

Vote and Aggregation in Combinatorial Domains with Structured Preferences

Jérôme Lang

IRIT - Université Paul Sabatier, Toulouse (France)

lang@irit.fr

Abstract

In many real-world collective decision problems, the set of alternatives is a Cartesian product of finite value domains for each of a given set of variables. The prohibitive size of such domains makes it practically impossible to represent preference relations explicitly. Now, AI has been developing languages for representing preferences on such domains in a succinct way, exploiting structural properties such as conditional preferential independence. Here we reconsider voting and aggregation rules in the case where voters' preferences have a common preferential independence structure, and address the decomposition a voting rule or an aggregation function following a linear order over variables.

1 Introduction

Researchers in social choice have extensively studied the properties of voting rules and aggregation functions, up to an important detail: candidates are supposed to be listed explicitly (typically, they are individuals or lists of individuals), which assumes that they are not too numerous. In this paper, we consider the case where the set of candidates has a *combinatorial structure*, i.e., is a Cartesian product of finite value domains for each of a finite set of variables.

Since the number of possible alternatives is then exponential in the number of variables, it is not reasonable to ask voters to rank all alternatives explicitly. Consider for example that voters have to agree on a common menu to be composed of a first course, a main course, a dessert and a wine, with a choice of 6 items for each. This makes 6^4 candidates. This would not be a problem if each of the four items to be chosen were independent from the other ones: in this case, this vote over a set of 6^4 candidates would come down to four independent votes over sets of 6 candidates each, and any standard voting rule could be applied without difficulty. Things become more complicated if voters express dependencies between items, such as “if the main course is meat then I prefer red wine, otherwise I prefer white wine”. Indeed, as soon as variables are not preferentially independent, it is generally a bad idea to decompose a vote problem with p variables into a set of p smaller problems, each one bearing on a single variable: “multiple election paradoxes” [Brams *et al.*, 1998]

show that such a decomposition leads to suboptimal choices, and give real-life examples of such paradoxes, including simultaneous referenda on related issues. They argue that the only way of avoiding the paradox would consist in “voting for combinations [of values]”, but they stress its practical difficulty without giving any hint for a practical solution.

Because the preference structure of each voter in such a case cannot reasonably be expressed by listing all candidates, what is needed is a compact *preference representation language*. Such languages have been developed within the Artificial Intelligence community so as to escape the combinatorial blow up of explicit representations. Many of these languages (including CP-nets and their extensions) are *graphical*: preferences are expressed locally (on small subsets of variables). The common feature of these languages is that they allow for a concise representation of the preference structure, while preserving a good readability (and hence a proximity with the way agents express their preferences in natural language).

Thus, AI gives a first answer to the problem pointed in [Brams *et al.*, 1998]. However, another problem arises: *once preferences have been elicited, and represented in some compact language, how is the voting rule computed?* The prohibitive number of candidates makes it practically impossible to compute these rules in a straightforward way.

When domains are not too large, it may still be reasonable to first generate the whole preference relations from their compact representations and then compute the outcome by a direct implementation of the voting rule. However, when domains become bigger, this naive method becomes too greedy and then we need to find a more sophisticated way of computing the outcome of the vote. Two methods come to mind: either (1) give up optimality and compute an *approximation* of the voting or aggregation rule, or (2) assume that the voters' preferences enjoy specific structural properties that can be exploited so as to *decompose* the problem into smaller, local subproblems. Here we address (2), and we focus on a specific restriction of preference profiles where all voters have a preference relation enjoying conditional preferential independencies compatible with a common acyclic graph G . After giving some background on preference relations over combinatorial domains and vote in Section 2, we introduce and study sequential voting rules in Section 3. Section 4 then considers preference aggregation over combinatorial domains, and Section 5 concludes.

2 Background

2.1 Preferences on combinatorial domains

Let $V = \{\mathbf{x}_1, \dots, \mathbf{x}_p\}$ be a set of variables. For each $\mathbf{x}_i \in V$, D_i is the value domain of \mathbf{x}_i . A variable v_i is binary if $D_i = \{x_i, \bar{x}_i\}$. Note the difference between the variable \mathbf{x}_i and the value x_i . If $X = \{\mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_m}\} \subseteq V$, with $i_1 < \dots < i_p$, then D_X denotes $D_{i_1} \times \dots \times D_{i_m}$.

$\mathcal{X} = D_1 \times \dots \times D_p$ is the set of all alternatives, or candidates. Elements of \mathcal{X} are denoted by \vec{x}, \vec{x}' etc. and represented by concatenating the values of the variables: for instance, if $V = \{\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3\}$, $x_1\bar{x}_2x_3$ assigns \mathbf{x}_1 to x_1 , \mathbf{x}_2 to \bar{x}_2 and \mathbf{x}_3 to x_3 . We allow concatenations of vectors of values: for instance, let $V = \{\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \mathbf{x}_4, \mathbf{x}_5\}$, $Y = \{x_1, x_2\}$, $Z = \{x_3, x_4\}$, $\vec{y} = x_1\bar{x}_2$, $\vec{z} = \bar{x}_3x_4$, then $\vec{y}.\vec{z}.\bar{x}_5$ denotes the alternative $x_1\bar{x}_2x_3x_4\bar{x}_5$.

A (strict) preference relation on \mathcal{X} is a strict order (an irreflexive, asymmetric and transitive binary relation). A linear preference relation is a complete strict order, i.e., for any \vec{x} and $\vec{y} \neq \vec{x}$, either $\vec{x} \succ \vec{y}$ or $\vec{y} \succ \vec{x}$ holds. If R is a preference relation, we generally note $\vec{x} \succ_R \vec{x}'$ instead of $R(\vec{x}, \vec{x}')$.

