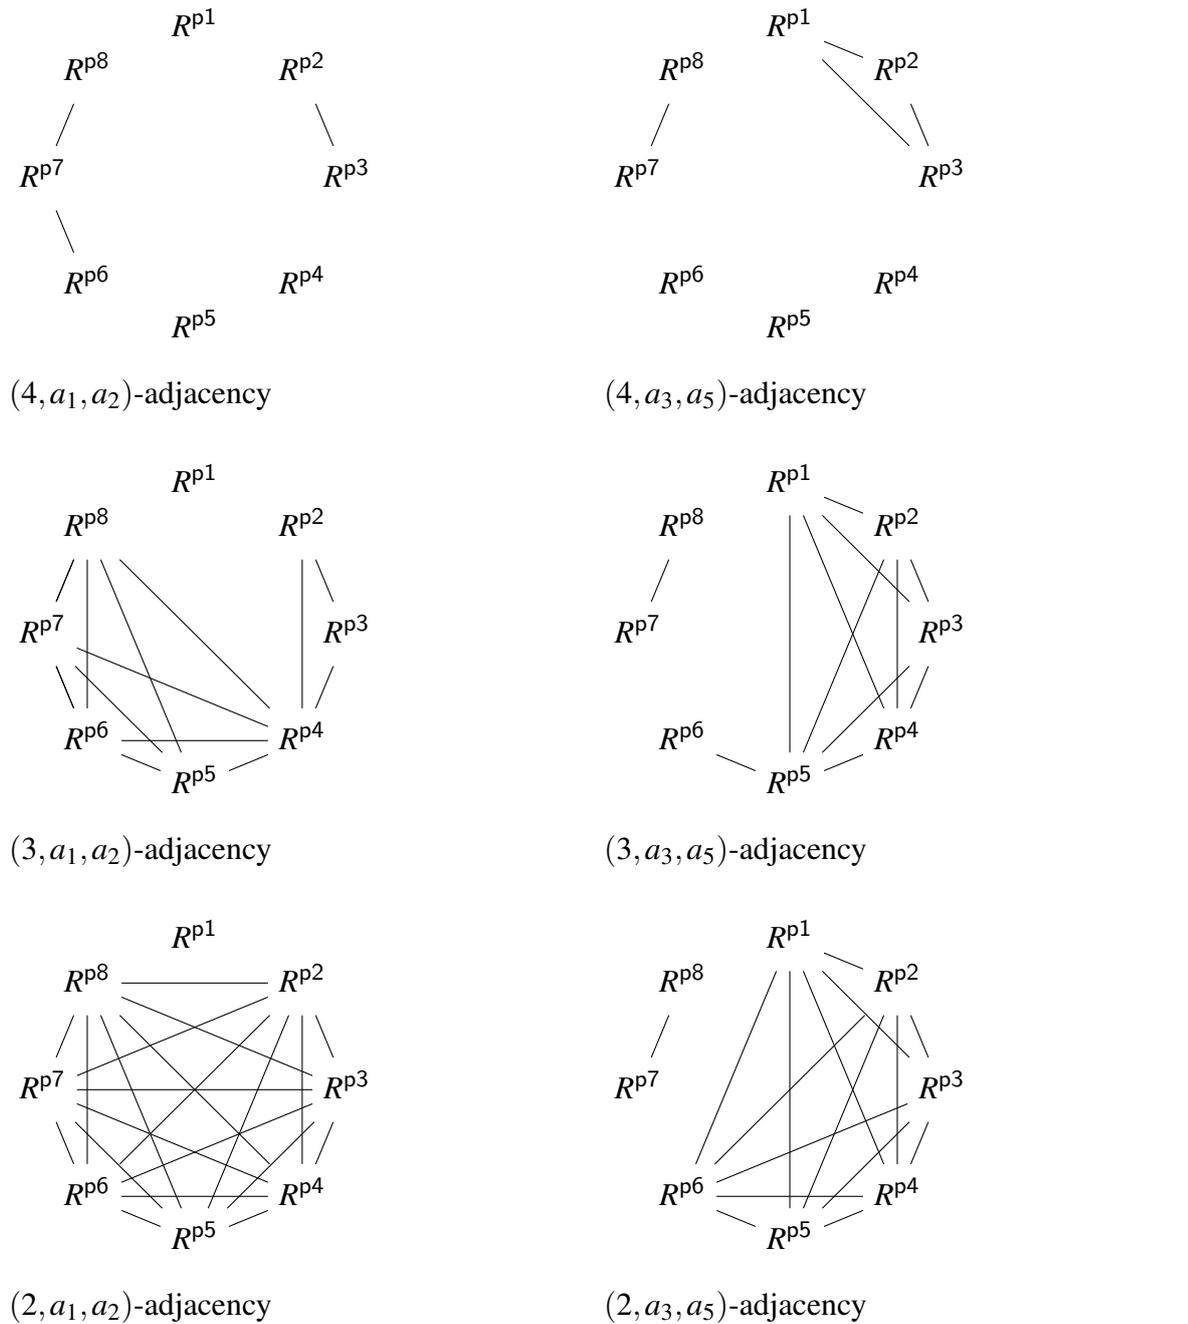
4.3 The Blau equivalence: when are different versions of independence equivalent?

Given a particular social welfare function defined on a particular domain, k -IND may or may not be satisfied. In this section we give a necessary and sufficient condition for this property. This will allow us to determine the domains for which k -IND implies l -IND, for arbitrary values of k and l . This generalises the question as posed by Blau, who showed that on the full domain, if a social welfare function satisfies k -IND then it also satisfies l -IND for $k, l < m$. We have already seen in Section 4.2 that there are domains where this general result does not hold. This section will allow us to determine when it does.

Definition 4.4. For a domain D , integer $k \in \{2, \dots, m\}$ and two alternatives $x, y \in A$, two preferences $R, R' \in D$ are (k, x, y) -adjacent if there is a set $X \subseteq A$ such that $|X| = k$, $\{x, y\} \subseteq X$ and $R \upharpoonright_X = R' \upharpoonright_X$.

Figure 4.1 displays some examples of adjacency on a domain that we will return to later in the chapter.

Just considering adjacency does not suffice; recall that Blau’s method uses a third profile to “connect” two profiles that are not directly adjacent. For domains that are proper subsets of the full domain, it may be necessary to use even more profiles to form this connection.



R^{p1}	R^{p2}	R^{p3}	R^{p4}	R^{p5}	R^{p6}	R^{p7}	R^{p8}
a_1	a_2	a_2	a_3	a_3	a_4	a_4	a_5
a_2	a_1	a_3	a_2	a_4	a_3	a_5	a_4
a_3	a_3	a_1	a_4	a_2	a_5	a_3	a_3
a_4	a_4	a_4	a_1	a_5	a_2	a_2	a_2
a_5	a_5	a_5	a_5	a_1	a_1	a_1	a_1

Figure 4.1: Graphs displaying (k, x, y) -adjacency for some example values on the single-peaked domain of Table 4.2 (which is also displayed above). Note that for $(x, y) = (a_1, a_2)$ or $(x, y) = (a_3, a_5)$, $(2, x, y)$ -reachability is the same as $(3, x, y)$ -reachability but not $(4, x, y)$ -reachability.

Definition 4.5. For a domain D , integer $k = 2, \dots, m$ and two alternatives $x, y \in A$, two preferences $R, R' \in D$ are (k, x, y) -reachable iff there is a sequence of preferences $R = S_0, S_1, \dots, S_t = R'$ such that S_s and S_{s+1} are (k, x, y) -adjacent. If $R, R' \in D$ are (k, x, y) -reachable we write $R \simeq_{xy}^k R'$.

Remark 4.6. For all $k \in \{2, \dots, m\}$ and $x, y \in A$, the relation \simeq_{xy}^k is an equivalence relation on D . Note that if $k < l$ then the relation \simeq_{xy}^l refines \simeq_{xy}^k .

In particular, the equivalence classes of \simeq_{xy}^m are as refined as possible: they are precisely the singleton subsets of D .

Lemma 4.7. A social welfare function F on D satisfies k -IND if for any pair of alternatives $x, y \in A$ and for any agent i and pair of profiles \mathbf{R} and (\mathbf{R}_{-i}, R'_i) , if R_i and R'_i are (k, x, y) -reachable then $F(\mathbf{R}) \upharpoonright_{\{x, y\}} = f(\mathbf{R}_{-i}, R'_i) \upharpoonright_{\{x, y\}}$.

Proof. (If.) Take a social welfare function F and suppose that the right hand side of the iff holds. Take X such that $|X| = k$ and $\mathbf{R}, \mathbf{R}' \in D$ such that $\mathbf{R}|_X = \mathbf{R}'|_X$. Now, for all $i \in N$, we have $R_i \upharpoonright_X = R'_i \upharpoonright_X$; thus for all $\{x, y\} \subseteq X$, for all $i \in N$, R_i, R'_i are (k, x, y) -adjacent, thus (k, x, y) -reachable. Thus by the right hand side of the iff, for all $\{x, y\} \subseteq X$ we have $F(\mathbf{R}) \upharpoonright_{\{x, y\}} = F(\mathbf{R}') \upharpoonright_{\{x, y\}}$, so clearly $F(\mathbf{R}) \upharpoonright_X = F(\mathbf{R}') \upharpoonright_X$ as required.

(Only if.) Suppose that k -IND is satisfied. Suppose R_i, R'_i are (k, x, y) -reachable. Thus there is a list $R_i = S_0, S_1, \dots, S_t = R'_i$ with each $S_i \in D$ such that for each $s < t$, there is a $X_s \supseteq \{x, y\}$ with $|X_s| = k$ and $S_s|_{X_s} = S_{s+1}|_{X_s}$. Thus by k -IND $F(\mathbf{R}_{-i}, S_s) \upharpoonright_{X_s} = F(\mathbf{R}_{-i}, S_{s+1}) \upharpoonright_{X_s}$, so in particular $F(\mathbf{R}_{-i}, S_s) \upharpoonright_{\{x, y\}} = F(\mathbf{R}_{-i}, S_{s+1}) \upharpoonright_{\{x, y\}}$. Thus

$$F(\mathbf{R}) \upharpoonright_{\{x, y\}} = F(\mathbf{R}_{-i}, S_1) \upharpoonright_{\{x, y\}} = \dots = F(\mathbf{R}_{-i}, R'_i) \upharpoonright_{\{x, y\}}$$

as required. □

This characterisation will allow us to determine when different versions of independence imply each other. Let us formally define this generalised concept.

Definition 4.8. For $2 \leq k, l \leq m$, a domain D is (k, l) -equivalent iff (any social welfare function on D is k -IND iff it is l -IND).

The main result of this section provides a necessary and sufficient condition on domains that determines when (k, l) -equivalence holds.

Theorem 4.9. A domain D is (k, l) -equivalent iff for all $x, y \in A$, $\simeq_{xy}^k = \simeq_{xy}^l$.

Proof. (If.) Suppose F satisfies l -IND, and suppose the right hand side of Theorem 4.9 holds. We want to show that F also satisfies k -IND: take arbitrary $x, y \in A$, agent $i \in N$ and pair of profiles \mathbf{R} and (\mathbf{R}_{-i}, R'_i) such that R_i, R'_i are (k, x, y) -reachable. By the right hand side

of Theorem 4.9, these are also (l, x, y) -reachable. Thus as F satisfies l -IND, by the only if direction of Lemma 4.7 we have $F(\mathbf{R}) \upharpoonright_{\{x,y\}} = F(\mathbf{R}_{-i}, R'_i) \upharpoonright_{\{x,y\}}$ as required.

(Only if.) If $k = l$ this is trivially satisfied, so assume $k < l$. We proceed by contraposition: suppose the negation of the right hand side of Theorem 4.9. To accord with Remark 4.6, \simeq_{xy}^l must properly refine \simeq_{xy}^k . Thus there are alternatives $x, y \in A$, agent $i \in N$, and $S, S' \in D$ such that $S \simeq_{xy}^k S'$ but not $S \simeq_{xy}^l S'$. We construct an F that satisfies l -IND but violates k -IND. First, let $\mathcal{S} = \{R \in D : R \simeq_{xy}^k S\}$ and $\mathcal{S}' = D \setminus \mathcal{S}$. Define F such that it

- 1) returns some fixed ordering over all pairs $\{z, w\}$ where $z, w \neq x, y$,
- 2) sets $x[F(\mathbf{R})]z$ and $y[F(\mathbf{R})]z$ for all $z \neq x$, and
- 3) returns $x[F(\mathbf{R})]y$ iff $\forall i \in N, R_i \in \mathcal{S}$.

We claim that F satisfies l -IND but not k -IND. First let us demonstrate the violation of k -IND: for \mathbf{R} where $\forall i \in N, R_i = S$, $F(\mathbf{R}) \upharpoonright_{\{x,y\}} \neq F(\mathbf{R}_{-1}, S') \upharpoonright_{\{x,y\}}$ but S, S' are (k, x, y) -reachable. Now we show the satisfaction of l -IND. First note for all pairs except $\{x, y\}$, F returns the same ordering over these pairs, thus the condition of Lemma 4.7 is trivially satisfied. It remains to check for the pair $\{x, y\}$. Take an arbitrary agent i and pair of profiles \mathbf{R} and (\mathbf{R}_{-i}, R'_i) such that R_i and R'_i are (k, x, y) -reachable. If for all $j \in N, R_j \in \mathcal{S}$, then as $R_i \in \mathcal{S}$ and $R_i \simeq_{xy}^k R'_i$ we also have $R'_i \in \mathcal{S}$, thus $F(\mathbf{R}) \upharpoonright_{\{x,y\}} = F(\mathbf{R}_{-i}, R'_i) \upharpoonright_{\{x,y\}}$. So suppose for some $j \in N, R_j \notin \mathcal{S}$. If $j \neq i$ then we still have $R_j \notin \mathcal{S}$ within (\mathbf{R}_{-i}, R'_i) , whereas if $j = i$ then $R'_i \notin \mathcal{S}$, thus $F(\mathbf{R}) \upharpoonright_{\{x,y\}} = F(\mathbf{R}_{-i}, R'_i) \upharpoonright_{\{x,y\}}$. \square

4.4 Violating the Blau equivalence: what Blau partitions are possible?

As the name suggests, for a given domain D , (k, l) -equivalence is an equivalence relation on the integers $\{2, \dots, m\}$. Given that we can determine when the different versions of independence imply one another, we now ask: what is the structure of these implications for a given domain?

Definition 4.10. For a given domain, the Blau partition is the partition of $\{2, \dots, m\}$ determined by (k, l) -equivalence. That is, given a domain D , integers k and l with $k < l$ are in the same equivalence class of the Blau partition iff D is (k, l) -equivalent.

Definition 4.11. For $p, q \in \mathbb{N}$, an integer interval is a set $\{x \in \mathbb{N} : p \leq x \leq q\}$. An interval partition is a partition whose equivalence classes are all integer intervals.

We use the following notation for integer intervals and interval partitions. An interval partition of $\{2, \dots, m\}$ can be concisely expressed by the maxima of the members of its equiva-

lence classes. So instead of writing the full interval partition

$$\{\{2, \dots, q_1\}, \{q_1 + 1, \dots, q_2\}, \dots, \{q_{t-1} + 1, \dots, q_t\}\}$$

where $q_1 < q_2 < \dots < q_t = m$, we write $\llbracket q_1, q_2, \dots, q_t \rrbracket$.

We now show that Blau partitions do not contain any gaps, and thus must be composed of intervals.

Proposition 4.12. *Every Blau partition is an interval partition.*

Proof. Suppose that we have (k, l) -equivalence. Take p such that $k < p < l$. It suffices to show that we have (k, p) -equivalence. By Proposition 4.1, if a social welfare function is k -IND, then it is p -IND. Similarly, if the social welfare function is p -IND it is also l -IND, thus by (k, l) -equivalence it is k -IND as required. \square

For example, the Blau partition under the full domain is $\llbracket m - 1, m \rrbracket$. The Blau partition under the domain of Example 4.3 is $\llbracket 2, 4 \rrbracket$.

Remark 4.13. A domain satisfies the Blau equivalence if and only if it has Blau partition $\llbracket m - 1, m \rrbracket$.

One way of violating the Blau equivalence is if the Blau partition only contains a single equivalence class. This only occurs for extremely restricted domains.

Proposition 4.14. *The Blau partition contains one equivalence class if and only if $D \subseteq \{R, R^{\text{rev}}\}$ for some linear order R , where $xR^{\text{rev}}y$ if and only if yRx .*

Proof. (If.) For any social welfare function, the property of k -IND is trivially satisfied for all $k = 2, \dots, m$.

(Only if.) Contraposition: there must be $R, R' \in D$ such that $R \upharpoonright_{\{x,y\}} = R' \upharpoonright_{\{x,y\}}$ for some x, y . We describe a social welfare function F that satisfies m -IND but not 2-IND. Indeed, as m -IND is trivially satisfied for any social welfare function, setting $xF(\mathbf{R})y$ if and only if no agents in \mathbf{R} have the preference R suffices (arbitrarily fix the ordering over the other alternatives). \square

Of course, more interesting violations are also possible. We show in the next theorem that any interval partition that contains m as a singleton is the Blau partition under some domain. Indeed, for an arbitrary such interval partition we provide an explicit construction of the required domain.

Theorem 4.15. *For any interval partition of $\{2, \dots, m\}$ which contains $\{m\}$ as an equivalence class, there is a domain D such that the Blau partition under this domain is this interval partition.*

Proof. We use some additional notation. Recall we label the alternatives a_1, \dots, a_m . Let R^\downarrow be the linear order over A such that $a_i R a_j$ iff $i < j$. Write $R^{\downarrow k}$ for the ordering that has $a_i R^* a_j$ for $i < j, i \neq k, j \neq k$ and $a_i R^* a_k$ for all $i \neq k$. Informally, $R^{\downarrow k}$ starts with R^\downarrow and sends the k th element to the bottom, thus $R^\downarrow = R^{\downarrow m}$.

Take a set $\{q_1, \dots, q_k\} = K \subseteq \{2, \dots, m\}$. Let $D' = \{R^{\downarrow k} : k \in K\}$. We will show that the Blau partition of the domain D' is $\llbracket q_1, \dots, q_k \rrbracket$ if $m-1, m \in K$. By transitivity, it suffices to show that

- 1) if $k \in K$, then $(k, k+1)$ -equivalence does not hold, and
- 2) if $k \notin K$, then $(k, k+1)$ -equivalence holds.

We show these one after the other.

- 1) Here $R^{\downarrow k} \in D'$. Consider $a_1 R^* a_2 \dots P a_k \in \mathcal{W}(\{a_1, \dots, a_k\})$. Both R^\downarrow and $R^{\downarrow k}$ are extensions of this within D' , which implies that $R^\downarrow \simeq_{a_1 a_k}^l R^{\downarrow k}$. Now, all preferences $R' \neq R^{\downarrow k}$ in the domain have $a_k R' a_i$ for the $m-k$ alternatives a_i where $i > k$. This implies that no linear order in the domain coincides with $R^{\downarrow k}$ on a superset of $\{a_{k-1}, a_k\}$ of size $k+1$. This implies that $R^\downarrow \not\simeq_{a_1 a_k}^{k+1} R^{\downarrow k}$. As $\simeq_{a_1 a_k}^k \neq \simeq_{a_1 a_k}^{k+1}$, by Theorem 4.9 we do not have $(k, k+1)$ -equivalence.
- 2) Here $R^{\downarrow k} \notin D'$. We first show the following lemma.

