Voting behavior in one-shot and iterative multiple referenda[1]

Umberto Grandi, Jérôme Lang, Ali I. Ozkes, and Stéphane Airiau

Abstract

We consider voters making a collective decision via simultaneous vote on two binary issues. Their preferences are captured by payoffs assigned to combinations of outcomes for each issue; they can be nonseparable. We pursue an experimental approach and investigate the impact of iterative voting in this context, repeating the voting process until a final outcome is reached. Our results from experiments run in the lab show that voters tend to have an optimistic rather than a pessimistic behaviour when casting a vote on a non-separable issue and that iterated voting may in fact improve the social outcome. We provide the first comprehensive empirical analysis of individual and collective behavior in the multiple referendum setting.

1 Introduction

Iterative voting is based on the idea that better outcomes can be reached if voters are allowed to change their vote a number of times, after seeing the distribution of the others' votes.

Beyond single-winner voting, a context where iterative voting can be helpful for voters to coordinate and reach better outcomes (and thus, in a sense, vote strategically in the positive sense of the term), is *multiple referenda*: voters have to come up with a collective yes/no decision on a number of binary issues, upon which they may have nonseparable preferences. Nonseparable preferences in multiple referenda sometimes lead to *multiple election paradoxes* [6]. Designing ways to avoid multiple election paradoxes is a very difficult problem. Perhaps the most convincing method we know of (at least in some contexts) is the one suggested by Bowman et al. [4]: *allowing for iterative voting*. They run computer simulations on a number of voting profiles with two and three issues giving rise to multiple election paradoxes, enabling the voters with the possibility of changing their votes in an iterated process. Their results show that iteration leads to a statistically significant improvement in social welfare.

However, the extent to which the results based on simulations carry to the real world is not clear. We propose lab experiments to investigate the behavior of human subjects in this context and the impact of iterative voting on individual and collective behavior. For the sake of simplicity, we focus on two binary issues. We want to test a number of hypotheses. The first series of hypotheses is independent from the dynamic aspects of the process. When a voter must choose a value (*yes* or *no*) for an issue, three cases may occur:

- when the issue is separable we expect the voters to choose the dominating value.
- when the issue is nonseparable, and when one value stochastically dominates the other one, we expect voters to choose the stochastically dominant value.
- when the issue is nonseparable and no value stochastically dominates the other one, voters have to choose between a value that leads to the best or the worst payoff (say,

[1]This paper is based on the following publication:

Grandi, U., Lang, J., Ozkes, Airiau, S.: Voting behavior in one-shot and iterative multiple referenda. *Social Choice and Welfare* (2022).

It has been slightly rewritten for a computational social choice readership (and shortened).

1 or 4) and a value that leads an intermediate payoff (say, 2 or 3). It was conjectured in [13] that a majority of voters would make the former choice (which we may call 'optimistic'), Testing this long-standing hypothesis is one of our primary motivations.

The second series of hypotheses concern the individual behaviour of voters in the iterated voting process: we postulate that voters change their vote about an issue from round t to round $t+1$ more frequently if they are pivotal at t and the issue is nonseparable. The last hypothesis we want to test is central in iterated voting: iteration increases social welfare as well as Condorcet efficiency. This has been already verified in other contexts (see [9] as well as [12] for social welfare, see [21] and [23] for Condorcet efficiency), and is the justification that Bowman et al. [4] put forward for justifying their use of an iterated process for multiple referenda. Verifying this hypothesis would confirm the validity of their process.

Our conclusions are twofold. First, we find that iterative voting does improve the quality of the outcome when this is measured as the average utility of the election outcome among the voters, confirming the conclusions of Bowman et al. [4]; second, we show that voters tend to behave rationally in terms of (standard and stochastic) domination, and to vote 'optimistically' more frequently than 'pessimistically'.

Section 2 discusses the related literature on multiple referenda and iterative voting. Section 3 presents the formal setting and Section 4 the details of our experiments. We then present results on 'optimistic' and 'pessimistic' behaviour (Section 5), on the individual voting dynamics in iteration (Section 6), and on the evolution of global parameters in iteration (Section 7). Section 8 concludes the paper. We sometimes refer to our full paper [10] for more details about the profiles, the experiments, and further experimental findings.

2 Related work

2.1 Iterative voting

Iterative voting has became an established setting on the research agenda in computational social choice. For a survey, see [16] (note that a number of papers have appeared since then). Most work on iterative voting consists of theoretical analyses, such as convergence and efficiency guarantees. A noticeable exception is the analysis of behavioural online experiments by Meir et al. [19], who investigate the distribution of types among voters based on a number of myopic heuristics, and compare one-shot voting with access to an information poll to iterative voting where voters can also see each others' votes.

