Issue-by-issue voting: an experimental evaluation

Hélène Fargier[1] and **Jérôme Lang**[2] and **Jérôme Mengin**[1] and **Nicolas Schmidt**[1]

Abstract. Multiple referenda consists in making a common decision about each of a set of binary issues, given the preferences of a set of voters. Asking voters for their preferences on all combinations of values is not feasible in practice, because of the exponentially large number of such combinations; on the other hand, voting issue-by-issue on each of the issues can lead to strongly paradoxical outcomes. This paper proposes to measure experimentally to which extent it is suboptimal to vote issue-by-issue voting, in function of the voting rule to be implemented, and the nature of the voters' preferences (arbitrary, separable or additive). For this we use randomly generated separable profiles, which turns out to be a difficult problem.

keywords. Computational social choice; preferences; voting; combinatorial domains; random generation.

1 Introduction

In many practical group decision making contexts, voters have to agree on a value to be given to each of a set of variables, or issues. The simplest approach – which is the one generally used in practice – consists in decomposing the voting process into a set of elementary voting processes, each bearing on a single variable, and processed simultaneously (*i.e.*, in issue-by-issue). When voters have preferential dependencies, however, such a decomposition may (in theory) lead to counterintuitive results [5, 15, 1, 19]. Consider a first example, with 500 voters and two Boolean issues A and B, whose value domains are respectively $\{a, \bar{a}\}$ and $\{b, \bar{b}\}$. The voters' preferences are:

200 voters: $a\bar{b} \succ \bar{a}b \succ \bar{a}\bar{b} \succ ab$
200 voters: $\bar{a}b \succ a\bar{b} \succ \bar{a}\bar{b} \succ ab$
100 voter: $ab \succ a\bar{b} \succ \bar{a}b \succ \bar{a}\bar{b}$

The last group of voters prefer a to $\bar{a}$, whatever the (fixed) value of B, and *vice versa*: their preference order is *separable*. In contrast, the other 400 voters have nonseparable preferences: for example, the first two prefer $A = a$ to $A = \bar{a}$ if $B = \bar{b}$ and prefer $A = \bar{a}$ to $A = a$ if $B = b$. If one asks voters to vote issue-by-issue, and if they majoritarily behave optimistically (*i.e*, the first group majoritarily chooses $A = a$), then we get a majority for $A = a$ and a majority for $B = b$: the final decision (a, b) is the worst for 80% of the voters.

As remarked by Lacy and Niou [15], issue-by-issue voting is much less of a problem when voters' preferences are separable[3]. However, separability is a highly demanding condition; moreover, it does not allow to avoid all paradoxes, even in its strongest version, as shown on the following example, that issue-by-issue voting does not satisfy *Pareto efficiency*[1, 17, 19].

Example 1. *We have three Boolean issues and three voters whose preferences are:*

v_1: $ab\bar{c} \succ a\bar{b}\bar{c} \succ \bar{a}b\bar{c} \succ \bar{a}\bar{b}\bar{c} \succ abc \succ a\bar{b}c \succ \bar{a}bc \succ \bar{a}\bar{b}c$
v_2: $a\bar{b}c \succ a\bar{b}\bar{c} \succ \bar{a}\bar{b}c \succ \bar{a}\bar{b}\bar{c} \succ abc \succ \bar{a}bc \succ ab\bar{c} \succ \bar{a}b\bar{c}$
v_3: $\bar{a}bc \succ \bar{a}\bar{b}c \succ \bar{a}b\bar{c} \succ \bar{a}\bar{b}\bar{c} \succ abc \succ ab\bar{c} \succ a\bar{b}c \succ a\bar{b}\bar{c}$

Issue-by-issue voting leads to the decision abc, which is Pareto-dominated by $\bar{a}\bar{b}\bar{c}$, that is to say all voters prefer $\bar{a}\bar{b}\bar{c}$ to abc.

These two problems raise concerns on the social acceptability of issue-by-issue voting. On the other hand, voting on combinations (or "bundles") of values, which may be the only way to escape them, is practically impossible to implement, because of the combinatorial nature of the problem. As a matter of fact, issue-by-issue voting, as imperfect as it may be, is used in quite many contexts, and in particular in multiple referenda, as they are held for example in California [5]. Other solutions have been suggested, such as sequential voting [16, 18, 10], using compact representation languages such as in [7]; each of them shows to have some benefits but also some pitfalls.

In this paper we stick to issue-by-issue voting. We would like to know to which extent this procedure can approximate a voting rule on a combinatorial domain, comparing the outcome of this procedure to the one that we would have got if preferences over bundles had been elicited and aggregated using a given voting rule. Some rather negative results have been given in [8], who consider a few rules based on scores, and for each of them, give worst-case approximation bounds of the ratio between the score of the alternative chosen by issue-by-issue voting and the score of the alternative chosen by the voting rule, for separable profiles. However, these worst-case negative results do not give much information about the *average-case* of issue-by-issue voting. Here we address this question experimentally, via a random generation of profiles; however, generating separable profiles appear to be a difficult problem, because the ratio between the number of separable preferences and the number of arbitrary preferences is very low [11]; we address the problem by several random generation methods in Section 3. In Section 4 we use these methods to assess the average quality of issue-by-issue voting.

