10

Voting in Combinatorial Domains

Jérôme Lang[a] and Lirong Xia[b]

10.1 Motivations and classes of problems

This chapter addresses preference aggregation and voting on domains which are the Cartesian product (or sometimes, a subset of the Cartesian product) of finite domain values, each corresponding to an issue, a variable, or an attribute.

As seen in other chapters of this handbook, voting rules map a profile (usually, a collection of rankings, see Chapter 1 (Zwicker, 2015)) into an alternative or a set of alternatives. A key question has to do with the structure of the set of alternatives. Sometimes, this set has a simple structure and a small cardinality (*e.g.*, in a presidential election). But in many contexts, it has a complex combinatorial structure. We give here three typical examples:

- *Multiple referenda*. On the day of 2012 US presidential election, voters in California had to decide whether to adopt each of eleven propositions.[1] Five referenda were categorized as budget/tax issues. Specifically, two of them (Propositions 30 and 38) both aimed to raise taxes for education, with different details on the type and rate of the tax. Similarly, in Florida voters had to vote on 11 propositions, eight of which were categorized as budget/tax issues.
- *Group configuration or group planning*. A set of agents sometimes has to make a common decision about a complex object, such as a common menu (composed for instance of a first course, a main course, a dessert and a wine, with a few possible values for each), or a common plan (for instance, a group of friends have to travel together to a sequence of possible locations, given some constraints on the possible sequences).
- *Committee elections* and more generally *multiwinner elections*. A set of agents has to choose a group of delegates or representatives of some given size, from a larger set of candidates. As another example, a group of friends wants to choose a set of DVDs to purchase collectively, from a larger set, subject to some budget constraints.

In these three examples, the set of alternatives has acombinatorial structure: it is a Cartesian product $A = D_1 \times \ldots \times D_p$, where for each i, D_i is a finite value domain for

[a] LAMSADE, CNRS – Université Paris-Dauphine, France
[b] Rensselaer Polytechnic Institute (RPI), USA
[1] http://en.wikipedia.org/wiki/California_elections,_November_2012

a variable X_i, or, in the third example, a subset of a Cartesian product (see further below). For the menu example, we may have for instance $D_1 = \{soup, salad, quiche\}$, $D_2 = \{beef, salmon, tofu\}$, etc. For the multiple referenda example, or more generally in the case of *binary* variables (which for the sake of simplicity we assume in most of the chapter), we write $D_i = \{0_i, 1_i\}$ for each i. Also, when all variables are binary, we usually drop indices and parentheses: for instance, $(1_1, 0_2, 1_3)$ is denoted simply by 101.

Each of these examples has specific properties that may call for specific ways of solving them, which we review in this chapter. Still, the major issue for all classes of problems mentioned, is the *tradeoff between expressivity and cost*. This is illustrated in the following example for multiple referenda by Lacy and Niou (2000):

Example 10.1 We have three issues, and three voters with the following preferences:

- Voter 1: $110 \succ 101 \succ 011 \succ 001 \succ 100 \succ 010 \succ 000 \succ 111$;
- Voter 2: $101 \succ 011 \succ 110 \succ 010 \succ 100 \succ 001 \succ 000 \succ 111$;
- Voter 3: $011 \succ 110 \succ 101 \succ 100 \succ 010 \succ 001 \succ 000 \succ 111$.

At one extreme, we can allow the voters to be fully expressive: each voter submits a full ranking over all 2^3 alternatives. The number of alternatives grows exponentially in the number of issues, which imposes a high cognitive cost on the voters to construct their rankings as well as a high communication cost to report these rankings to the central authority that has to gather the votes and compute the outcome, cf. Chapter 10 (Boutilier and Rosenschein, 2015).

At the other extreme, we could ask each voter to report only her top-ranked alternative. This approach is almost cost-free, but the lack of expressivity can cause serious problems. Applying plurality voting (see Chapter 1 (Zwicker, 2015)) for winner selection is quite arbitrary, since three alternatives are tied in the first place by receiving a single vote. Applying the majority rule (see Chapter 1 (Zwicker, 2015)) to each issue separately, as commonly done for multiple referenda, leads to an even worse outcome: the winner, 111, is ranked last by *all* voters!

We consider separately the case where the common decision to be taken consists of choosing the members of a committee. Benoît and Kornhauser (1991) consider two classes of committee elections: *designated post committees*, and *at-large committees*. In designated post committees, candidates run for a specific post (and the size of the committee is the number of posts); in at-large committees, they don't run for a specific post, and the size of the committee is specified explicitly. Designated post committee elections are naturally expressed as elections on a combinatorial domain: variables correspond to posts, and the domain of each variable is the set of candidates applying for the post. The case of at-large elections is more subtle. An obvious choice consists in having binary variables corresponding to candidates, but then the cardinality constraint restricts the set of feasible committees: we are here in a case of *constrained voting on a combinatorial domain*, where the set of alternatives is not simply the Cartesian product of the domains but a subset of it. Voting for at-large committees takes this cardinality constraint into account for restricting the set of admissible outputs[2] (and, sometimes, the set of admissible inputs), and gives rise to widely used voting rules for *multiwinner elections* (Brams and Fishburn, 2004).

Consider again Example 10.1. At one extreme, one could view all these domains as ordinary domains, and proceed as usual by eliciting voters' preferences over the set of alternatives A and then applying a given voting rule. Because the number of alternatives grows exponentially with the number of variables, this is unrealistic as soon as one has more than a few variables; we can definitely not expect individuals to spend hours or days expressing rankings explicitly on thousands of alternatives. At the other extreme, one may think of considering each variable or issue separately, and then organizing votes in parallel for each of them. (This is the way it is usually done in multiple referenda, where each voter has to cast a yes/no ballot for each of the variables simultaneously.) This is much less expensive in terms of communication and computation, but amounts to making the very strong assumption that voters have *separable preferences*, that is, voters' preferences for the value of any variable do not depend on the values of other variables. This assumption is patently unrealistic in many contexts. In multiple referenda, it is likely that a voter's preference over some of these referenda depends on the outcomes of the other referenda, especially when budget/tax issues are concerned, because voters typically have some maximal budget or tax amount they are willing to pay. In group configuration, the value taken by a variable (such as the main course) may have a dramatic influence on a voter's preferences on other variables (such as the wine). In a committee election, it is often the case that a voter's preference for having A over B in the committee depends on whether C is also in the committee, because for instance she wants some balance between genders or between members of different communities.

There are several criteria on which we may assess the practical implementability of a method for voting in combinatorial domains. Perhaps the most important one is the *communication cost* necessary to elicit the votes. Since the communication burden is borne by the voters, making sure that it is reasonably low is a crucial requirement. A second criterion is the *computational cost* needed to compute the outcome. A third criterion is the *generality* of the approach, that is, its applicability to a large variety of profiles: some are widely applicable, whereas some rely on strong domain restrictions. Lastly, and crucially, is the *quality of the outcome*: as we shall soon see, some approaches may lead to extremely controversial, sometimes absolutely unacceptable, outcomes, while others may satisfy desirable social choice axiomatic properties such as Pareto Optimality that give a guarantee about the quality of the solution.

Each of the following sections focuses on families of methods for implementing elections on combinatorial domains. Section 10.2 considers simultaneous voting. As we shall see, simultaneous voting may perform extremely poorly when separability does not hold (and may perform poorly—although much less so—even when separability holds); more precisely, we will list a few important criteria for evaluating methods for implementing elections in combinatorial domains, and will show that simultaneous voting performs

[2] Designated post committees also need constraints if some candidates apply for more than one post.

poorly on all of them except communication and computation cost. In Section 10.3 we discuss methods that assume that voters' preferences are partially specified and then completed automatically using some *completion principle*. Various completion principles are discussed in subsections: using a distance between alternatives is discussed in Section 10.3.1; using a preference extension from singletons to subsets is discussed in Section 10.3.2; and more generally, using a language for compact preference representation such as CP-nets (Section 10.3.3), lexicographic preferences trees (Section 10.3.4), and languages for cardinal preference representation (Section 10.3.5). In Section 10.4 we present methods based on sequential voting, where variables (or groups of variables) are voted on one after another. Section 10.5 concludes by discussing the respective merits and drawbacks of different classes of methods, and briefly addresses related problems.

10.2 Simultaneous voting and the separability issue

10.2.1 Preliminaries

In this chapter, $\mathcal{X} = \{X_1, \ldots, X_p\}$ is a set of *variables*, or *issues*, where each issue X_i takes a value in a finite *local domain* D_i. The set of alternatives, or the *domain*, is $A = D_1 \times \ldots \times D_p$. For $\vec{x} = (x_1, \ldots, x_p) \in A$, and $I \subseteq \{1, \ldots, p\}$, we denote $\vec{x}_I = (x_i)_{i \in I}$. We also make use of the notational convention $-i = \{1, \ldots, p\} \setminus \{i\}$.

Let $\succ$ be a linear order (a transitive, irreflexive and complete preference relation) on A. We say that $\succ$ is *separable* (Debreu, 1954) if and only if for all $i \leq p$, $x_i, y_i \in D_i$ and $(\vec{x}_{-i}, \vec{y}_{-i}) \in D_{-i}$ we have $(x_i, \vec{x}_{-i}) \succ (y_i, \vec{x}_{-i})$ if and only if $(x_i, \vec{y}_{-i}) \succ (y_i, \vec{y}_{-i})$. When $\succ$ is separable, the $\succ^i$ is defined by $x_i \succ^i y_i$ if and only if $(x_i, \vec{x}_{-i}) \succ (y_i, \vec{x}_{-i})$ for an arbitrary $\vec{x}_{-i}$.

Given n voters, a *profile* is a collection $R = \langle \succ_1, \ldots, \succ_n \rangle$ of linear orders on A. A profile R is separable when each of $\succ_i$ is separable. Given a separable profile over a domain composed of binary variables, the *simultaneous*[3] *majority outcome* $m(R)$ is defined by $m(R) = (x_1^*, \ldots, x_p^*)$, where a majority of voters prefer x_i^* to the opposite value $1 - x_i^*$ (for the sake of simplicity we assume an odd number of voters, so that there are no ties and the majority outcome is uniquely defined). When variables are not binary, simultaneous voting uses a specific voting rule for each variable. *In the rest of this section, for the sake of simplicity we focus on binary variables.*

In simultaneous voting, each voter only has to report a ranking over D_i for each i, therefore the communication requirement of simultaneous voting is $O(n \sum_i |D_i| \log |D_i|)$. Since all variables are binary, each voter has only to report a ballot consisting of a (preferred) value for each variable, hence the requirement complexity is $O(np)$. For instance,

[3] The terminology "simultaneous voting" is used by Lacy and Niou (2000). It is also called *standard voting* by Brams et al. (1997), *propositionwise aggregation* by Brams et al. (1998), and *seat-by-seat voting* by Benoît and Kornhauser (2010). We choose the terminology 'simultaneous voting' although it is a little bit ambiguous: it does not only mean that voters vote simultaneously, but also that they vote *simultaneously and separately on all issues*. Approaches reviewed in Section 10.3 do not satisfy that, although, in some sense, they may also be considered as being 'simultaneous' in the sense that all voters vote simultaneously.

if a voter prefers 1_1 over 0_1, 0_2 over 1_2 and 1_3 over 0_3, then she reports the ballot 101, which represents $1_1 0_2 1_3$. In this case, separability implies that this ballot also corresponds to her most preferred alternative: in other words, simultaneous voting is a *tops-only* voting rule.