Lemma 4.16. *Take integers k', k'' such that $2 \leq k' < k'' \leq m$. Suppose for any $S, S' \in D$ and set $X \subset A$ such that $S \upharpoonright_X = S' \upharpoonright_X$ and $|X| = k'$, there are $Y, Y' \supset X$ and $S'' \in D$ such that $S \upharpoonright_Y = S'' \upharpoonright_Y$, $S'' \upharpoonright_{Y'} = S' \upharpoonright_{Y'}$ and $|Y| = |Y'| = k''$. Then we have (k', k'') -equivalence.*

Proof of Lemma 4.16. To see this, take arbitrary $\{x, y\} \subset A$ and $R, R' \in D$ such that $R \simeq_{xy}^{k'} R'$. We want to show that $R \simeq_{xy}^{k''} R'$. So, suppose there is a list $R = S_1, S_2, \dots, S_t = R'$ such that for each $i = 1, \dots, t-1$ there is a set $X_i \supset \{x, y\}$, with $|X_i| = k'$, and such that $S_i \upharpoonright_{X_i} = S_{i+1} \upharpoonright_{X_i}$. We want to find a list $R = S'_1, S'_2, \dots, S'_t = R'$ such that for each $i = 1, \dots, t-1$ there is a set $Y_i \supset \{x, y\}$, with $|Y_i| = k''$, and such that $S'_i \upharpoonright_{Y_i} = S'_{i+1} \upharpoonright_{Y_i}$. We expand the first list, adding in a single new linear order between each adjacent pair. First, for $1 \leq i \leq t$, set $S_i = S'_{2i-1}$, with in particular $S_t = S'_{2t-1} = S'_t$. Then, for $1 \leq i < t$, take the two linear orders required by the condition as $S = S_i$ and $S' = S_{i+1}$, and define each $S'_{2i} = S''$. ■

We now show that the condition of Lemma 4.16 is satisfied for $k' = k$ and $k'' = k+1$. Consider an arbitrary set $X = \{a_{s_1}, \dots, a_{s_k}\}$, with $s_i > s_j$ if $i > j$ and $|X| = k$, and an arbitrary ordering S over this set.

Suppose $S \neq R^\downarrow \upharpoonright_X$, then there is at most one extension of S in D^* , so there is nothing to prove. So suppose $S = R^\downarrow \upharpoonright_X$, and consider R and R' in the domain that extend S . Note these must be of the form $R^{\downarrow p}$ and $R^{\downarrow q}$ for $p, q \neq s_1, s_2, \dots, s_{k-1}$, as otherwise $S \neq R^\downarrow \upharpoonright_X$. Without loss of generality suppose $p < q$.

Case 1. If $p, q \neq s_k$, $R^{\downarrow p} \upharpoonright_{X \cup \{q\}} = R^\downarrow \upharpoonright_{X \cup \{q\}}$ and $R^{\downarrow q} \upharpoonright_{X \cup \{p\}} = R^\downarrow \upharpoonright_{X \cup \{p\}}$ as required.

Case 2. Suppose $p = s_k$.

Subcase a. Suppose $s_k < m$. Then there is $r \notin K$ such that $r > s_k$. As $s_k > k$, there is some $r' \notin K$ such that $r' < s_k$. We have $R^{\downarrow p}|_{X \cup \{r'\}} = R^{\downarrow}|_{X \cup \{r'\}}$ and $R^{\downarrow q}|_{X \cup \{r\}} = R^{\downarrow}|_{X \cup \{r\}}$ as required.

Subcase b. Suppose $s_k = m$. As $m - 1, m \in K$, $k < m - 1$, thus there are distinct $r, r' \notin K$ such that $r, r' < s_k$. Without loss of generality suppose $r \neq p$ and $r' \neq q$, we then have $R^{\downarrow p}|_{X \cup \{r\}} = R^{\downarrow}|_{X \cup \{r\}}$ and $R^{\downarrow q}|_{X \cup \{r'\}} = R^{\downarrow}|_{X \cup \{r'\}}$ as required. \square

4.5 Blau's impossibility: do k -IND social welfare functions exist on Arrovian domains?

We now have a better idea of when l -IND implies k -IND for $l > k$. The fact that for some domains the Blau equivalence fails opens up the possibility that on a domain where the Arrovian impossibility holds, i.e. a domain where 2-IND and the Pareto criterion imply dictatorship, we may nonetheless have a social welfare function that satisfies k -IND, the Pareto criterion and non-dictatorship with $2 < k < m$. This may be thought of as a *successful* weakening of independence; such a weakening overcomes what we call Blau's impossibility.

In fact, we have already seen a somewhat trivial successful weakening of independence. For the domain of Example 4.3 ranking by Borda scores satisfies 3-IND, but as every triple is free this domain is subject to Arrow's impossibility. This example is trivial in that 3 belongs to the "top" equivalence class in the Blau partition, thus any social welfare function on this domain satisfies 3-IND. We now describe an example of a social welfare function that is not trivial in this sense, defined on the union of single-peaked and single-dipped domains. Of course, single-peakedness is well known as an escape from the Arrovian impossibility.³

Definition 4.17. A domain is **single-peaked** if there is a linear order \succ on A , called the **axis**, such that $(x \succ y \succ z \text{ or } z \succ y \succ x)$ implies $(xRy \text{ implies } yRz)$.

Note that single-peakedness, considered as a property, is satisfied by multiple domains, however reference is often made to "the single-peaked domain", which could have multiple interpretations: for a fixed axis \succ there is a unique largest (by cardinality and inclusion) single-peaked domain that contains all other single-peaked domains with the same axis.⁴ In fact, this version of the largest single-peaked domain cannot be expressed within our current framework, because it is not a Cartesian power domain.

³ Black (1948) provides the classic reference concerning single peaked domains.

⁴Another (incompatible) interpretation of the largest single-peaked domain is the domain that consists of those profiles for which there is *some* axis along which the candidates are single-peaked.

Definition 4.18. A **maximal single-peaked domain** is a largest single-peaked domain with respect to some given axis.

Maximal single-peaked domains are examples of domains that satisfy the Blau equivalence but where the Arrovian impossibility does not apply.

Proposition 4.19. For $m > 2$, any maximal single-peaked domain satisfies the Blau equivalence, i.e. it has Blau partition $\llbracket m - 1, m \rrbracket$.

Proof. As for $m > 2$ a maximal single-peaked domain has more than two elements, by Proposition 4.14 we know that there is not one single partition. Thus it suffices to show $(l, l + 1)$ -equivalence for $l = 2, \dots, m - 2$. We proceed by showing that the condition of Lemma 4.16 holds. Thus suppose that $S \upharpoonright_X = S' \upharpoonright_X$ for $S, S' \in D$ and $X \subset A$ with $|X| = l$. Note as $|X| = l \leq m - 2$ there are $a, b \in A \setminus X$. For $x \in A$, $Y \subseteq A$ and $R \in D$, define

$$\mu(x, Y, R) = |\{y \in Y : x \succ y, yRx\}| - |\{y \in Y : y \succ x, yRx\}|.$$

This may be thought of as a measure of “how far anti-clockwise” x is in R with respect to some set of alternatives Y . Now, there is a $S'' \in D$ such that $S'' \upharpoonright_X = S \upharpoonright_X$, $\mu(a, X, S) = \mu(a, X, S'')$ and $\mu(b, X, S') = \mu(b, X, S'')$. Then, as required, $S \upharpoonright_{X \cup \{a\}} = S'' \upharpoonright_{X \cup \{a\}}$ and $S' \upharpoonright_{X \cup \{b\}} = S'' \upharpoonright_{X \cup \{b\}}$. \square

By removing preferences from a maximal single-peaked domain we can create a domain to which neither the Blau equivalence nor the Arrovian impossibility applies. For instance, consider the domain D_p displayed in Table 4.2.⁵

Table 4.2: A single-peaked domain.

R^{p1}	R^{p2}	R^{p3}	R^{p4}	R^{p5}	R^{p6}	R^{p7}	R^{p8}
a_1	a_2	a_2	a_3	a_3	a_4	a_4	a_5
a_2	a_1	a_3	a_2	a_4	a_3	a_5	a_4
a_3	a_3	a_1	a_4	a_2	a_5	a_3	a_3
a_4	a_4	a_4	a_1	a_5	a_2	a_2	a_2
a_5	a_5	a_5	a_5	a_1	a_1	a_1	a_1

Let us draw attention to the pair (a_1, a_2) . Because R^{p2} and R^{p3} are $(4, a_1, a_2)$ -adjacent,⁶

⁵ Preferences were not removed arbitrarily to create this domain. For the remaining preferences, note that the alternatives on either side of the “peak” are balanced; that these alternatives are interspersed as far as possible. Compare to the idea of **equidistantly single-peaked** domains described by Ozdemir and Sanver (2007).

⁶The full list of pairs of rankings that are $(4, a_1, a_2)$ -adjacent is (R^{p2}, R^{p3}) , (R^{p6}, R^{p7}) and (R^{p7}, R^{p8}) . See Figure 4.1.

therefore

(4.i) $\{5\}$ is a member of the Blau partition of D_p .

However, none of R^{p6} , R^{p7} or R^{p8} are $(4, a_1, a_2)$ -adjacent to R^{p4} , as the rankings, firstly, of a_1 and a_5 and, secondly, of a_2 and a_4 are different in each pair of preferences. It can similarly be seen that R^{p5} is not $(4, a_1, a_2)$ -adjacent to any of the other rankings in D_p . Therefore it is not $(4, a_1, a_2)$ -reachable to any other ranking, although it is certainly $(2, a_1, a_2)$ -reachable to other rankings. Thus 2 and 4 are in different sets of the Blau partition for this domain, so

(4.ii) D_p does not satisfy the Blau equivalence.

The inverse of a single-peaked domain is a single-dipped domain. Such a single dipped domain, D_d , is displayed in Table 4.3.

Table 4.3: A single-dipped domain.

R^{d1}	R^{d2}	R^{d3}	R^{d4}	R^{d5}	R^{d6}	R^{d7}	R^{d8}
a_1	a_1	a_1	a_1	a_5	a_5	a_5	a_5
a_2	a_2	a_2	a_5	a_1	a_4	a_4	a_4
a_3	a_3	a_5	a_2	a_4	a_1	a_3	a_3
a_4	a_5	a_3	a_4	a_2	a_3	a_1	a_2
a_5	a_4	a_4	a_3	a_3	a_2	a_2	a_1

We now consider the union domain $D_u = D_p \cup D_d$. This satisfies the free triple condition and thus is subject to the Arrovian impossibility: there is no 2-IND and non-dictatorial social welfare function that satisfies the Pareto criterion defined on D_u . As the only rankings in the domain with a_2 preferred to a_1 are in D_p (because $R^{d8} = R^{p8}$), the statements (4.i) and (4.ii) equally apply if we substitute D_u for D_p . This opens up the possibility for a successful weakening of independence. Define $D_L = \{D^{pi}, D^{di} : i = 1, 2, 3, 4\}$ and $D_R = \{D^{pi}, D^{di} : i = 5, 6, 7, 8\}$.

Definition 4.20. Selective pairwise majority is the function on D_u that proceeds as follows: divide the voters into those who have preferences in D_L and those who have preferences in D_R , then perform *pairwise majority*⁷ on the larger of these two groups. Break ties arbitrarily in a consistent and deterministic manner.

Theorem 4.21. *Selective pairwise majority is non-dictatorial social welfare function on D_u that satisfies the Pareto criterion and 4-IND.*

⁷Pairwise majority determines the ranking over pairs of alternatives based upon which is preferred by more voters.

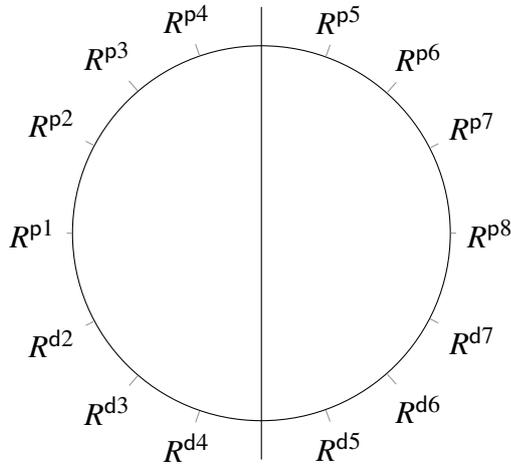


Figure 4.2: The alternatives in the domain that is the union of the single-peaked domain of Table 4.2 and the single-dipped domain of Table 4.3 domain arranged on a circle. The single-crossing line for D_L is the semicircle to the left; that for D_R is the semicircle to the right.

Proof. First note that D_L and D_R are each **single-crossing domains**⁸ (see Figure 4.2), thus this function does actually return a complete preorder as required. As pairwise majority on D_L —or indeed on D_R —is non-dictatorial, so to is selective pairwise majority. If all agents prefer x to y , then a majority also do so for whichever of D_L or D_R has more voters; therefore selective pairwise majority satisfies the Pareto criterion.

To finish we show that selective pairwise majority is 4-IND. Take arbitrary $R \in D_L$ and $R' \in D_R$. Note that $a_1 R a_5$ but $a_5 R' a_1$; and also $a_2 R a_4$ but $a_4 R' a_2$. As there are only five alternatives, R and R' never coincide on a set of four alternatives. Thus R and R' are not (k, x, y) -adjacent for any pair of alternatives x, y . As R and R' were chosen arbitrarily, they are not (k, x, y) -reachable either.⁹ So for any R'' such that $R \simeq_{xy}^4 R''$, $R'' \in D_L$, by the definition of the social welfare function substituting this R'' for R in a profile will not change the outcome. The same point applies to R' . Thus by Lemma 4.7 we have 4-IND. \square

⁸ A domain is single-crossing if its preferences can be listed R^1, R^2, \dots, R^t —or placed on a line—such that, for all x and y , if $x R^1 y$ and $y R^s x$, then $x R^t y$ for $s \leq i \leq t$. Gans and Smart (1996) provide some economic applications of this property. Rothstein (1991) shows that for any profile on a single-crossing domain there is a *representative voter* whose (strict) preferences coincide with the (strict) majority relation (though note that he does not use the term single-crossing).

⁹For all alternatives $x \neq y$, the set of equivalence classes of \simeq_{xy}^4 on D_u is the following refinement of $\{D_L, D_R\}$: $\{\{R^{p1}, R^{p2}, R^{p3}, R^{d2}, R^{p3}\}, \{R^{p6}, R^{p7}, R^{p8}, R^{d6}, R^{p7}\}, \{R^{p4}\}, \{R^{p5}\}, \{R^{d4}\}, \{R^{d5}\}\}$.

4.6 Final remarks on domain restrictions and independence of irrelevant alternatives

We have focused on the Blau (1971) equivalence, which concerns the impossibility of weakening binary independence by considering independence over sets of higher cardinality. We started by observing that this equivalence may vanish under domain restrictions, i.e., for some domains k -IND diverges from 2-IND for $2 < k < m$. Next, we provided, for any given domain D and any given value of k , a necessary and sufficient condition for a social welfare function to satisfy k -IND. We used this result to identify a necessary and sufficient condition for a domain to be D (k, l) -equivalent—a result which enables the determination of cases where the Blau equivalence holds. The Blau equivalence itself was defined in terms of Blau partitions; we also saw that for almost any such partition there is a domain that exemplifies it. Finally, in the section preceding this one we defined a particular five alternative domain that is subject to the Arrovian impossibility but violates the Blau equivalence, and a non-dictatorial social welfare function on this domain that satisfies the Pareto criterion and 4-IND.

What results can we draw from our work? There is no logical dependence between the Blau equivalence and the Arrovian impossibility. We have seen domains where both hold (the full domain), where just the Blau equivalence holds (any maximal single-peaked domain), where just the Arrovian impossibility holds (the domain of Example 4.3 and the domain D_u in Section 4.5), and also where neither hold (the domains D_p and D_d in Section 4.5).

Let us restrict our attention to domains that are subject to Arrow's impossibility. Here the Blau equivalence implies that the impossibility holds even for weaker versions of independence. Hence, when the Blau equivalence fails, a potential escape from Arrow's impossibility arises: there may be domains where Arrow's result holds, i.e. where 2-IND, the Pareto criterion and non-dictatoriality are logically incompatible; but also where there exist non-dictatorial social welfare functions that satisfy the Pareto criterion and some version of k -IND.

As was already noted by Blau (1971), it is not interesting to weaken independence to m -IND, because this is trivially satisfied by any social welfare function. For example, on any domain, ranking by Borda scores is non-dictatorial and satisfies m -IND and the Pareto criterion, though this escape from Arrow's impossibility is somewhat unsatisfying. We consider this social welfare function to be similarly unsatisfying on the domain of Example 4.3. Although on this four alternative domain ranking by Borda scores satisfies 3-IND, so too will any other social welfare function: 3 and 4 are in the same equivalence class of the Blau partition of this domain. Blau's impossibility, when elaborated in line with the above strand of thought, may be properly stated: there is no non-dictatorial and k -IND social welfare

function that satisfies the Pareto criterion, for any k not in the equivalence class of the Blau partition that contains m . Thus Blau's impossibility holds on domains that are subject to Arrow's impossibility where there are only two equivalence classes in the Blau partition, though Section 4.4 shows that there are many domains with more than two such equivalence classes.

Section 4.5 gives an example of a *non-trivial* weakening of independence: a non-dictatorial social welfare function defined on a five alternative domain that is subject to Arrow's impossibility that satisfies the Pareto criterion and 4-IND, where 4 is in a different equivalence class to 5 in the Blau partition. Thus Blau's impossibility is extensionally a stronger result than Arrow's impossibility; it applies to fewer domains. Of course, the domain and social welfare function we describe are (somewhat) designed to provide the necessary example. However, it is interesting to note that during their construction we started with single-peaked and single-dipped domains and ended up with a single-crossing condition. These are all Condorcet domains, where the pairwise majority relation is transitive, but we see no particular reason why successful weakenings must necessarily involve domains of this type. We have found at least one other (less simply expressed) example on the domain D_u ; we conjecture that there are many different types of successful weakenings.

Chapter 5

Combining ordinal and evaluative information in social choice

The traditional approach in social choice theory involves aggregating ordinal preferences over alternatives. We¹ will refer to this as the **ranking approach**. This can be contrasted with what we call the **evaluative approach**, which involves aggregating evaluations of the alternatives made by the agents. For example, approval voting, as pioneered by Brams and Fishburn (1978), is evaluative; it affords two possible evaluations for each of the alternatives: approved or disapproved. In this chapter we prove an impossibility in combining the ranking and evaluative approaches. Our result implies that there is an incompatibility between the two.

Imagine that a committee is tasked with selecting a candidate for a position that does not need to be filled. The position may be offered to the best candidate. If this candidate rejects the offer, the position can then be offered to the second best candidate, and so on until it would be preferable to leave the position unfilled rather than offer it to the next best available candidate. The task can be divided into two: creating a ranking over the candidates—the ordinal part; and determining which candidates are suitable for the position—the evaluative part. It is these two different parts that are the subject of this chapter.