2 Background

2.1 Voting

Let $\mathcal{X}$ be a finite set of m alternatives. A *vote* over $\mathcal{X}$ is a linear order $\succ$ over $\mathcal{X}$. A *profile* $P = (\succ_1, \dots, \succ_n)$ is a collection of n votes over $\mathcal{X}$, where $\succ_i$ is the vote of voter i. The vote of i represents her preferences, assuming that votes are sincere. A *voting rule* is a function r that associates to each profile P an alternative $r(P) \in \mathcal{X}$.

Several classical voting rules have been extensively studied (see e.g. [4] for a panorama of voting rules). In this paper, we are mainly interested in two groups of voting rules. The first group is that of scoring voting rules, that associate a score with each alternative, based on the ranks of the alternative in the votes. More precisely, given a vote $\succ$ and an alternative $x \in \mathcal{X}$, let $rk(\succ, x) \in \{1, \dots, m\}$

[1] IRIT – University of Toulouse; {lastname}@irit.fr
[2] LAMSADE – University of Paris-Dauphine; lang@irit.fr
[3] Three forms of separability will be defined in Section 2; here it is enough to say that the argument holds for each form.

denote the rank of x in $\succ$. A scoring voting rule is defined by a vector of scores $\langle s_1, \ldots, s_m \rangle$, such that $s_1 \geq \ldots \geq s_m$, so that s_i is the score associated with rank i. Every time an alternative x is ranked ith for some voter, this vote contributes s_i to the overall score of x. Given a profile $P = \langle \succ_1, \ldots, \succ_n \rangle$, the score of x for P is therefore $s(P, x) = \Sigma_{i=1}^{m} s_{rk(\succ_i, x)}$. The alternatives are then ranked according to their global score $s(P, x)$, and $r(P)$ is the alternative x that maximizes $s(P, x)$. Three distinguished scoring rules that will be considered in this paper are:

Borda: $s_1 = m - 1, s_2 = m - 2, \ldots, s_m = 0$;
Plurality: $s_1 = 1, s_2 = \ldots = s_m = 0$;
$\frac{m}{2}$-approval: $s_1 = s_2 = \ldots = s_{\frac{m}{2}} = 1, s_{\frac{m}{2}+1} = \ldots = s_m = 0$

Given a profile P, x is a *Condorcet winner* if for every $y \neq x$, a majority of voters rank x ahead of y. A voting rule r is said to be *Condorcet-consistent* if, for every profile P for which there is Condorcet winner x, $r(P) = x$. It is well-known that no scoring rule is Condorcet-consistent. The second group of rules that we study in the sequel contains Condorcet-consistent rules. For most of these rules, the winner can be determined from the *majority graph* associated with profile P: this graph contains an oriented edge from alternative x to y if a majority of voters ranks x ahead of y. More generally, a weighted majority graph associated with P indicates, for every pair of alternatives (x, y), the number of voters $N_P(x, y)$ that rank x ahead of y. In particular, we will consider two Condorcet-consistent voting rules in the sequel:

Copeland: the alternative that wins the most duels wins;
Maximin: the alternative x that maximizes $\min_{y \neq x} N_P(x, y)$ wins.

2.2 Combinatorial domains

We consider a set $\mathcal{I} = \{A, B, C, ...\}$ of p *issues*, each issue being associated with a binary domain (the possible answers): $D(A) = D(B) = D(C) = \cdots = \{0, 1\}$. Then $\mathcal{X} = D(A) \times D(B) \times D(C) \times \ldots$ is the set of the possible alternatives, or, using voting terminology, *candidates*. The number of alternatives is thus $m = 2^p$. The elements of $\mathcal{X}$ are vectors $\vec{x}$, $\vec{x'}$; we will often concatenate the answers to describe a particular alternative. For instance, if $\mathcal{I} = \{A, B, C\}$, $(1, 0, 1)$ denotes the alternative that has answer 1 for issue A, answer 0 for issue B, and answer 1 for issue C. We will also use concatenation for vectors of answers for disjoint sequences of issues: for instance, if $\mathcal{I} = \{A, B, C, D\}$, $Y = (A, B)$, $Z = (C, D)$, $\vec{y} = (1, 0)$, $\vec{z} = (0, 1)$, then $\vec{y}.\vec{z}$ denotes the alternative $(1, 0, 0, 1)$. Lastly, for every $X \subseteq \mathcal{I}$, D_X is the set of assignments $\vec{x}$ of elements of X in their respective domains.