Let $\{X, Y, Z\}$ be a partition of the set V of variables and \succ a linear preference relation over D_V . X is (conditionally) preferentially independent of Y given Z (w.r.t. \succ) if and only if for all $\vec{x}_1, \vec{x}_2 \in D_X, \vec{y}_1, \vec{y}_2 \in D_Y, \vec{z} \in D_Z$,

$$\vec{x}_1.\vec{y}_1.\vec{z} \succ \vec{x}_2.\vec{y}_1.\vec{z} \text{ iff } \vec{x}_1.\vec{y}_2.\vec{z} \succ \vec{x}_2.\vec{y}_2.\vec{z}$$

Unlike probabilistic independence, preferential independence is a directed notion: X may be independent of Y given Z without Y being independent of X given Z .

A CP-net \mathcal{N} [Boutilier et al., 2004] over V is a pair consisting of a directed graph G over V and a collection of conditional preference tables $CPT(\mathbf{x}_i)$ for each $x_i \in V$. Each conditional preference table $CPT(\mathbf{x}_i)$ associates a total order¹ $\succ_{\vec{u}}$ with each instantiation \vec{u} of \mathbf{x}_i 's parents $Pa(\mathbf{x}_i) = U$. For instance, let $V = \{\mathbf{x}, \mathbf{y}, \mathbf{z}\}$, all three being binary, and assume that preference of a given agent over 2^V can be defined by a CP-net whose structural part is the directed acyclic graph $G = \{(x, y), (y, z), (x, z)\}$; this means that the agent's preference over the values of \mathbf{x} is unconditional, preference over the values of \mathbf{y} (resp. \mathbf{z}) is fully determined given the value of \mathbf{x} (resp. the values of \mathbf{x} and \mathbf{y}).

The conditional preference statements contained in these tables are written with the usual notation, that is, $x_1\bar{x}_2 : x_3 \succ \bar{x}_3$ means that when \mathbf{x}_1 is true and \mathbf{x}_2 is false then $\mathbf{x}_3 = x_3$ is preferred to $\mathbf{x}_3 = \bar{x}_3$. In this paper we make the classical assumption that G is acyclic. A CP-net \mathcal{N} induces a preference ranking on \mathcal{X} : $\vec{x} \succ_{\mathcal{N}} \vec{y}$ iff there is a sequence of improving flips from \vec{y} to \vec{x} , where an improving flip is the flip of a single variable \vec{x}_i "respecting" the preference table $CPT(\mathbf{x}_i)$ (see [Boutilier et al., 2004]). Note that the preference relation induced from a CP-net is generally not complete.

Let G be a directed graph over V , and \succ a linear preference relation. \succ is said to be compatible with G iff for each $\mathbf{x} \in V$, \mathbf{x} is preferentially independent of $V \setminus (\{\mathbf{x}\} \cup Par(\mathbf{x}))$ given $Par(\mathbf{x})$. The following fact is obvious, but important:

¹More generally, the entries of conditional preference tables may contain partial orders over the domains of the variables (see [Boutilier et al., 2004]), but we don't need this here.

Observation 1 A linear preference relation \succ is compatible with G if and only if there exists a CP-net \mathcal{N} whose associated graph is G and such that \succ extends $\succ_{\mathcal{N}}$.

Let G be an acyclic graph over V and let $\mathcal{O} = \mathbf{x}_1 > \dots > \mathbf{x}_p$ be a linear order on V . G is said to follow \mathcal{O} iff for every edge $(\mathbf{x}_i, \mathbf{x}_j)$ in G we have $i < j$. A linear preference relation \succ is said to follow \mathcal{O} iff it is compatible with some acyclic graph G following \mathcal{O} . Clearly, \succ follows $\mathcal{O} = \mathbf{x}_1 > \dots > \mathbf{x}_p$ if and only if for all $i < p$, \mathbf{x}_i is preferentially independent of $\{\mathbf{x}_{i+1}, \dots, \mathbf{x}_p\}$ given $\{\mathbf{x}_1, \dots, \mathbf{x}_{i-1}\}$ with respect to \succ . If \succ follows \mathcal{O} then the projection of \succ on \mathbf{x}_i given $(x_1, \dots, x_{i-1}) \in D_1 \times \dots \times D_{i-1}$, denoted by $\succ_{\mathbf{x}_i | \mathbf{x}_1=x_1, \dots, \mathbf{x}_{i-1}=x_{i-1}}$, is the linear preference relation on D_i defined by: for all $x_i, x'_i \in D_i$, $x_i \succ_{\mathbf{x}_i | \mathbf{x}_1=x_1, \dots, \mathbf{x}_{i-1}=x_{i-1}} x'_i$ iff $x_1 \dots x_{i-1} x_i x_{i+1} \dots x_p \succ x_1 \dots x_{i-1} x'_i x_{i+1} \dots x_p$ holds for all $(x_{i+1}, \dots, x_p) \in D_{i+1} \times \dots \times D_p$.

Due to the fact that \succ follows \mathcal{O} and that \succ is a linear order, $\succ_{\mathbf{x}_i | \mathbf{x}_1=x_1, \dots, \mathbf{x}_{i-1}=x_{i-1}}$ is a well-defined linear order as well. Note also that if \succ follows both $\mathcal{O} = \mathbf{x}_1 > \dots > \mathbf{x}_p$ and $\mathcal{O}' = \mathbf{x}_{\sigma(1)} > \dots > \mathbf{x}_{\sigma(k-1)} > \mathbf{x}_i (= \mathbf{x}_{\sigma(k)}) > \dots > \mathbf{x}_{\sigma(p)}$, then $\succ_{\mathbf{x}_i | \mathbf{x}_1=x_1, \dots, \mathbf{x}_{i-1}=x_{i-1}}$ and $\succ_{\mathbf{x}_i | \mathbf{x}_{\sigma(1)}=x_{\sigma(1)}, \dots, \mathbf{x}_{\sigma(k-1)}=x_{\sigma(k-1)}}$ coincide. In other words, the local preference relation on \mathbf{x}_i depends only on the values of the variables that precede \mathbf{x}_i in \mathcal{O} and in \mathcal{O}' .