In the above task, candidates are individually evaluated as being suitable, or not, for a position. A less interpreted presentation of the evaluation occurs in the preference-approval framework—originally conceived of by Brams and Sanver (2009), and investigated further by Erdamar et al. (2017). The smallest non-trivial application of this framework involves two alternatives, a and b , which are ranked with respect to each other and also individually evaluated as approved or disapproved. There are natural restrictions upon how rankings and evaluations should be combined. For example, if a is approved and b is disapproved, it would be unreasonable to also rank the two alternatives as indifferent. A consistent combination of ranking and approvals is known as a **preference-approval**.

A preference-approval can be represented as a complete preorder over the alternatives that is also equipped with a **zero-line** to demarcate the approval level. For the two alternatives a and b there are eight possible preference-approvals. These are displayed in Figure 5.1.

¹This chapter is based upon joint work with Remzi Sanver

<i>a</i>	<i>a</i>	<i>b</i>	<i>b</i>	<i>ab</i>			
<i>b</i>	<i>b</i>	<i>a</i>	<i>a</i>	<i>b</i>	<i>a</i>	<i>ab</i>	
1 st	2 nd	3 rd	4 th	5 th	6 th	7 th	8 th

Figure 5.1: The possible preference approvals for two alternatives. An alternative is strictly preferred to the other if and only if it is higher on the page, and alternatives are approved if and only if they are above the horizontal line.

The following example of a collection of preference-approvals may set off alarm bells.

Example 5.1. Suppose that there are only three agents and that they have, respectively, the first, third and fifth preference-approvals displayed in Figure 5.1. Amongst these agents *a* is ranked above *b* by a majority—if we aggregate the ranking according to the majority method of May (1952), then *a* will be ranked above *b* in the social ordering. At the same time, *a* is disapproved by a majority and *b* is approved by a majority—if we aggregate the evaluations by majority, thereby complying with Sanver’s (2010) axiom of majoritarian approval, *b* should be uniquely approved in the output. Thus applying majority on the decomposed components of these preference-approvals results in an inconsistent outcome.

Example 5.1 effectively recreates the Condorcet paradox with the zero-line as an implicit third alternative. We extend the example into a more general result, mirroring the connection between Condorcet’s paradox and Arrow’s (1950) impossibility theorem. So the Condorcet-like example presages a more general phenomenon, which we prove as an impossibility theorem. Our result states that if aggregation of preference-approvals is performed in a decomposed manner and the approval aggregation part satisfies unanimity, then the only possibility is a dictatorship. Unanimity is a weak condition; the important condition is that of decomposability, which takes the place of the well-known independence condition of Arrow (1950): decomposability requires that the social evaluation for each alternative only depends upon the individuals’ evaluations for that alternative, and that the social ranking between the alternatives only depends upon the individuals’ rankings over the alternatives. Decomposability can be weakened in a variety of ways. For example, we identify various possibilities when decomposability is weakened within the evaluative part, but the ordinal and evaluative parts are kept separate—these possibilities are stated as Propositions 5.6, 5.7 and 5.8. We also identify four specific possibilities that arise when the decomposability between the ordinal and evaluative parts is relaxed.

To justify the importance of decomposability, let us expand upon its analogy to Arrow’s ax-

iom of independence. To determine social ranking between two alternatives, one can very well use information about how agents compare these to a third alternative. In fact, we know by Arrow's Theorem that this must be the case (if one wants a sensible ranking aggregation). One may like or dislike Arrow's axiom as a principle of ranking aggregation, but it is important to know that ranking aggregation is not decomposable into pairs. In a similar vein, one may like or dislike decomposability between ranking and evaluation aggregation as a principle, but it is worth knowing that such a decomposition is not possible.

To give a brief overview of the rest of the chapter: in Section 5.1 we define the basic model of preference-approvals. We show our central impossibility concerning the aggregation of preference-approvals in Section 5.2. Possibilities arise if one does not require that preference-approvals are aggregated in a decomposed manner, which we explore in Section 5.3. We prove an extended version of the impossibility, applicable for more than two evaluation levels, in Section 5.4. Final remarks are provided in Section 5.5.

5.1 The preference-approval model

In this chapter agents not only express preferences in terms of complete preorders (see the introduction for definitions) but also in terms of sets of alternatives: we write $B_i \in 2^A$ for the set of alternatives approved by $i \in N$.

Sets of alternatives form part of the definition of preference-approvals, which were first conceived by Brams and Sanver (2009). A **preference-approval** of agent $i \in N$ is a pair $p_i = (R_i, B_i) \in \mathcal{W}(A) \times 2^A$, composed of a complete preorder and set of approved alternatives, with the extra requirement that

$$(5.i) \quad (\forall x, y \in A) ((xR_i y \text{ and } y \in B_i) \Rightarrow x \in B_i).$$

We write Π for the set of all preference-approvals.

Example 5.2. We have already seen one example of a possible interpretation of a preference-approval, that we here restate alongside two other interpretations.

- 1) Alternatives are candidates for a position that need not be filled, and the zero-line represents the point at which it is preferred that the position be left empty.
- 2) Alternatives are time-slots at which a meeting may be held. Attendees have preferences over dates, and are unavailable at those dates that are below the line.
- 3) Alternatives are possible venues for an event, but transport restrictions only allow attendees to come to certain of the venues.

Arguably, item (3) in Example 5.2 may not belong to the preference-approval framework: an agent may well recognise that a venue that they cannot attend is nonetheless a better venue

than one that they can attend. Our analysis requires that the evaluations and rankings cannot be divorced in this manner. Another possibility that we do not account for is indifference in ranking between two alternatives that are evaluated into different categories.²

A profile of preference approvals is some $\mathbf{p} = (p_1, p_2, \dots, p_n) \in \Pi^N$. As a general rule we write vectors in boldface.

A **preference-approval aggregator** is a function $\pi: D^N \rightarrow \Pi$, where $D \subseteq \Pi$. A preference-approval aggregator π is **dictatorial** if there is an agent d whose strict preference and approval line are reproduced in the output preference-approval; i.e. $d \in N$ is a dictator if for all $\mathbf{p} = ((R_1, B_1), \dots, (R_n, B_n))$ we have $\pi(\mathbf{p}) = (R, B_d)$ where R is some ranking such that xR^*y if xR_d^*y .

Example 5.3. There are somewhat natural examples of preference-approval aggregators for each of our example interpretations of the preference approval framework.

- 1) Shortlist the candidates using the evaluations, and then perform your preferred ordinal aggregation on the remaining candidates.
- 2) Supposing that we want as many people as possible to attend the meeting: approve only those slots where a maximum of attendees can attend. For the remaining dates, rank first in terms of weak dominance, then in terms of time (earlier dates preferred to later dates).
- 3) Chose amongst the venues that most people can attend, but also take into account the ranking of those who cannot even attend the chosen venue.

The issues raised by the latter two examples are interesting, but will not be treated further in this chapter.

In this chapter we consider **social welfare functions** with restricted domains; functions $F: D^N \rightarrow \mathcal{W}(A)$, where $D \subseteq \mathcal{W}(A)$. We write $F^*(\mathbf{R})$ for the strict part of $F(\mathbf{R})$. We say that f satisfies the **Pareto criterion** if, for all $\mathbf{R} \in \mathcal{W}(A)$, xP_iy for all $i \in N$ implies $xF^*(\mathbf{R})y$. We say that F is **dictatorial** if there is a $d \in N$ such that for all $\mathbf{R} \in \mathcal{W}(A)$, xR_d^*y implies $xF^*(\mathbf{R})y$.

We also consider aggregation in the approval voting tradition. For $x \in A$, an **elementary approval aggregator** α_x is a function $\alpha_x: \{0, 1\}^N \rightarrow \{0, 1\}$. We refer to the elementary approval aggregator of a subscripted alternative, such as a_i , directly with the subscript, i.e. $\alpha_i = \alpha_{a_i}$. An elementary approval aggregator should be interpreted as a map from a profile where each agent either approves x (assigns it 1) or disapproves x (assigns it 0) into a situation where either x is globally approved (the result is 1) or x is globally disapproved (the result is 0). Such a function satisfies **unanimity** if $\alpha_x(0, \dots, 0) = 0$ and $\alpha_x(1, \dots, 1) = 1$. An

²Only allowing for this weakening could be seen as analogous to moving from total preorders to quasi-orders.

approval aggregator α is a function $\alpha: (2^A)^N \rightarrow 2^A$. We say that α satisfies **(alternative-wise) unanimity** if, for any $x \in A$, (1) if $x \in B_i$ for all $i \in N$ then $x \in \alpha(\mathbf{B})$ and (2) if $x \in B_i$ for no $i \in N$ then $x \notin \alpha(\mathbf{B})$.³

Our final definitions concern splitting up various aggregators into sub-functions. An approval aggregator is **decomposable** if, for each alternative $x \in A$ there is an elementary approval aggregator α_x such that:

$$x \in \alpha(\mathbf{B}) \Leftrightarrow \alpha_x(1_{B_1}(x), \dots, 1_{B_n}(x)) = 1$$

where the indicator function of $X \subseteq A$, is the function $1_X: A \rightarrow \{0, 1\}$ defined by

$$1_X(x) = \begin{cases} 1 & \text{if } x \in X \\ 0 & \text{otherwise.} \end{cases}$$

In such a case we write $\alpha = (\alpha_x)_{x \in A}$. Similarly, π can be decomposed into (F, α) if

$$\pi((R_1, B_1), \dots, (R_n, B_n)) = (F(R_1, \dots, R_n), \alpha(B_1, \dots, B_n));$$

and further into $(F, (\alpha_x)_{x \in A})$ if $\alpha = (\alpha_x)_{x \in A}$. In such cases we write, respectively, $\pi = (F, \alpha)$ and $\pi = (F, (\alpha_x)_{x \in A})$.

5.2 An impossibility in aggregating preference-approvals

Theorem 5.4. *If $D = \Pi$ then any $\pi = (F, \alpha_1, \dots, \alpha_m)$ such that F satisfies the Pareto criterion and α_x satisfies unanimity for all $x \in A$ is dictatorial.*

Proof. Take a preference-approval aggregator $\pi = (F, \alpha_1, \dots, \alpha_m): \Pi^N \rightarrow \Pi$ such that π satisfies the Pareto criterion and each α_i satisfies unanimity. To show that π is dictatorial, we define a social welfare function \bar{F} over an extended set of alternatives $A \cup \{z\} = X$.

The function \bar{F} corresponds to π , while the alternative z represents the zero-line. Using two results from the literature on social welfare functions, we show that \bar{F} is dictatorial, which in turn implies that π is dictatorial

We say that a social welfare function $g: D^N \rightarrow \mathcal{W}(X)$ satisfies **binary independence** if for every pair $\{x, y\} \subseteq X$, for all $\mathbf{R}, \mathbf{R}' \in D^N$, if $\mathbf{R} \upharpoonright_{\{x, y\}} = \mathbf{R}' \upharpoonright_{\{x, y\}}$ then $g(\mathbf{R}) \upharpoonright_{\{x, y\}} = g(\mathbf{R}') \upharpoonright_{\{x, y\}}$. Note that in order for binary independence to not be trivially satisfied it is necessary

³ The term ‘‘alternative-wise’’ indicates a distinction from a (weaker) version of unanimity which requires that if all agents have exactly the same evaluations for *all* the alternatives, then the output evaluations are identical. We do not consider the weaker version of unanimity in this chapter.

that $|X| > 2$. We say a domain is **Arrovian** if any social welfare function that satisfies the Pareto criterion and binary independence must also be dictatorial. Kelly (1994) gives some conditions that imply that a domain is Arrovian.

Given a preference-approval $p = (R, B)$ define a complete preorder $\bar{p} \in \mathcal{W}(B \cup \{z\})$ by, for $\{a, b\} \subseteq B$, $a\bar{p}b \Leftrightarrow aRb$, $b\bar{p}a \Leftrightarrow bRa$, $a\bar{p}z \Leftrightarrow a \in B$, and $z\bar{p}a \Leftrightarrow a \notin B$. It can be verified that \bar{p} is complete and transitive. Write \bar{p}^* for the strict part of \bar{p} ; the above definitions imply that, for any $\{a, b\} \subseteq A$,

$$(5.ii) \quad aPb \Leftrightarrow a\bar{p}^*b$$

and

$$(5.iii) \quad a \in A \Leftrightarrow a\bar{p}^*z.$$

Define

$$D = \{\bar{p} : p \in \Pi\} \subset \mathcal{W}(A \cup \{z\}) .$$

By definition $p \mapsto \bar{p}$ is a bijection from Π to D . This means there is also a bijection from preference approval profiles $\mathbf{p} \in \Pi^N$ to $\bar{\mathbf{p}} = (\bar{p}_i)_{i \in N} \in D^N$. The domain D^N is a (Cartesian) product domain that contains all possible linear orders and yet does not contain the complete preorder where all pairs of alternatives are considered indifferent. Thus this domain is Arrovian according to Kelly's (1994) "Theorem 1".

Define the social welfare function $\bar{F} : D^N \rightarrow \mathcal{W}(A \cup \{z\})$ as, for an arbitrary $\bar{\mathbf{p}} \in D^N$,

$$(5.iv) \quad \bar{F}(\bar{\mathbf{p}}) = \overline{\pi(\mathbf{p})} .$$

We now show that \bar{F} satisfies the Pareto criterion and satisfies binary independence.

For the Pareto criterion, take an arbitrary $\bar{\mathbf{p}}$. We split into two cases. First we consider pairs $\{a, b\} \subseteq A$. Suppose for all $i \in N$ $a\bar{p}_i^*b$. The result follows directly from the fact that F satisfies the Pareto criterion. Second, we consider pairs $\{a, z\}$ with $a \in A$. As a first subcase, suppose for all $i \in N$, $a\bar{p}_i^*z$. This implies that all agents approve a in \mathbf{p} . By unanimity of α_a , a is approved in $\pi(\mathbf{p})$, thus $a\bar{F}(\bar{\mathbf{p}})z$ as required. For the subcase where for all $i \in N$ we have $z\bar{p}_i^*a$, replace approve with disapprove in the preceding sentence.

For the proof of binary independence, we utilise results due to Campbell and Kelly (2000), who establish the notion of **relevant sets**. A relevant set is defined with respect to social welfare function $g : D^N \rightarrow \mathcal{W}(X)$, where $D \subseteq \mathcal{W}(X)$ and g satisfies the Pareto criterion. A relevant set for a pair $\{x, y\} \subseteq X$ is the minimal (by inclusion) set of alternatives $Y \supseteq \{x, y\}$, $Y \subseteq X$, such that for all profiles $\mathbf{R}, \mathbf{R}' \in D \subseteq \mathcal{W}(X)$, if $\mathbf{R} \upharpoonright_Y = \mathbf{R}' \upharpoonright_Y$ then $g(\mathbf{R}) \upharpoonright_{\{x, y\}} = g(\mathbf{R}') \upharpoonright_{\{x, y\}}$. Campbell and Kelly's "Theorem 2 part (I)", paraphrased, states: for any social

welfare function g that satisfies the Pareto criterion, and distinct alternatives x, y, w , if the relevant set of $\{x, y\}$ is $\{x, y\}$ then the relevant set of $\{y, w\}$ either (1) contains x or (2) is identical to $\{y, w\}$.

We know already that \bar{F} satisfies the Pareto criterion. Take arbitrary $\{a, b\} \subseteq A$. By the decomposability of π , (1) the relevant set of $\{z, a\}$ is $\{z, a\}$ and (2) the relevant set of $\{a, b\}$ is a subset of $A \setminus z$. Because the relevant set of $\{a, b\}$ does not contain z , by ‘‘Theorem 2 part (I)’’, the relevant set of $\{a, b\}$ must be $\{a, b\}$. Thus we have binary independence.

Altogether this implies that

(5.v) \bar{F} is dictatorial, with a dictator d .

We claim that d must also be a dictator for π : take an arbitrary profile $\mathbf{p} = (\mathbf{R}, \mathbf{B})$ and write $\pi(\mathbf{p}) = (R, B)$.⁴ The definitions of \bar{p} and \bar{F} and the dictatoriality of \bar{F} lead to the following chains of implications, for $a, b \in A$.

(5.vi)

$$aP_d b \stackrel{(5.ii)}{\Leftrightarrow} a\bar{p}_d^* b \stackrel{(5.v)}{\Leftrightarrow} a\bar{F}^*(\bar{\mathbf{p}})b \stackrel{(5.iv)}{\Leftrightarrow} a\pi(\mathbf{p})^* b \stackrel{(5.ii)}{\Leftrightarrow} aPb$$

(5.vii)

$$a \in B_d \stackrel{(5.iii)}{\Leftrightarrow} a\bar{p}_d^* z \stackrel{(5.v)}{\Leftrightarrow} a\bar{F}^*(\bar{\mathbf{p}})z \stackrel{(5.iv)}{\Leftrightarrow} a\pi(\mathbf{p})^* z \stackrel{(5.iii)}{\Leftrightarrow} a \in B$$

Thus the preference-approval aggregator π is dictatorial. \square

There are two subtle requirements in this proof, concerning properties of the created domain and the informational restrictions of the created social welfare function. Concerning the first, the presence of all linear orders and the absence of the preference of complete indifference is sufficient for a domain to be Arrovian. This means that for a possibility based on a restricted domain it will be necessary to remove at least one preference-approval that corresponds to a linear order. Concerning the second, it was crucial that the relevant set of an alternative and the zero-line alternative was exactly that pair: if we do not decompose the aggregation of approvals, then this need no longer be the case, and the theorem of Campbell and Kelly (2000) no longer applies.

The impossibility vanishes for certain interesting restricted domains of preference-approvals. Denote by Π^- the domain where all alternatives below the line must be considered indifferent. For $A = \{a, b\}$, Π^- contains exactly the six preference-approvals displayed in Figure 5.2.

⁴Note that to ease presentation, we transpose the Cartesian product without comment, i.e. when we write (\mathbf{R}, \mathbf{B}) , where $\mathbf{R} = (R_1, \dots, R_n)$ and $\mathbf{B} = (B_1, \dots, B_n)$, we are technically referring to the profile $((R_1, B_1), (R_2, B_2), \dots, (R_n, B_n))$.

a	a	b	b	ab
b	b	a	a	ab

Figure 5.2: The six preference-approvals in Π^- .

The restriction to the domain Π^- is reasonable for many applications of the preference-approval framework: in particular, for the interpretation of alternatives as time-slots for a meeting, as in Example 5.2 (2), it seems reasonable to assume that time-slots that cannot be attended are considered indifferent. However, for the selection of candidates for a position that may be left empty as in Example 5.2 (1), although it seems reasonable to consider candidates who are considered unsuitable as indifferent in the *output*, in the *profile* it may be useful information to know one committee member ranks two candidates that she considers unsuitable.

Let α_x^{ud} be the “unanimous disapproval rule” that disapproves x iff x is disapproved by all agents in N . Note that this elementary approval aggregator satisfies unanimity. We strengthen the Pareto criterion to the **strong Pareto criterion** by requiring of a social welfare function F that if for $x, y \in A$ and $\mathbf{R} \in D^N$ such that $xR_i y$ for all $i \in N$ and $xR^*_i y$ for some $i \in N$, then $xF^*(\mathbf{R})y$.