When considering orderings over combinatorial domains, there exist three definitions of separability.
Let $\succ$ be a vote over $\mathcal{X}$, it is:

weakly separable if for every variable $A \in \mathcal{I}$, every $v, v' \in D_A$, $\vec{x}, \vec{x}' \in D_{\mathcal{I} \setminus \{A\}}$: $v.\vec{x} \succ v'.\vec{x} \iff v.\vec{x'} \succ v'.\vec{x'}$
strongly separable if for every partition $\{X, Y\}$ of $\mathcal{I}$ and every $\vec{x}, \vec{x}' \in D_{\mathcal{X}}, \vec{y}, \vec{y}' \in D_{\mathcal{Y}}$: $\vec{x}.\vec{y} \succ \vec{x'}.\vec{y} \iff \vec{x}.\vec{y'} \succ \vec{x'}.\vec{y'}$
additively separable if for every issue $X_i \in \mathcal{I}$ there exists a function $u_i : D_i \to \mathbb{R}^+$ such that for every $\vec{x}, \vec{y} \in \mathcal{X}$ we have $\vec{x} \succ \vec{y}$ if and only if $\sum_i u_i(x_i) > \sum_i u_i(y_i)$.

Example 1. (cont.) *The voters on example 1 have strongly separable preferences. For instance, we have $a \succ_1 \bar{a}, \forall B, C$.*

Strong separability is sometimes called mutual preferential independence, as in [13]. It requires that preferences over combinations of values of any subset of variables do not depend on the values of other variables. Weak separability only requires that preferences over values of a single variable do not depend on the fixed values of other variables; it is met by preference relations associated with a CP-net with no edge in the dependency graph [2]. Additive separability implies strong separability, which in turn implies weak separability. As soon as $p \geq 5$, additive separability is strictly stronger than strong separability [14], whereas both notions coincide for $p \leq 4$ [3]. Therefore, we have only one notion of separability for $p = 2$, two distinct notions for $p = 3$ and $p = 4$. On continuous domains, strong and additive separability are equivalent [9].

2.3 Problematics

Our main objective is to compare the results of the strict application of a voting rule over a combinatorial domain with the results obtained when issue-by-issue voting is used. It is known that as soon as $p \geq 3$, issue-by-issue voting does not satisfy neutrality nor efficiency [1, 19], which implies that any voting rule on a combinatorial domain with more than two variables that satisfies any of these two properties (and all commonly used voting rules do) does not coincide with the issue-by-issue voting rule. For instance, in Example 1, the issue-by-issue winner is abc, whereas the Plurality cowinners are $ab\bar{c}$, $a\bar{b}c$ and $\bar{a}bc$, and the Borda winner is $\bar{a}\bar{b}\bar{c}$, which is also a Condorcet winner.

However, this impossibility theorem does not exclude the possibility that the outcomes do coincide *in general*. The rest of the paper addresses this question by trying to determine the probability that the outcomes coincide. Because of the difficulty of an analytical approach, we estimate this probability experimentally. Ideally it would be interesting to conduct these experiments with real-world data, but such voting data are rarely available, and, to our knowledge, there exist no available data of significant size for multiple referenda (in the data on multiple referenda used in [5], only the combination of the preferred values of each voter is given, not their entire preference relations). Therefore we choose to generate random samples.

3 Generation of separable profiles

The generation of random profiles in social choice is not a new problem: the *impartial culture assumption*, which assumes a uniform distribution over the set of possible profiles, is often made. In our case, we have $m = 2^p$ alternatives, thus there are $2^p!$ possible votes for each voter. Generating uniformly distributed profiles over such domains is not a problem as long as the number p of issues is low. However, the probability of obtaining a weakly separable (or, *a fortiori*, a strongly or additively separable) profile among the $2^p!$ possibilities is extremely weak. For instance, when $p = 4$, of the 16! possible orderings only 5376 are strongly (and additively) separable (a ratio of around $1/10^8$), and 26886144 are weakly separable (a ratio of around $1/10^6$). The exact numbers of orderings that are weakly (resp. additively, resp strongly) separable for $p > 5$ (resp. $p > 6$, resp. $p > 7$) are currently unknown – to our knowledge.

We have investigated several ways of generating random separable orders. A first, naive method, consists in picking random orders uniformly distributed and keeping only the (weakly/strongly/addtively) separable ones. Given the very low probability of picking a separable order, this method is practically not feasible as soon as $p > 4$.