Lastly, for any acyclic graph G over V , we say that two linear preference relations R_1 and R_2 are G -equivalent, denoted by $R_1 \sim_G R_2$, if and only if R_1 and R_2 are both compatible with G and for any $\mathbf{x} \in V$, for any $\vec{y}, \vec{y}' \in Dom(par(\mathbf{x}))$ we have $R_1^{\mathbf{x} | par(\mathbf{x})=\vec{y}} = R_2^{\mathbf{x} | par(\mathbf{x})=\vec{y}}$.

Observation 2 For any linear preference relations R and R' , $R \sim_G R'$ if and only if there exists a CP-net \mathcal{N} whose associated graph is G and such that R and R' both extend $\succ_{\mathcal{N}}$.

Example 1 Let $V = \{\mathbf{x}, \mathbf{y}, \mathbf{z}\}$, all three being binary. and let R and R' be the following linear preference relations:

$$R : xyz \succ xy\bar{z} \succ x\bar{y}\bar{z} \succ x\bar{y}z \succ \bar{x}y\bar{z} \succ \bar{x}\bar{y}\bar{z} \succ \bar{x}yz \succ \bar{x}\bar{y}z$$

$$R' : xyz \succ xy\bar{z} \succ \bar{x}y\bar{z} \succ x\bar{y}\bar{z} \succ \bar{x}y\bar{z} \succ \bar{x}\bar{y}\bar{z} \succ x\bar{y}z \succ \bar{x}\bar{y}z$$

Let G the graph over V whose set of edges is $\{(\mathbf{x}, \mathbf{z}), (\mathbf{y}, \mathbf{z})\}$. R and R' are both compatible with G . Moreover, $R \sim_G R'$, since all local preference relations coincide: $x \succ_R^{\bar{x}} \bar{x}$ and $x \succ_{R'}^{\bar{x}} \bar{x}$; $z \succ_R^{x, \bar{x}, y=y} \bar{z}$ and $z \succ_{R'}^{x, \bar{x}, y=y} \bar{z}$; etc. The CP-net \mathcal{N} such that R and R' both extend $\succ_{\mathcal{N}}$ is defined by the following local conditional preferences: $x \succ \bar{x}$; $y \succ \bar{y}$; $xy : z \succ \bar{z}$; $x\bar{y} : \bar{z} \succ z$; $\bar{x}y : \bar{z} \succ z$; $\bar{x}\bar{y} : \bar{z} \succ z$.

2.2 Voting rules and correspondences

Let $\mathcal{A} = \{1, \dots, N\}$ be a finite set of voters and \mathcal{X} a finite set of candidates. A (collective) preference profile w.r.t. \mathcal{A} and \mathcal{X} is a collection of N individual linear preference relations over \mathcal{X} : $P = (\succ_1, \dots, \succ_N) = (P_1, \dots, P_N)$. Let $\mathcal{P}_{\mathcal{A}, \mathcal{X}}$ be the set of all preference profiles for \mathcal{A} and \mathcal{X} .

A voting correspondence $C : \mathcal{P}_{\mathcal{A}, \mathcal{X}} \rightarrow 2^{\mathcal{X}} \setminus \{\emptyset\}$ maps each preference profile P of $\mathcal{P}_{\mathcal{A}, \mathcal{X}}$ into a nonempty subset $C(P)$ of \mathcal{X} . A voting rule $r : \mathcal{P}_{\mathcal{A}, \mathcal{X}} \rightarrow \mathcal{X}$ maps each preference profile P of $\mathcal{P}_{\mathcal{A}, \mathcal{X}}$ into a single candidate $r(P)$. A rule can be obtained from a correspondence by prioritization over candidates (for more details see [Brams and Fishburn, 2004]).

To give an example, consider the well-known family of *positional scoring rules and correspondences*. A positional scoring correspondence is defined from a *scoring vector*, that is, a vector $\vec{s} = (s_1, \dots, s_m)$ of integers such that $s_1 \geq s_2 \geq \dots \geq s_m$ and $s_1 > s_m$. Let $rank_i(x)$ be the rank of x in \succ_i (1 if it is the favorite candidate for voter i , etc.). The score of x is defined by $S(x) = \sum_{i=1}^N s_{rank_i(x)}$. The correspondence defined from \vec{s} selects all candidates maximizing S . A positional voting rule is defined as a positional scoring correspondence plus a tie-breaking mechanism. Well-known examples are the *Borda* rule, given by $s_k = m - k$ for all k , and the *plurality* rule, by $s_1 = 1$, and $s_k = 0$ for all k .

We also recall the definition of a *Condorcet winner* (CW). Given a profile $P = (\succ_1, \dots, \succ_N)$, $x \in \mathcal{X}$ is a *Condorcet winner* iff it is preferred to any other candidate by a strict majority of voters, that is, for all $y \neq x$, $\#\{i, x \succ_i y\} > \frac{N}{2}$. It is well-known that there are some profiles for which no CW exists. Obviously, when a CW exists then it is unique.

3 Sequential voting

Given a combinatorial set of alternatives and a compact representation (in some preference representation language \mathcal{R}) of the voters' preferences, how can we compute the (set of) winner(s)? The naive way consisting in "unfolding" the compactly expressed preferences (that is, generating the whole preference relations on $D_1 \times \dots \times D_p$ from the input), and then applying a given voting rule, is obviously unfeasible, except if the number of variables is really small. We can try to do better and design an algorithm for applying a given voting rule r on a succinctly described profile P without generating the preferences relations explicitly. However, we can't be too optimistic, because it is known that the latter problem is computationally hard, even for simple succinct representation languages and simple rules (see [Lang, 2004]).

A way of escaping this problem consists in *restricting the set of admissible preference profiles* in such a way that computationally simple voting rules can be applied. A very natural restriction (that we investigate in the next Section) consists in assuming that preferences enjoy some specific structural properties such as conditional preferential independencies.