Proposition 5.5. *If F is a social welfare function with domain $\mathcal{W}(A)^N$ that satisfies the strong Pareto criterion, then $(F, (\alpha_x^{\text{ud}})_{x \in A})$ is a preference-approval with domain $(\Pi^-)^N$.*

Proof. We need to verify that for all $(\mathbf{R}, \mathbf{B}) \in (\Pi^-)^N$,

$$(f(\mathbf{R}), (\alpha_x^{\text{ud}})_{x \in A}(\mathbf{B})) = (R, B) \in \Pi^- .$$

For $x, y \in A$, suppose that xRy and $x \notin B$, we only need to show that $y \notin B$ (i.e. we show the contrapositive of the consequent of (5.i)). Because $x \notin B$, by the definition of α_x^{ud} there are only two possible preference approvals in Π^- that each agent can have: each agent can either approve y and disapprove x , or disapprove both and rank them as indifferent. But no agent can have the first, otherwise yR^*x by the strong Pareto criterion, contradiction xRy . Thus all agents disapprove y , and thus $y \notin B$ by the definition of α_y^{ud} , as required. \square

5.3 Possibilities from relaxing decomposability

Our notion of decomposing a preference-approval aggregator takes place at two different levels. Theorem 5.4 requires:

- 1) that determining whether a specific alternative is approved is independent of the approvals of the other alternatives in the profile.
- 2) that the aggregation of rankings and approvals are independent of each other.

Relaxing either of these requirements leads to possibilities.

5.3.1 Relaxing the decomposability of the approval aggregation

We characterize two families of rules that arise when we relax the internal decomposability of the approval part of a preference-approval aggregator. Our characterisations apply to the two alternative case: *well-behaved* approval aggregators that work alongside any social welfare function that satisfies the Pareto criterion; and *maximally discriminating* approval aggregators that distinguish between alternatives as often as possible. However, for three or more alternatives, these characterisations can no longer be applied; we nevertheless show that there are still possibilities.

5.3.1.1 Well-behaved approval aggregators

There are approval aggregators over $\{a, b\}$ that create a preference-approval aggregator no matter what social welfare function that satisfies the Pareto criterion that they are paired with. Let \mathcal{N} be the set of non-empty proper subsets of N . Take any function $g: \mathcal{N} \times \mathcal{N} \rightarrow \{0, 1\}$. Define the approval aggregator α_g as follows:

- 1) If a is approved by all agents and b approved by none, approve only a ,
- 2) Else if b is approved by all agents and a approved by none, approve only b .
- 3) Else:
 - 3.1) if a is approved by no agents or b is approved by none, approve neither.
 - 3.2) if a is approved by all agents or b is approved by all, approve both.
 - 3.3) otherwise, let N_a be the set of agents that approve a and N_b be the set of agents that approve b ; if $g(N_a, N_b) = 1$ approve both alternatives, otherwise disapprove both.

Note that α_g satisfies unanimity. Let G be the set of functions of type $g: \mathcal{N} \times \mathcal{N} \rightarrow \{0, 1\}$. One may consider that an approval aggregator α is *well-behaved* if (f, α) is well-defined for any f that satisfies the Pareto criterion. Under such an interpretation, Theorem 5.6 characterises the well-behaved approval aggregators.

Proposition 5.6. *Suppose $m = 2$. If f satisfies the Pareto criterion and has domain $\mathcal{W}(A)$, then, for any $g \in G$, (f, α_g) defines a preference-approval aggregator with domain Π^N . In the*

other direction, if $\alpha \neq \alpha_g$ for some $g \in G$, then there is some f that satisfies the Pareto criterion with domain $\mathcal{W}(A)^N$ such that (f, α) does not define a preference-approval aggregator with domain Π^N .

Proof. Take f and α_g as described above, let $\pi = (f, \alpha_g)$. Consider a profile $\mathbf{p} = (\mathbf{R}, \mathbf{B})$. Suppose that $x \in \pi_2(\mathbf{p})$ and $y \notin \pi_2(\mathbf{p})$, for $\{x, y\} = \{a, b\}$. It suffices to show that $x\pi_1^*(\mathbf{p})y$. As we must be in case 1 or case 2, x is approved by all agents and y is approved by none, thus by consistency $xP_i y$ for all $i \in N$, thus by Pareto optimality $x\pi_1^*(\mathbf{p})y$ as required.

Now take an arbitrary approval aggregator $\alpha' \neq \alpha_g$ that satisfies unanimity. This implies that the outputs given in cases 1 and 2 still apply, thus the output in one or other of 3 (a), (b) or (c) must be different. Suppose that (a) is not the case; the proof for the others is similar. This implies that there is a non-empty proper subset N_x of agents such that when x is approved only by all agents in N_x and y is approved by no agents, x is approved in the output and y is disapproved in the output. Consider the function f' such that $wf'(\mathbf{R})z$ if and only if $wP_i z$ for any $i \in N$. It is not hard to see that f' a well-defined social welfare function that satisfies the Pareto criterion, however (f', α') is not admissible. \square

5.3.1.2 Maximally discriminating approval aggregators

The approval aggregators of the previous subsection are very undiscriminating, which is precisely what is needed to make them compatible with many different social welfare functions that satisfy the Pareto criterion. On the other hand, if we start with a social welfare function, we can choose an approval aggregator that tends to discriminate between the alternatives, in terms of approvals, as much as possible.

We start with a social welfare function f over $\{a, b\}$ that satisfies the Pareto criterion. Consider a linear order profile over the two alternatives with a reduced electorate $N' \subseteq N$, $\mathbf{R}' \in \mathcal{L}(A)^{N'}$. We say that \mathbf{R}' is *sufficient* for $x \in A$ if for all profiles $\mathbf{R} \in \mathcal{L}(A)^N$ such that for all $i \in N'$ $R'_i = R_i$, $xf^*(\mathbf{R})y$, where $y \in A, y \neq x$.

We now define an aggregator α_f in terms of f . Take an arbitrary approval profile \mathbf{B} . The *total approval count* is $\sum_{i \in N} |B_i|$. Let I be the set of agents who approve a and disapprove b and let J be the set of agents who approve b and disapprove a . We write $\mathcal{L}(A)$ for the set of linear orders over A . Let $\mathbf{R}' \in \mathcal{L}(A)^{I \cup J}$ be such that $aP_i b$ if $i \in I$ and $bP_i a$ if $i \in J$. For $x \in A$, if \mathbf{R}' is sufficient for x , $\alpha_f(\mathbf{B}) = \{x\}$. Otherwise, if the total approval count is greater than or equal to $n = |N|$, approve both alternatives; if the total approval count is less than n disapprove both alternatives.

One may consider an approval aggregator *maximally discriminating* if it classifies different

alternatives into different classes as often as possible.⁵ Theorem 5.7 shows that α_f is maximally discriminating if one wants to be consistent with f .

Proposition 5.7. *Suppose $m = 2$. If f satisfies the Pareto criterion and has domain $\mathcal{W}(A)^N$, then (f, α_f) is a preference-approval aggregator defined on the domain Π^N . Further, if α is an approval aggregator such that for some profile \mathbf{B} and alternative $x \in \{a, b\}$, $\alpha(\mathbf{B}) = \{x\}$ but $\alpha_f(\mathbf{B}) = \emptyset$ or $\alpha_f(\mathbf{B}) = B$, then (f, α) does not define a preference-approval aggregator with domain Π^N .*

Proof. We first show that, for an arbitrary profile (\mathbf{R}, \mathbf{B}) , the output of

$$(f, \alpha_f)(\mathbf{R}, \mathbf{B}) = (R, B)$$

is a preference-approval. Supposing that $x \in B$ and $y \notin B$, we need to show that xPy . Because $\alpha_f(A) = \{x\}$, this means that $\mathbf{R}' \in \mathcal{L}(A)^{N'}$, $N' \subseteq N$ is sufficient for x , where for all $i \in N$, $x \in B_i$ and $y \notin B_i$ implies that $i \in N'$ and $xP'y$, and $y \in B_i$ and $x \notin B_i$ implies that $i \in N'$ and $yP'_i x$. Because each (R_i, B_i) is a preference-approval, this means that $R_i = R'_i$ for all $i \in N'$. Because \mathbf{R}' is sufficient for x , this implies that xPy as required.

Now suppose that α is an approval aggregator such that for some profile \mathbf{B} and alternative $x \in \{a, b\}$, $\alpha(\mathbf{B}) = \{x\}$ but $\alpha_f(\mathbf{B}) = \emptyset$ or $\alpha_f(\mathbf{B}) = A$. Suppose that $\alpha_f(\mathbf{B}) = \emptyset$; the case with $\alpha_f(\mathbf{B}) = A$ is very similar. This means that the total approval count is less than n , and that $\mathbf{R}' \in \mathcal{L}(A)^{N'}$, $N' \subseteq N$ is not sufficient for x . This means that there is some \mathbf{R} such that for all $i \in N'$, $R_i = R'_i$ but $yf(\mathbf{R})x$. Applying (f, α) to the profile (\mathbf{R}, \mathbf{B}) will not produce a preference-approval. \square

5.3.1.3 Relaxing the decomposability of the approval aggregation with three or more alternatives

For three or more alternatives, our analysis is complicated by the fact that there are social welfare functions that satisfy the Pareto criterion which cannot be used to create a consistent preference-approval aggregator. Ranking by the sum of Borda scores provides an example—in this setting involving indifferences, we take the Borda score of an alternative to be the number of other alternatives it is weakly preferred to. Consider the two preference-approval profiles in Figure 5.3, each profile containing two agents and three alternatives. Note that in both a must be approved and c must be disapproved by unanimity. Using Borda scoring, b

⁵Note that this precise formulation of discriminability has strong implications for cases with more than three alternatives. Let $A = \{a, b, c, d\}$. Obviously, if all alternatives in A are approved, then there is no discrimination between the alternatives. If only one alternative is approved, then *three pairs* of alternatives have different evaluations. If two alternatives are approved, then *four pairs* of alternatives have different evaluations. Thus one may distinguish between **discriminability** and **decisiveness**.

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

Figure 5.3: Two preference-approval profiles that show that if there are three alternatives, there are social welfare functions that satisfy the Pareto criterion that cannot be used to create a consistent preference-approval aggregator

is considered indifferent to c in the left profile, and indifferent to a in the right profile. This then means that b must be disapproved in the left profile and approved in the right profile, however b has the same approvals for both, thus no approval aggregator is consistent with ranking by Borda scores.

This means that for three or more alternatives there are no well-behaved approval aggregators. Nor can we find a maximally discriminating approval aggregators for any given ranking function, because no approval aggregator may exist at all. So the two characterisations given above cannot be directly extended to the case of three or more alternatives.

Nonetheless, there are still possibilities when there are three or more alternatives. Define a disjunctive social welfare function f_{db} as ranking by Borda scores if any agent is completely indifferent between all the alternatives and as a copy of the first agent's preferences otherwise. Similarly, define α^{db} such that it approves all alternatives if each alternative is approved at least once, otherwise it disapproves all alternatives if each alternative is disapproved at least once, otherwise it copies the evaluations of the first agent.

Proposition 5.8. *For $m \geq 2$, (f_{db}, α^{db}) is a preference-approval aggregator defined on the domain Π^N .*

Proof. We need to check the consistency of the output. So suppose that for two alternatives $x, y \in A$, x is approved and y is disapproved. We need to show that xPy . Because x and y are in different approval classes, we must be in a profile where the first agent must have had her approval levels copied. As such, no agent can be indifferent between all the alternatives, so the first agent also has their preferences copied, thus xPy as required. \square

5.3.2 Relaxing decomposability between rankings and approvals

There are also possibilities if the output ranking is allowed to depend upon the approvals in the profile. In this subsection we describe four such preference-approval aggregators.

5.3.2.1 Shortlist by elementary approval aggregators then rank

We first give an example of a preference-approval aggregator π^I , where the approval aggregation is internally decomposed and independent of the ranking aggregation while the output ranking depends upon the approvals in the profile. Under π^I , the approvals are used to select an approved shortlist, and then the ranking is performed separately upon the approved and non-approved alternatives. For $x \in A$, we define α_x^{maj} for arbitrary $\mathbf{z} \in \{0, 1\}^N$ as

$$\alpha_x^{\text{maj}}(\mathbf{z}) = \begin{cases} 1 & \text{if } |\{i \in N : z_i = 1\}| \geq n/2, \\ 0 & \text{otherwise.} \end{cases}$$

Define the approval part of π^I such that $\pi_2^I(\mathbf{R}, \mathbf{B}) = (\alpha_x^{\text{maj}})(\mathbf{B})$. For the ranking part, we use a local version of Borda where the score for an alternative is calculated using either only the approved or only the disapproved alternatives. First, for each $X \subseteq A$, let

$$\text{bscore}_X : \mathcal{W}(A)^N \rightarrow X \rightarrow \mathbb{N}$$

be defined as

$$\text{bscore}_X(\mathbf{R})(x) = \sum_{i \in N} |\{y \in X : xR_i y\}|.$$

Define $\text{borda}_X : \mathcal{W}(A)^N \rightarrow \mathcal{W}(X)$ by

$$(x, y) \in \text{borda}_X(\mathbf{R}) \text{ if and only if } \text{bscore}_X(\mathbf{R})(x) \geq \text{bscore}_X(\mathbf{R})(y).$$

Returning to the definition of π^I , and writing $B = \pi_2^I(\mathbf{R}, \mathbf{B})$, let $R = \pi_1^I(\mathbf{R}, \mathbf{B})$ be defined by xPy for $x \in B$ and $y \notin B$; xRy if and only if $(x, y) \in \text{borda}_B(\mathbf{R})$ for $x, y \in B$; and xRy if and only if $(x, y) \in \text{borda}_{A \setminus B}(\mathbf{R})$ for $x, y \notin B$.

If we suppose that an approval aggregator is composed of elementary approval aggregators such as α_x^{maj} , there will be many cases where either all alternatives are approved or all alternatives are disapproved. One may desire instead to try to approve half of the alternatives. It is impossible to always do so without violating unanimity, but we can create an approval aggregator that only approves all alternatives when every alternative is approved by all agents and only disapproves all alternatives when every alternative is disapproved by all agents.

5.3.2.2 Shortlist by non-decomposable approval aggregator then rank

We now describe a preference-approval aggregator π^II that typically approves around half the alternatives; under π^II , the approval aggregation is not anymore internally decomposed but is still independent of the ranking aggregation while the output ranking depends upon the approvals in the profile.

We informally describe the approval aggregation part as an iterative process. First approve all alternatives that are each approved by all the agents. If more than half of the alternatives are approved, disapprove the remaining alternatives. Otherwise, from the remaining set of alternatives, approve all those alternatives with maximal approval support—repeat this step until at least half the alternatives are approved, and afterwards disapprove the remaining alternatives. Formally, define $\gamma: \mathbb{N} \rightarrow (A \rightarrow \mathbb{N}) \rightarrow 2^A$ by $x \in \gamma_j(g)$ if and only if $g(x) \geq j$.⁶ Define $\text{ascore}: (2^A)^N \rightarrow A \rightarrow \mathbb{N}$ by $\text{ascore}(\mathbf{B})(x) = |\{i \in N : x \in B_i\}|$. Note that for $j < k$ we have $\gamma_k(\text{ascore}(\mathbf{B})) \subseteq \gamma_j(\text{ascore}(\mathbf{B}))$ and that in particular $\gamma_0(\text{ascore}(\mathbf{B})) = A$ and $\gamma_{n+1}(\text{ascore}(\mathbf{B})) = \emptyset$; these facts mean that the following definition is well-formed. Let $\alpha^{\text{half}}(\mathbf{B}) = \gamma_j(\text{ascore}(\mathbf{B}))$ for the maximal $j \in \mathbb{N}$ such that $|\gamma_j(\text{ascore}(\mathbf{B}))| \geq m/2$.

Our second preference-approval aggregator is like the first, π^{I} , only using the non-decomposable approval aggregator defined above. Define the approval part of π^{II} such that $\pi_2^{\text{II}}(\mathbf{R}, \mathbf{B}) = \alpha^{\text{half}}(\mathbf{B})$. For the ranking part, writing $B = \pi_2^{\text{II}}(\mathbf{R}, \mathbf{B})$ let $R = \pi_1^{\text{II}}(\mathbf{R}, \mathbf{B})$ be defined by xPy for $x \in B$ and $y \notin B$; xRy if and only if $(x, y) \in \text{borda}_B(\mathbf{R})$ for $x, y \in B$; and xRy if and only if $(x, y) \in \text{borda}_{A \setminus B}(\mathbf{R})$ for $x, y \notin B$.

5.3.2.3 Perform ranking aggregation then approve according to the ranking

We now consider a function π^{III} whose ranking aggregation is independent of approvals while the output approvals depend upon the rankings in the profile. There is a somewhat trivial family of examples here: perform ranking aggregation by any desired method, and approve some fixed number or proportion of alternatives. This method completely ignores the input approvals, so the output ranking is obviously independent of the approvals in the profile. For an example social welfare function, we consider Borda ranking. Formally, define the ranking part of π^{III} such that $\pi_1^{\text{III}}(\mathbf{R}, \mathbf{B}) = \text{borda}_A(\mathbf{R})$. For the approval part, let $\pi_2^{\text{III}} = \gamma_j(\text{bscore}(\mathbf{R}))$ for the maximal $j \in \mathbb{N}$ such that $|\gamma_j(\text{bscore}(\mathbf{R}))| \geq m/2$. Note that \mathbf{B} occurs nowhere in the definition of π^{III} ; the output completely ignores the input approvals.

5.3.2.4 Borda with a movable zero determined by the zero-line

For completeness, we also describe a natural function π^{IV} for which both the ranking and the approvals are interdependent. Suppose that each agent assigns a score to each alternative depending upon how many alternatives there are in between the alternative and the zero-line, with a positive or negative score respectively if the alternative is above or below the zero-line; for each alternative, sum their scores; rank the alternatives according to their sum,

⁶Our definitions are given in functional style, i.e. functions are used as arguments. Note that we assume left associativity of expressions, i.e. $h(i)(x)$ is $(h(i))(x)$, indeed we will typically write $h_i(x)$; and that we assume right associativity of type definitions, i.e. $h: X \rightarrow Y \rightarrow Z$ is implicitly $h: X \rightarrow (Y \rightarrow Z)$.

and disapprove any and only alternatives with a sum strictly less than zero. This may be thought of as *Borda ranking with a movable zero*—note that although the idea behind this is natural, the detailed definition requires some arbitrary decisions. For π^{IV} , first define $\text{mscore}: \Pi \rightarrow A \rightarrow \mathbb{N}$ by

$$\text{mscore}(R, B)(x) = \begin{cases} |\{y \in B : xRy\}| & \text{if } x \in B, \\ 0 - |\{y \in A \setminus B : yRx\}| & \text{if } x \notin B, \end{cases}$$

Define $\pi^{\text{IV}}(\mathbf{R}, \mathbf{B}) = (R, B)$ by

$$xRy \text{ if and only if } \sum_{i \in N} \text{mscore}(R_i, B_i)(x) \geq \sum_{i \in N} \text{mscore}(R_i, B_i)(y)$$

and $x \in B$ if and only if $\text{mscore}(R_i, B_i) \geq 0$.

5.4 Extending the impossibility to more evaluation levels

The impossibility of combining the ordinal and evaluative approaches also applies to the case where there are more than two evaluation levels. In order to consider multiple evaluation levels we must extend the preference-approval framework, which requires more definitions that immediately follow.