How good is simultaneous voting? We know already that it has a low communication cost, as well as a very low computation cost when variables are binary (and more generally for most commonly used voting rules, when variables are not binary). Things become much worse when we turn to the *quality* of the outcome. Even though there is no single way of measuring the quality of the outcome, in most cases a popular type of negative results is to show that simultaneous voting is prone to *paradoxes*, called *multiple election paradoxes*, or *paradoxes of multiple referenda* (see next subsection). Positive results, on the other hand, proceed by showing that some desirable axiomatic properties are satisfied.

A key issue in assessing the quality of the outcome is whether we assume voters to have separable preferences or not. We start with the general case.

10.2.2 Simultaneous voting with nonseparable preferences

When preferences are not separable, a first problem that arises is that if a voter's preferred alternative is $\vec{x} = (x_1, \ldots, x_p)$, then there is no guarantee that she will report x_1 as her preferred value for X_1. For example, if her preference relation is $111 \succ 000 \succ 001 \succ 010 \succ 100 \succ 110 \succ 101 \succ 011$, then for three of the four combinations of values of X_2 and X_3, 0_1 is preferred to 1_1, and similarly for X_2 and X_3; therefore, even though the value of X_1 in her preferred alternative is 111, she might well report 0_1 as her preferred value for X_1, as well as 0_2 and 0_3 as her preferred values for X_2 and X_3. A voter whose preferred value for X_i is always the value of X_i in her preferred alternative will be called *optimistic*, because reporting in such a way comes down to assuming that the outcome over all other issues will be the most favorable one. In our example, if the voter is optimistic then she should vote for 1_1, for 1_2 and for 1_3. More generally, choosing a preferred value to report for an issue depends on the voter's beliefs about the outcomes of the other issues, which in turn depends on her beliefs about the other voters' behavior. A game-theoretic analysis of this complex phenomenon is given by Ahn and Oliveros (2012).

The multiple election paradoxes studied by Brams et al. (1998) and Scarsini (1998) occur when the winner of simultaneous voting receives the fewest votes.

Example 10.2 (Brams et al., 1998) There are 3 issues and 3 voters voting respectively for 110, 101 and 011. Simultaneous voting outputs 111, whereas 111 receives support from none of the voters.

An even more striking paradox, again due to Brams et al. (1998), is obtained with 4 issues, with the outcome being 1111 whereas 1111 is the *only* alternative that receives no vote and 0000 is the only alternative that receives the most votes.

Whether these are paradoxical outcomes or not depends on the voters' preferences over the whole domain. The implicit assumption in these examples is that voters have *plurality-based preferences*: each voter i submits her preferred alternative $\vec{x}^i$, prefers $\vec{x}^i$ to all other alternatives, and is indifferent between any two alternatives different from $\vec{x}^i$. Such dichotomous, plurality-based preferences are not separable. Under this assumption, in the three-issue example above, 111 is Pareto-dominated by 110, 101 and 011; in the four-issue example, 1111 is Pareto-dominated by *all other alternatives*, which is clearly a very undesirable outcome.

The assumption that preferences are plurality-based is very demanding and is very often not plausible. A weaker assumption is *top-consistency*: which only states that each voter prefers her reported alternative $\vec{x}^i$ to all other alternatives. If, instead of assuming plurality-based preferences, we only assume top-consistency, the quality of the outcome can be even worse, as it can be seen on the following example.

Example 10.3 We have two issues: building a swimming pool or not (1_S or 0_S), and building a tennis court or not (1_T or 0_T). We have $2k+1$ voters:

- k voters: $1_S0_T \succ 0_S1_T \succ 0_S0_T \succ 1_S1_T$;
- k voters: $0_S1_T \succ 1_S0_T \succ 0_S0_T \succ 1_S1_T$;
- 1 voter: $1_S1_T \succ 0_S1_T \succ 1_S0_T \succ 0_S0_T$.

It is unclear what the first k voters will report when choosing between 1_S and 0_S. Indeed, their preferences are non-separable: they prefer the swimming pool to be built if the tennis course is not, and *vice versa*. Now, if they vote for 1_S, their vote, when it is a decisive, leads to either 1_S0_T or 1_S1_T, that is, to the voter's best alternative or to her worst alternative. On the other hand, voting for 0_S, again when it is a decisive vote, leads to either 0_S0_T or 0_S1_T, that is, to one of the voter's 'intermediate' alternatives. This shows why the first k voters may be hesitant to vote for 1_S or for 0_S. They may also be hesitant to vote for 1_T or for 0_T, although the situation here is a bit different (a decisive vote for 1_T leads to the second-ranked or to the worst alternative, while a decisive vote for 1_T leads to the best or to the third-ranked alternative). If we assume that these first k voters do not have any knowledge about the others' preferences (or even if they do, but do not use this information for voting strategically), then these voters will feel ill at ease when voting and may experience regret once they know the final outcome (*e.g.*, if they vote for 1_S, wrongly believing that the group will decide not to build the tennis court). The case for the next k voters is symmetric (with the roles of S and T being swapped). Only the last voter, who has separable preferences, has no problem voting for 1_S and for 1_T and does not experience regret after the election. The analysis of the paradox by Lacy and Niou (2000) assumes that voters choose to vote optimistically (thus the first k voters would vote for 1_S)[4]: under this assumption, the simultaneous voting outcome 1_S1_T *is ranked last by all but one voters*.

Take the profile in Example 10.1 as another example. Assuming again that voters vote optimistically, the simultaneous voting outcome (111) is ranked *last* by *all voters*, which is, arguably, a very bad decision.

[4] This assumption is often reasonable, even if it has a certain level of arbitrariness.

These paradoxes are partly due to the implicit assumption that voters do not have any knowledge about other votes. However, even if voters' preferences are common knowledge, and voters vote strategically, strong paradoxes can still arise (Lacy and Niou, 2000) (see also Section 10.4.2). As argued by Saari and Sieberg (2001), the source of these paradoxes is the loss of information that occurs when separating the input profile into smaller profiles for single issues.

10.2.3 Simultaneous voting with separable preferences

Assuming separability allows us to avoid some of the paradoxes described above. First, when all voters have separable preferences, they can vote safely for their preferred value, for each one of the issues, and without any risk of experiencing regret (this is called *simple voting* by Benoît and Kornhauser (1991)). Second, under the separability assumption, simultaneously voting enjoys some desirable properties, including the election of a Condorcet winner when there is one (Kadane, 1972).[5]

However, some paradoxes still remain. In particular, the outcome may be Pareto-dominated by another alternative (Özkal-Sanver and Sanver, 2006; Benoît, 2010), as shown in the following example.

Example 10.4 (Özkal-Sanver and Sanver, 2006) We have three issues, and three voters whose preferences are as follows:

- Voter 1: $111 \succ 011 \succ 101 \succ 001 \succ 110 \succ 010 \succ 100 \succ 000$;
- Voter 2: $100 \succ 000 \succ 101 \succ 001 \succ 110 \succ 010 \succ 111 \succ 011$;
- Voter 3: $010 \succ 011 \succ 000 \succ 001 \succ 110 \succ 111 \succ 100 \succ 101$.

Note that these preferences are separable: voter 1 prefers 1_1 to 0_1, whichever the values of X_2 and X_3 (that is, she prefers 100 to 000, 101 to 001, 110 to 010 and 111 to 011), prefers 1_2 to 0_2, whichever the values of X_1 and X_3, and 1_3 to 0_3, whichever the values of X_1 and X_2. Similar reasoning shows that preferences for voters 2 and 3 are also separable. The outcome of simultaneous voting is 110, which is Pareto-dominated by 001, that is, all three voters prefer 001 to 110.

Benoît and Kornhauser (2010) prove a more general result. One may wonder whether there could be rules other than issue-wise majority that would escape the paradox. Unfortunately this is not the case: as soon as there are at least three issues, or when there are exactly two issues, one of which has at least three possible values, then simultaneous voting is efficient if and only if it is dictatorial. This result was generalized to irresolute voting rules by Xia and Lang (2009).

[5] This holds both in the assumption that voters vote sincerely and in the assumption that voters' preferences are common knowledge and voters vote strategically.

10.2.4 Discussion

Evaluating simultaneous voting on the criteria evoked in the introduction (Section 10.1), it is now clear that simultaneous voting has a low communication cost, and has also a low computation cost, provided that the "local" rules used to determine the outcome for each variable are easy to compute (which is obviously the case if variables are binary). Then, there are two possibilities: either we are able to assume separability, and in that case the outcome has some quality guarantees (even in this case it remains prone to some paradoxes, see Section 10.2.3); or we do not assume separability, and then the quality of the outcome can be extremely bad. We note that separability is a very strong assumption: the proportion of preferences on a combinatorial domain that are separable is very low (Bradley et al., 2005), and there are many domain-specific arguments (such as budget constraints) showing that in many domains, it is almost hopeless to expect that voters' preferences are separable.

10.3 Approaches based on completion principles

One way of escaping the paradoxes of simultaneous voting, discussed in particular by Brams et al. (1997) and Lacy and Niou (2000), consists in having voters *vote on combinations* (or *bundles*) or values. This section discusses various ways of implementing this. Before addressing several classes of more complex methods, we mention three very simple solutions, which are relevant in some cases.

1. Voters rank all alternatives (*i.e.*, all combinations) and a classical voting rule, such as Borda, is used.
2. Voters give only their top alternatives, and the plurality rule is used.
3. Voters rank a small number of pre-selected alternatives, and use a classical voting rule.

The first way is clearly the best method when the set of alternatives is small (say, up to four of five binary variables). It becomes inapplicable when the number of issues becomes more than a few, since asking voters to rank explicitly more than a few dozens of alternatives is already hopeless. The second way, advocated by Brams et al. (1997), has the obvious advantage that it is relatively inexpensive both in terms of communication and computation; it is feasible provided that the set of alternatives is small enough with respect to the set of voters; when this is not the case, the plurality votes are likely to be completely dispersed (for instance, with 10 binary variables and 100 voters, the number of alternatives (2^{10}) is ten times larger than the number of voters and it may plausibly happen that each alternative will get no more than one vote), which does not help much making a decision. The third way avoids both problems, but the arbitrariness of the preselection phase can make the whole process very biased, and gives too much power to the authority who determines the preselected alternatives.

Ideally, methods should avoid the arbitrariness of methods 2 and 3 and the communication requirement of method 1. Recall that simultaneous voting has a low elicitation cost, at the price of considering all issues independently. One way of introducing links between issues while keeping the low communication cost of simultaneous voting consists in asking voters to specify a small part of their preference relations and then complete them into full (or, at last, more complete) preference relations using a fixed *completion* (or *extension) principle*. After this completion has been performed, we may apply a classical voting rule, or a voting rule specifically designed for this extension principle. We consider several families of completion principles, in increasing order of sophistication.

- *Top-based input*: the voters submit their preferred alternative and the completion principle makes use of a predefined distance over alternatives (typically, the Hamming distance); the completion principle ranks all alternatives according to their proximity to the preferred alternative.
- *Singleton ranking-based input*: this completion principle works only for binary variables; voters specify a ranking over single issues; the completion principle then extends it into a preference relation over all alternatives. This class of methods is often used for the selection of a set of items (typically, a committee).
- *Hypercube-based input*: the input consists of a compact representation of each voter's preference between all pairs of alternatives that are identical on all issues but one (this set of pairs of alternatives is also called the hypercube associated with A).
- *Inputs based on more sophisticated inputs*, such as conditionally lexicographic preference trees, weighted or prioritized logical formulas, generalized additive independence networks, or weighted constraints.