Generating additively separable orders by random utility generation

When the value of p becomes too large for an explicit enumeration, we shall rely on multiattribute utility theory as a way of gen-

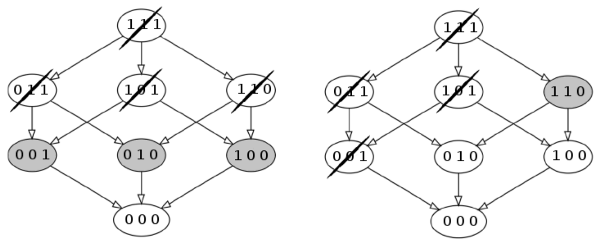

Figure 1. The lattice of options for $p = 3$, after the first four options have been chosen: (a) three possible choices, (b): one possible choice.

erating additively separable orders. The method consists in doing the following: for each variable X_i, we generate a random utility $u_i(x_i) \in (0, 1]$ for each of the possible values $x_i \in D(X_i)$; this results in a utility function on alternatives $u : \mathcal{X} \to \mathbf{R}$, and then we rank alternatives according to u. This method is a simplified version of the one developed in [6]. The distribution we get reaches every additively separable order with a positive probability, but is not uniform. It nevertheless has the advantage of being based on a well characterized model of rational decision makers.

Storing all normalized orders (weak and strong separability)

An ordering is normalized [11] if (i) the best alternative is $(1, ..., 1)$ and (ii) $(0, 1, 1, ..., 1) \succ (1, 0, 1, ..., 1) \succ \cdots \succ (1, 1, 1, ..., 0)$. Every separable order can be obtained from some normalized, separable order by permutation of some issues and inversion of some answers. Thus, in order to generate separable orders, we can first build a table of all normalized, separable orders, and then pick some of them at random and apply some permutations and inversions at random.

This method can be used for generating weakly separable orders if there are no more than $p = 4$ issues: for $p = 5$ issues our first experiments show that one would need a table of at least 5 terabytes (probably much more). For strongly separable orders, this method works well in practice for up to $p = 6$ issues

Generating weakly separable orders by lattice exploration

Normalized, weakly separable orders can be directly generated by considering the partial order $\succ^0$ that contains all pairs of the form $(1.\vec{x}, 0.\vec{x})$ for every $\vec{x} \in D_{\mathcal{I}\setminus\{A\}}$, for every $A \in \mathcal{I}$. All completions of $\succ^0$ are then normalized, weakly separable orders (and random permutations of issues and random inversions of answers will generate any weakly separable order). Such completions can be generated by random, top-down traversal of the graph that corresponds to $\succ^0$. The traversal must be such that no alternative is reached before preferred alternatives. The order in which the alternatives is reached is then a weakly separable order. More precisely, the first alternative $\vec{x}_1$ is the one that has the preferred value for every issue, $11 \dots 1$ if we build a normalized order. We can then pick at random the second preferred alternative $\vec{x}_2$ among those that are dominated by $\vec{x}_1$ only. The third preferred alternative is then picked at random among those that are only dominated by $\vec{x}_1$ and $\vec{x}_2$, and so on.

However, a uniform distribution over all possible alternatives every time a new alternative must be picked does not guarantee that the resulting distribution over complete, normalized, weakly separable orders will be uniform. This is due to the fact that the number of possible alternatives at a given step depends on the already chosen alternatives. Figure 3 depicts two different possibilities to pick the first four alternatives when there are $p = 3$ issues: it shows that the number of possible alternatives at this step is not constant. The probabilities of the most frequent and of the least frequent complete orders generated in this way can be calculated. For $p = 4$ issues, they are $1/331776$ and $1/238878720$ respectively.

A better idea to complete the partial order $\succ^0$ is to proceed level by level, top-down. After all the alternatives above a given level in the lattice have been ranked, we consider all alternatives $\vec{x}$ at the next level: we determine the possible ranks for each of them in the current order, so as to respect the weak separability property: the rank of $\vec{x}$ will then be picked randomly among these possible ranks. The alternative $\vec{x}$ which has the lowest number of possible ranks is ranked first – since ranking the other alternatives will not change its set of possible ranks, whereas ranking an alternative highly constrained may diminish the number of possible ranks for alternatives less constrained. Once all alternatives at this level have been ranked, we proceed to the next level of the lattice of $\succ^0$.

For instance, consider the following order obtained after adding the two first levels; $111 \succ 011 \succ 101 \succ 110$. At the next level, the least restricted alternative is 001, since there are two possible choices for it (and only one for 100 and for 010). If we insert first 001, it will have a probability one half to be in the fourth position, and one half to be between the positions 5 and 7 (since 100 and 010 will be inserted afterwards). If 001 is inserted last, there is one $\frac{1}{4}$ to have it finally ranked at each of its possible locations (4, 5, 6 and 7).

Generating weakly separable orders by reparation of non-separable orders

Another method to generate weakly separable preferences consists in randomly generating (not necessarily separable) orders $\vec{x}_1 \succ \vec{x}_2 \succ \dots \succ \vec{x}_m$ over $\mathcal{X}$ with a uniform distribution, and "reparing" each such order so as to make it weakly separable. The reparation works issue by issue, as follows: for every issue $A \in \mathcal{I}$, we first record the value that $\vec{x}_1$ has for A as being the preferred value for A; then we scan the entire ordering, making permutations every time we encounter pairs of alternatives that violate the weak separability condition with respect to A and the chosen preferred value. It can be proved that the resulting order, when it has been repaired w.r.t. all issues, is weakly separable. The order in which the issues are repaired has to be generated at random, since it has an effect on the resulting order over the alternatives.