Denote by E a set of possible evaluations, $|E| \geq 2$. A **preference-evaluation** is a pair $v = (R, s)$ where $s: A \rightarrow E$ is a **sorting function** into $|E|$ categories. We suppose that there is linear order \succeq on E . A preference-evaluation $v = (R, s)$ is **consistent** with \succeq if xRy implies $s(x) \succeq s(y)$. Let Ω be the set of consistent preference evaluations.

The existence of a linear order on E is central to our interpretation of the evaluative approach. For $e, e' \in E$, the interpretation of $e \succ e'$ is that e is a better evaluative category than e' . Consistency with ordinal preference states, roughly speaking, that it is impossible to prefer a worse category to a better category.

A **preference-evaluation aggregator** is a function $\omega: \Omega^N \rightarrow \Omega$. We say that a preference-evaluation aggregator ω is dictatorial if there is an agent d whose strict preference and evaluations are reproduced in the output preference-evaluation; i.e. d is a dictator if for all profiles of preference-evaluations $\mathbf{v} = ((R_1, s_1), \dots, (R_n, s_n))$, we have $\omega(\mathbf{v}) = (R, s_d)$ such that for all $x, y \in A$, xR_d^*y implies xR^*y .

An **elementary evaluation aggregator** for an alternative x is a function

$$\eta_x: E^N \rightarrow E .$$

For $j \in \{1, \dots, m\}$, we typically shorten η_{a_j} to η_j . Such a function satisfies **(evaluation) unanimity** if $\eta_x(t, \dots, t) = t$ for all $t \in E$. We write $\omega = (F, \eta_1, \dots, \eta_m)$ if⁷

$$\omega((R_1, s_1), \dots, (R_n, s_n)) = (F(R_1, \dots, R_n), \lambda x. \eta_x(s_1(x), \dots, s_n(x))).$$

Theorem 5.9. *Any preference-evaluation aggregator $\omega = (f, (\eta_x)_{x \in A})$ such that η_x satisfies unanimity for all $x \in A$ is dictatorial.*

Proof. Take a preference-evaluation aggregator ω as in the statement of the theorem. We first show that f satisfies the Pareto criterion. We then show that every pair of evaluation categories have a “local” dictator. Such a local dictator must be a dictator on the ranking function—this implies that all such local dictators coincide. Finally, we argue that this agent is also an evaluation dictator.

Suppose that \mathbf{R} is such that for all $i \in N$, $aP_i b$, we want to show that $af^*(\mathbf{R})b$. Take $e_1, e_2 \in E$ such that $e_1 \succ e_2$. Define \mathbf{s} such that for all $i \in N$,

$$s_i(x) = \begin{cases} e_1 & \text{for all } x \text{ such that } xR_i a \\ e_2 & \text{otherwise} \end{cases}$$

For each $i \in N$, each (R_i, s_i) is a preference-evaluation consistent with \succeq . By unanimity, $\omega_2(\mathbf{R}, \mathbf{s})_1(a) = e_1$ and $\omega_2(\mathbf{R}, \mathbf{s})(b) = e_2$. By consistency of the output, $af^*(\mathbf{R})b$, thus f satisfies Pareto optimality.

Let $\omega = (f, (\eta_x)_{x \in A}): \Omega^N \rightarrow \Omega$ such that f satisfies the Pareto criterion and η_x satisfies unanimity for all $x \in A$. Take an arbitrary pair of evaluation categories $e, e' \in E, e \neq e'$. Write $\Omega |_{e, e'} = \Omega \cap (\mathcal{W}(A) \times \{e, e'\}^A)$. There is an agent $d_{e, e'} = d$ such that for all $(\mathbf{R}, \mathbf{s}) \in (\Omega |_{e, e'})^N$, for $x \in A$ and $\omega(\mathbf{v}) = (R, s)$ we have $s_d(x) = s(x)$, otherwise we can translate ω with its domain restricted to $(\Omega |_{e, e'})^N$ into a preference-approval aggregator that violates Theorem 5.4.

We now argue that d is a dictator over F . Take arbitrary $x, y \in A$ and $\mathbf{R} \in \mathcal{W}(A)^N$ such that $xR_d^* y$. Without loss of generality suppose that $e \succ e'$, consider a profile \mathbf{s} such that for all $z \in A$ such that $zR_d x$, $s_d(z) = e$, for all $z \in A$ such that $xP_d z$, $s_d(z) = e'$, and for all $z \in A$, for all $i \in N \setminus \{d\}$, $s_i(z) = e'$. Note that this \mathbf{s} forms a consistent preference-evaluation profile when combined with \mathbf{R} . By consistency, for $\omega(\mathbf{R}, \mathbf{s}) = (R, s)$, xPy . Thus $xF^*(\mathbf{R})y$ if and only if $xR_d^* y$. Note that this implies that $d_{e, e'} = d_{e'', e''}$ for any $e, e', e'', e''' \in E$.

Finally we argue that d is also an evaluation dictator. Take arbitrary $x \in A$ and $\mathbf{s} \in (E^A)^N$. We want to show that $\eta_x(\mathbf{s}) = s_d(x)$. Write $e = s_d(x)$, and pick some $y \in A, y \neq x$. Consider some profile \mathbf{s}' where for all $i \in N$, $s'_i(x) = s_i(x)$ and $s'_i(y) = e$. Consider any profile \mathbf{R} such that $xR_d^* y$ and that is consistent with \mathbf{s}' , note that such a profile exists. By unanimity, for

⁷ We use lambda notation in the following manner: $\lambda x. F(x) = F$.

$\omega(\mathbf{R}, \mathbf{s}') = (R, s)$, $s(y) = e$. By dictatorship of d on F , xR^*y . By consistency, $s'(x) \succeq s'(y) = e$. Similar to above, consider any profile \mathbf{R}' such that yR_d^*x and that is consistent with \mathbf{s}' . By a similar chain of arguments, $s'(y) \succeq s'(x)$, thus $s'(x) = e$. Because for each $i \in N$, $s'_i(x) = s_i(x)$, $\eta_x(\mathbf{s}) = s_d(x)$. \square

Theorem 5.9 is a significant strengthening of Theorem 5.4, as it removes the Pareto optimality of the social welfare function and allows for more than two evaluation levels. We nevertheless present the two theorems separately. One reason for this is because we do not see an obvious direct proof of Theorem 5.9 that does not use Theorem 5.4. Moreover, it is less obvious how the possibilities expressed in Section 5.3 (which hold even when the social welfare function is required to satisfy the Pareto criterion) apply when there are more than two evaluation levels—while positive results are possible, the extra technicalities are cumbersome and would clutter the results. In a similar vein, we do not discuss domain restrictions when there are more than two evaluation levels.

5.5 Final remarks

There is a view that social choice should be performed using evaluations rather than rankings. In fact, the literature contains several examples of social choice procedures that use evaluations, including approval voting (Brams and Fishburn, 1978), threshold aggregation involving three-graded rankings (Aleskerov et al., 2007; Alcantud and Laruelle, 2014), utilitarian voting (Hillinger, 2005) and range voting (Gaertner and Xu, 2012; Pivato, 2014; Zahid and De Swart, 2015; Macé, 2018). A further example is majority judgment, introduced by Balinski and Laraki (2011), which selects the alternative with the highest median evaluation. To be sure, the median has earlier usages as a social choice rule, for example by Bassett and Persky (1999) who apply it within the traditional Arrovian ranking framework, however it should be noted that majority judgment can choose between alternatives with tied highest median, which (depending on the setting) may be an important improvement in decisiveness. Balinski and Laraki's contribution goes beyond this extra tiebreaking step: they provide a whole framework within which evaluative methods can be analysed.⁸ Part of their defense of majority judgment consists of a defense of the evaluative approach as a whole, for example they claim that

the central problem becomes *how to transform many individual grades of a common language into a single collective grade* where the individuals may have unknown preferences that are too complex to be formulated. Sharing a common

⁸Later work picks up this approach and compares different evaluative methods within an evaluative framework, for example work by Brams and Potthoff (2015).

language of grades makes no assumptions about an agent's or a judge's utilities or preferences. (Balinski and Laraki, 2011, p. xiii, their italics).

The debate between defenders of ranking aggregation and defenders of evaluation aggregation goes back to the early days of approval voting⁹ and has been reactivated by the work of Balinski and Laraki.¹⁰ This literature not only notes that particular evaluative methods are incompatible with majoritarianism conditions defined on social welfare functions (such as our example in the introduction) but reflects an incommensurability between the two approaches which is substantiated by our results. In fact, Theorem 5.9 shows that when unanimous evaluation aggregators are used, dictatorship is the only consistent way to independently aggregate rankings and evaluations. Even dictatorship fails to work when the social welfare function is assumed to satisfy mild conditions such as near unanimity, where if everyone except one agent ranks an alternative top, this alternative should be ranked top in the output. In brief, caution is advised in searching for a compromise that combines the two approaches by allowing individuals to have both rankings and evaluations.¹¹

The impossibilities expressed by Theorems 5.4 and 5.9 depend on decomposing the ranking and evaluation aggregators from each other (which embodies the incommensurability between the two approaches) but also on the decomposability of the evaluation aggregation itself. This latter decomposability is satisfied by every evaluative method that we are aware of; indeed, Balinski and Laraki (2011) argue that this decomposability is the correct interpretation of the independence property.

When decomposability is relaxed there are various possibilities, as described in Section 5.3. In particular, relaxing the internal decomposability of the approval aggregation leads to the possibilities given by Proposition 5.6 and Proposition 5.7. More generally, we can structure the space of possible preference-approval aggregators along two axes. One axis imposes restrictions upon the ranking part of the aggregation: the strongest restriction here is a version of binary independence applied to preference-approval aggregators, which requires that the

⁹For example, see the back and forth between Saari and Van Newenhizen (1988), Brams et al. (1988b), Saari and Van Newenhizen (1988) and again Brams et al. (1988a).

¹⁰For example Balinski and Laraki (2007), Felsenthal and Machover (2008), Brams (2011), Edelman (2012), Balinski and Laraki (2016) and Laslier (2017).

¹¹The preference-approval framework that we treat in this chapter can, as Sanver (2010) discusses, be mathematically placed within the traditional literature of social welfare functionals (Sen, 1977) where cardinal or interpersonal comparisons are allowed. Preference-approvals present a weak version of ordinal level comparisons (OLC) which are explored by Roberts (1980). The closest previous work of this type that we are aware of is due to List (2001), who considers a narrow informational addition that he calls OLC+0, which only allows a single level of ordinal comparability. This is almost equivalent to the preference-approval framework, but it allows for alternatives to be on the zero-line, thus there is a third evaluative category within which indifference is forced. Also, List's results concern functions that produce choice sets or ordinal rankings, not functions that output preference-approvals or their equivalent.

Table 5.1: Possible preference-approval aggregators within the space where the vertical axis imposes restrictions on the evaluative part of aggregation and the horizontal axis imposes restrictions on the ranking part of the aggregation. We implicitly suppose that the evaluative part respects unanimity, and that the preference-approval aggregator is non-dictatorial.

Approval aggregation	Ranking aggregation		
	Satisfies binary independence	Independent of approvals	No restriction
Independent of rankings and completely decomposable	Impossible for any number of alternatives by Theorem 5.9		Shortlist by majority approvals then rank, π^I
Independent of rankings	Possible for two alternatives (by Proposition 5.6 or 5.7 and May (1952)), impossible for more than two (Arrow, 1950)	Possible by Proposition 5.8	Shortlist a given proportion by approvals then rank, π^{II}
No restriction		Approve k highest ranked, π^{III}	Borda with movable zero, π^{IV}

output ranking of each pair only depends upon the input rankings of these pairs. Binary independence implies the following weaker condition: that the ranking aggregation only depends upon the ranking part of the profile. Finally, there are aggregators with no restriction on the ranking aggregation. The second axis imposes restrictions on the evaluative part of the aggregation: the evaluations of each alternative can be computed independently of any other information; or can depend upon the evaluations of all the alternatives; or can be calculated using all the information from the profile. This structure, and various possibilities and impossibilities within, is drawn in Table 5.1—note that the results of the table only apply to preference-approvals, not to preference-evaluations.

As a final note, it may be wondered: given that we consider a particular form of aggregation, can we embed our work into a more general aggregation frameworks? Wilson (1975), Nehring (2003), or Maniquet and Mongin (2016) provide possibilities, but perhaps the most promising general framework would be that of *judgment aggregation*, for which List and Puppe (2007) provide a survey. The linear order version of Arrow's theorem is relatively easy to embed into judgement aggregation (Dietrich and List, 2007), however it required much more technical consideration of Dietrich (2015) for the full generality of Arrow's theorem over weak preferences to be thus embedded. Without the extra generality afforded by

weak preferences, much of our results become extremely simple: thus although it seems possible to perform an embedding into judgement aggregation, it requires the heavy machinery of Dietrich (2015) to be done in an interesting manner.

Conclusion

The first three chapters investigate the concept of decisiveness in increasingly abstract terms. The concept arises in the first chapter as a slightly under-defined, but intuitively desirable property, that seems to be in potential conflict with other properties when we consider a real-life application of social choice theory.

The second chapter abstracts away from any particular application, but still considers concrete choice functions. It provides a general method that makes specific choice functions output smaller choice sets, and may be viewed as an attempt to further understand the concept of decisiveness. The method is a success insofar as it does certainly make the rules more decisive, however the costs of applying it are unclear. The method requires a utilisation of more information from the profile, and it is not clear that this use is justified. We possibly lose the character of the original social choice function, as is evidenced by the fact that it is hard to guarantee the inheritance of properties from the original function to their more decisive versions. Worse, there are distinct possibilities as to how the function can be made more decisive, and it is not obvious which of these possibilities is better. Perhaps this is because the decisive versions attempt to use too much information from the profile—so perhaps we should attempt to restrict the amount of information that is used from a profile.

The fourth chapter takes an even more abstract view of the problem: we now move to considering families of choice functions. It provides a definition of maximal decisiveness with respect to the information basis of a rule. However, it turns out that very few rules actually used in practice are maximally decisive with respect to our intuitive informational bases. There is simply too much global information to make choices. Further, although we thereby can class some rules as more and less decisive than each other, this classification has many incompatibilities.

It is, thus, not clear that we have the correct definition of decisiveness, nor even that such a definition is possible given a global take on information. This suggests that a *local* look at the information in the profile may be more useful. The third chapter gives a result concerning the classic method of enforcing that only local information is used in determining the output of a social welfare function: independence of irrelevant alternatives. This property generally leads to strong impossibility results. We investigate a method of escaping this impossibility through a combination of means: slightly weakening the strict imposition that requires that only extremely local information be relevant, and also applying domain restrictions.

The final chapter returns, in part, to the problems of the first chapter. It describes a different paradigm for social choice: the evaluative paradigm. Perhaps the problems of social choice are better served by such a paradigm? Our results may be interpreted as showing that if

we want to adopt the evaluative paradigm, we should do so *in whole*, rather than trying to combine the two different approaches.

So, what is the upshot of this for social choice theory? The issues surrounding decisiveness seem to dissolve in an evaluative framework. However, as we see in the first chapter, there are applications of social choice problems where it truly is the ordinal framework that is desired. Under such a framework, the problem of decisiveness remains unsolved: taking a global look of profiles forces distinctions between alternatives which should probably not be distinguished, whereas the local look leads to strong impossibilities or difficult technical results.

Glossary

(k, l) -equivalent a domain D is such if any social welfare function on D is k -IND iff it is l -IND. 89–92, 94, 97, 122

(k, x, y) -adjacent requires that there is a set $X \subseteq A$ such that $|X| = k$, $\{x, y\} \subseteq X$ and $R \upharpoonright_X = R' \upharpoonright_X$. 87, 89, 92, 94–96

(k, x, y) -reachable (\simeq_{xy}^k) requires that there is a set $X \subseteq A$ such that $|X| = k$, $\{x, y\} \subseteq X$ and $R \upharpoonright_X = R' \upharpoonright_X$. 89, 90, 92, 95, 96

agent (i, j, k, \dots) 7, 26, 40, 65, 67–69, 84–86, 89, 90, 101, 128, 129, 131

agents, set of (N) a set such that $|N| = n$. 7, 26, 50, 64, 65, 67–69, 84, 85, 89, 90, 101, 121, 124, 127–129, 131

alternative ($a_1, a_2, \dots, a_m, a, b, c, \dots, x, y, z, \dots$) 6, 7, 26, 40, 43–45, 47, 50–52, 57, 65–67, 69, 74, 84–87, 89–96, 99, 100, 121, 123–126, 129–131

alternatives, set of (A) a set such that $|A| = m$. 6, 7, 15, 26, 46–48, 50–52, 56, 64–69, 84–87, 89, 90, 92–94, 101, 113, 121–131

alternative-wise unanimity requires that for any $x \in A$, (1) if $x \in B_i$ for all $i \in N$ then $x \in \alpha(\mathbf{A})$ and (2) if $x \in B_i$ for no $i \in N$ then $x \notin \alpha(\mathbf{A})$. 103

altruistic manipulation manipulation that aids some other athlete. 19, *see* discipline manipulation condition

anonymity treat agents symmetrically. 16, 40, 129

approval aggregator (α) of type $\alpha: (2^A)^N \rightarrow 2^A$. 103, 130

approval voting social choice procedure in the social choice theory, evaluative approach where alternatives are either approved or disapproved by the agents, alternatives with the most approvals are selected. 99

argmax tree function (AM) given $g_A: \mathcal{L}(A) \rightarrow \tau_A$, define

$$g_A^{\text{AM}}(T) = \operatorname{argmax}_{x \in A} |\{R \in \mathcal{L}(A) : g_A(R)(T) = x\}|. \quad 50, 57–59, 61, 133$$

argmax universe concept run a procedure in all possible universes and return those outcomes that happen in *most* universes, refines the parallel universe concept. 44

argument max function (argmax) elements of X that return the largest value for some function g : $\operatorname{argmax}_{x \in X} g(x) = \{x \in X : \forall y \in X, g(x) \geq g(y)\}$. 50, 51, 69, 121, 129, 131

Arrovian domain domain on which Arrow's impossibility holds. 104

athlete 8, 14–21, 24–27, 30, 31, 40, 129, *alias for* alternative

athlete-neutrality treat athletes symmetrically. 16, 40, *alias for* neutrality

athletes, set of (A) a set such that $|A| = m$. 14–16, 18–20, 24, 26, 27, 30, 129, *alias for* alternatives, set of

axiom in social choice theory, refers to desirable properties of a social choice procedure. 9

axis (\succ) linear order on A used for single-peaked or single-dippedness. 93, 94, 129