3.1 Sequential voting rules and correspondences

Now comes the central assumption to the sequential approach: *there exists an acyclic graph G such that the linear preference relation of every voter is compatible with G* . This assumption is not as restrictive as it may appear at first look: suppose indeed that the linear preference relations $(\succ_1, \dots, \succ_N)$ are compatible with the acyclic graphs G_1, \dots, G_N , whose sets of edges are E_1, \dots, E_N . Then they are a fortiori compatible with the graph G^* whose set of edges is $E_1 \cup \dots \cup E_N$. Therefore, if G^* is *acyclic*, then sequential voting will be applicable to $(\succ_1, \dots, \succ_N)$ (of course, this is no longer true if G^* has cycles). Moreover, in many real-life domains it may be deemed reasonable to assume that preferential dependencies between variables coincide for all agents.

Sequential voting consists then in applying "local" voting rules or correspondences on single variables, one after the other, in such an order that the local vote on a given variable

can be performed only when the local votes on all its parents in the graph G have been performed.

We define $Comp_G$ as the set of all collective profiles $P = (\succ_1, \dots, \succ_N)$ such that each \succ_i is compatible with G .

Definition 1 Let G be an acyclic graph on V ; let $P = (P_1, \dots, P_N)$ in $Comp_G$, $\mathcal{O} = \mathbf{x}_1 > \dots > \mathbf{x}_p$ a linear order on V following G , and (r_1, \dots, r_p) a collection of deterministic voting rules (one for each variable \mathbf{x}_i). The sequential voting rule $Seq(r_1, \dots, r_p)$ is defined as follows:

- $x_1^* = r_1(P_1^{\mathbf{x}_1}, \dots, P_N^{\mathbf{x}_1});$
- $x_2^* = r_2(P_1^{\mathbf{x}_2 | \mathbf{x}_1 = x_1^*}, \dots, P_N^{\mathbf{x}_2 | \mathbf{x}_1 = x_1^*});$
- ...
- $x_p^* = r_p(P_1^{\mathbf{x}_p | \mathbf{x}_1 = x_1^*, \dots, \mathbf{x}_{p-1} = x_{p-1}^*}, \dots, P_N^{\mathbf{x}_p | \mathbf{x}_1 = x_1^*, \dots, \mathbf{x}_{p-1} = x_{p-1}^*})$

Then $Seq(r_1, \dots, r_p)(P) = (x_1^*, \dots, x_p^*)$.

Example 2 Let $N = 12$, $V = \{\mathbf{x}, \mathbf{y}\}$ with $Dom(\mathbf{x}) = \{x, \bar{x}\}$ and $Dom(\mathbf{y}) = \{y, \bar{y}\}$, and $P = (P_1, \dots, P_{12})$ the following 12-voter profile:

$$\begin{aligned} P_1, P_2, P_3, P_4: & \quad xy \succ \bar{x}y \succ x\bar{y} \succ \bar{x}\bar{y} \\ P_5, P_6, P_7: & \quad \bar{x}y \succ xy \succ \bar{x}\bar{y} \succ x\bar{y} \\ P_8, P_9, P_{10}: & \quad \bar{x}y \succ \bar{x}\bar{y} \succ xy \succ x\bar{y} \\ P_{10}, P_{11}: & \quad \bar{x}y \succ \bar{x}\bar{y} \succ x\bar{y} \succ xy \end{aligned}$$

All these linear preference relations are compatible with the graph G over $\{\mathbf{x}, \mathbf{y}\}$ whose single edge is (\mathbf{x}, \mathbf{y}) ; equivalently, they follow the order $\mathbf{x} > \mathbf{y}$. Hence, $P \in Comp_G$. The corresponding conditional preference tables are:

voters 1,2,3,4	voters 5,6,7	voters 8,9,10	voters 11,12
$x \succ \bar{x}$	$x \succ \bar{x}$	$\bar{x} \succ x$	$\bar{x} \succ x$
$x : y \succ \bar{y}$	$x : \bar{y} \succ y$	$x : y \succ \bar{y}$	$x : \bar{y} \succ y$
$\bar{x} : y \succ \bar{y}$	$\bar{x} : y \succ \bar{y}$	$\bar{x} : y \succ \bar{y}$	$\bar{x} : y \succ \bar{y}$

Take r_x and r_y both equal to the majority rule, together with a tie-breaking mechanism which, in case of a tie between x and \bar{x} (resp. between y and \bar{y}), elects x (resp. y). The projection of P on \mathbf{x} is composed of 7 votes for x and 5 for \bar{x} , that is, P_i^x is equal to $x \succ \bar{x}$ for $1 \leq i \leq 7$ and to $\bar{x} \succ x$ for $8 \leq i \leq 12$. Therefore $x^* = r_x(P_1^x, \dots, P_{12}^x) = x$: the \mathbf{x} -winner is $x^* = x$. Now, the projection of P on \mathbf{y} given $\mathbf{x} = x$ is composed of 7 votes for y and 5 for \bar{y} , therefore $y^* = y$, and the sequential winner is now obtained by combining the \mathbf{x} -winner and the conditional \mathbf{y} -winner given $\mathbf{x} = x^* = x$, namely $Seq_{r_x, r_y}(P) = xy$.

In addition to sequential voting rules, we also define *sequential voting correspondences* in a similar way: if for each i , C_i is a correspondence on D_i , then $Seq(C_1, \dots, C_p)(P)$ is the set of all outcomes (x_1, \dots, x_p) such that $x_1 \in C_1(P_1^{\mathbf{x}_1}, \dots, P_N^{\mathbf{x}_1})$, and for all $i \geq 2$, $x_i \in C_i(P_i^{\mathbf{x}_i | \mathbf{x}_1 = x_1, \dots, \mathbf{x}_{i-1} = x_{i-1}}, \dots, P_N^{\mathbf{x}_i | \mathbf{x}_1 = x_1, \dots, \mathbf{x}_{i-1} = x_{i-1}})$. Due to the lack of space, we give results for voting rules only.