However, using this algorithm to repair randomly, uniformly generated orders does not give a uniform distribution over weakly separable orders: this reparation procedure can be seen as a local search in the neighbourhood of the initial order; so if a weakly separable order has more not separable neighbours than another one, it has more chances of being obtained.

Evaluating the random generators

We thus get the utility-based generator (for that seems natural for additively separable preferences), the storing generator (for weakly separable preferences up to $p = 4$ and strongly separable preferences up to $p = 6$), and the reparation- and exploration- based generators, that have been designed to generate weakly separable preference for higher values of p in a way close to equiprobability.

In order to compare the quality of the random generators described above, we have tested them in the case of $p = 4$. We know in

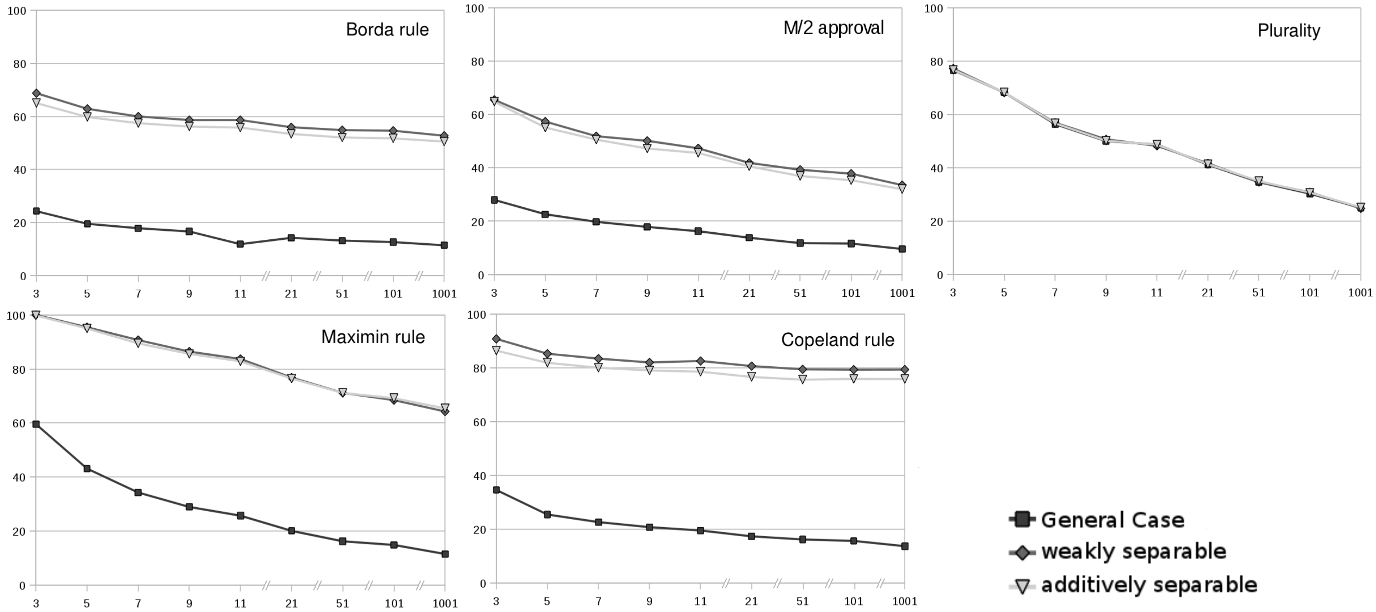

Figure 2. Success rate of issue-by-issue voting w.r.t. the number of voters; 4 issues

this case that there are 70016 normalized weakly separable orders (and therefore $70016 \times 2^4 \times 4! = 26886144$ weakly separable orders), and 14 strongly separable normalized orders (and therefore $14 \times 2^4 \times 4! = 5376$ strongly separable orders). Each generator was run at least 10^7 times so as to generate a normalized order, and we counted apparition of each preference order. The following table gives the frequencies of occurrence of the least frequent and most frequent orders, the ratio between these two frequencies, as well as an estimation as the entropy of each generator — the entropy measures the closeness of a probability distribution to equiprobability, and varies from 1 (equiprobability) to 0 (determinism).

	Storing	Utility	Reparation	Exploration
Weak sep.	X		X	X
Strong sep.	X			
additive sep.		X		
Max. freq.		4/25	1/3900	1/44050
Min. freq.		1/25	1/1400000	1/133333
max. freq. / min. freq.	1	4	360	3
Entropy	1	0,92	0,97	0,9994

Table 1. Entropy, minimal and maximal frequencies of apparition of the generated orders, by generator

Generating by storing is the only perfectly equiprobable generator. It can be used when the number of issues remains low: up to 4 for weakly separable orders, and up to 6 for strongly separable orders.