Backus-Naur method for describing context free grammars, e.g. $A ::= a \mid aA$ generates all strings containing only 'a'. 47

binary independence can restrict attention to pairs of alternatives: $\mathbf{R} \upharpoonright_{\{x,y\}} = \mathbf{R}' \upharpoonright_{\{x,y\}}$ implies $F(\mathbf{R}) \upharpoonright_{\{x,y\}} = F(\mathbf{R}') \upharpoonright_{\{x,y\}}$. 6, 26, 40, 83, 85, 103, 125

binary relation (Q) set of ordered pairs, subset of Cartesian square. 15, 123, 125

binary relations, properties of given $R \subseteq X \times X$:

antisymmetric (xRy and yRx) implies $x = y$, all $x, y \in X$. 126

asymmetric (not xRy) or (not yRx), all $x, y \in X$. 130

complete xRy or yRx , all $x, y \in X$. 123, 125, 126

connected xRy or yRx or $x = y$, all $x, y \in X$. 15, 126, 130

irreflexive not xRx , all $x \in X$. 126

symmetric if xRy then yRx , all $x, y \in X$. 125

transitive (xRy and yRz) implies xRz , all $x, y, z \in X$. 123, 125, 126, 130, 131

binary tree (τ) structure containing nodes and leaves, each node has two children, each child is a node or leaf. 44, 46–48, 50, 56, 57, 129

binary trees, families of defined recursively for all lists of alternatives:

balanced tree (ft) $\text{ft}(12\dots m) = \text{ft}(1\dots \lfloor \frac{m}{2} \rfloor) (\text{ft}(\lceil \frac{m}{2} \rceil \dots m))$. 48, 58

Banks tree (bn) $\text{bn}(12\dots m) = \text{bn}(13\dots m)(\text{bn}(23\dots m))$. 47, 48, 53, 58, 59

iterative Condorcet tree (ic) $\text{ic}(12\dots m) = 12\dots m \text{ic}(2\dots m)$. 48, 53, 54, 57–59, 133

simple tree (st) $\text{st}(12\dots m) = 1(\text{st}(2\dots m))$. 48, 53, 57–59, 133

two-leaf tree (tt) $\text{tt}(12\dots m) = 12$. 48, 53, 57–59

binary trees, properties of

complete every alternative $a \in A$ appears at least once as a leaf. 48

repetitive some alternative $a \in A$ appears more than once as a leaf. 48, 55

binary trees, set of all (τ_A) for fixed A , generated by $\tau ::= a \in A \mid (\tau\tau)$. 47, 48, 50, 121, 127

Blau equivalence for $k \neq l$, $k, l < m$, k -IND and l -IND are equivalent. 83, 93, 94, 97

Blau impossibility for $k < m$, k -IND and the Pareto criterion and together imply dictatorship. 83, 93

Blau partition equivalence class on integers determining where k -IND is equivalent to l -IND, the partition of $\{2, \dots, m\}$ determined by (k, l) -equivalence. 84, 90–95, 97, 98

Borda scoring vector the vector $(m-1, m-2, \dots, 0)$ or, in canonical form,

$$(1, (|A|-2)/(|A|-1), \dots, 1/(|A|-1), 0). 15, 87, 97, 122$$

Borda social choice function scoring function using the Borda scoring vector. 69

Cartesian domain (D) D^N , where, for example, $D \in \{\mathcal{W}(A), \mathcal{L}(A), \Pi\}$. 84–87, 89–92, 94–98, 121, 123, 125

clear winner condition an athlete who always defeats another should be considered better. 15, 40, *alias for* Pareto criterion

- complete preorder** (R) transitive and complete binary relation. 6, 7, 14–16, 18–20, 26, 27, 30, 84–87, 89–96, 106, 113, 121, 123, 124, 126–130
- coincide** R and R' coincide on X if $R \upharpoonright_X = R' \upharpoonright_X$. 85, 86
- extend** $R \in \mathcal{W}(X)$ extends $R' \in \mathcal{W}(Y)$, $Y \subseteq X$, if R and R' coincide on X . 85, 86, 125
- ordered construction** (\downarrow) given A as context, $R^\downarrow = a_1 a_2 \dots a_m$. 92, 93, 123
- ordered construction, alternative to bottom** ($\downarrow \cdot$) as ordered construction and send alternative to bottom, $R^{\downarrow k} = (R^\downarrow \cup \{(a_i, a_k) : i \neq k\}) \setminus \{(a_k, a_i) : i \neq k\}$. 92, 93, 123
- complete preorder, asymmetric part of** (R^*) for $R \in \mathcal{W}(X)$ the set $R^* = R \setminus \{(x, x) : x \in X\}$. 6, 14, 16, 85, 92, 102, 106, 114, 115, 127, 128, 130
- complete preorder, list notation of** elements in parentheses are indifferent, otherwise leftwards elements are preferred, e.g. $(a_1 a_3) a_2$ is the complete preorder $\{(a_1, a_3), (a_3, a_1), (a_1, a_3), (a_3, a_2)\}$. 48, 67, 123
- complete preorders, set of** (\mathcal{W}) over X , $\mathcal{W}(X)$. 6, 7, 14, 15, 27, 30, 84–86, 92, 123, 128, 129
- composition consistency** for any $T = (T'; T_1, \dots, T_m)$, index i and alternative $a \in A_i$, $a \in f(T)$ iff $a_i \in f(T')$ and $a \in f(T_i)$. 52, *see also* weak composition consistency
- computational complexity** study of resources required to execute algorithms. 3, 123
- computational social choice theory** interdisciplinary field, adds theoretical computer science notions to social choice theory, e.g. computational complexity. 3
- Condorcet consistency** if Condorcet winner exists, must be unique output. 52
- Condorcet domain** Cartesian domain for which for all profiles, the majority relation is always transitive. 98
- Condorcet independence of irrelevant alternatives** if arbitrary top pairs are ranked the same in two profiles, then one of the pair cannot not be selected for one profile and selected for the other. 27, 125
- Condorcet loser** unique minimum of majority relation, if exists. 13
- Condorcet paradox** the observation that in a profile of three voters such that $aP_1bP_1P_1c$, $bP_2cP_2P_2a$, and $cP_3aP_3P_3b$, comparing the alternatives by pairwise majorities creates a cycle. 100
- Condorcet winner** unique maximum of majority relation, if exists. 13, 123
- consistent preference-evaluation** (v) a preference-evaluation such that xRy implies $s(x) \succeq s(y)$. 113
- consistent preference-evaluations, set of** (Ω) set equipped with a linear order \succeq . 113, 114, 128
- covering relation** given a tournament T , a covers b if aTb and $(bTx$ implies $aTx)$. 51, 55
- decision theory** formal study of individual rational choice. 18
- decisiveness** how (small/close to linear) is output of social choice procedure, informal. 4, 5, 8, 10, 31, 44, 62

- decomposability** a social choice procedure aggregating a composite structure can be split into multiple social choice procedures, which combine to produce a consistent output. 100
- decomposability of approval aggregator** $\alpha = (\alpha_x)_{x \in A}$ if $\forall x \in A \exists \alpha_x$ such that $x \in \alpha(\mathbf{B}) \Leftrightarrow \alpha_x(1_{B_1}(x), \dots, 1_{B_n}(x)) = 1$. 103
- determined** some discipline's top rank is always reproduced. 16, 40, *see also* top dictator
- dictator** (d) privileged agent $d \in N$. 113, 124, 129, 130
- dictatorial** function has dictator of given type. 85–87
- dictatorial preference-approval aggregator** if for $\mathbf{p} = ((R_1, B_1), \dots, (R_n, B_n))$, $\omega(\mathbf{p}) = (R, B_d)$ and $R \subseteq R_d$. 102
- dictatorial preference-evaluation aggregator** if for $\mathbf{v} = ((R_1, s_1), \dots, (R_n, s_n))$, $\omega(\mathbf{p}) = (R, s_d)$ and $R \subseteq R_d$. 113
- dictatorial social welfare function** if $F(\mathbf{R}) \subseteq R_d$. 102, 103
- discipline** single competition between athletes. 8, 14–18, 20, 24–27, 30, 31, 40, 129, *alias for agent*
- discipline aggregator** (F) method for ranking athletes given results in disciplines. 9, 14–16, 18, 20, 24, 40, 124, *alias for social welfare function*
- discipline manipulation condition** athlete a can manipulate from sincere profile \mathbf{P} to \mathbf{P}' , $R = F(\mathbf{P})$, $R' = F(\mathbf{P}')$ if $r_{R'}(a) < r_R(a)$. 18
- egoistic manipulation** requires $r_{R'}(b) < r_R(b)$. 18, 40
- altruistic manipulation** for some b , requires $r_{R'}(b) < r_R(b)$. 19, 40
- without sacrifice** requires $r_{R'}(a) \leq r_R(a)$. 19, 20, 40
- completely without sacrifice** requires $|\{x \in A : xR'a\}| \geq |\{x \in A : xRa\}|$. 19, 20
- spoiler** requires $r_{R'}(b) < r_R(a)$, a spoils result for better competitors. 20, 40
- discipline-neutrality** treat disciplines symmetrically. 16, 40, *alias for anonymity*
- disciplines, set of** (N) a set such that $|N| = n$. 14–18, 20, 24, 27, 30, 129, *alias for agents, set of*
- domain** (\mathcal{D}) set of possible inputs, e.g. $\mathcal{W}(A)^N$, $\mathcal{L}(A)^N$, $(2^A)^N$, $\mathcal{L}(A)^{\mathcal{A}}$, Π^N . 64, 69, 72, 73, 125, 129
- domain restriction** a restriction applied to the domain of a social choice function, social welfare function, preference-approval aggregator. 5, 83
- egoistic manipulation** manipulation that aids the manipulator. 18, *see discipline manipulation condition*
- elementary approval aggregator** (α_x) for $x \in A$, $\alpha_x: \{0, 1\}^N \rightarrow \{0, 1\}$. 102
- elementary evaluation aggregator** (η_x) of type $\eta_x: E^N \rightarrow E$. 113–115, 125
- empty set** (\emptyset) contains no elements, $\{\}$. 7, 33, 37, 38, 46, 51, 52, 68, 69, 72, 73, 109, 112, 127, 129, 131
- equivalence class** given an equivalence relation $\sim \subseteq X \times X$, the equivalence class of $x \in X$ is $\{y \in X : x \sim y\}$. 89–91, 93, 97, 98, 126

- equivalence relation** (\sim) a symmetric, transitive and complete binary relation. 89
- evaluation unanimity** $\eta_x(t, \dots, t) = t$ for all $t \in E$. 114
- ex ante** before the resolution of a chance event. 44
- ex post** after the resolution of a chance event. 44
- exogenous tiebreaker** method selecting from tied alternatives external to choice procedure. 13, 63
- expressible (by basis)** f expressible by \sim if for all $\mathbf{P}, \mathbf{P}' \in \mathcal{D}$, $\mathbf{P} \sim \mathbf{P}'$ implies that $f(\mathbf{P}) = f(\mathbf{P}')$. 64
- finite initial sequences of natural numbers** (\mathcal{N}) set of initial subsets of \mathbb{N} , $\mathcal{N} = \{\{0\}, \{0, 1\}, \{0, 1, 2\}, \dots\}$. 64–69, 71, 127, 129–131
- free triple** $\{a, b, c\} \in A$ such that every possible (linear order/complete preorder) over the three alternatives extend to some member in a Cartesian domain D . 86, 93
- game theory** formal study of rational choice amongst multiple agents. 18
- Gibbard-Satterthwaite impossibility** resolute social choice functions are manipulable or dictatorial, and variations. 18
- Gibbard-Satterthwaite manipulation** changing a single agents ranking in a preference profile results in a better outcome for this agent, according to their original ranking. 40
- hardness result** e.g. if an oracle can solve problem, then can solve NP problem in polynomial time, unfeasible on large inputs if $P \neq NP$. 4
- impartial culture** generate preferences uniformly at random. 21
- impossibility result** conditions that cannot simultaneously hold. 4, 10, 23
- independence of irrelevant alternatives (IIA)** can restrict attention to any subset of alternatives: $\mathbf{R} \upharpoonright_X = \mathbf{R}' \upharpoonright_X$ implies $F(\mathbf{R}) \upharpoonright_X = F(\mathbf{R}') \upharpoonright_X$. 5, 10, 83, 125
- independence properties** family concerning local information, includes binary independence, Condorcet independence of irrelevant alternatives, independence of irrelevant alternatives (IIA), k -independence. 83
- indicator function** (1_X) of $X \subseteq A$, the function $1_X : A \rightarrow \{0, 1\}$ defined by $1_X(x) = 1$ if $x \in X$; 0 otherwise. 103, 124
- infimum** (\wedge) given Y and reference binary relation R over $X \supseteq Y$, maximal by R element $x \in X$ such that yRx for all $y \in Y$. 73, 78, 134
- informational basis** (\sim) equivalence class of profiles. 64–66, 69, 73–76, 125, 126
- informational restriction** ignore some information from a profile, treat different profiles as equivalent. 64
- integers** (\mathbb{Z}) whole numbers, 0, 1, 2, ..., -1, -2... 66, 67, 77, 130
- interpersonal comparison of utility** measure preferences, well-being across different agents.

interval partition ($\llbracket \cdot \rrbracket$) partition of $\{2, \dots, m\}$ whose equivalence classes are integer intervals, concisely expressed by maxima of equivalence classes: for $\{\{2, \dots, q_1\}, \{q_1 + 1, \dots, q_2\}, \dots, \{q_{t-1} + 1, \dots, q_t\}\}$ where $q_1 < q_2 < \dots < q_t = m$, we write $\llbracket q_1, q_2, \dots, q_t \rrbracket$. 90–92, 94, 126

k -independence can restrict attention to sets of alternatives of cardinality k : for X such that $|X| = k$, $\mathbf{R} \upharpoonright_X = \mathbf{R}' \upharpoonright_X$ implies $F(\mathbf{R}) \upharpoonright_X = F(\mathbf{R}') \upharpoonright_X$. 26, 85–87, 89–91, 93, 95–98, 121, 125

lambda notation $\lambda x.f(x) = f$. 114

left associativity if parentheses missing, apply first to leftmost elements, we assume left associativity of algebraic formulas, $xyz = (xy)z$. 47, 48, *see also* right associativity

linear order (R) transitive, complete and antisymmetric relation. 48, 50, 56, 93, 121, 127, 130, *see also* complete preorder

linear orders, set of (\mathcal{L}) over X , $\mathcal{L}(X)$. 6, 48, 50, 56, 121, 127, 128

majority cycle loop in majority relation e.g. $aTbTcTa$, equivalently set with no Condorcet loser and no Condorcet winner. 9, 12

majority relation (T) connected and irreflexive $T \subset A \times A$. 12, 123, 130, *see also* majority relation of a profile

majority relation of a profile (T) $T: \mathcal{W}^*(A)^{\mathcal{N}} \rightarrow 2^{\text{pairs}}$ defined by $(x, y) \in T(\mathbf{P})$ iff $|\{i \in N : xP_i y\}| \geq |\{i \in N : yP_i x\}|$. 15, 16, 46, 65, 66, 68, 123, *see also* majority relation

manipulation distortion of input of social choice procedure for gain. 4, 9, 129

maxima (max) undominated elements of a (normally transitive) binary relation, for $R \in \mathcal{W}(A)$, $\max(R) = \{x \in A : \forall y \in A, yRx \Rightarrow xRy\}$. 16, 123, 129, 130

maxima of set (max) overload notation, add restriction to set e.g. $\max(R)(X) = \max(R \upharpoonright_X)$, if no binary relation take natural ordering on numbers e.g. $\max X = \max(>)(X)$. 90, 126, 130

maximal single-peaked domain the largest (by cardinality and inclusion) single-peaked domain with respect to some axis. 94

maximally decisive for basis \sim if for any profile \mathbf{R} such that $a, b \in f(\mathbf{R})$, there is a permutation σ such that $\sigma(a) = b$ and $\mathbf{R} \sim \sigma(\mathbf{R})$. 5, 74

membership (\in) $x \in X$ if x is in X . 6, 7, 14–21, 24–27, 30, 31, 33–38, 47, 48, 50–52, 56, 64–74, 76, 77, 79, 81, 84–87, 89–94, 96, 101–115, 121–131

minima (min) undominating elements of a (normally transitive) binary relation, e.g. $\min(R) = \{x \in A : \forall y \in A, xPy\}$. 123

monotonic increasing the input in some manner cannot decrease the output in some manner. 18

monotonicity for tournament functions if $a \in f(T)$ and bTa , then $a \in f(T_{\langle a, b \rangle})$. 52, 53, 57, 59, 60

monotonicity, of ranks increasing an alternative's rank in the profile, with no other changes, cannot decrease this alternative's rank in the output. 18, 40, *see also* egoistic manipulation

Nash product take the product of multiple utilities. 15

natural numbers (\mathbb{N}) numbers used for counting, we include 0, then 1, 2, ... 64, 65, 67–70, 73, 125, 128, 131

neutrality treat alternatives symmetrically. 5, 16, 40, 44, 129

non-reducible tournament tournament with a cycle through all alternatives. 59

null representation of a profile (null) of type $\mathcal{L}^*(A)^{\mathcal{N}} \rightarrow \{\emptyset\}$. 68

odds description of probabilities using weights, e.g. $w_1 : w_2 : w_3 \dots$ where probability of the event associated with w_i is $w_i / (\sum_j w_j)$. 21

ordered pairs distinct of alternatives (pairs) the set $\{(a, b) : a, b \in A, a \neq b\}$. 65–67, 126, 130, 131

parallel universe concept run a procedure in all possible universes and take the union of the results. 43, 44, 121

parallel universe function given a set of possible universes X and function $g : X \rightarrow Y \rightarrow Z$, the function $g^{\text{PU}} : Y \rightarrow 2^Z$ is . 44

parallel universe tree function (PU) given $g_A : \mathcal{L}(A) \rightarrow \tau_A$, define $g_A^{\text{PU}}(T) = \{g_A(R)(T) : R \in \mathcal{L}(A)\}$. 50, 53, 56, 57, 59, 61

Pareto criterion prefer Pareto optimal outcomes, for F, \mathbf{R} , if $\forall i \in N, aR^*i; b$ then $aF^*(\mathbf{R})b$. 16, 40, 42, 85–87, 93, 95, 97, 98, 102

permutation onto function with identical domain and range. 7, 16

permutation of agents (ρ) we abuse notation:

$$\rho((R_1, \dots, R_n)) = (R_{\rho(1)}, \dots, R_{\rho(n)}) . 7, 16, 127$$

permutation of alternatives (σ) we abuse notation: $\sigma(R) = \{(\sigma(x), \sigma(y)) : (x, y) \in R\}$ and $\sigma((R_1, \dots, R_n)) = (\sigma(R_1), \dots, \sigma(R_n))$. 7, 16, 74, 126, 127