10.3.1 Top-based inputs and distance-based completion principles

One way to express a small part of the agents' preferences and to complete them automatically consists of asking each voter to specify *her top alternative* $\vec{x}^*$, and then applying the following intuitive completion principle: the closer to $\vec{x}^*$ with respect to a predefined distance d between alternatives, the more preferred. Formally: given a voter's top alternative $\vec{x}^*$, $\succ$ is d-*induced* if for all $\vec{y}, \vec{z} \in A$, $\vec{y} \succ \vec{z}$ iff $d(\vec{y}, \vec{x}^*) < d(\vec{z}, \vec{x}^*)$, and $\succ$ is d-*consistent* iff for all $\vec{y}, \vec{z} \in A$, $d(\vec{y}, \vec{x}^*) < d(\vec{z}, \vec{x}^*)$ implies $\vec{y} \succ \vec{z}$.

A trivial choice of a distance is the Dirac (or drastic) distance, defined by $d(\vec{x}, \vec{y}) = 0$ if $\vec{x} = \vec{y}$ and $d(\vec{x}, \vec{y}) = 1$ if $\vec{x} \neq \vec{y}$. We recover here the plurality-based extension principle discussed in Section 10.2.2, which can thus be seen as a distance-based extension.

While many choices of a nontrivial distance can be made, the most obvious one is perhaps the *Hamming distance* d_H: for all $\vec{x}, \vec{y} \in A$, $d_H(\vec{x}, \vec{y})$ is the number of issues on which $\vec{x}$ and $\vec{y}$ disagree. We say that $\succ$ is *Hamming-induced* (resp. *Hamming-consistent*) iff it is d_H-*induced* (resp. d_H-*consistent*).

Once such a preference extension principle has been fixed, we can apply a voting rule to select the winner. A prominent example of such a rule is *minimax approval voting*, defined by Brams et al. (2007) in the context of committee elections (although there is no reason not to apply it in more general contexts); for this reason, we describe the rule in a committee election setting, thus, with binary variables (also, it is not entirely trivial to extend minimax approval voting to nonbinary domains). There are n voters, p candidates, $k \leq p$ positions to be filled; each voter casts an approval ballot $V_i = (v_i^1, \ldots, v_i^p) \in \{0,1\}^p$, where $v_i^j = 1$ if voter i approves candidate j. Then for every subset S of k candidates, let $d_H(S, (V_1, \ldots, V_n)) = \max_{i=1,\ldots,n} d_H(S, V_i)$ be the largest Hamming distance between S and a ballot. Minimax approval voting selects a committee S minimizing $d_H(S, (V_1, \ldots, V_n))$. Minimax approval voting makes sense if there are few voters, but much less so in large electorates, because a single voter can have a huge influence, even if everyone else agrees. Note that minimizing the *sum* of Hamming distances would be equivalent to outputting the candidates with the k largest approval scores (see Section 10.3.2).

Example 10.5 Let $n = 4$, $p = 4$, $k = 2$. The ballots are defined as follows, together with the computation of Hamming distance between the votes and any subset S composed of 2 candidates (there are 6 such candidates):

	$V_1 : 1110$	$V_2 : 1101$	$V_3 : 1010$	$V_4 : 1010$	max
1100	1	1	2	2	**2**
1010	1	3	0	0	3
1001	3	1	3	3	3
0110	1	3	2	2	3
0101	3	1	4	4	4
0011	3	3	2	2	3

The winning committee under minimax approval voting is 1100. Minimizing the *sum* of Hamming distances would lead to selecting 1010.

Because there are $\binom{p}{k}$ possible committees, winner determination for minimax approval voting is computationally intractable: finding a winning committee is NP-hard (Frances and Litman, 1997). LeGrand et al. (2007) give a polynomial-time 3-approximation algorithm; a better approximation (with ratio 2) is given by Caragiannis et al. (2010).

Another line of research that makes use of preference extensions based on the Hamming distance is that of Laffond and Lainé (2009) and Cuhadaroğlu and Lainé (2012). Recall from Section 10.2 that even if voters' preferences are separable, the simultaneous voting outcome can be Pareto-dominated. If furthermore voters' preferences are Hamming-consistent, then two positive results arise: (a) the simultaneous voting outcome cannot be Pareto-dominated, (b) the simultaneous voting outcome is in the top cycle (*a fortiori*, simultaneous voting is Condorcet-consistent). However, weaker negative results remain: not only may the outcome be majority-defeated but it can also fail to be in the uncovered set (Laffond and Lainé, 2009). To which extent are the positive results specific to the Hamming extension principle? An answer is given by Cuhadaroğlu and Lainé (2012), who show that under some mild conditions, the largest set of preferences for which the simultaneous voting outcome is Pareto-efficient is the set of Hamming-consistent preferences.

Distance-based approaches have a lot in common with *belief merging* (see a recent survey by Konieczny and Pérez (2011)), which aggregates several propositional formulas K_1, ..., K_n into a collective propositional formula $\Delta(K_1, \ldots, K_n)$. The set of alternatives corresponds to the set of propositional valuations (or interpretations). Perhaps the most well-studied family of belief merging operators is the class of *distance-based* merging operators: there is a predefined, integer-valued, agent-independent distance d over propositional valuations (typically, the Hamming distance), and a symmetric, non-decreasing aggregation function $\star$ over integers, and the output is a formula whose models minimize $\star\{d(., K_i)|i = 1 \ldots n\}$, where $d(\vec{x}, K_i) = \min_{\vec{y} \models K_i} d(\vec{x}, \vec{y})$. The complexity of distance-based belief merging is addressed by Konieczny et al. (2004). Although coming from a different area, distance-based belief merging shares a lot with combinatorial voting with distance-based preference extensions (especially minimax approval voting). Two important differences are that in belief merging: (a) the input may consist of a *set* of equally most preferred alternatives, rather than a single one; and (b) the input is represented compactly by a logical formula.

10.3.2 Input based on rankings over single variables

In this section we focus specifically on the selection of a *collective set of items* S by a group of agents. The meaning we give here to "items" is extremely general and can cover a variety of situations, with two typical examples being *committee elections*, where the "items" are representatives, and *group recommendations*, where items are objects such as books, movies, etc.

Formally, this can be cast as a combinatorial domain where the set of binary issues is $\mathcal{X} = \{X_1, \ldots, X_p\}$, with $D_i = \{0_i, 1_i\}$ for each i. These binary issues correspond to a set of items $C = \{c_1, \ldots, c_p\}$, where $X_i = 1_i$ (resp. 0_i) means that item c_i is (resp. is not) in the selection S. Because of the focus on the selection of a subset of items, we change the notational convention by denoting an alternative $\vec{x} \in A = \{0_1, 1_1\} \times \ldots \times \{0_p, 1_p\}$ as a *subset of issues* S composed of items c_i's with $X_i = 1_i$. Thus, alternatives are elements of 2^C.

In most cases, the set of feasible subsets is a proper subset of 2^C, defined by a *constraint* Γ restricting the set of feasible or allowed subsets. In committee elections, the most common constraints are *cardinality constraints* that restrict the size of a committee, by specifying an exact size k, or a lower and/or an upper bound. More generally, Lu and Boutilier (2011) consider budget constraints, defined by a price for each item and a maximum total cost – hence the terminology *budgeted social choice*.

The approaches discussed in this section proceed by first eliciting from each agent some preference information (typically, a ranking) over single *items*, then extending these preferences over single items to preferences over sets of items, and finally selecting a set of items S.

The most obvious way of doing so is *multiwinner approval voting* (which can, to some extent, be seen as the multiwinner version of simultaneous voting): each voter approves as many candidates as she wants, and the winners are the k candidates approved most often. In *single nontransferable vote (SNTV)* and *bloc voting*, there is an additional restriction on the number of candidates approved: 1 in SNTV and k in bloc voting (these rules are thus multiwinner versions of plurality and k-approval, respectively). Finally, in *cumulative voting*, voters distribute a fixed number of points among the candidates, and the winners are the k candidates maximizing the number of points. The common point of all these rules is that voters' preferences are assumed to be separable; reformulated in terms of preference extensions, each input defines a score over single candidates, and the total score of a candidate is the *sum* of the scores it gets from the voters. Computational aspects of strategic behavior (manipulation by a single voter and control by the chair) for these multiwinner voting rules have been studied by Meir et al. (2008).

In the remainder of this section, we focus on classes of methods where the input consists of *rankings* over single items.

In committee elections, where the items are individuals supposed to represent the voters, the rationale for the last step is that a committee election is used to elect an assembly whose members will make decisions on behalf of the society. As argued by Betzler et al. (2013), finding a committee of representatives should satisfy two criteria: *representativity* (the composition of the committee should globally reflect the will of the voters), and *accountability* (each voter should be represented by a given member of the committee). In consensus recommendations, the rationale for the last step is that each user will benefit from the best option according to her own preferences (Lu and Boutilier, 2011); in this case, the "representative" of a voter is her most preferred item in S.[6] [7] The latter interpretation leads to an obvious choice for defining representative items for voters: if the set of items S is chosen, then the representative item of voter i is $c \in S$ if $c \succ_i c'$ for all $c' \in C \setminus \{c\}$. Alternatively, we say that each agent is *represented* by an item in S. In committee elections, this principle is the basis of the *Chamberlin and Courant* multiwinner election scheme (discussed later).

We now describe these multiwinner election schemes (grouped under the terminology "fully proportional representation") more formally. For each voter i and each item c there is a *misrepresentation value* $\mu_{i,c}$, representing the degree to which item c misrepresents voter i. *A positional misrepresentation function* makes use of a scoring vector $\vec{s} = \langle s_1, \ldots, s_p \rangle$ such that $s_1 \leq \ldots \leq s_p$. In particular, the Borda scoring vector $\vec{s_B}$ is defined by $s_k = k$ for all k. By $pos_i(c)$ we mean the position of item c in i's preference ranking (from 1 for the most preferred item to p for the least preferred one). The misrepresentation function induced by $\vec{s}$ is $\mu_{i,c} = s_{pos_i(c)}$. Intuitively, s_i is the amount of dissatisfaction that a voter derives from being represented by an alternative that she ranks in position i. Another simple

6 As discussed by Lu and Boutilier (2011), this can also be seen as a segmentation problem (Kleinberg et al., 2004), where one more generally seeks k solutions to some combinatorial optimization problem that will be used by $n \geq k$ different users, each with a different objective value on items; optimization requires segmenting users into k groups depending on which of the k items gives them the greatest benefit.

7 Skowron et al. (2015) generalize this scheme by taking into account more than one item by agent, but still giving more importance to an agent's most preferred item than to her second best preferred item, etc.

way of defining a *misrepresentation based on approval ballots* is: every voter submits a subset of candidates that she approves, and $\mu_{i,c}$ is 0 if i approves c, and 1 otherwise.

An *assignment function* π maps every voter to an item in the selected subset S. The misrepresentation of voter i under π is $\mu_{i,\pi(i)}$. Once individual misrepresentation has been defined, we need to define the *global misrepresentation* of the society when selecting a subset S of items. There are two traditional ways of doing so: *utilitarianism* (global misrepresentation is the sum of all individual misrepresentation) and *egalitarianism* (global misrepresentation is the misrepresentation of the least well-represented agent). Formally, the global misrepresentation of assignment π is defined as:

- (utilitarianism) $\mu_U(\pi) = \sum_{i \leq n} \mu_{i,\pi(i)}$.
- (egalitarianism) $\mu_E(\pi) = \max_{i \leq n} \mu_{i,\pi(i)}$.