As we could expect, the utility-based generator has a bad entropy in the case of 4 issues. Nevertheless, it has the advantage of being simple, fast, and based on a realistic voter preference model, therefore this is the one we shall use to generate additively separable preferences for more than 6 issues. We do not currently know how to efficiently do this for more than 6 issues; therefore, this generator will also be used for strong separability, although it gives a zero probability to any strongly separable order which is not additively separable.

As to generating weakly separable preferences, our experiments suggest that the generator by lattice exploration is better than the reparation-based generator; we will use this argument to choose this generator for our experiments reported in Section 4.

4 Experimental study

The aim of the following experiments is to evaluate the interest of issue-by-issue voting for multiple referenda as an approximation of the application of a specific voting rule, applied to a profile over a combinatorial domain. We consider five voting rules: Borda, $\frac{m}{2}$-approval, Plurality, Maximin and Copeland, and we study the influence, on the quality of the approximation, of parameters such as the number of issues, the number of voters and the type of preference (weakly separable, strongly separable, additively separable).

Since the issues are binary, issue-by-issue voting leads to applying majority voting on each issue. When the profiles are not separable, we suppose that the voters adopt an optimistic attitude and prefer the values as prescribed by their preferred alternative. In order to limit the occurrence of ties, we assume the number of voters to be odd. The outcome of the issue-by-issue voting is then compared to the alternative chosen by the application of each specific voting rule r.

As for the generation of general profiles, without any assumption of separability, we use a uniform distribution over all profiles. For the generation of weakly/strongly/additively separable profiles, we use the storing-based generator for $p \leq 4$. When $p > 4$, we use the exploration-based generator and the utility based generator.

In the first experiment (Figures 2 to 5), we count the percentage of profiles that lead to the success of issue-by-issue voting (that is, the proportion of the generated profiles for which the original rule and issue-by-issue voting elect the same alternative; in case the application of the original voting rule gives a tie, we consider that issue-by-issue succeeds as soon as it elects one of the tied winners). Each of the following experiments repeats the test 10000 times (each point in the curve is computed on 10000 profiles).

4.1 Influence of separability

Unsurprisingly, the experimental results (cf. Figure 2; notice that for $p = 4$, strong and additive separability are equivalent, hence the figures draw only one curve for both concepts) are consistent with the theoretical results of [15] and [12]: issue-by-issue is sounder on separable profiles. For all the voting rules considered but Plurality, the success rate is better on separable samples than on the purely random samples. The result keeps holding when the number of voters increases (Figure 2) and when the number of issues increases (see Figure 3 for Borda; for the other rules studied, except Plurality, we get a similar behaviour). Notice that the success rate seems slightly better for weakly separable profiles than for strongly separable ones; we do not have a clear interpretation of this fact. For Plurality, the same results are obtained, whether or not separability is assumed. This can be easily explained by the fact that Plurality, like issue-by-issue voting, is tops-only, *i.e.*, the outcome is determined from the top of the votes, which implies that separability has no influence.

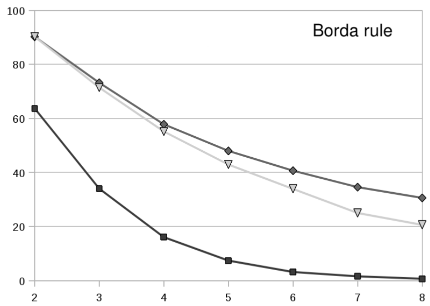

Figure 3. Borda's rule: success rate of issue-by-issue voting w.r.t. the number of issues; 11 voters.

From now on, we conduct the experiments on separable samples only: first, because purely random preferences lead to a quite bad success rate; and second, because a voter with non separable preferences can hardly give her preference issue by issue; we made the assumption that they report votes optimistically, which seems to be observed in practice, but this assumption can be questioned.

4.2 Comparison of voting rules

Let us first look at the case $p = 4$ (Figure 2; Figure 4 summarizes the 5 rules on the weakly separable sample).

We first observe that the three scoring rules are badly approximated; as soon a the number of voters reaches 7, the success rate of the approximation goes below the 60%, that is, for at least 6 cases among 10, issue-by-issue voting elects a winner that is different than the one designated by the application of the original voting rule. This rate is especially bad for Plurality and for $m/2$-approval.

Some rules, like plurality, $m/2$-approval or maximin generate many ties when the number of voters is low; this boosts the success rate of these rules for a few voters samples. That's why, for example, the success rate of maximin is 100% with 3 voters but decreases quickly when the number of voters increase.

It can moreover be noticed that the success rate gets worse as the number of voters increase. In a second series of tests, we measure the success rate for samples of 7 voters, letting the number of issues increase from 2 to 10, for the rules that were not too badly approximated according to the first experiment, namely Borda, Copeland and Maximin. Figure 5 reports our results for weakly separable profiles, generated by exploration (similar ones have been obtained for separable profiles and additively separable profiles): the success rate clearly depends on the number of issues (the more issues, the worst).