Plackett-Luce model generate linear orders according to each alternative's probability of being selected/ranked first: write p_i for the probability that a_i is selected/ranked first, given that a_j for $j \in J \subset \{1, \dots, m\}, i \notin J$ were selected/ranked in some specific order before a_i then a_i has a $p_i / (1 - \sum_{i \in J} p_j)$ chance of being selected/ranked next. 21

possible universe one given instance among a set of different possibilities. 121, 127

power set (2) for X , 2^X contains all subsets of X . 7, 46, 50, 51, 65, 69, 101, 121, 129

preference-approval (p) object $p = (R, B) \in \mathcal{W}(A) \times 2^A$ such that $(\forall x, y \in A) ((xR_iy \text{ and } y \in B_i) \Rightarrow x \in B_i)$. 99, 101

preference-approval aggregator (π) function $\pi : D^N \rightarrow \Pi$, where $D \subseteq \Pi$. 102–105, 124, 130

preference-approvals, set of (Π) all consistent preference-approvals for fixed A and N . 101–106, 127, 128

preference-evaluation (v) a pair (R, s) . 113, 123

preference-evaluation aggregator (ω) of type $\omega: \Omega^N \rightarrow \Omega$. 113–115, 124

profile vector with indices in N . 3, 125, 129

profile of complete preorders (\mathbf{R}) member of $\mathcal{W}(A)^N$. 3, 7, 26, 74, 84–86, 89–91, 106, 122, 124–130

profile of preference-approvals (\mathbf{p}) member of Π^N . 104, 124

profile of preference-evaluations (\mathbf{v}) member of Ω^N . 113, 124

profile of strict linear orders (\mathbf{P}) member of $\mathcal{L}(A)^N$. 14–16, 20, 24–27, 30, 33, 64–73, 76, 77, 124, 125, 128–131

projected image of a profile (proj) where $\text{proj}_i: \mathcal{L}^*(A)^N \rightarrow \mathcal{L}^*(A)$ is $\text{proj}_i(\mathbf{P}) = P_i$. 68, 80

proper refinement a binary relation R properly refines R' if $R' \subset R$. 90

properly maximally decisive for basis maximally decisive and expressible by a basis. 75

quota rule fix $q \in \mathbb{N}$, (select/approve) $a \in A$ if more than q agents (rank a first/approve a). 30

rank function (r) the number of alternatives ranked above an alternative plus one, $r(R)(a) = r_R(a) = |\{x \in A: xPa\}| + 1$, if context clear $r_i = r(R_i) = r(R^*_i)$. 14–16, 18–20, 24, 27, 30, 69, 124, 129

real numbers (\mathbb{R}) rational numbers and friends (irrational, transcendental). 68, 69

refinement a binary relation R refines R' if $R' \subseteq R$. 66, 73, 89

relevant set defined with respect to social welfare function that satisfies the Pareto criterion $g: D^N \rightarrow \mathcal{W}(X)$, where $D \subseteq \mathcal{W}(X)$; a relevant set for a pair $\{x, y\} \subseteq X$ is the minimal (by inclusion) set of alternatives $Y \supseteq \{x, y\}$, $Y \subseteq X$, such that for all profiles $\mathbf{R}, \mathbf{R}' \in D \subseteq \mathcal{W}(X)$, if $\mathbf{R} \upharpoonright_{\{x, y\}} = \mathbf{R}' \upharpoonright_{\{x, y\}}$ then $g(\mathbf{R}) \upharpoonright_{\{x, y\}} = g(\mathbf{R}') \upharpoonright_{\{x, y\}}$. 104

resolute for a social choice function output is always singleton, for a social welfare function output is always a linear order. 11, 31, 42, 44, 62

restrictions given a subset of some set, produces subsets of sets of structures.

restrictions applied to function domain (\upharpoonright) e.g. for $Y \subseteq X$, $Z \subseteq W^X$, we have

$$Z \upharpoonright_Y = \{(x, w) \in Z : x \in Y\}. \quad 124, 128$$

restrictions applied to function range (\downharpoonright) e.g. for $Y \subseteq X$, and $Z \subseteq X^k$, $Z \downharpoonright_Y = \{(x_1, \dots, x_k) \in Z : x_1, \dots, x_k \in Y\}$. 52, 84–87, 89–94, 121–123, 125, 126, 128, 129

restrictions applied to structure (\upharpoonright) context dependent: for profile $\mathbf{R} = (R_1, \dots, R_n)$ and $X \subseteq A$, $\mathbf{R} \upharpoonright_X = (R_1 \upharpoonright_X, \dots, R_n \upharpoonright_X)$; for set $D \subseteq \mathcal{W}(A)$, $D \upharpoonright_X = \{R \upharpoonright_X : R \in D\}$. 27, 84–86, 89, 90, 92, 93, 103, 104, 114, 126, 128, 130

right associativity if parentheses missing, apply first to rightmost elements, we assume right associativity of types, $f: X \rightarrow Y \rightarrow Z = f: X \rightarrow (Y \rightarrow Z)$. *see also* left associativity

- scoring function** (f_x) given $\mathcal{D} \in \{\mathcal{L}^*(A)^N, \mathcal{L}^*(A)^{\mathcal{N}}\}$, a function $f_x: \mathcal{D} \rightarrow 2^A \setminus \{\emptyset\}$ defined by $f_x(\mathbf{P}) = \operatorname{argmax}_{b \in A} \sum_{i \in N} x_{r_i(a)}$ where $r_i(a) = |\{b \in A : bP_i a\}| + 1$. . 69, 122
- scoring rule** assigns points based upon positions in a neutral and anonymous fashion, better alternatives have higher sums of points. 15
- scoring vector** (\mathbf{x}) vector of real numbers used to assign points for a scoring function. 68, 69, 122, 129
- set of categories** (E) set equipped with a linear order \succeq . 113–115, 123–125, 130
- sets of binary relations** typeset using calligraphic symbols, e.g. \mathcal{W} . 1
- simulation** artificial (computer generated) model of (social) state. 10
- sincere profile** non-manipulated profile. 17, 124
- single-crossing domain** preferences can be listed R^1, R^2, \dots, R^t such that if xR^1y and $yR^s x$, then $xR^i y$ for $s \leq i \leq t$. 96, 98
- single-dipped domain** set of complete preorders that are single-dipped with respect to some axis. 93, 95, 98
- single-peaked domain** set of complete preorders that are single-peaked with respect to some axis. 83, 93, 98
- single-peaked with respect to** \succ if $(x \succ y \succ z \text{ or } z \succ y \succ x)$ implies $(xRy \text{ implies } yRz)$. 93
- singleton set** contains a unique element. 89, 91
- singleton tree function** (τ) function based on binary tree, of type $\tau: \mathcal{T}(A) \rightarrow A$, defined $x: T \mapsto x \mid xy: T \mapsto (\text{if } xTy: x; \text{ otherwise } y) \mid \tau\tau': T \mapsto \tau(T)\tau'(T)(T)$. 50
- social choice function** (f) selects alternatives, $f: \mathcal{D} \rightarrow 2^A \setminus \{\emptyset\}$. 7, 46, 64, 74, 124–126, 128–130
- social choice functions, examples of**
- constant function** $\exists a, \forall \mathbf{R}, f(\mathbf{R}) = a$. 27
- dictatorship** $\exists d, \forall \mathbf{R}, f(\mathbf{R}) = \max(R_d)$. 27, 68
- instant runoff voting** (irv) successively eliminate alternatives with least first places. 5, 30, 43, 68
- inverse-Borda-Nash** (ibn) order alternatives increasingly (smaller is better) by product of inverse Borda scores, breaking ties by majorities if possible, i.e. define $Q \subseteq A \times A$ by xQy iff $\prod_{i \in n} r_i(x) > \prod_{i \in n} r_i(y)$ or $(\prod_{i \in n} r_i(x) = \prod_{i \in n} r_i(y) \text{ and } xT_p y)$; and define $\text{ibn}(\mathbf{P}) = Q^+ \cup \{(x, x) : x \in A\}$. 4, 15, 21, 23
- iterative first place elimination** (ifpe) let $\text{localwin}(\mathbf{P}) = (\text{if } \exists a \in A, |\{i \in N : r_i(a) = 1\}| > n/2 \mapsto \{a\}; \text{ otherwise } \{x \in A : \exists i \in N, r_i(x) = 1\})$, let $\mathbf{P}^1 = \mathbf{P}$, for $t \geq 1$, recursively define $\mathbf{P}^{t+1} = \mathbf{P}^t \upharpoonright_{A \setminus \text{localwin}(\mathbf{P}^t)}$, define $x(\text{ifpe}(\mathbf{P}))y$ iff there are integers s, t such that $s \leq t$ and $x \in \text{localwin}(\mathbf{P}^s)$ and $y \in \text{localwin}(\mathbf{P}^t)$. 27, 30
- iterative successive last removal** (isr) For an arbitrary profile \mathbf{P} , let $\text{lose}_t(\mathbf{P}) = \{a \in A : r_t(a) \text{ is maximal}\}$. Let $\mathbf{P}^1 = \mathbf{P}$, and for $t \geq 1$ recursively define \mathbf{P}^{t+1} as the restriction of \mathbf{P}^t to $A \setminus \text{lose}_t \text{ mod } n(\mathbf{P}^t)$. Writing $P = \text{isr}(\mathbf{P})$, for $x, y \in A$, define xRy iff there are integers $s, t \leq m$ such that $s \geq t$ and $x \in \text{lose}_s \text{ mod } n(\mathbf{P}^s)$ and $y \in \text{lose}_t \text{ mod } n(\mathbf{P}^t)$. 27

- social choice procedure** formal method from social state to structure over alternatives, e.g. social choice function, social welfare function, approval aggregator, preference-approval aggregator. 3, 123, 126, 130
- social choice theory** formal study of voting methods, choice procedures. 3, 8, 123
- social choice theory, evaluative approach** aggregate evaluations expressed as classifications. 6, 99, 121
- social choice theory, ordinal/traditional/Arrovian approach** aggregate preferences expressed as binary relations. 3, 8, 83, 99
- social choice theory, paradox of** undesirable behaviour of social choice procedure. 9
- social welfare function** (F) aggregates complete preorders $F: \mathcal{W}(A)^N \rightarrow \mathcal{W}(A)$. 7, 26, 40, 83–86, 89–91, 93, 97, 98, 102–106, 114, 115, 121, 122, 124–128, 130, 131
- social welfare function, asymmetric part of** (F^*) as for complete preorder, asymmetric part of. 84, 85
- social welfare function, winners of** alternatives ranked top in the output, $\max(F(\mathbf{R}))$. 15
- sorting function** (s) function from alternatives to categories, type $s: A \rightarrow E$. 113, 123, 128
- stability** given a variable agenda social choice function and sensible notion of restriction of a profile structure \mathbf{Q} , requires that for all nonempty $X, Y \subseteq A$ and $Z \subseteq X \cap Y$, $Z = f(\mathbf{Q}|_X) = f(\mathbf{Q}|_Y)$ if and only if $Z = f(\mathbf{Q}|_{X \cup Y})$. 52
- strict linear order** (P) transitive, connected and asymmetric relation, asymmetric part of some linear order. 7, 14–20, 24, 27, 65, 67–69, 128, 129, 131, *see also* strict weak order
- strict linear orders, set of** (\mathcal{L}^*) over X , $\mathcal{L}^*(X)$. 7, 14, 15, 27, 30, 64–69, 71, 80, 127–131
- strict weak orders, set of** (\mathcal{W}^*) over X , $\mathcal{W}^*(X)$. 7, 126
- supremum** (\vee) given Y and reference binary relation R over $X \supseteq Y$, minimal by R element $x \in X$ such that xRy for all $y \in Y$. 73
- symmetric weighted majority relation of a profile** ($\text{sym}W$) $\text{sym}W: \mathcal{L}^*(A)^{\mathcal{N}} \rightarrow \mathbb{Z}^{\text{pairs}}$ is $\text{sym}W(\mathbf{P})(x, y) = W(\mathbf{P})(x, y) - W(\mathbf{P})(y, x)$. 66–68, 71, 72
- tabular presentation of profile of linear orders** Columns correspond to agents, higher alternatives are more preferred. 43, 65, 75
- tie** two or more alternatives that cannot be chosen between. 9, 63
- top dictator** d with top preferences reproduced, $\max F(\mathbf{R}) = \max R_d$. 40
- tournament** (T) asymmetric majority relation. 45–47, 50–52, 54, 56, 57, 59, 121, 123, 127, 130, 131
- tournament function** (f) non-variable agenda method to select alternatives from tournaments, a type of social choice function. 51, 52, 123, 131
- tournament solution** variable agenda method to select alternatives from tournaments. 5, 44
- tournament solutions, examples of** defined in terms of T :
- Banks set** the maximal elements (by T) of maximal (by inclusion) subsets of A that are transitive according to T . 51

- Copeland set** $\operatorname{argmax}_a |\{x \in A : aTx\}|$. 51
- Markov solution concept** alternatives with maximally probability in the stationary distribution of the Markov chain with transition probabilities determined by T . 59
- set of Condorcet non-losers** all alternatives that dominate at least one other alternative according to T . 51, 53, 59
- top cycle** smallest set $B \subseteq A$ such that bTa for all $b \in B$ and $a \in A \setminus B$. 51, 53
- tournament solutions based on binary trees** 50, *see* binary tree, functions based on
- uncovered set** contains every a such that there is no element b such that bTa and for all x such that aTx, bTx . 51
- tournament with one inversion** $(\langle \cdot, \cdot \rangle) T_{\langle a, b \rangle}$ is T with relation between a, b inverted. 51, 52, 56, 131
- tournament, summary of** Let $T' \in \mathcal{T}(\{a_1, \dots, a_m\})$, and, for each i , $T_i \in \mathcal{T}(A_i)$ for pairwise disjoint sets of alternatives A_i . we write $T = (T'; T_1, \dots, T_m)$ to refer to the tournament over $\bigcup A_i$ defined by, for each $a \in A_i$ and $b \in A_j$, if $i = j$ then aTb iff $aT_i b$ and if $i \neq j$ then aTb iff $a_i T' a_j$. 51, 123
- tournaments, set of** (\mathcal{T}) over A , $\mathcal{T}(A)$. 46, 50, 51, 129, 131
- transitive closure** $(+)$ given a relation R , the smallest transitive relation R^+ such that $R \subseteq R^+$. 12, 15, 129
- unanimity** if all agents are in agreement, this should be reflected in the social outcome. 100, 102
- undominated element** given binary relation R over X , element $x \in X$ such that $\forall y \in X$ if yRx then xRy . 126
- variable agenda** allow the set of alternatives to vary. 50, 51, 130
- variable electorate** allow the set of agents to vary. 46, *see also* finite initial sequences of natural numbers
- vector** values in coordinates, given in bold type, e.g. \mathbf{x} . 69, 128
- vector value substitution** the vector (\mathbf{x}_{-i}, x') is \mathbf{x} with x_i replaced by x' . 7
- voting situation** an anonymous profile, where the names of the voters are not recorded. 67
- voting situation representation of a profile** (vs) where $\text{vs}: \mathcal{L}^*(A)^{\mathcal{N}} \rightarrow \mathbb{N}^{\mathcal{L}^*(A)}$ and $\text{vs}(\mathbf{P})(P) = |\{i \in N : P_i = P\}|$. 67, 68
- weak composition consistency** for all $\{x, y\} \subseteq Y \subseteq X$, (1) $f(T) \cap (X \setminus Y) = f(T_{\langle x, y \rangle}) \cap (X \setminus Y)$ and (2) $f(T) \cap Y \neq \emptyset \Leftrightarrow f(T_{\langle x, y \rangle}) \cap Y \neq \emptyset$. 52, *see also* composition consistency
- weakly determined** some discipline's top rank is always partially reproduced. 16, *alias for* weak top dictator of F
- weighted majority relation of a profile** (W) $W: \mathcal{L}^*(A)^{\mathcal{N}} \rightarrow \mathbb{N}^{\text{pairs}}$ is defined by $W(\mathbf{P})(x, y) = |\{i \in N : xP_i y\}|$. 65–68, 76, 77, 130

zero-line marker in a ranking, alternatives above which are approved, below which disapproved. 99

List of Figures

1.1	The majority relation between six athletes	13
1.2	An example profile and output	24
1.3	Two profiles involving a change for one athlete but not for the output winner	24
1.4	Profiles showing a discipline whose ranking affects the winner	25
1.5	Profiles showing a discipline that is important in deciding the winner	25
1.6	Profiles showing a discipline that fully determines the winner	26
1.7	A ranking corresponding to a potential manipulation and its cycles	35
1.8	A ranking before and after a manipulation	36
2.1	Possible outcomes for an instance of instant runoff voting	45
2.2	Two representations of a tournament.	46
2.3	An example binary tree	47
2.4	An example of the reduction of a tree	47
2.5	Binary trees with four alternatives.	49
2.6	Illustration of equivalence of the iterative Condorcet rule with the tree ic	54
2.7	A tournament for which a covered alternative is the output of a balanced tree.	55
2.8	Two tournaments and balanced trees.	55
2.9	Tournaments and a tree demonstrating a violation of monotonicity	56
2.10	A tournament for which st^{AM} and ic^{AM} are distinct.	57
3.1	Representations of a profile as a tournament and weighted tournament	65
3.2	Representations of a profile as a symmetric weighted tournament	66
3.3	A weighted tournament showing Borda is not maximally decisive for $symC2$	75
3.4	The partial weighted tournament of a permutation of the profile in Table 3.7.	77
3.5	An almost symmetric tournament	79
3.6	Informational bases ordered by refinement	81
4.1	Examples of adjacency on a single-peaked domain	88
4.2	Circle of alternatives in a single-peaked domain and a single-dipped domain	96
5.1	The possible preference approvals for two alternatives.	100
5.2	The six preference-approvals in Π^-	106
5.3	Two preference-approval profiles over three alternatives	110

List of Tables

1.1	Possible ranks of seven of the athletes after the first round of competition	12
1.2	Predicted results for six of the athletes in the final round	13
1.3	Profile that shows that inverse-Borda-Nash is not weakly determined	17
1.4	Number of randomly generated profiles that were tied or manipulable	22
1.5	An example of iterative successive last removal	28
1.6	An example of iterative first place elimination	29
1.7	Profile showing that iterative first place elimination is not weakly determined	31
1.8	Comparison of terminology in discipline aggregation and social choice theory	40
2.1	An example profile in tabular form	43
2.2	Outcomes of argmax rules on non-reducible tournaments of size six	58
2.3	Summary of identities of rules based on trees.	61
2.4	Known properties for rules based on trees.	61
3.1	Profile for which resoluteness is undesirable	62
3.2	Profile with two alternatives that it is difficult to choose between.	63
3.3	The profile of Figure 3.2 with three agents with cyclic preferences added.	67
3.4	A profile in which all alternatives are ranked third twice.	70
3.5	A profile showing Borda is not maximally decisive for scoring vectors.	76
3.6	A profile whose symmetric weighted tournament only contains zeros.	76
3.7	A profile whose scaled positional matrix is the same as the profile in Table 3.6.	77
3.8	A profile demonstrating Borda is not maximally decisive for $C2 \wedge npM$	78
4.1	A domain for which every triple is free.	87
4.2	A single-peaked domain.	94
4.3	A single-dipped domain.	95
5.1	Preference-approval aggregators along two axes	117

Bibliography

- Alcantud, J. C. R. and Laruelle, A. (2014). Disapproval voting: a characterization. *Social Choice and Welfare*, 43(1):1–10.
- Aleskerov, F., Yakuba, V., and Yuzbashev, D. (2007). A ‘threshold aggregation’ of three-graded rankings. *Mathematical Social Sciences*, 53:106–110.
- Altman, A. and Kleinberg, R. (2010). Nonmanipulable randomized tournament selections. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence*.
- Anscombe, G. E. (1976). On frustration of the majority by fulfilment of the majority’s will. *Analysis*, 36(4):161–168.
- Arrow, K. J. (1950). A difficulty in the concept of social welfare. *Journal of Political Economy*, 58(4):328–346.
- Balinski, M. and Laraki, R. (2007). A theory of measuring, electing, and ranking. *Proceedings of the National Academy of Sciences*, 104(21):8720–8725.
- Balinski, M. and Laraki, R. (2011). *Majority Judgment: Measuring, Ranking, and Electing*. MIT press.
- Balinski, M. and Laraki, R. (2016). Majority Judgment vs Majority Rule. <https://hal.archives-ouvertes.fr/hal-01304043>. Preprint.
- Banks, J. S. (1985). Sophisticated voting outcomes and agenda control. *Social Choice and Welfare*, 1(4):295–306.
- Barberà, S., Bossert, W., and Pattanaik, P. K. (2004). Ranking sets of objects. In *Handbook of utility theory*, pages 893–977. Springer.
- Bartholdi, J. J., Tovey, C. A., and Trick, M. A. (1989). The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3):227–241.
- Bassett, G. W. and Persky, J. (1999). Robust voting. *Public Choice*, 99(3):299–310.
- Benson, A. (2012). Formula 1 bosses clarify blocking rule. *BBC*. <https://www.bbc.com/sport/formula1/16416932> (accessed 12 April 2018).
- Black, D. (1948). On the rationale of group decision-making. *Journal of political economy*, 56(1):23–34.
- Blau, J. H. (1957). The existence of social welfare functions. *Econometrica*, 25(2):302–313.