An important property of such sequential voting rules and correspondences is that the outcome does not depend on \mathcal{O} , provided that G follows \mathcal{O} . This can be expressed formally:

Observation 3 Let $\mathcal{O} = (\mathbf{x}_1 > \dots > \mathbf{x}_p)$ and $\mathcal{O}' = (\mathbf{x}_{\sigma(1)} > \dots > \mathbf{x}_{\sigma(p)})$ be two linear orders on V such that G follows both \mathcal{O} and \mathcal{O}' . Then

$Seq(r_1, \dots, r_p)(P) = Seq(r_{\sigma(1)}, \dots, r_{\sigma(p)})(P)$
and similarly for voting correspondences.

Note that when all variables are binary, all “reasonable” neutral voting rules (we have no space to comment on what “reasonable” means – and this has been debated extensively in the social choice literature) coincide with the majority rule when the number of candidates is 2 (plus a tie-breaking mechanism). Therefore, if all variables are binary and the number of voters is odd (in which case the tie-breaking mechanism is irrelevant), then the only “reasonable” sequential voting rule is $Seq(r_1, \dots, r_n)$ where each r_i is the majority rule.

It is important to remark that, in order to compute $Seq(r_1, \dots, r_p)(P)$, we do not need to know the linear preference relations P_1, \dots, P_N entirely: everything we need is the local preference relations: for instance, if $V = \{\mathbf{x}, \mathbf{y}\}$ and G contains the only edge (\mathbf{x}, \mathbf{y}) , then we need first the unconditional linear preference relations on \mathbf{x} and then the linear preference relations on \mathbf{y} conditioned by the value of \mathbf{x} . In other words, if we know the conditional preference tables (for all voters) associated with the graph G , then we have enough information to determine the sequential winner for this profile, even though some of the preference relations induced from these tables are incomplete. This is expressed more formally by the following fact (a similar result holds for correspondences):

Observation 4 Let $V = \{\mathbf{x}_1, \dots, \mathbf{x}_p\}$, G an acyclic graph over V , and $P = (P_1, \dots, P_N)$, $P' = (P'_1, \dots, P'_N)$ two complete preference profiles such that for all $i = 1, \dots, N$ we have $P_i \sim_G P'_i$. Then, for any collection of local voting rules (r_1, \dots, r_p) , we have

$$Seq(r_1, \dots, r_p)(P) = Seq(r_1, \dots, r_p)(P').$$

This, together with Observation 2, means that applying sequential voting to two collections of linear preference relations corresponding to the same collection of CP-nets gives the same result. This is illustrated on the following example.

Example 3 Everything is as in Example 2, except that we don't know the voters' complete preference relations, but only their corresponding conditional preference tables. These conditional preferences contain strictly less information than P , because some of the preference relations they induce are not complete: for instance, the induced preference relation for the first 4 voters is $xy \succ \bar{x}\bar{y} \succ \bar{x}y$, $xy \succ x\bar{y} \succ \bar{x}\bar{y}$, with $x\bar{y}$ and $\bar{x}y$ being incomparable. However, we have enough information to determine the sequential winner for this profile, even though some of the preference relations are incomplete. For instance, taking again the majority rule for r_x and r_y , the sequential winner is xy for any complete profile $P' = (P'_1, \dots, P'_{12})$ extending the incomplete preference relations induced by the 12 conditional preference tables above.

3.2 Sequential decomposability

We now consider the following question: given a voting rule r , is there a way of computing r sequentially when the linear preference relations enjoy common preferential independencies?

Definition 2 A voting rule r on $X = D_1 \times \dots \times D_p$ is **decomposable** if and only if there exist n voting rules r_1, \dots, r_p on D_1, \dots, D_p such that for any linear order $\mathcal{O} = \mathbf{x}_1 \succ \dots \succ \mathbf{x}_p$ on V and for any preference profile $P = (P_1, \dots, P_N)$ such

that each P_i follows \mathcal{O} , we have $Seq(r_1, \dots, r_p)(P) = r(P)$. The definition is similar for correspondences.

An interesting question is the following: for which voting rules r does the sequential winner (obtained by sequential applications of r) and the “direct” winner (obtained by a direct application of r) coincide? The following result shows that this fails for the whole family of scoring rules (and similarly for correspondences).

Proposition 1 No positional scoring rule is decomposable.

Proof sketch: We give a proof sketch for the case of two binary variables (this generalizes easily to more variables, as well as to non-binary variables). Let r be a decomposable scoring rule on $2^{\{\mathbf{x}, \mathbf{y}\}}$: there exist two local rules r_x and r_y such that whenever P follow $\mathbf{x} \succ \mathbf{y}$, we have $Seq(r_x, r_y)(P) = r(P)$. Then we show that r_x and r_y are both the majority rule (this follows easily from the fact that some properties of r , including monotonicity, carry on to r_x and r_y .) Now, consider the same preference profile P as in Example 3. P follows the order $\mathbf{x} \succ \mathbf{y}$. Now, let $s_1 \geq s_2 \geq s_3 \geq s_4 = 0$ (with $s_1 > 0$) the weights defining r . The score of xy is $4s_1 + 3s_2 + 3s_3$; the score of $\bar{x}\bar{y}$ is $5s_1 + 4s_2 + 3s_3$, which is strictly larger than the score of xy , therefore xy cannot be the winner for r , whatever the values of s_1, s_2, s_3 . ■

Such counterexamples can be found for many usual voting rules outside the family of scoring rules (we must omit the results due to the lack of space), including the whole family of voting rules based on the majority graph. Positive results, on the other hand, seem very hard to get. Obviously, dictatorial rules (electing the preferred candidate of some fixed voter) and constant rules (electing a fixed candidate whatever the voters' preferences) are decomposable. But the latter rules are of course not reasonable, and we conjecture that the answer to the above question is negative as soon as a few reasonable properties are required.

A particular case of preferential independence is when all variables are preferentially independent from each other, which corresponds to a dependency graph G with no edges. In this case, the preference profile follows any order on the set of variables, and the sequential winner is better called a *parallel winner*, since the local votes on the single variables can be performed in any order. We might then consider the following property of *separability*:

Definition 3 A deterministic rule r is separable if and only if for any preference profile $P = (\succ_1, \dots, \succ_N)$ such that the variables are pairwise conditionally preferentially independent, the parallel winner of r w.r.t. P is equal to $r(P)$.