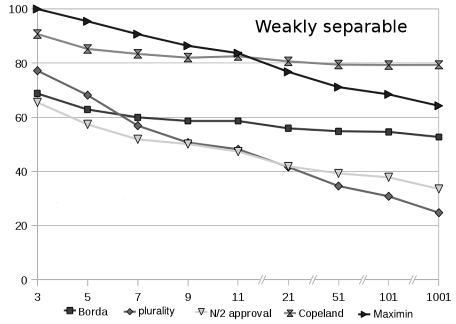

Figure 4. Success rate for Borda, Plurality, $m/2$-approval, Copeland, Maximin w.r.t. number of voters; weakly separable profiles, 4 issues

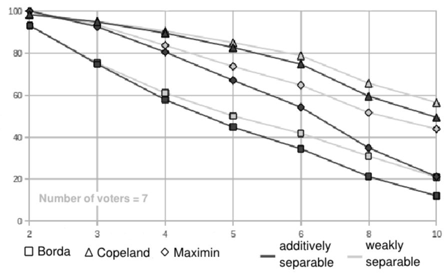

Figure 5. Success rate for Borda, Copeland, Maximin w.r.t. the number of issues, 7 voters

In summary, the success rate of issue-by-issue voting thus gets worse as the number of voters increases and, to a larger extent, as the number of issues increases. It quickly becomes very bad for Plurality and $\frac{m}{2}$-approval (once again). Results are better for Borda, but nevertheless falls below 50% for 5 issues (for 7 voters) or 10 voters (for 4 issues), which is disappointing. Finally, it is much better for Copeland and (to a lesser extent) Maximin; in both cases, a closer look to the sample reveals that bad results are highly correlated with the absence of a Condorcet winner.

4.3 Quality of the approximation

For 5 issues, the success rate of issue-by-issue falls below 50% for all considered scoring rules, and below 80% for Copeland and Maximin.However, the probability that both winners coincide is perhaps not the best way of measuring the approximation of a voting rule; as in [8], we may consider instead, for any voting rule r based on the maximization of some numerical score, the ratio between the score of the alternative elected by the issue-by-issue rule and the score of the winner of r. The following table gives these ratios. The first three lines of the table gives the average ratio taking all generated profiles into account, and the last one the average ratio when only the 'unsuccessful' profiles, that is, those for which the issue-by-issue winner and the winner for r differ. (Note that using these ratios for comparing different rules should be done with care, as these ratios depend heavily on the definition of the score.) For the sake of completeness we also give the success rate for separable profiles.

Once again, the best results are obtained for Borda, Copeland and Maximin: for these rules, the average approximation ratio is above 97%). The result are not as good for Plurality and $\frac{m}{2}$-approval.

$rule$:	Borda	plur.	$\frac{m}{2}$-app.	Cop.	maxim.
General case					
Average score of the real winner	72.71	1.9	5.73	14.22	3.47
Average score of the issue-by-issue winner	70.98	1.3	5.08	13.94	3.38
Ratio	0.976	0.684	0.887	0.98	0.974
Success rate	0.608	0.570	0.527	0.838	0.907
Unsuccessful profiles					
Average score of the real winner	70.59	1.99	5.77	13.36	3
Average score of the issue-by-issue winner	66.18	0.6	4.39	11.65	1.97
Ratio	0.938	0.302	0.761	0.872	0.657
Maximal distance	21	3	4	8	2
Minimal distance	1	1	1	1	1

Figure 6. Distance in the scores of the real and issue-by-issue winners , 7 voters and 4 issues, weakly separable profiles.

4.4 Pareto Efficiency

Recall that one of the drawbacks of issue-by-issue voting is its failure to satisfy efficiency as soon as $p \geq 3$. We give here the probability that the issue-by-issue winner is Pareto-dominated. An analytical study for $p = 3, m = 3$ gives a probability of $\frac{1}{2304}$ for strongly separable profiles and $\frac{1}{18432}$ for weakly separable profiles. For more than 3 issues or 3 voters, computing the probability analytically seems difficult, therefore we have once again proceeded to experiments with randomly generated profiles. For each couple $(\#issue, \#voters)$ we generated a set of 10^6 additively profiles, and for each of them, checked if the outcome is Pareto-dominated. Figure 7 give our results for additive separable profiles.

	$m = 3$	$m = 5$	$m = 7$	$m = 9$	$m = 11$
$p = 3$	436	20	4	0	0
$p = 4$	1699	111	6	0	0
$p = 5$	4211	329	23	4	0
$p = 6$	8268	671	34	4	0
$p = 7$	14052	1360	86	3	0
$p = 8$	21984	2217	150	14	0
$p = 9$	32324	3378	284	16	0

Figure 7. Pareto-dominated profiles for 10^6 profiles; additive separability

We can see that even if the probability of having a Pareto-dominated outcome is non-negligible when the number of voters is low (up to 3,2% for 3 voters and 9 issues), this probability decreases very quickly as the number of voters increases, and becomes negligible from 9 voters on – which is of course very unsurprising. Similar tests on weakly separable profiles give a probability of getting a Pareto-dominated outcome 10 to 100 times smaller than with additively separable profiles, the shape of the graphics being similar.