Finally, let $\mathcal{F}$ be the set of feasible subsets of items; typically, if k items are to be elected then $\mathcal{F}$ is the set $\mathcal{S}_k$ of all subsets of C of size k.

The *Chamberlin and Courant* scheme (Chamberlin and Courant, 1983) simply outputs the committee of size k that minimizes μ_U. Because there is no constraint on the assignment function, every voter is assigned to her preferred item in the selected subset S. That is, $\pi(i) = \arg\min_{c \in S} \mu_{i,c}$. Then, her misrepresentation when selecting the feasible subset S is equal to $\mu_{i,S} = \min_{c \in S} \mu_{i,c}$. The best committee is then the feasible subset S minimizing $\mu_U(\pi)$ (under utilitarianism) or S that minimizes $\mu_E(\pi)$ (under egalitarianism). The *Monroe* scheme (Monroe, 1995) additionally requires that the assignment π is balanced: each candidate in S must be assigned to at least $\lfloor n/k \rfloor$ voters.[8] Formally, the Monroe scheme selects the allocation π minimizing $\mu_U(\pi)$ subject to the constraints $|\pi^{-1}(s)| \geq \lfloor \frac{n}{k} \rfloor$ for all $s \in Range(\pi)$.

Budgeted social choice (Lu and Boutilier, 2011) generalizes Chamberlin-Courant by redefining feasibility via a budget constraint: each item c has a fixed cost (to be counted if x is selected) and a unit cost (to be counted k times if k agents are represented by the item), the maximum budget is K, and $\mathcal{F}$ is the set of all assignments with total cost $\leq K$.

The egalitarian version of multiwinner schemes is due to Betzler et al. (2013). Elkind et al. (2014) discuss some properties of multiwinner voting schemes.

Example 10.6 Let $C = \{c_1, c_1, c_3, c_4\}$, $K = 2$, and the following 4 agents' preferences:

$$\begin{array}{ll} \langle & c_1 \succ c_2 \succ c_3 \succ c_4, \\ & c_1 \succ c_2 \succ c_3 \succ c_4, \\ & c_1 \succ c_3 \succ c_2 \succ c_4, \\ & c_1 \succ c_3 \succ c_2 \succ c_4, \\ & c_2 \succ c_4 \succ c_3 \succ c_1, \\ & c_4 \succ c_3 \succ c_2 \succ c_1 \quad \rangle \end{array}$$

For the Borda misrepresentation function, the optimal Chamberlin-Courant committee of

[8] In indirect democracy, that is, when the set of representatives has to make a decision on behalf of the society, it may be a good idea to give more power to people who represent more people than to those who represent less people; for instance, Chamberlin and Courant suggested to give to each committee member a weight equal to the number of voters she represents.

2 items is $\{c_1, c_4\}$, whereas for Monroe it is $\{c_1, c_2\}$. For the egalitarian versions, both $\{c_1, c_4\}$ and $\{c_2, c_3\}$ are optimal for Chamberlin-Courant and $\{c_2, c_3\}$ is optimal for Monroe.

Because the set of feasible subsets is generally exponentially large, finding the optimal subset is highly nontrivial. Brams and Potthoff (1998) were the first to discuss the computation of the Chamberlin-Courant and the Monroe voting schemes, showing that the optimal committee can be determined using integer programming. This provides a method that works in practice when the number of voters and items are small, but may not scale up well. They formulate an improved integer program for settings where the number of agents is large, but this modified integer program is still too large to be solved when the number of items is large.

One cannot really do better in the general case; indeed, we have the following hardness results:

- Winner determination for the Chamberlin-Courant and the Monroe schemes with approval ballots are both NP-complete (Procaccia et al., 2008);
- Winner determination for the Chamberlin-Courant scheme with the Borda misrepresentation function is NP-complete (Lu and Boutilier, 2011);
- Winner determination for the minimax versions of the Chamberlin-Courant and Monroe schemes is NP-complete (Betzler et al., 2013).

Some slightly more positive results are obtained:

- *Parameterized complexity*. Procaccia et al. (2008) show that winner determination for Chamberlin-Courant and Monroe is tractable for small committees: if the size of the subset to be selected is constant, then winner determination is polynomial for both voting schemes. Betzler et al. (2013) investigate further the parameterized complexity of fully proportional representation by establishing a mixture of positive and negative results: they mainly prove fixed-parameter tractability with respect to the number of candidates or the number of voters, but fixed-parameter intractability with respect to the number of winners.
- *Approximation*: Lu and Boutilier (2011) give a polynomial algorithm with approximation ratio $1 - \frac{1}{e}$ for Chamberlin-Courant with the Borda misrepresentation function. Skowron et al. (2013a,c) give further approximability results.
- *Domain restrictions*: for single-peaked profiles, most multi-winner problems discussed above become polynomial; the only rule that remains NP-hard for single-peaked electorates is the classical Monroe rule (Betzler et al., 2013). These results are extended by Cornaz et al. (2012) to profiles with bounded *single-peaked width*, and by Yu et al. (2013) who consider profiles that are more generally *single-peaked on a tree*. Skowron et al. (2013b) address the case of single-crossing profiles.

Finally, the generalization of full proportional representation schemes to incomplete preferences was considered by Lu and Boutilier (2013) (see also Chapter 10 (Boutilier and Rosenschein, 2015)).

The notion of *Condorcet winning set*(Elkind et al., 2015) also evaluates a subset according to a best item in it. The criteria for selecting a "best" subset does not use a misrepresentation function but is simply based on the Condorcet principle: $S \subseteq C$ is a Condorcet winning set if for every $z \notin S$, a majority of voters prefers *some* $s \in S$ to z. For every m-candidate profile, there is a Condorcet winning set of size at most $\log_2 m + 1$, therefore, finding a Condorcet winning set can be done by enumerating all subsets of candidates of size $\lfloor \log_2 m \rfloor + 1$, *i.e.*, in *quasipolynomial* time. It may actually be even easier: it is an open issue whether for all k there exists a profile for which the smallest Condorcet winning set has size k.

10.3.3 Hypercube-based inputs

Specifying top-based inputs (respectively, rankings over variables) needs $O(np)$ (respectively, $O(np \log p)$) space, hence the communication requirement of the two previous subclasses of methods is low: each agent needs only to report $O(np)$ (respectively, $O(np \log p)$) bits to the central authority. On the other hand, their applicability is very weak, because only a tiny fraction of preference relations comply with the required domain restrictions. We now consider more expressive approaches that are based on *compact representations*: the votes, or a significant part of the votes, are not given extensively but are described in some formal language that comes with a function mapping any input of the language to a (partial or complete) vote (Lang, 2004). Formally, a *compact preference representation language* is a pair $L = \langle \Sigma_L, I_L \rangle$ where Σ_L is a formal language, and I_L is a function from Σ_L to the set of preference relations over A. $I_L(\Sigma_L)$ is the set of all preference relations expressible in L. A language L_1 is more expressive than L_2 if $I_{L_1}(\Sigma_{L_1}) \supset I_{L_2}(\Sigma_{L_2})$ and more succinct than L_2 if there is a function $f : \Sigma_{L_2} \to \Sigma_{L_1}$ and a polynomial function *pol* such that for all $\sigma \in \Sigma_{L_2}$, we have (i) $|f(\sigma)| \leq pol(|\sigma|)$ and (ii) $I_{L_1}(f(\sigma)) = I_{L_2}(\sigma)$. Conditions (i) and (ii) together mean that any preference relation expressible in L_2 can also be expressed in L_1 without a superpolynomial increase in the size of expression.

If a language L is totally expressive (*i.e.*, $I_L(\Sigma_L)$ is the set of *all* rankings over D) then the worst-case size necessary for expressing a ranking is exponentially large in the number of variables.[9] Therefore, there is a tradeoff to be made between having a fully expressive language which, at least for some preference relations, will not be compact at all, or making a domain restriction that will allow for a compact input in all cases.

Some of the solutions advocated in the previous sections were, to some extent, making use of very rough compact preference representation languages. Expressing only the top

[9] A simple proof of this fact in the case of binary variables: for p variables there are $(2^p)!$ possible rankings, and the best we can do to express a ranking is to use $\log((2^p)!)$ bits in the worst case; and $\log((2^p)!)$ is exponential in p.

alternative, say 111, is a compact representation of the partial preference relation

$$111 \succ A \setminus \{111\}$$

or, in the case of the Hamming distance completion, of the complete preorder

$$111 \succ \begin{array}{c} 110 \\ 101 \\ 001 \end{array} \succ \begin{array}{c} 100 \\ 010 \\ 001 \end{array} \succ 000$$

Expressing a ranking over single items is a compact representation of a partial or complete preorder over committees: for 2-committees and the Chamberlin-Courant scheme, for instance, $1 \succ 2 \succ 3 \succ 4$ is a compact representation of

$$\begin{array}{c} \{1,4\} \\ \{1,2\} \\ \{1,3\} \end{array} \succ \begin{array}{c} \{2,3\} \\ \{2,4\} \end{array} \succ \{3,4\}$$

As said above, these first two compact representation languages are admittedly very compact, but also very inexpressive. We now give some examples of more expressive languages.

The first compact representation language we consider is that of *conditional preference networks* (CP-nets). CP-nets (Boutilier et al., 2004) allow for a compact representation of the *preference hypercube* associated with a preference relation over D. Given a preference relation $\succ$ over $D = \prod_{i=1}^{n}\{0_i, 1_i\}$, the preference hypercube $\succ_H$ is the restriction of $\succ$ to the set of pairs of alternatives $\vec{x}$, $\vec{y}$ differing on only one variable (such as, for instance, 0101 and 0111). CP-nets are based on the notion of *conditional preferential independence* (Keeney and Raiffa, 1976): given a strict preference relation $\succ$, $X_i \in \mathcal{X}$, $Y \subseteq \mathcal{X} \setminus \{X_i\}$ and $Z = \mathcal{X} \setminus (\{X_i\} \cup Y)$, we say that X_i is preferentially independent of Y given Z with respect to $\succ$ if for any $x_i, x_i' \in D_i, \vec{y}, \vec{y}' \in D_Y$, and $\vec{z} \in D_Z$, we have $(x_i, \vec{y}, \vec{z}) \succ (x_i', \vec{y}, \vec{z})$ if and only if $(x_i, \vec{y}', \vec{z}) \succ (x_i', \vec{y}', \vec{z})$. A CP-net $\mathcal{N}$ over A consists of two components.

- The first component is a directed graph G expressing preferential independence relations between variables: if $Par_G(X_i)$ denotes the set of the parents of X_i in G, then every variable X_i is preferentially independent of $\mathcal{X} \setminus (Par(X_i) \cup \{X_i\})$ given $Par(X_i)$.
- The second component is, for each variable X_i, a set of linear orders $\succ_{\vec{u}}^{i}$ over D_i, called conditional preferences, for each $\vec{u} \in D_{Par_G(X_i)}$. These conditional preferences form the *conditional preference table* for issue X_i, denoted by $CPT(X_i)$.

The preference relation $\succ_{\mathcal{N}}$ induced by $\mathcal{N}$ is the transitive closure of

$$\{(a_i, \vec{u}, \vec{z}) \succ (b_i, \vec{u}, \vec{z}) \mid i \leq p; \vec{u} \in D_{Par_G(X_i)}; a_i, b_i \in D_i, a_i \succ_{\vec{u}}^{i} b_i; \vec{z} \in D_{-(Par_G(X_i) \cup \{X_i\})}\}$$

When all issues are binary, $\succ_{\mathcal{N}}$ is equivalent to a preference hypercube and $\mathcal{N}$ is a compact representation of this preference hypercube, whose size is the cumulative size of all its conditional preference tables.