Obviously, any decomposable rule is separable. Are there any separable rules? Focusing on positional scoring rules, we find a rather intriguing result (the proof of which is omitted):

Proposition 2 Let $V = \{\mathbf{x}_1, \dots, \mathbf{x}_p\}$ (with $p \geq 2$).

- if $p = 2$ and both variables are binary, exactly one positional scoring rule is separable: the rule associated with the weights $s_1 = 2s_2 = 2s_3$ (and $s_4 = 0$).
- in all other cases ($p \geq 3$ or at least one variable has more than 2 possible values), then no positional scoring rule is separable.

3.3 Sequential Condorcet winners

We may now wonder whether a Condorcet winner (CW), when there exists one, can be computed sequentially. Sequential Condorcet winners (SCW) are defined similarly as for sequential winners for a given rule: the SCW is the sequential combination of “local” Condorcet winners.

Definition 4 Let G be an acyclic graph and $P = (\succ_1, \dots, \succ_N)$ a profile in $Comp_G$. Let $\mathcal{O} = \mathbf{x}_1 > \dots > \mathbf{x}_p$ be a linear order on V following G . (x_1^*, \dots, x_p^*) is a sequential Condorcet winner for P if and only if

- $\forall x'_1 \in D_1, \#\{i, x_1^* \succ_i x'_1\} > \frac{N}{2}$;
- for every $k > 1$ and $\forall x'_k \in D_k$,
 $\#\{i, x_k^* \succ_i^{x_k | x_1=x_1^*, \dots, x_{k-1}=x_{k-1}^*} x'_k\} > \frac{N}{2}$.

This definition is well-founded because we obtain the same set of SCWs for any \mathcal{O} following G . The question is now, do SCWs and CWs coincide? Clearly, the existence of a SCW is no more guaranteed than that of a CW, and there cannot be more than one SCW. We have the following positive result, whose proof (which we omit) does not present any difficulty:

Proposition 3 Let G be an acyclic graph and $P = (\succ_1, \dots, \succ_N)$ in $Comp_G$. If $(x_1^*, x_2^*, \dots, x_n^*)$ is a Condorcet winner for P , then it is a sequential Condorcet winner for P .

The following example shows that the converse fails: 2 voters have the preference relation $x\bar{y} \succ \bar{x}\bar{y} \succ xy \succ \bar{x}y$, one voter $xy \succ x\bar{y} \succ \bar{x}y \succ \bar{x}\bar{y}$, and 2 voters $\bar{x}y \succ \bar{x}\bar{y} \succ xy \succ x\bar{y}$. \mathbf{x} and \mathbf{y} are mutually preferentially independent in all relations, therefore the SCW is the combination of the locals CW for $\{\mathbf{x}\}$ and for $\{\mathbf{y}\}$, provided they exist. Since 3 voters unconditionally prefer x to \bar{x} , x is the $\{\mathbf{x}\}$ -CW; similarly, 3 voters unconditionally prefer y to \bar{y} and y is the $\{\mathbf{y}\}$ -CW. Therefore, xy is the SCW for the given profile; but xy is not a CW for this profile, because 4 voters prefer $\bar{x}\bar{y}$ to xy .

Note that Proposition 3 extends to weak Condorcet winners (where strict majority is replaced by weak majority).

We now give a condition on the preference relations ensuring that SCWs and CWs coincide. Let $\mathcal{O} = \mathbf{x}_1 \triangleright \dots \triangleright \mathbf{x}_p$ be a linear order on V . We say that a linear preference relation \succ on D_X is *conditionally lexicographic* w.r.t. \mathcal{O} if there exist p local conditional preference relations $\succ^{x_i | x_1=x_1, \dots, x_{i-1}=x_{i-1}}$ x'_i for $i = 1, \dots, p$, such that $\bar{x} \succ \bar{y}$ if and only if there is a $j \leq p$ such that (a) for every $k < j$, $x_k = y_k$ and (b) $x_j \succ^{x_j | x_1=x_1, \dots, x_{j-1}=x_{j-1}} y_j$. A profile $P = (\succ_1, \dots, \succ_N)$ is *conditionally lexicographic* w.r.t. \mathcal{O} if each \succ_i is conditionally lexicographic w.r.t. \mathcal{O} . Such preference relations can be represented by TCP-nets [Brafman et al., 2006] or conditional preference theories [Wilson, 2004].

Proposition 4 Let \mathcal{O} be a linear strict order over V . If $P = (\succ_1, \dots, \succ_N)$ is conditionally lexicographic w.r.t. \mathcal{O} , then \bar{x} is a sequential Condorcet winner for P if and only if it is a Condorcet winner for P .

4 Arrow’s theorem and structured domains

We end this paper by considering decomposable domains from the point of view of preference aggregation. A preference aggregation function maps a profile to an aggregated

profile representing the preference of the group. Arrow’s theorem [Arrow, 1963] states that any aggregation function defined on the set of all profiles and satisfying unanimity and independence of irrelevant alternatives (IIR) is dictatorial. An Arrow-consistent domain \mathcal{D} is a subset of \mathcal{P} allowing for unanimous, IIR and nondictatorial aggregation functions.

It is easy to see that for any acyclic graph G , $Comp(G)$ is Arrow-consistent. Indeed, consider the preference aggregation function defined as follows:

- reorder the variables in an order compatible with G , i.e., w.l.o.g., assume that there is no edge (x_i, x_j) in G with $i \geq j$. Such an order exists because G is acyclic.
- let $h : V \rightarrow \mathcal{A}$ be a *non-constant* mapping from variables to voters (such a mapping exists because $|V| \geq 2$ and $|\mathcal{A}| \geq 2$).
- for any \bar{x} and $\bar{y} \neq \bar{x}$, let $k(\bar{x}, \bar{y}) = \min\{j, x_j \neq y_j\}$.
- for any collective profile $\langle \succ_1, \dots, \succ_N \rangle$, define $\succ_* = f_h(\succ_1, \dots, \succ_N)$ by: for all \bar{x} and \bar{y} , $\bar{x} \succ_* \bar{y}$ if $x_k \succ_{h(k)}^{x_k | x_1=x_1, \dots, x_n=x_n} y_k$, where $k = k(\bar{x}, \bar{y})$.