- Blau, J. H. (1971). Arrow's theorem with weak independence. *Economica*, 38(152):413–420.
- Brams, S. J. (2011). Grading candidates. *American Scientist*, 99(5):426.
- Brams, S. J. and Fishburn, P. C. (1978). Approval voting. *American Political Science Review*, 72(03):831–847.
- Brams, S. J., Fishburn, P. C., and Merrill, S. (1988a). Rejoinder to Saari and Van Newenhizen. *Public Choice*, 59(2):149.
- Brams, S. J., Fishburn, P. C., and Merrill, S. (1988b). The responsiveness of approval voting: Comments on Saari and Van Newenhizen. *Public Choice*, 59(2):121–131.
- Brams, S. J., Kilgour, D. M., and Zwicker, W. S. (1998). The paradox of multiple elections. *Social Choice and Welfare*, 15(2):211–236.
- Brams, S. J. and Potthoff, R. F. (2015). The paradox of grading systems. *Public Choice*, 165(3-4):193–210.
- Brams, S. J. and Sanver, M. R. (2009). Voting systems that combine approval and preference. In Brams, S. J., Gehrlein, W. V., and Roberts, F. S., editors, *The Mathematics of Preference, Choice and Order*, pages 215–237. Springer.
- Brandt, F. (2011). Group-strategyproof irresolute social choice functions. In *Proceedings of the 22nd International Joint Conference On Artificial Intelligence*, volume 22, page 79. Citeseer.
- Brandt, F. and Brill, M. (2011). Necessary and sufficient conditions for the strategyproofness of irresolute social choice functions. In *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge*, pages 136–142. ACM.
- Brandt, F., Brill, M., and Harrenstein, P. (2016). Tournament solutions. In Brandt, F., Conitzer, V., Endriss, U., Lang, J., and Procaccia, A. D., editors, *Handbook of Computational Social Choice*, chapter 3. Cambridge University Press.
- Brill, M. and Conitzer, V. (2015). Strategic voting and strategic candidacy. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*, volume 15, pages 819–826.
- Brill, M. and Fischer, F. (2012). The price of neutrality for the ranked pairs method. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence*, pages 1299–1305.
- Bubboloni, D. and Gori, M. (2014). Anonymous and neutral majority rules. *Social Choice and Welfare*, 43(2):377–401.
- Campbell, D. E. and Kelly, J. S. (2000). Information and preference aggregation. *Social Choice and Welfare*, 17(1):3–24.

- Campbell, D. E. and Kelly, J. S. (2002). Impossibility theorems in the arrovian framework. In Arrow, K. J., Sen, A. K., and Suzumura, K., editors, *Handbook of Social Choice and Welfare*, volume 1, chapter 1, pages 35–94. Elsevier.
- Chevaleyre, Y., Lang, J., Maudet, N., and Ravilly-Abadie, G. (2009). Compiling the votes of a subelectorate. In *Proceedings of the 21st International Joint Conference On Artificial Intelligence*, pages 97–102.
- Condorcet, N. d. (1785). *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. Imprimerie Royale.
- Conitzer, V., Rognlie, M., and Xia, L. (2009). Preference functions that score rankings and maximum likelihood estimation. In *Proceedings of the 21st International Joint Conference On Artificial Intelligence*, pages 109–115.
- Conitzer, V. and Sandholm, T. (2002). Complexity of manipulating elections with few candidates. In *Proceedings of the 18th AAAI Conference on Artificial Intelligence*, pages 314–319.
- Conitzer, V. and Sandholm, T. (2005). Communication complexity of common voting rules. In *Proceedings of the 6th ACM Conference on Electronic Commerce*, pages 78–87. ACM.
- Debord, B. (1987). Caractérisation des matrices de préférences nettes et méthodes d'agrégation associées. *Mathématiques et Sciences humaines*, 97(5):17.
- Dietrich, F. (2015). Aggregation theory and the relevance of some issues to others. *Journal of Economic Theory*, 160:463–493.
- Dietrich, F. and List, C. (2007). Arrow's theorem in judgment aggregation. *Social Choice and Welfare*, 29(1):19–33.
- Dutta, B., Jackson, M. O., and Le Breton, M. (2001). Strategic candidacy and voting procedures. *Econometrica*, 69(4):1013–1037.
- Edelman, P. H. (2012). Review of majority judgment: measuring, ranking and electing. *Public Choice*, 151(3/4):807–810.
- Eğecioğlu, Ö. and Giritligil, A. E. (2013). The impartial, anonymous, and neutral culture model: A probability model for sampling public preference structures. *The Journal of Mathematical Sociology*, 37(4):203–222.
- Elkind, E., Faliszewski, P., and Slinko, A. (2009). Swap bribery. In *International Symposium on Algorithmic Game Theory*, pages 299–310. Springer.

- Erdamar, B., Sanver, M. R., and Sato, S. (2017). Evaluationwise strategy-proofness. *Games and Economic Behavior*, 106:227–238.
- Faliszewski, P., Hemaspaandra, E., and Hemaspaandra, L. A. (2006). The complexity of bribery in elections. In *Proceedings of the 21st AAAI Conference on Artificial Intelligence*, pages 641–646.
- Felsenthal, D. S. and Machover, M. (2008). The majority judgement voting procedure: a critical evaluation. *Homo Oeconomicus*, 25(3):319–334.
- Fine, B. and Fine, K. (1974). Social choice and individual ranking I. *The Review of Economic Studies*, 41(3):303–322.
- Fischer, F., Procaccia, A. D., and Samorodnitsky, A. (2009). A new perspective on implementation by voting trees. In *Proceedings of the 10th ACM conference on Electronic commerce*, pages 31–40. ACM.
- Fishburn, P. C. (1977). Condorcet social choice functions. *SIAM Journal on applied Mathematics*, 33(3):469–489.
- Freeman, R., Brill, M., and Conitzer, V. (2014). On the axiomatic characterization of runoff voting rules. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence*, pages 675–681.
- Gaertner, W. (2001). *Domain Conditions in Social Choice Theory*. Cambridge University Press.
- Gaertner, W. (2002). Domain restrictions. In *Handbook of Social Choice and Welfare*, volume 1, pages 131–170. Elsevier.
- Gaertner, W. and Xu, Y. (2012). A general scoring rule. *Mathematical Social Sciences*, 63(3):193–196.
- Gans, J. S. and Smart, M. (1996). Majority voting with single-crossing preferences. *Journal of public Economics*, 59(2):219–237.
- García-Lapresta, J. L. and Martínez-Panero, M. (2017). Positional voting rules generated by aggregation functions and the role of duplication. *International Journal of Intelligent Systems*, 32(9):926–946.
- Garman, M. B. and Kamien, M. I. (1968). The paradox of voting: Probability calculations. *Behavioral Science*, 13(4):306–316.
- Gehrlein, W. V. and Fishburn, P. C. (1976). Condorcet’s paradox and anonymous preference profiles. *Public Choice*, 26(1):1–18.

- Gibbard, A. (1973). Manipulation of voting schemes: a general result. *Econometrica*, 41(4):587–601.
- Goldsmith, J., Lang, J., Mattei, N., and Perny, P. (2014). Voting with rank dependent scoring rules. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence*, pages 698–704.
- Gpupdate (2008). Barrichello admits Ferrari team orders. <https://www.gpupdate.net/en/f1-news/203773/barrichello-admits-ferrari-team-orders/>. Accessed: 28 February 2018.
- Grofman, B. and Feld, S. L. (2004). If you like the alternative vote (aka the instant runoff), then you ought to know about the Coombs rule. *Electoral studies*, 23(4):641–659.
- Henderson, D. A., Kirrane, L. J., et al. (2018). A comparison of truncated and time-weighted Plackett–Luce models for probabilistic forecasting of Formula One results. *Bayesian Analysis*, 13(2):335–358.
- Hillinger, C. (2005). The case for utilitarian voting. Discussion Paper 2005-11, University of Munich.
- Horan, S. (2013). Implementation of majority voting rules. *preprint*.
- Hudry, O. (2004). A note on “banks winners in tournaments are difficult to recognize” by gj woeginger. *Social Choice and Welfare*, 23(1):113–114.
- International Federation of Sports Climbing (2017). Combined 2017 format for IFSC events and olympic games. https://www.ifsc-climbing.org/images/World_competitions/Event_regulations/170314_Combined-Olympic-Format.pdf. Accessed: 19 February 2018.
- Jimenez, T. (2012). Cycling - I crashed on purpose, says victorious Briton. <https://uk.reuters.com/article/uk-oly-cycl-ctmspr-hindes-day7/cycling-i-crashed-on-purpose-says-victorious-briton-idUKBRE87201X20120803>. Accessed: 28 February 2018.
- Kelly, J. S. (1994). The Bordes-LeBreton exceptional case. *Social Choice and Welfare*, 11(3):273–281.
- Konczak, K. and Lang, J. (2005). Voting procedures with incomplete preferences. In *Proceedings of the IJCAI-05 Multidisciplinary Workshop on Advances in Preference Handling*, pages 124–129.
- Kruger, J. and Airiau, S. (2017). Permutation-based randomised tournament solutions. In *Multi-Agent Systems and Agreement Technologies*, pages 235–250. Springer.

- Kruger, J. and Sanver, M. R. (2018). Restricting the domain allows for weaker independence. *Social Choice and Welfare*, 51(3):563–575.
- Laffond, G., Lainé, J., and Laslier, J.-F. (1996). Composition-consistent tournament solutions and social choice functions. *Social Choice and Welfare*, 13(1):75–93.
- Lang, J., Maudet, N., and Polukarov, M. (2013). New results on equilibria in strategic candidacy. In *International Symposium on Algorithmic Game Theory*, pages 13–25. Springer.
- Lang, J., Pini, M. S., Rossi, F., Venable, K. B., and Walsh, T. (2007). Winner determination in sequential majority voting. In *Proceedings of the 20th International Joint Conference On Artificial Intelligence*, pages 1372–1377.
- Laslier, J.-F. (1997). *Tournament solutions and majority voting*. Springer Verlag.
- Laslier, J.-F. (2017). L'étrange "jugement majoritaire". Working papers 2017-23, PSE.
- Laslier, J.-F. and Sanver, M. R. (2010). *Handbook on approval voting*. Springer Science & Business Media.
- Lev, O. and Rosenschein, J. S. (2012). Convergence of iterative voting. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems*, pages 611–618.
- List, C. (2001). A note on introducing a "zero-line" of welfare as an escape route from Arrow's theorem. *Pacific Economic Review*, 6(2):223–238.
- List, C. and Puppe, C. (2007). Judgment aggregation: a survey. <http://faculty.econ.ucdavis.edu/faculty/nehrling/papers/Survey-List-Puppe-Judgment%20Aggregation-2007.pdf>. Accessed: 24 October 2018.
- Luce, D. R. (1959). *Individual Choice Behavior*. Wiley.
- Macé, A. (2018). Voting with evaluations: characterizations of evaluative voting and range voting. *Journal of Mathematical Economics*, 79:10–17.
- Maniquet, F. and Mongin, P. (2016). A theorem on aggregating classifications. *Mathematical Social Sciences*, 79:6–10.
- May, K. O. (1952). A set of independent necessary and sufficient conditions for simple majority decision. *Econometrica*, 20(4):680–684.
- May, K. O. (1954). Intransitivity, utility, and the aggregation of preference patterns. *Econometrica*, 22(1):1–13.

- McGarvey, D. C. (1953). A theorem on the construction of voting paradoxes. *Econometrica*, pages 608–610.
- Meir, R., Polukarov, M., Rosenschein, J. S., and Jennings, N. R. (2010). Convergence to equilibria in plurality voting. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence*, pages 823–828.
- Meyer, J. (2018). Interviewed by Sebastian Schneckenburger. Personal communication with the Sport Director of the International Federation of Sports Climbing.
- Moon, J. W. (1968). *Topics on Tournaments in Graph Theory*. Holt, Rinehart and Winston.
- Moulin, H. (1983). *The Strategy of Social Choice*. North-Holland.
- Moulin, H. (1986). Choosing from a tournament. *Social Choice and Welfare*, 3(4):271–291.
- Nehring, K. (2003). Arrow’s theorem as a corollary. *Economics Letters*, 80(3):379–382.
- Obraztsova, S. and Elkind, E. (2012). Optimal manipulation of voting rules. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems*, pages 619–626.
- Obraztsova, S., Elkind, E., and Hazon, N. (2011). Ties matter: Complexity of voting manipulation revisited. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems*, pages 71–78.
- Obraztsova, S., Markakis, E., Polukarov, M., Rabinovich, Z., and Jennings, N. R. (2015). On the convergence of iterative voting: how restrictive should restricted dynamics be? In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*, pages 993–999.
- Ozdemir, U. and Sanver, M. R. (2007). Dictatorial domains in preference aggregation. *Social Choice and Welfare*, 28(1):61–76.
- Pivato, M. (2014). Formal utilitarianism and range voting. *Mathematical Social Sciences*, 67:50–56.
- Plackett, R. L. (1975). The analysis of permutations. *Applied Statistics*, pages 193–202.
- Rae, D. W. and Daudt, H. (1976). The Ostrogorski paradox: a peculiarity of compound majority decision. *European Journal of Political Research*, 4(4):391–398.
- Reny, P. J. (2001). Arrow’s theorem and the gibbard-satterthwaite theorem: a unified approach. *Economics Letters*, 70(1):99–105.
- Roberts, K. W. (1980). Possibility theorems with interpersonally comparable welfare levels. *The Review of Economic Studies*, 47(2):409–420.
- Rothstein, P. (1991). Representative voter theorems. *Public Choice*, 72(2-3):193–212.

- Russell, T. and Van Beek, P. (2011). An empirical study of seeding manipulations and their prevention. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*, pages 350–356.
- Saari, D. G. and Van Newenhizen, J. (1988). The problem of indeterminacy in approval, multiple, and truncated voting systems. *Public choice*, 59(2):101–120.
- Sanver, M. R. (2010). Approval as an intrinsic part of preference. In Laslier, J.-F. and Sanver, M. R., editors, *Handbook on Approval Voting*, pages 469–481. Springer.
- Sanver, M. R. and Zwicker, W. S. (2009). One-way monotonicity as a form of strategy-proofness. *International Journal of Game Theory*, 38(4):553–574.
- Sato, S. (2009). Informational requirements of social choice rules. *Mathematical Social Sciences*, 57(2):188–198.
- Sato, S. (2016). Informational requirements of social choice rules to avoid the condorcet loser: A characterization of the plurality with a runoff. *Mathematical Social Sciences*, 79:11–19.
- Satterthwaite, M. A. (1975). 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.
- Sen, A. (1977). On weights and measures: informational constraints in social welfare analysis. *Econometrica*, pages 1539–1572.
- TheGuardian (2012). Olympic badminton players charged with trying to lose their games. <https://www.theguardian.com/sport/2012/aug/01/olympic-badminton-players-charged-lose>. Accessed: 28 February 2018.
- Tsetlin, I., Regenwetter, M., and Grofman, B. (2003). The impartial culture maximizes the probability of majority cycles. *Social Choice and Welfare*, 21(3):387–398.
- Wilson, R. (1972). Social choice theory without the Pareto principle. *Journal of Economic Theory*, 5:478–486.
- Wilson, R. (1975). On the theory of aggregation. *Journal of Economic Theory*, 10:89–99.
- Woeginger, G. J. (2003). Banks winners in tournaments are difficult to recognize. *Social Choice and Welfare*, 20(3):523–528.
- Wright, M. (2014). Or analysis of sporting rules—a survey. *European Journal of Operational Research*, 232(1):1–8.
- Xia, L. and Conitzer, V. (2010). Compilation complexity of common voting rules. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence*.

-
- Yu, N. N. (2015). A quest for fundamental theorems of social choice. *Social Choice and Welfare*, 44(3):533–548.
- Zahid, M. A. and De Swart, H. (2015). The Borda majority count. *Information Sciences*, 295:429–440.
- Zwicker, W. S. (2016). Introduction to the theory of voting. In *Handbook of Computational Social Choice*, chapter 2. Cambridge University Press.

Résumé

Le choix social repose sur l'agrégation de préférences, parfois conflictuelles, afin de produire une décision collective pour le groupe. Nous examinons dans quelle mesure une procédure de choix social peut être décisive—c'est-à-dire renvoyer une décision sans *ex-æquo*. Pour introduire cette question, nous prenons l'exemple d'un concours sportif, pour lequel on ne veut qu'un seul gagnant.

De manière générale, nous décrivons une méthode pour rendre des procédures plus décisives. En particulier nous étudions des procédures basées sur les tournois, une représentation spécifique des préférences.

Avec une interprétation globale des informations sur les préférences, nous définissons un concept abstrait, saisissant le caractère décisif d'une procédure. D'un autre côté, en nous basant sur une interprétation locale des informations disponibles, nous rendons la propriété célèbre d'«indépendance des alternatives non pertinentes» plus faible en utilisant certaines restrictions de domaine sur les préférences.

Enfin nous combinons deux cadres d'agrégation: l'approche «Arrovian» qui s'appuie sur des préférences relatives, où les alternatives sont comparées par paire; et l'approche que nous nommons «évaluative», où les alternatives sont jugées dans l'absolu, indépendamment les uns des autres.

Mots Clés

théorie du choix social; tournois, caractère décisif

Abstract

The problem of social choice concerns aggregating multiple, perhaps conflicting preferences into a collective preference or outcome. We consider how the information utilised affects how decisive the procedure can be—that is, to what extent ties can be avoided. As an introduction to these problems, we consider the problem of selecting the winners of a sports competition, for which a single winner is a desirable property.

We describe a general method for making social choice procedures more decisive. In particular, we study procedures based upon preferences represented as structures called tournaments.

Taking a global view of available information, we define an abstract concept which captures the notion of the decisiveness of a procedure. On the other hand, taking a local view of available information, we successfully weaken the property known as “(Arrovian) independence of irrelevant alternatives” in conjunction with domain restrictions.

Finally we combine two frameworks of aggregation; the “Arrovian” approach wherein preferences are taken to be binary relations, for which alternatives are compared in pairs; and the approach that we call “evaluative”, where the alternatives are given absolute evaluations, each alternative evaluated independently of the others.

Keywords

social choice theory; tournaments, decisiveness