Example 10.7 Let $p = 3$. The following figure represents a CP-net $\mathcal{N}$ together with its induced preference hypercube $\succ_{\mathcal{N}}$. For the sake of simplicity, 000 represents the alternative $0_1 0_2 0_3$, etc.

$$X_1 \longrightarrow X_2 \longrightarrow X_3$$

$$0_1 \succ 1_1 \qquad \begin{array}{l} 0_1 : 0_2 \succ 1_2 \\ 1_1 : 1_2 \succ 0_2 \end{array} \qquad \begin{array}{l} 0_2 : 0_3 \succ 1_3 \\ 1_2 : 1_3 \succ 0_3 \end{array}$$

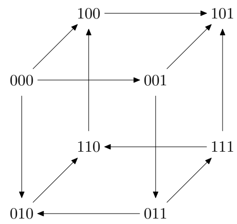

Group decision making in multi-issue domains via CP-net aggregation has been considered in a number of papers, which we briefly review in a nonchronological order. We will discuss in Section 10.4 the role of CP-nets in sequential voting (Lang and Xia, 2009; Airiau et al., 2011): this way of proceeding sequentially leads to interleave elicitation and aggregation, and elicits only a small part of the voters' CP-nets. Another way of proceeding consists in first eliciting the voters' CP-nets entirely, then proceeding to aggregation. Then, two ways are possible.

The first aggregation consists of mapping each of the individual CP-nets to its associated preference relation, and aggregating these into a collective preference relation. This method was initiated by Rossi et al. (2004), who consider several such aggregation functions, and was studied further by Li et al. (2010), who give algorithms for computing Pareto-optimal alternatives with respect to the preference relations induced by the CP-nets, and fair alternatives with respect to a cardinalization of these preference relations.

A second technique, considered by Xia et al. (2008), Li et al. (2011), and Conitzer et al. (2011), consists of aggregating the individual CP-nets into a collective CP-net, and then outputting the nondominated alternatives of this collective CP-net. No domain restriction is made on the individual CP-nets. For every set of "neighboring" alternatives (differing only in the value of one issue), a local voting rule (typically majority if domains are binary) is used for deciding the common preferences over this set, and finally, optimal outcomes are defined based on the aggregated CP-net.

Example 10.8 We have two issues, X_1 and X_2, and the following three CP-nets.

voter 1	voter 2	voter 3
$X_1 \longrightarrow X_2$	$X_1 \longleftarrow X_2$	$X_1 \quad X_2$
$0_1 \succ 1_1$; $0_1 : 1_2 \succ 0_2$, $1_1 : 0_2 \succ 1_2$	$0_2 : 1_1 \succ 0_1$, $1_2 : 0_1 \succ 1_1$; $0_2 \succ 1_2$	$1_1 \succ 0_1$; $1_2 \succ 0_2$

The majority aggregation of $\mathcal{N}_1$, $\mathcal{N}_2$ and $\mathcal{N}_3$ is the following CP-net, depicted with its induced preference relation.

$$X_1 \longleftrightarrow X_2$$

$0_2 : 1_1 \succ 0_1$	$0_1 : 1_2 \succ 0_2$
$1_2 : 0_1 \succ 1_1$	$1_1 : 0_2 \succ 1_2$

$$\begin{tikzcd} 0_1 0_2 & 1_1 0_2 \arrow[l] \arrow[d] \\ 0_1 1_2 \arrow[u] \arrow[r] & 1_1 1_2 \end{tikzcd}$$

The dependency graph of this collective CP-net contains an edge from X_1 to X_2 (resp., from X_2 to X_1) because the dependency graph of voters 1 (resp., 2) CP-net does. In the preference table for X_1, we have $0_2 : 1_1 \succ 0_1$ because voters 2 and 3 (unlike voter 1) prefer 1_1 to 0_1 when $X_2 = 0_2$.

Once the CP-nets from agents have been aggregated to a common CP-net $\mathcal{N}^*$, the next task consists of finding a set of solutions. Because $\mathcal{N}^*$ only specifies pairwise preferences between neighbour alternatives, usual solution concepts are not directly applicable. In particular, there is generally no way of checking whether a Condorcet winner exists; however, we can check if there are *hypercubewise Condorcet winners* (HCW), that is, alternatives that dominate all of their neighbours in $\mathcal{N}^*$. Unlike Condorcet winners, a profile may possess no HCW, one HCW, or several HCW (in Example 10.8 there are two, namely $0_1 1_2$ and $1_1 0_2$). The notion of a HCW was first defined by Xia et al. (2008) and studied further by Li et al. (2011), who study some of its properties and propose (and implement) a SAT-based algorithm for computing them, whereas the probability of existence of a HCW is addressed by Conitzer et al. (2011). More solution concepts (such as the top cycle, Copeland, maximin, or Kemeny) can also be generalized to profiles consisting of preference hypercubes (Xia et al., 2008; Conitzer et al., 2011), while new solution concepts, based on distances between alternatives in the hypercube, have been proposed by Xia et al. (2010).

10.3.4 Conditionally lexicographic preferences

A *conditionally lexicographic preference* can be represented compactly by a *lexicographic preference tree (LP-tree)* (Booth et al., 2010), consisting of (i) a *conditional importance tree*, where each node is labelled by a variable X_i and has either one child, or two children associated with the values 0_i and 1_i taken by X_i; (ii) and, for each node v of the tree, labelled by X_i, a *conditional preference table* expressing a preference order on D_i for all possible combination of values of (some of) the ancestor variables that have not yet assigned a value in the branch from the root to v.

Example 10.9 An LP-tree with $p = 3$ is illustrated in the figure below. The most important variable is X_1, and its preferred value is 0_1; when $X_1 = 0_1$ then the second most important variable is X_2, with preferred value 0_2, then X_3 with preferred value 0_3; when $X_1 = 1_1$ then the second most important variable is X_3, with preferred value 1_3, then X_2, with preferred value 0_2 if $X_3 = 0_3$ and 1_2 if $X_3 = 1_3$. The preference relation induced by an LP-tree compares two alternatives by looking for the first node (starting from the root) that discriminates them: for instance, for $\vec{x} = 111$ and $\vec{y} = 100$, this is the node labelled by X_3 in the branch associated with $X_1 = 1_1$. Because $1_3 \succ 0_3$ at that node,

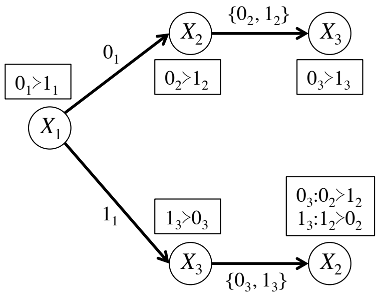

$\vec{x}$ is preferred to $\vec{y}$. The complete preference relation associated with the above LP-tree is $001 \succ 000 \succ 011 \succ 010 \succ 111 \succ 101 \succ 100 \succ 110$.

Assuming preferences are conditionally lexicographic imposes an important domain restriction (as does separability), but for some voting rules, determining the outcome is efficient in communication and computation (Lang et al., 2012). We give an example with 2^{p-2}-approval. Given an LP-tree T compactly expressing a ranking $\succ_T$, an alternative is one of the 2^{p-2} best alternatives (*i.e.*, in the top quarter) if and only if it gives the preferred value to the most important variable (in the example above, $X_1 = 0_1$) and the preferred value to the second most important variable given this value ($X_2 = 0_2$). This gives, for every voter, a conjunction of two literals (here $\neg X_1 \wedge \neg X_2$); the 2^{p-2}-approval winners are exactly those who satisfy a maximal number of such formulas, thus the winner determination problem can be solved using a MAXSAT solver. Note that, although the problem is NP-hard, there are efficient MAXSAT solvers (and a MAXSAT track of the SAT competition). Results about other rules can be found in (Lang et al., 2012).

Example 10.10 Let $n = 3$, $p = 3$, and consider the three LP-trees in Figure 10.1. The first LP-tree is the same as that in Example 10.9, and an alternative is ranked in its top $2^{p-2} = 2$ positions if and only if $\neg X_1 \wedge \neg X_2$ is satisfied. In the second LP-tree, the top 2^{p-2} alternatives can be represented as $\neg X_1 \wedge X_2$. In the third LP-tree, the top 2^{p-2} alternatives can be represented as $X_2 \wedge X_3$. These are the formulas in the MAXSAT instance and the winner for 2-approval is 011.

10.3.5 Cardinal preferences

In general, voting rules are using ordinal inputs. Allowing for a numerical representation of preferences (and possibly assuming interpersonal comparison of preference) opens the door to a different class of approaches, based on the maximization of an aggregation function. Many languages for compact preference representation of numerical preferences have been defined and equipped with efficient algorithms, especially valued CSPs (Bistarelli et al., 1999) and GAI-nets (Bacchus and Grove, 1995; Gonzales and Perny, 2004). In both cases, local utility functions are defined over small (and possibly intersecting) subsets of variables $S_1, \ldots, S_q$, and the global utility function is the sum (or more gen-

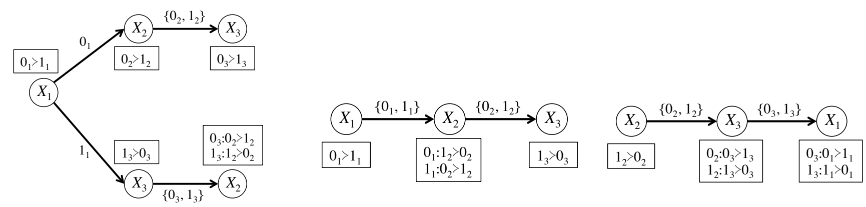

Figure 10.1 Three LP-trees.

erally the aggregation, for some suitable aggregation function) of the local utilities obtained from the local tables by projecting the alternatives on each of the S_i's. Gonzales et al. (2008) use such a representation based on GAI-nets and study algorithms for finding a Pareto-optimal alternative. Lafage and Lang (2000) and Uckelman (2009) assume that individual preferences are compactly represented using weighted propositional formulae, and that a collectively optimal alternative is defined through the maximization of a collective utility function resulting in the aggregation of individual utilities, for some suitable aggregation function (which requires not only that preferences be numerical but also that they be interpersonally comparable). See also the work by Dalla Pozza et al. (2011), discussed in Section 10.4.

10.4 Sequential Voting

The basic principle of sequential voting is that at each step voters' preferences over the values of a single issue are elicited, the decision about this variable is taken using a local voting rule, and the outcome is communicated to the voters before they vote on the next variable. Formally, a sequential voting protocol on A is defined by (1) an order $\mathcal{O}$ over $\mathcal{X}$ – without loss of generality, we let $\mathcal{O} = X_1 \rhd X_2 \rhd \cdots \rhd X_p$; and (2) for each $i \leq p$, a resolute voting rule r_i over D_i. The sequential voting protocol $Seq_{\mathcal{O}}(r_1, \ldots, r_p)$ (Lang and Xia, 2009) is defined below.

Algorithm 1: Sequential Voting Protocol

Input: An order $\mathcal{O} = X_1 \rhd X_2 \rhd \cdots \rhd X_p$ over $\mathcal{X}$; p local voting rules $r_1, \ldots, r_p$.
Output: The winner $(d_1, \ldots, d_p)$

```
for t = 1 to p do
    Ask every agent i to report her preferences ≻_t^i over D_t given (d_1, ..., d_{t−1}).
    Let P_t = ⟨≻_t^1, ..., ≻_t^n⟩ and d_t = r_t(P_t).
    Communicate d_t to the voters.
end
return (d_1, ..., d_p)
```

The above definition can easily be extended to irresolute voting rules. In the remainder of this section, we will use $Seq_{\mathcal{O}}$ as a shorthand notation for $Seq_{\mathcal{O}}(r_1, \ldots, r_p)$.