Proposition 5 f_h is a nondictatorial aggregation function on $Comp(G)$ satisfying unanimity and IIR.

Thus, restricting the set of admissible profiles to $Comp(P)$ is sufficient for escaping Arrow’s theorem: in other terms, $Comp(G)$ is Arrow-consistent². f_h is easier understood when it is turned into a voting rule: voter $h(x_1)$ first chooses his preferred value for variable x_1 , then voter $h(x_2)$ comes into play and chooses his preferred value for variable x_2 given the value assigned to x_1 , and so forth. Note that f_h is truly nondictatorial, because no agent imposes her own preference to the group; however f_h however has p local dictators (one for each variable), since voter $h(i)$ imposes his preference on the domain of x_i . We now prove that there must exist such local dictators, provided that the following three properties are satisfied:

(P1) preservation of the independence structure

f is a mapping from $Comp(G)^N$ to $Comp(G)$.

(P2) independence of irrelevant values and variables

For any $\mathbf{x}_i \in V$, $\bar{z} \in D_{Par(\mathbf{x}_i)}$, and $P = \langle P_1, \dots, P_N \rangle$, $Q = \langle Q_1, \dots, Q_N \rangle$ in $Comp(G)^N$ such that for every j and all $x, x' \in D_i$, $x \succ_{P_j}^{x_i | \bar{z}} x'$ iff $x \succ_{Q_j}^{x_i | \bar{z}} x'$, we have $x \succ_{f(P)}^{x_i | \bar{z}} x'$ if and only if $x \succ_{f(Q)}^{x_i | \bar{z}} x'$.

(P3) local unanimity

For any $P \in Comp(G)$, $\mathbf{x}_i \in V$ and $\bar{z} \in D_{Par(\mathbf{x}_i)}$, if $P_1^{x_i | \bar{z}} = \dots = P_N^{x_i | \bar{z}}$ then $f(P)^{x_i | \bar{z}} = P_1^{x_i | \bar{z}}$.

(P1) expresses that the preferential independencies expressed in the graph G should be transferred to the aggregated preference relation. Therefore, under (P1), for any preference relation \succ_* resulting from the aggregation of N preferences relations in $Comp(G)$ there exist p local conditional preference relations $\succ_*^{x_i | Par(\mathbf{x}_i)}$, for $i = 1, \dots, p$.

²As a referee pointed out to me, this domain restriction does not preclude the existence of Condorcet triples, which suggests that it is not as strong as it may seem at first glance.

(P2) is a local version of independence of irrelevant alternatives: whether the society prefers a value x_i to another value y_i of \mathbf{x}_i given an assignment \vec{z} of the parent variables of \mathbf{x}_i depends only on the voters' preferences between these two values given \vec{z} (and not on their preferences on other values of \mathbf{x}_i nor on their preferences on other variables.)

(P3) tells that if all voters have the same local preference relation over the values of a variable \mathbf{x}_i given a fixed value \vec{z} of its parents, then the local collective preference on $D_{\mathbf{x}_i}$ given \vec{z} should be equal to this local preference relation.

Importantly, note that the way (P2) and (P3) are written depends on the assumption that (P1) holds – otherwise we would not have been allowed to write $\succ_{f(P)}^{\mathbf{x}_i|\vec{z}}$.

Proposition 6 *Let G be an acyclic graph over a set of variables $V = \{\mathbf{x}_1, \dots, \mathbf{x}_p\}$ with domains D_1, \dots, D_p , with $|D_i| \geq 3$ for each i . An aggregation function f on $\text{Comp}(G)$ satisfies (P1), (P2) and (P3) if and only if there exists a local dictator $d(\mathbf{x}_i, \vec{z})$ for each \mathbf{x}_i and each $\vec{z} \in D_{\text{Par}(\mathbf{x}_i)}$, such that for each $\vec{t} \in D_V \setminus (\text{Par}(\mathbf{x}_i) \cup \{\mathbf{x}_i\})$ we have*

$$\vec{t}x_i \succ_{f(\succ_1, \dots, \succ_N)} \vec{t}x'_i \Leftrightarrow \vec{t}x_i \succ_{d(\mathbf{x}_i, \vec{z})} \vec{t}x'_i.$$