5 Conclusion

Although the initial motivation of this paper was an experimental comparison between issue-by-issue voting and common voting rules applied to separable profiles, it turned out that a surprisingly difficult issue that had to be addressed first was the random generation of separable profiles. As soon as there are at least 5 issues, we do not know how to generate weakly separable, separable nor additively separable profiles with an equiprobable distribution; we remedied this to some extent, by proposing two methods that generate distributions 'not too far' from equiprobability, and finally chose an algorithm based on the exploration of the lattice of alternatives for the generation of weakly separable profiles, and an algorithm relying on a utility-based representation of preferences, for strongly separable profiles. The first one performs well in terms of entropy (it is close to be equiprobable) and the second one has the advantage of being based on a well known (and well characterized) model of rational decision makers.

Concerning issue-by-issue voting, our result are rather negative, confirming the theoretical results in [8]. Although they show that the violation of efficiency is rare (which was expected), they also show that issue-by-issue voting is a bad approximation of scoring rules, in particular plurality. For r = plurality, Borda or $\frac{m}{2}$-approval, the issue-by-issue winner is different of the one elected by r for more than 70% of the profiles, even for only 4 issues and 5 voters. These results become worse when the number of issues or the number of voters increase. Copeland and (to a lesser extent) maximin do better: for instance, their winners coincide with the issue-by-issue winner for 80%of the profiles, for 4 issues. Unsurprisingly, our results also confirm that issue-by-issue behaves better on separable preference than on purely random profiles.

REFERENCES

[1] J.-P. Benoit and L.A. Kornhauser, 'Only a dictatorship is efficient or neutral', Technical report, NYU Law School, (2006).

[2] C. Boutilier, R. Brafman, C. Domshlak, H. Hoos, and D. Poole, 'CP-nets: a tool for representing and reasoning with conditional ceteris paribus preference statements', *JAIR*, **21**, 135–191, (2004).

[3] W. Bradley, J. Hodge, and M Kilgour, 'Separable discrete preferences', *Mathematical Social Sciences*, **49**(3), 335–353, (2005).

[4] S. Brams and P. Fishburn, 'Voting procedures', in *Handbook of Social Choice and Welfare, Volume 1*, eds., K. Arrow, A. Sen, and K. Suzumura, 173–236, Elsevier, (2002).

[5] S. Brams, W. Zwicker, and M. Kilgour, 'The paradox of multiple elections', *Social Choice and Welfare*, **15**(2), 211–236, (1998).

[6] D. Braziunas and C. Boutilier, 'Local utility elicitation in gai models', in *Proceedings of UAI'05*, pp. 42–49, (2005).

[7] V. Conitzer, J. Lang, and L. Xia, 'Hypercubewise preference aggregation on multi-issue domains', in *IJCAI*, (2011).

[8] V. Conitzer and L. Xia, 'Approximation of common voting rules by sequential voting rules on multi-issue domains', in *KR*, (2012). to appear.

[9] W.M. Gorman, 'The structure of utility functions', *Review of Economic Studies*, **35**, 367–390, (1968).

[10] U. Grandi and U. Endriss, 'Aggregating dependency graphs into voting agendas in multi-issue elections', in *IJCAI*, pp. 18–23, (2011).

[11] J. Hodge, 'Permutations of separable preference orders', *Discrete Appl. Math.*, **154**, 1478–1499, (2006).

[12] J. Hodge and P. Schwallier, 'How does separability affect the desirability of referendum election outcomes?', *Theory and Decision*, **61**(3), 251–276, (2006).

[13] R. L. Keeney and H. Raiffa, *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*, John Wiley and Sons, 1976.

[14] C. Kraft, J. Pratt, and A. Seidenberg, 'Intuitive probability on finite sets', *The Annals of Mathematical Statistics*, **30**(2), 408–419, (1959).

[15] D. Lacy and E. Niou, 'A problem with referenda', *J. of Theoretical Politics*, **12**(1), 5–31, (2000).

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

[17] I. Oskal-Sanver and R. Sanver, 'Ensuring pareto optimality by referendum voting', *Social Choice and Welfare*, **27**(1), 211–219, (Aug. 2006).

[18] G. Dalla Pozza, F. Rossi, M.-S. Pini, and B. Venable, 'Multi-agent soft constraint aggregation via sequential voting', in *IJCAI*, pp. 172–177, (2011).

[19] L. Xia and J. Lang, 'A dichotomy theorem on the existence of efficient or neutral sequential voting correspondences', in *IJCAI*, (2009).