We have not discussed yet agents' behavior in each step. The main complication is that a preference for one issue may depend on the results for other issues, hence the difficulty for a voter to decide her local preferences to report.

Example 10.11 Let X_1 and X_2 be two binary issues and let P denote the following 3-voter profile:

$$1_1 1_2 \succ 0_1 1_2 \succ 1_1 0_2 \succ 0_1 0_2$$
$$1_1 0_2 \succ 1_1 1_2 \succ 0_1 1_2 \succ 0_1 0_2$$
$$0_1 1_2 \succ 0_1 0_2 \succ 1_1 0_2 \succ 1_1 1_2$$

If the order $\mathcal{O}$ is $X_2 \rhd X_1$, then voters 2 and 3 cannot unambiguously report their preferences over X_2, because they depend on the value of X_1 (for instance, voter 2 prefers 0_2 to 1_2 when $X_1 = 1_1$ and 1_2 to 0_2 when $X_1 = 0_1$), which has not been fixed yet. In other terms, marginal (or local) preference over X_2 does not have a precise meaning here.

10.4.1 Safe sequential voting

The condition that ensures that voters can report their preferences unambiguously is $\mathcal{O}$-*legality*: given $\mathcal{O} = X_1 \rhd X_2 \rhd \cdots \rhd X_p$, a preference relation $\succ$ over A is *O-legal* if for every $k \leq p$, X_k is preferentially independent of $X_{k+1}, \cdots X_p$ given $X_1, \cdots, X_{k-1}$; or, equivalently, $\succ$ extends the preference relation $\succ_{\mathcal{N}}$ induced by a CP-net whose dependency graph is compatible with $\mathcal{O}$ (*i.e*, does not contain any edge from X_i to X_j such that $X_j \rhd X_i$). Let $Legal(\mathcal{O})$ denote the set of all $\mathcal{O}$-legal profiles.

Example 10.11, continued. P is not $(X_2 \rhd X_1)$-legal, because voters 2 and 3 have preferences over X_2 that depend on X_1. On the other hand, P is $(X_1 \rhd X_2)$-legal, because all voters have unconditional preferences on X_1: voters 1 and 2 prefer 1_1 to 0_1, and voter 3 prefers 0_1 to 1_1, independently of the value of X_2.

In presence of the O-legality domain restriction, we say that sequential voting is *safe*. In this section we assume $\mathcal{O}$ to be fixed and apply sequential voting to the domain of $\mathcal{O}$-legal profiles only. A crucial property of simultaneous voting under the separability restriction carries over to safe sequential voting: it makes sense for a voter to report her local preferences on the current issue given the value of earlier issues, without having to wonder about the values of issues that have not been decided yet.

When all agents' preferences are $\mathcal{O}$-legal, a sequential voting protocol can also be considered *as a voting rule*, because a voter's preference on the values of X_i given the values of earlier variables is unambiguously defined given her preference relation over A. More precisely, given any $\mathcal{O}$-legal profile P, $Seq_{\mathcal{O}}(P)$ is defined to be the output of the sequential voting protocol where in step 2, $P_t = \langle \succ_t^1, \ldots, \succ_t^n \rangle$ where $\succ_t^i$ represents local preferences of agent i over D_t given $X_1 = d_1, \ldots, X_{t-1} = d_{t-1}$.

Example 10.12 Suppose there are two binary issues X_1 and X_2. Let P denote the same profile as in Example 10.11. Let $\mathcal{O} = X_1 \rhd X_2$. As we discussed above, P is $\mathcal{O}$-legal. To apply $Seq_{\mathcal{O}}(maj, maj)$, where maj denote the majority rule, in step 1 the voters are asked to report their (unconditional) preferences on X_1, which gives $P_1 = \langle 1_1 \succ 0_1, 1_1 \succ 0_1, 0_1 \succ 1_1 \rangle$. Therefore, $d_1 = maj(P_1) = 1_1$. In step 2, the voters report their preferences over D_2 given $X_1 = 1_1$, which leads to $P_2 = \langle 1_2 \succ 0_2, 0_2 \succ 1_2, 0_2 \succ 1_2 \rangle$, and then $d_2 = maj(P_2) = 0_2$. Therefore, $Seq_{\mathcal{O}}(maj, maj)(P) = 1_1 0_2$.

Normative Properties

We recall from Chapter 2 that one classical way to assess voting rules is to study whether they satisfy certain normative properties. In this subsection we examine the normative properties of safe sequential voting.

Classical normative properties are defined for voting rules where the input is composed of linear orders over the alternatives. We note that $Seq_{\mathcal{O}}(P)$ is only defined for $\mathcal{O}$-legal profiles. Some normative properties can be easily extended to $Seq_{\mathcal{O}}(P)$, for example *anonymity* and *consistency*, while others need to be modified. For example, the classical *neutrality* axiom states that for any profile P and any permutation M of the alternatives, $r(M(P)) = M(r(P))$. However, even if P is $\mathcal{O}$-legal, $M(P)$ might not be $\mathcal{O}$-legal. Therefore, we will focus on a weaker version of neutrality that requires $r(M(P)) = M(r(P))$ for all P and M such that both P and $M(P)$ are $\mathcal{O}$-legal. *Monotonicity* can be modified in a similar way.

Whether $Seq_{\mathcal{O}}$ satisfies a specific normative property often depends on whether the local voting rules satisfy the same property. Some properties are inherited by sequential compositions from their local rules r_i; for others, satisfaction of the property by the local voting rules is merely a necessary but not sufficient condition.

Theorem 1 (Lang and Xia, 2009)

- $Seq_{\mathcal{O}}(r_1, \ldots, r_p)$ satisfies *anonymity* (respectively, *consistency*) if and only if r_i satisfies anonymity (*consistency*) for all $i = 1, \ldots, p$;

- if $Seq_{\mathcal{O}}(r_1, \ldots, r_p)$ satisfies *neutrality* (respectively, *Condorcet-consistency*, *participation*, and *Pareto-efficiency*) then r_i satisfies neutrality (respectively, Condorcet-consistency, participation, and Pareto-efficiency) for all $i = 1, \ldots, p$;
- $Seq_{\mathcal{O}}(r_1, \ldots, r_p)$ satisfies *monotonicity* if and only if r_p satisfies monotonicity.

Proofs for the normative properties mentioned in Theorem 1 follow a similar pattern. Take consistency, for example. Recall that a voting rule r satisfies consistency if for all disjoint profiles P_1 and P_2 such that $r(P_1) = r(P_2)$, we have $r(P_1 \cup P_2) = r(P_1)$ (see see Chapter 1 (Zwicker, 2015)). We first prove that consistency can be lifted from all local rules to their sequential composition $Seq_{\mathcal{O}}(r_1, \ldots, r_p)$. For any disjoint sets of profiles P_1 and P_2 such that $Seq_{\mathcal{O}}(r_1, \ldots, r_p)(P_1) = Seq_{\mathcal{O}}(r_1, \ldots, r_p)(P_2)$, let $\vec{d}$ denote the outcome. We then prove that $Seq_{\mathcal{O}}(r_1, \ldots, r_p)(P_1 \cup P_2) = \vec{d} = (d_1, \ldots, d_p)$ by induction on the round t of sequential voting. When $t = 1$, let P_1^1 and P_2^1 denote the agents' preferences over X_1, which are well-defined because P_1 and P_2 are $\mathcal{O}$-legal. Due to consistency of r_1, we have $r_1(P_1^1 \cup P_2^1) = d_1$. Suppose the outcome of sequential voting is the same as in $\vec{d}$ up to round $k-1$. It is not hard to verify that in round k, the winner is d_k by considering agents' preferences over X_t conditioned on previous issues taking $d_1, \ldots, d_{k-1}$. This proves that $Seq_{\mathcal{O}}(r_1, \ldots, r_p)$ satisfies consistency.

Conversely, if $Seq_{\mathcal{O}}(r_1, \ldots, r_p)$ satisfies consistency and for the sake of contradiction suppose that a local rule does not satisfy consistency. Let t denote the smallest number such that r_t is not consistent, and let P_1^t, P_2^t denote the profiles over X_t with $r_t(P_1^t) = r_t(P_2^t) \neq r_t(P_1^t \cup P_2^t)$. Let $r_1, \ldots, r_{t-1}, r_{t+1}, \ldots, r_p$ be rules that always output the same winner regardless of the local profile. We can extend P_1^t and P_2^t to profiles P_1, P_2 over the whole combinatorial domain so that $Seq_{\mathcal{O}}(r_1, \ldots, r_p)(P_1) = Seq_{\mathcal{O}}(r_1, \ldots, r_p)(P_2)$, and agents' local preferences over X_t are P_1^t and P_2^t. It is not hard to verify that $Seq_{\mathcal{O}}(r_1, \ldots, r_p)(P_1) \neq Seq_{\mathcal{O}}(r_1, \ldots, r_p)(P_1 \cup P_2)$, which contradicts the assumption that $Seq_{\mathcal{O}}(r_1, \ldots, r_p)$ satisfies consistency.

In fact, for neutrality and Pareto-efficiency, a stronger result has been proved for *irresolute* sequential voting rules: Xia and Lang (2009) show that except in the case where the domain is composed of two binary issues, the only neutral irresolute sequential voting rules are dictatorships, anti-dictatorships,[10] and the trivial irresolute rule that always outputs the whole set of alternatives; and the only Pareto-efficient irresolute sequential voting rules are dictatorships and the trivial irresolute rule. When the domain is composed of two binary issues, sequential majority is neutral and Pareto-efficient.

Strategic behavior

In the previous subsection, when we talked about normative properties, it was implicitly assumed that agents were truthful. However, in practice an agent may misreport her preferences at step 2 of the sequential voting protocol (Algorithm 1). If some variables are non-binary, then sequential voting will inherit manipulability from the local rules, even if the profile is separable. However, in case all variables are binary, it is not immediately clear

[10] A rule is an anti-dictatorship if there exists an agent such that the winner is always her least preferred alternative.

if a sequential voting rule defined over $\mathcal{O}$-legal profiles is strategy-proof (see Chapters 2 and 6), since the Gibbard-Satterthwaite theorem is not directly applicable. The following example shows that sequential voting is not strategy-proof, even when all variables are binary, and the agents' preferences are $\mathcal{O}$-legal for some $\mathcal{O}$.

Example 10.13 Let P be the profile defined in Example 10.11. If agent 1 knows the preferences of agent 2 and agent 3, then she has no incentive to vote truthfully on issue X_1, *even though her preference relation is separable*: if she votes for 1 sincerely, then the outcome is 10. If she votes for 0 instead, then the outcome is 01, which is better to her.

The problem of characterizing strategy-proof voting rules in binary multi-issue domains has received some attention. Barbera et al. (1991) characterize strategy-proof voting rules when the voters' preferences are separable, and each issue is binary. Ju (2003) characterizes all strategy-proof voting rules on binary multi-issue domains (satisfying a mild additional condition) where each issue can take three values: "good", "bad", and "null". Le Breton and Sen (1999) prove that if the voters' preferences are separable, and the restricted preference domain of the voters satisfies a richness condition, then a voting rule is strategy-proof if and only if it is a simultaneous voting rule for which each local voting rule is strategy-proof over its respective domain.