Proof sketch: The \Leftarrow direction is straightforward. For the \Rightarrow direction, let f satisfying (P1), (P2) and (P3). (P1) guarantees that for every \mathbf{x}_i and $\vec{z} \in D_{\text{Par}(\mathbf{x}_i)}$, \mathbf{x}_i is independent of $V \setminus (D_{\text{Par}(\mathbf{x}_i)} \cup \{\mathbf{x}_i\})$ given \vec{z} , therefore there exists a local collective preference relation $\succ_{\mathbf{x}_i|\vec{z}}$ such that for all $\vec{t} \in D_V \setminus (\text{Par}(\mathbf{x}_i) \cup \{\mathbf{x}_i\})$ and for all $x_i, x'_i \in D_{\mathbf{x}_i}$, $x_i \succ_{\mathbf{x}_i|\vec{z}} x'_i$. (P2) implies that $\succ_{\mathbf{x}_i|\vec{z}}$ is fully determined by voters' preferences on the values of \mathbf{x}_i given \vec{z} . Therefore, there exists a local aggregation function $f^{\mathbf{x}_i|\vec{z}}$ such that $\succ_{\mathbf{x}_i|\vec{z}} = f^{\mathbf{x}_i|\vec{z}}(\succ_1^{\mathbf{x}_i|\vec{z}}, \dots, \succ_N^{\mathbf{x}_i|\vec{z}})$. $f^{\mathbf{x}_i|\vec{z}}$ is easily shown to satisfy the conditions of Arrow' theorem. Applying Arrow's theorem to each $f^{\mathbf{x}_i|\vec{z}}$ enables us to conclude to the existence of a dictator $d(\mathbf{x}_i, \vec{z})$ for each \mathbf{x}_i and each $\vec{z} \in D_{\text{Par}(\mathbf{x}_i)}$, such that $x_i \succ_{f^{\mathbf{x}_i|\vec{z}}(\succ_1^{\mathbf{x}_i|\vec{z}}, \dots, \succ_N^{\mathbf{x}_i|\vec{z}})} x'_i \Leftrightarrow x_i \succ_{d(\mathbf{x}_i, \vec{z})} x'_i$. Now, the fact that $\succ_{\mathbf{x}_i|\vec{z}} = f^{\mathbf{x}_i|\vec{z}}(\succ_1^{\mathbf{x}_i|\vec{z}}, \dots, \succ_N^{\mathbf{x}_i|\vec{z}})$ allows us to conclude that $x_i \succ_{d(\mathbf{x}_i, \vec{z})} x'_i$ is equivalent to $x_i \succ_{f(\succ_1, \dots, \succ_N)} x'_i$, which in turn is equivalent to: for all $\vec{t} \in D_V \setminus (\text{Par}(\mathbf{x}_i) \cup \{\mathbf{x}_i\})$ and for all $x_i, x'_i \in D_{\mathbf{x}_i}$, $\vec{t}x_i \succ_{\ast} \vec{t}x'_i$, where $\succ_{\ast} = f(\succ_1, \dots, \succ_N)$. ■

Note that the local dictator for a given variable may depend on the values of its parents. For instance, with two variables \mathbf{x} and \mathbf{y} and a dependency graph with the edge (\mathbf{x}, \mathbf{y}) , we have a single dictator for \mathbf{x}_i and up to $|D_{\mathbf{x}}|$ dictators for \mathbf{y} .

Corollary 1 *Let $\text{Sep}(V)$ be the domain of fully separable linear preference relations on D_V . An aggregation function f on $\text{Sep}(V)$ satisfies (P1), (P2) and (P3) iff there exists p local dictators $d(\mathbf{x}_1), \dots, d(\mathbf{x}_p)$ such that for each $\vec{t} \in D_V \setminus \{\mathbf{x}_i\}$,*

$$\vec{t}x_i \succ_{f(\succ_1, \dots, \succ_N)} \vec{t}x'_i \Leftrightarrow \vec{t}x_i \succ_{d(\mathbf{x}_i)} \vec{t}x'_i.$$

Finally, note that knowing the local dictators does not fully determine f . For instance, let $V = \{\mathbf{x}, \mathbf{y}\}$, and G with no edge. Assume voter 1 prefers x_1 to x_2 and voter 2 prefers y_1 to y_2 . $\succ_{\ast} = f(\succ_1, \succ_2)$ satisfies $x_1 \succ_{\ast} x_2$ and $y_1 \succ_{\ast} y_2$, but this does not tell whether $x_1y_2 \succ_{\ast} x_2y_1$ or $x_2y_1 \succ_{\ast} x_1y_2$.

5 Conclusion

As far as we know, aggregating structured preferences on combinatorial domains exploiting preferential independence properties has never been considered neither in social choice nor in AI. [Rossi *et al.*, 2004] define a multi-agent extension to CP-nets and propose various semantics for aggregating preferences; but they do not address computational issues.

This paper contains several negative results. But one important question is left unanswered: *what are the sequentially decomposable voting rules?* Answering this question (by finding a small set of properties implying that a rule cannot be decomposable) seems much more difficult than we thought, and this is of course an issue for further research.

Next, we identified a domain for which direct and sequential Condorcet winners coincide. Clearly, lexicographic preferences are very specific, so that we would like to find more reasonable restrictions for the latter property to hold.

Another important issue stems from the fact that in combinatorial domains with structural properties (independencies), direct (global) voting rules are generally not computable by a sequential application of local rules: so, what should we favor? Global voting rules, which are well studied in social choice but which take no advantage of preferential independencies, or sequential local rules, which are based on the dependency graph, thereby being more intuitive and easier to compute? A theoretical comparison between global voting and sequential local voting is a highly promising issue.

Acknowledgements

I wish to thank Christophe Gonzales, Martin Meier, Patrice Perny and Lirong Xia for many helpful discussions.

References

- [Arrow, 1963] K. Arrow. *Social Choice and Individual Values*. Wiley, 1963. 2nd edition.
- [Boutilier *et al.*, 2004] C. Boutilier, R. Brafman, C. Domshlak, H. Hoos, and D. Poole. CP-nets: a tool for representing and reasoning with conditional ceteris paribus statements. *JAIR*, 21:135–191, 2004.
- [Brafman *et al.*, 2006] R. Brafman, C. Domshlak, and E. Shimony. On graphical modeling of preference and importance. *JAIR*, 25:389–424, 2006.
- [Brams and Fishburn, 2004] S. Brams and P. Fishburn. Voting procedures. In *Handbook of Social choice and Welfare*, chapter 4. Elsevier, 2004.
- [Brams *et al.*, 1998] S. Brams, D. Kilgour, and W. Zwicker. The paradox of multiple elections. *Social Choice and Welfare*, 15:211–236, 1998.
- [Lang, 2004] J. Lang. Logical preference representation and combinatorial vote. *Annals of Mathematics and Artificial Intelligence*, 42(1):37–71, 2004.
- [Rossi *et al.*, 2004] F. Rossi, K. Venable, and T. Walsh. mCP-nets: representing and reasoning with preference of multiple agents. In *Proceedings of AAAI-04*, 2004.
- [Wilson, 2004] N. Wilson. Consistency and constrained optimisation for conditional preferences. In *Proceedings of ECAI-04*, pages 888–892, 2004.