We may wonder whether this extends to safe sequential voting. However, the following impossibility theorem of Xia and Conitzer (2010) answers the question negatively: there is no strategy-proof sequential voting rule on $Legal(\mathcal{O})^n$ that satisfies non-imposition, except a dictatorship. Xia and Conitzer (2010) also prove a positive result in the further restricted case of $\mathcal{O}$-legal conditionally lexicographic preferences: essentially, the strategy-proof rules on this domain are generalized sequential voting rules, where the choice of the local rule to apply on a given issue may depend on the values taken by more important issues.

10.4.2 Sequential voting: the general case

In the absence of $\mathcal{O}$-legality, sequential voting suffers from the same problem as simultaneous voting in the absence of separability: there is no clear way for voters to report their local preferences on the domain of an issue, since it may depend on the value of issues yet to be decided. Moreover, choosing the agenda (the order on which the issues are decided) can be tricky: What is a good agenda? Who chooses it? This problem is raised by Airiau et al. (2011), who suggest designing a voting procedure for choosing the agenda: each voter reports its dependency graph and these graphs are aggregated into an acyclic graph, for instance using a distance-based aggregation function. Another approach to unrestricted sequential voting is described by Dalla Pozza et al. (2011), who assume that at each stage, voters report their local preferences according to the projection of a weighted CSP on the current variable.

Another challenge in the analysis of sequential voting without the $\mathcal{O}$-legality assumption is that the outcome may depend on the order in which the issues are decided. This can give the chair (or whoever chooses the order) an effective way of controlling the election (see Chapter 7). This can be seen in the following example:

Example 10.3, continued. Suppose the voters report their preferences optimistically, which is known to the chair. If S is decided before T, then we get $k+1$ votes for 1_S, k votes for 0_S, leading to the local outcome 1_S; then, given $S = 1_S$, we have $2k$ votes for 0_T and 1 vote for 1_T, therefore the final outcome is $1_S 0_T$. Symmetrically, if T is decided before S, the final outcome is $0_S 1_T$. Thefore, the chair's strategy can be to choose the order S $\rhd$ T if she prefers $1_S 0_T$ to $0_S 1_T$, and the order T $\rhd$ S otherwise. (Note that $1_S 1_T$ and $0_S 0_T$ cannot be obtained).

This shows that the chair can sometimes, and to some extent, control the election by fixing the agenda (see also Chapter 7). This drawback of sequential voting is however tempered by the fact that under some reasonable assumptions about the way the voters's behaviors are represented, in most cases, most of these agenda control problems are NP-hard (Conitzer et al., 2009).

We mentioned above that in the absence of $\mathcal{O}$-legality, there is no clear way for voters to report their local preferences on the domain of a variable. However, there is a case where voters may in fact be able to determine valid reports of their local preferences. When voters' preferences are assumed to be common knowledge, the sequential voting protocol can be framed as an extensive-form game, called a *strategic sequential voting process*, denoted by $\text{SSV}_\mathcal{O}$ (Xia et al., 2011). We assume that all variables are binary. Without loss of generality, let $\mathcal{O} = X_1 \rhd \cdots \rhd X_p$. The game is defined as follows:

- The players are the voters; their preferences are linear orders over A; their possible actions at stage $t \leq p$ are 0_t and 1_t.
- In each stage t, all voters vote on X_t simultaneously, r_t is used to choose the winning value d_t for X_t, and d_t is reported back to the voters.
- We assume complete information: all voters know the other voters' preferences, the local voting rules $r_1, \ldots, r_p$ and the order $\mathcal{O}$.

When all issues are binary, $\text{SSV}_\mathcal{O}$ can be solved by backward induction where in each stage all voters move simultaneously and perform a dominant strategy, as illustrated in the following example.

Example 10.14 Let P be the profile defined in Example 10.11. If the outcome of the first stage of sequential voting is 1_1, then in the second stage it is voter 1's dominant strategy to vote for 1_2 because $1_1 1_2 \succ_1 1_2 0_2$ and the majority rule is strategyproof. Similarly, in this case voters 2 and 3 will vote for 0_2. Therefore, by the majority rule, the winner will be $1_1 0_2$. Similarly, if the outcome of the first stage of sequential voting is 0_1, then the votes at the second stage will be unanimously 0_2, and the winner will be $0_1 1_2$. Given the above reasoning, in the first stage, each agent is comparing $1_1 0_2$ to $0_1 1_2$, and will vote 1_1 if he prefers $1_1 0_2$ to $0_1 1_2$ and 0_1 if he prefers $0_1 1_2$ to $1_1 0_2$. Again we have two alternatives, and the majority rule is strategy-proof. This means that voter 1 will vote for 0_1, voter 2 will vote for 1_1, and voter 3 will vote for 0_1. Hence, the winner for X_1 is 0_1, and the overall winner is $0_1 1_2$. This backward induction process is shown in Figure 10.2.

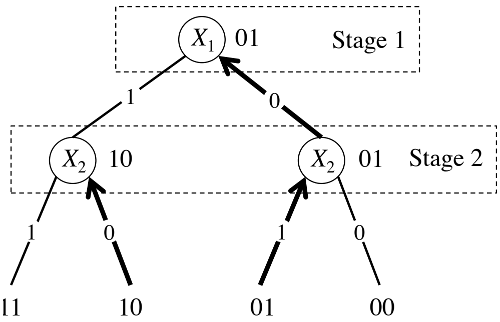

Figure 10.2 Backward induction tree for Example 10.14.

In Example 10.14, the backward induction winner is unique. This observation can be extended to an arbitrary number of binary issues. Let $SSV_{\mathcal{O}}(P)$ denote the backward induction winner when the voters' true preferences are P.[11] Despite being unique, this outcome is extremely undesirable in the worst case: Xia et al. (2011) prove that for any $p \in \mathbb{N}$ and any $n \geq 2p^2 + 1$, there exists a profile P such that (1) $SSV_{\mathcal{O}}(P)$ is ranked within the bottom $\lfloor p/2 + 2 \rfloor$ positions in *every* voter's true preferences, and (2) $SSV_{\mathcal{O}}(P)$ is Pareto-dominated by $2^p - (p+1)p/2$ alternatives. We note that when p is not too small, $\lfloor p/2 + 2 \rfloor$ and $(p + 1)p/2$ are much smaller than $|A| = 2^p$. Therefore, the SSV winner can indeed be extremely undesirable. A stronger form of the theorem and a similar negative result for $\mathcal{O}$-legal profiles are given by Xia et al. (2011).

10.4.3 Discussion

The key property of sequential voting is that it interleaves preference elicitation and winner determination, whereas the approaches outlined in Sections 10.2 and 10.3 proceed in a more usual fashion, by eliciting preferences in one round, and determining the winner afterwards. As a result, sequential voting can save a lot in communication costs, but is applicable only when it is realistic to elicit preferences step by step. Now, the quality of the outcome obtained by sequential voting primarily depends on whether it is realistic to assume that voters' preferences are $\mathcal{O}$-legal for some common order $\mathcal{O}$: if so, then sequential

[11] This should be distinguished from the classical social choice setting, where the input consists of the reported preferences.

voting enjoys some good properties; if not, it offers far fewer quality guarantees. Therefore, one criticism of sequential voting is that it still needs a strong domain restriction to work well (Xia et al., 2008); but still, when compared to the separability restriction needed for simultaneous voting, $\mathcal{O}$-legality is much weaker (Lang and Xia, 2009). Conitzer and Xia (2012) evaluate the quality of the outcome of sequential voting w.r.t. a scoring function, which can be seen as numerical versions of multiple election paradoxes.

There could be other ways of interleaving elicitation and winner determination. A completely different way of proceeding was proposed very recently by Bowman et al. (2014), who propose an iterative protocol that allows voters to revise their votes based on the outcomes of previous iterations.

10.5 Concluding discussion

After reviewing several classes of methods for voting in combinatorial domains, we are left with the (expected) conclusion that none of them is perfect. More precisely, when choosing a method, we have to make a tradeoff between generality, communication (and, to a lesser extent, computation) costs, and the quality of the outcome, evaluated with respect to classical social choice criteria. If specific domain restrictions such as separability or, more generally, $\mathcal{O}$-legality, are realistic for the case at hand, then many of the methods discussed in this chapter work reasonably well. Otherwise, one has to be prepared to make some tradeoff between communication requirements and the quality of the outcome. One possibility that has not really been developed yet is to choose some intermediate method that requires some weak domain restriction, some reasonable communication and computation costs, and offers some reasonable guarantees about the quality of the outcome.

Voting in combinatorial domains is related to several other issues studied in this book and elsewhere:

- *Incomplete information and communication* (Chapter 10, (Boutilier and Rosenschein, 2015)): as some communication saving can be made by eliciting only a part of the voters' preferences, winner determination in combinatorial domains can benefit from approaches to winner determination from incomplete preferences as well as from the design of communication-efficient voting protocols.
- *Judgment aggregation* (Chapter 17, (Endriss, 2015)) is also concerned with making common decisions about possibly interrelated issues. There are interesting parallels between judgment aggregation and voting in combinatorial domains. Simultaneous voting corresponds to some extent to propositionwise voting: while the first works well when preferences are separable, the second outputs a consistent outcome if the agenda enjoys an independence property that resembles separability. Note that in judgment aggregation, difficulties are often caused by the logical relations between the elements of the agenda, while in voting in combinatorial domains, they are mainly due to preferential depen-

dencies. Relating both areas is a promising research direction; see Grandi and Endriss (2011) for some first steps in this direction.

- *Fair division of indivisible items* (Chapter 12, (Bouveret et al., 2015)) is another field where a common decision has to be made on a combinatorial space of alternatives (the set of all allocations). In the settings we reviewed in this chapter, we assumed that all agents were equally concerned with all issues (which is patently false in fair division, where agents are primarily – sometimes even exclusively – concerned by their share); but in some settings, some issues concern some (subsets of) voters more than others, which will call for introducing fairness criteria into multi-issue voting.

Acknowledgements

We would like to thank Craig Boutilier, Vince Conitzer, Ulle Endriss and Umberto Grandi for very helpful feedback.

References

Ahn, D. S., and Oliveros, S. 2012. Combinatorial Voting. *Econometrica*, **80**(1), 89–141.

Airiau, S., Endriss, U., Grandi, U., Porello, D., and Uckelman, J. 2011. Aggregating Dependency Graphs into Voting Agendas in Multi-Issue Elections. Pages 18–23 of: *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*.

Bacchus, F., and Grove, A. 1995. Graphical models for preference and utility. Pages 3–10 of: *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*. Morgan Kaufmann.

Barbera, S., Sonnenschein, H., and Zhou, L. 1991. Voting by Committees. *Econometrica*, **59**(3), 595–609.

Benoît, J.-P. 2010. Only a dictatorship is efficient. *Games and Economic Behavior*, **70**(2), 261–270.

Benoît, J.-P., and Kornhauser, L. A. 2010. Only a dictatorship is efficient. *Games and Economic Behavior*, **70**(2), 261–270.

Benoît, J.-P., and Kornhauser, L.A. 1991. *Voting simply in the election of assemblies*. Tech. rept. 91-32. C.V. Starr Center for Applied Economics.

Betzler, N., Slinko, A., and Uhlmann, J. 2013. On the Computation of Fully Proportional Representation. *Journal of Artificial Intelligence Research*, **47**, 475–519.

Bistarelli, S., Fargier, H., Montanari, U., Rossi, F., Schiex, T., and Verfaillie, G. 1999. Semiring-Based CSPs and Valued CSPs: Frameworks, Properties and Comparison. *Constraints*, **4**(3), 199–240.

Booth, R., Chevaleyre, Y., Lang, J., Mengin, J., and Sombattheera, C. 2010. Learning conditionally lexicographic preference relations. Pages 269–274 of: *Proceedings of the European Conference on Artificial Intelligence (ECAI)*.

Boutilier, C., and Rosenschein, J. 2015. Incomplete Information and Communication in Voting. Chap. 10 of: Brandt, F., Conitzer, V., Endriss, U., Lang, J., and Procaccia, A. D. (eds), *Handbook of Computational Social Choice*. Cambridge University Press.

Boutilier, C., Brafman, R., Domshlak, C., Hoos, H., and Poole, D. 2004. CP-nets: a tool for representing and reasoning with conditional ceteris paribus statements. *Journal of Artificial Intelligence Research*, **21**, 135–191.

Bouveret, S., Chevaleyre, Y., and Maudet, N. 2015. Fair Allocation of Indivisible Goods. Chap. 12 of: Brandt, F., Conitzer, V., Endriss, U., Lang, J., and Procaccia, A. D. (eds), *Handbook of Computational Social Choice*. Cambridge University Press.

Bowman, C., Hodge, J., and Yu, A. 2014. The potential of iterative voting to solve the separability problem in referendum elections. *Theory and Decision*, **77**(1), 111–124.

Bradley, W. J., Hodge, J. K., and Kilgour, D. 2005. Separable discrete preferences. *Mathematical Social Science*, **49**(3), 335–353.

Brams, S., and Fishburn, P. 2004. Voting procedures. Chap. 4 of: Arrow, K., Sen, A., and Suzumura, K. (eds), *Handbook of Social Choice and Welfare*. Elsevier.

Brams, S., and Potthoff, R. F. 1998. Proportional Representation: Broadening the Options. *Journal of Theoretical Politics*, **10**, 147–178.

Brams, S., Kilgour, D., and Zwicker, W. 1997. Voting on Referenda: the Separability Problem and Possible Solutions. *Electoral Studies*, **16**(3), 359–377.

Brams, S., Kilgour, D., and Zwicker, W. 1998. The Paradox of Multiple Elections. *Social Choice and Welfare*, **15**(2), 211–236.

Brams, S., Kilgour, D., and Sanver, R. 2007. A Minimax Procedure for Electing Committees. *Public Choice*, **3-4**(132), 401–420.

Caragiannis, I., Kalaitzis, D., and Markakis, E. 2010. Approximation Algorithms and Mechanism Design for Minimax Approval Voting. In: *Proceedings of the National Conference on Artificial Intelligence (AAAI)*.

Chamberlin, J. R., and Courant, P. N. 1983. Representative deliberations and representative decisions: proportional representation and the Borda rule. *American Political Science Review*, **77**(3), 718–733.

Conitzer, V., and Xia, L. 2012. Paradoxes of Multiple Elections: An Approximation Approach. Pages 179–187 of: *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning (KR)*.

Conitzer, V., Lang, J., and Xia, L. 2009. How Hard Is It to Control Sequential Elections via the Agenda? Pages 103–108 of: *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*.

Conitzer, V., Lang, J., and Xia, L. 2011. Hypercubewise Preference Aggregation in Multi-Issue Domains. Pages 158–163 of: *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*.

Cornaz, D., Galand, L., and Spanjaard, O. 2012. Bounded Single-Peaked Width and Proportional Representation. Pages 270–275 of: *Proceedings of the European Conference on Artificial Intelligence (ECAI)*.

Cuhadaroğlu, T., and Lainé, J. 2012. Pareto efficiency in multiple referendum. *Theory and Decision*, **72**(4), 525–536.

Dalla Pozza, G., Pini, M. S., Rossi, F., and Venable, K. B. 2011. Multi-Agent Soft Constraint Aggregation via Sequential Voting. Pages 172–177 of: *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*.

Debreu, G. 1954. Representation of a preference ordering by a numerical function. *Decision processes*, **3**, 159–165.

Elkind, E., Faliszewski, P., Skowron, P., and Slinko, A. 2014. Properties of Multiwinner Voting Rules. Pages 53–60 of: *Proceedings of the International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*.

Elkind, E., Lang, J., and Saffidine, A. 2015. Condorcet Winning Sets. *Social Choice and Welfare*.

Endriss, U. 2015. Judgment Aggregation. Chap. 17 of: Brandt, F., Conitzer, V., Endriss, U., Lang, J., and Procaccia, A. D. (eds), *Handbook of Computational Social Choice*. Cambridge University Press.

Frances, M., and Litman, A. 1997. On covering problems of codes. *Theory of Computing Systems*, **30**, 113–119.

Gonzales, C., and Perny, P. 2004. GAI networks for utility elicitation. Pages 224–233 of: *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning (KR)*.

Gonzales, C., Perny, P., and Queiroz, S. 2008. Preference Aggregation with Graphical Utility Models. Pages 1037–1042 of: *Proceedings of the National Conference on Artificial Intelligence (AAAI)*.

Grandi, U., and Endriss, U. 2011. Binary Aggregation with Integrity Constraints. Pages 204–209 of: *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*.

Ju, B.-G. 2003. A characterization of strategy-proof voting rules for separable weak orderings. *Social Choice and Welfare*, **21**(3), 469–499.

Kadane, J. 1972. On Division of the Question. *Public Choice*, **13**, 47–54.

Keeney, R., and Raiffa, H. 1976. *Decision with Multiple Objectives: Preferences and Value Tradeoffs*. Wiley and Sons.

Kleinberg, J. M., Papadimitriou, C. H., and Raghavan, P. 2004. Segmentation problems. *Journal of the ACM*, **51**(2).

Konieczny, S., and Pérez, R. Pino. 2011. Logic Based Merging. *Journal of Philosophical Logic*, **40**(2), 239–270.

Konieczny, S., Lang, J., and Marquis, P. 2004. DA^2 merging operators. *Artificial Intelligence*, **157**(1-2), 49–79.

Lacy, D., and Niou, E. 2000. A problem with referenda. *Journal of Theoretical Politics*, **12**(1), 5–31.

Lafage, C., and Lang, J. 2000. Logical representation of preferences for group decision making. Pages 457–468 of: *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning (KR)*.

Laffond, G., and Lainé, J. 2009. Condorcet choice and the Ostrogorski paradox. *Social Choice and Welfare*, **32**(2), 317–333.

Lang, J. 2004. Logical preference representation and combinatorial vote. *Annals of Mathematics and Artificial Intelligence*, **42**(1), 37–71.

Lang, J., and Xia, L. 2009. Sequential composition of voting rules in multi-issue domains. *Mathematical Social Sciences*, **57**(3), 304–324.

Lang, J., Mengin, J., and Xia, L. 2012. Aggregating Conditionally Lexicographic Preferences on Multi-issue Domains. Pages 973–987 of: *Proceedings of the International Conference on Principles and Practice of Constraint Programming (CP)*.

Le Breton, M., and Sen, A. 1999. Separable preferences, strategyproofness, and decomposability. *Econometrica*, **67**(3), 605–628.

LeGrand, R., Markakis, E., and Mehta, A. 2007. Some results on approximating the minimax solution in approval voting. Pages 198:1–198:3 of: *Proceedings of the International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*.

Li, M., Vo, Q. B., and Kowalczyk, R. 2010. An Efficient Majority-Rule-Based Approach for Collective Decision Making with CP-Nets. In: *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning (KR)*.

Li, M., Vo, Q. B., and Kowalczyk, R. 2011. Majority-rule-based preference aggregation on multi-attribute domains with CP-nets. Pages 659–666 of: *Proceedings of the International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*.

Lu, T., and Boutilier, C. 2011. Budgeted Social Choice: From Consensus to Personalized Decision Making. In: *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*.

Lu, T., and Boutilier, C. 2013. Multiwinner Social Choice with Incomplete Preferences. In: *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*.

Meir, R., Procaccia, A. D., Rosenschein, J. S., and Zohar, A. 2008. Complexity of Strategic Behavior in Multi-Winner Elections. *Journal of Artificial Intelligence Research*, **33**, 149–178.

Monroe, B. L. 1995. Fully Proportional Representation. *American Political Science Review*, **89**, 925–940.

Özkal-Sanver, İ., and Sanver, R. 2006. Ensuring Pareto-optimality by referendum voting. *Social Choice and Welfare*, **27**, 211–219.

Procaccia, A.D., Rosenschein, J., and Zohar, A. 2008. On the Complexity of Achieving Proportional Representation. *Social Choice and Welfare*, **30**(3), 353–362.

Rossi, F., Venable, K., and Walsh, T. 2004. mCP nets: representing and reasoning with preferences of multiple agents. Pages 729–734 of: *Proceedings of the National Conference on Artificial Intelligence (AAAI)*.

Saari, D., and Sieberg, K. 2001. The sum of the parts can violate the whole. *The American Political Science Review*, **95**(2), 415–433.

Scarsini, M. 1998. A strong paradox of multiple elections. *Social Choice and Welfare*, **15**(2), 237–238.

Skowron, P., Faliszewski, P., and Slinko, A. 2013a. Achieving fully proportional representation is easy in practice. Pages 399–406 of: *Proceedings of the International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*.

Skowron, P., Yu, L., Faliszewski, P., and Elkind, E. 2013b. The Complexity of Fully Proportional Representation for Single-Crossing Electorates. Pages 1–12 of: *Proceedings of International Symposium on Algorithmic Game Theory (SAGT)*.

Skowron, P., Faliszewski, P., and Slinko, A. 2013c. Fully Proportional Representation as Resource Allocation: Approximability Results. In: *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*.

Skowron, P., Faliszewski, P., and Lang, J. 2015. Finding a Collective Set of Items: From Proportional Multirepresentation to Group Recommendation. In: *Proceedings of the National Conference on Artificial Intelligence (AAAI)*.

Uckelman, J. 2009. *More Than the Sum of Its Parts: Compact Preference Representation Over Combinatorial Domains*. Ph.D. thesis, University of Amsterdam.

Xia, L., and Conitzer, V. 2010. Strategy-Proof Voting Rules over Multi-issue Domains with Restricted Preferences. Pages 402–414 of: *Proceedings of International Conference on Web and Internet Economics (WINE)*.

Xia, L., and Lang, J. 2009. A Dichotomy Theorem on the Existence of Efficient or Neutral Sequential Voting Correspondences. Pages 342–347 of: *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*.

Xia, L., Conitzer, V., and Lang, J. 2008. Voting on Multiattribute Domains with Cyclic Preferential Dependencies. Pages 202–207 of: *Proceedings of the National Conference on Artificial Intelligence (AAAI)*.

Xia, L., Conitzer, V., and Lang, J. 2010. Aggregating preferences in multi-issue domains by using maximum likelihood estimators. Pages 399–408 of: *Proceedings of the International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*.

Xia, L., Conitzer, V., and Lang, J. 2011. Strategic Sequential Voting in Multi-Issue Domains and Multiple-Election Paradoxes. Pages 179–188 of: *Proceedings of the ACM Conference on Electronic Commerce (EC)*.

Yu, L., Chan, H., and Elkind, E. 2013. Multiwinner Elections Under Preferences thar are Single-Peaked on a Tree. In: *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*.

Zwicker, William S. 2015. Introduction to the Theory of Voting. Chap. 2 of: Brandt, F., Conitzer, V., Endriss, U., Lang, J., and Procaccia, A. D. (eds), *Handbook of Computational Social Choice*. Cambridge University Press.

Index