Iterative voting methods can be designed as voting rules *per se* [9, 20, 2], or can be used for modelling voter response to polls preceding an election. Under this interpretation, convergence issues have been investigated [see, *e.g.*, 17, 18, 15], the quality of the resulting equilibria analysed [7, 12], as well as the possibility of manipulation by the central authority [23, 3]. A further subdivision of this research stream is obtained by distinguishing models that consider voters' uncertainty from those that do not. Uncertainty comes either from an exogenous poll or by making the votes of the other voters visible in iterated elections [19].

Modelling voter responses to polls starts with Reijngoud and Endriss [21] who formulated a first family of heuristics. This is developed further by Fairstein et al. [8]: voters are given a cardinal utility function over a set of candidates, and have access to a poll giving the expected number of votes that each candidate will get.

2.2 Multiple referenda

A multiple referendum consists of taking a collective decision over each of a set of binary variables, called *issues*.[2] Nonseparability is a huge problem in multiple referenda: on the one hand, letting voters vote individually on each variable ignores their preferential dependencies between issues and may lead to bad collective decisions; on the other hand, letting voters express their preferences over all combinations of values is not feasible due to the size of the domain, as soon as the number of issues exceeds a few units. This dilemma has been identified for a long time [5, 6, 13]. Since then, a number of solutions has been suggested and studied, such as simultaneous voting, sequential voting, compact preference representation, or expressing preferences as distance to the top alternative. They are reviewed by Lang and Xia [14]. Another solution, suggested by Bowman et al. [4], stands out as structurally different than the ones cited above. It consists of an iterative voting protocol that allows voters to revise their votes based on the outcome of previous iterations. It allows to reach a solution of reasonably good quality, while keeping the communication and cognitive burden at a reasonable level. Focusing on three-issue multiple referenda, their computer simulations show that iterative voting help voters get out of multiple election paradoxes.

3 Context and setting

Consider a domain made of two (binary) issues, called Row and Col, with domains $D_{\text{Row}} = \{\text{Top}, \text{Bottom}\}$ and $D_{\text{Col}} = \{\text{Left}, \text{Right}\}$.[3] The set of alternatives, called *cells*, is $A = \{TL, TR, BL, BR\}$, where T, B, L, R stand for Top, Bottom, Left, Right.

A (strict) preference relation $\succ$ over A is a transitive and asymmetric relation. If it is moreover complete, it is said to be *linear*. We say that Col is *separable* for $\succ$ if ($TR \succ TL$ if and only if $BR \succ BL$) and ($TL \succ TR$ if and only if $BL \succ BR$).[4]; and similarly for Row. If both Col and Row are separable for $\succ$ then $\succ$ is said to be *separable*; if only one of Col and Row is separable for $\succ$ then $\succ$ is said to be *semi-separable*; if none of Col and Row is separable for $\succ$ then $\succ$ is said to be *fully nonseparable*. When we say that a preference is non-separable we mean that it is either semi-separable or fully non-separable.

The restriction to two binary variables makes it very easy to count the number of linear preference relations in each class:[5]

- *separable*: $TR \succ TL \succ BR \succ BL$, and the seven other linear preference relations obtained by exchanging the role of Col and Row and/or permuting one of the values, or both, of Col and Row.
- *semi-separable*: $TR \succ TL \succ BL \succ BR$, and the seven other linear preference relations obtained by applying a permutation as above.
- *fully non-separable*: $TR \succ BL \succ BR \succ TL$, and the seven other linear preference relations obtained by applying a permutation as above.

[2]Examples from the real world are ubiquitous, including the commonly cited decision-making processes in Switzerland and California. For the latter, for instance, see: `https://ballotpedia.org/California_2020_ballot_propositions`.

[3]Choosing to restrict the study to two variables is, admittedly, a loss of generality. We made this choice so as to limit as much as possible the burden of cognitive effort: three binary variables involve reasoning with eight values, which could be difficult for many subjects.

[4]If $\succ$ is a linear order then ($TR \succ TL$ if and only if $BR \succ BL$) implies ($TL \succ TR$ if and only if $BL \succ BR$).

[5]The restriction to two variables is crucial: counting the number of separable preferences on an arbitrary number of variables is an open problem [11].

With each of the three types of preference relations, there are a number of possible associated behaviours:

1. Assume variable V is separable for $\succ$, with x_V as preferred value. Then x_V is a dominating action in the sense that a voter voting for x_V can never be better off by changing her vote. We say that such a voter, in a given situation, is *V-rational* if she votes for x_V. For instance, the agent with the separable preference $TR \succ TL \succ BR \succ BL$ who votes for (T, L) is Row-rational but Col-irrational.

2. An agent with a separable preference is said to be *rational* if she is both Col-rational and Row-rational, *half-rational* if she is rational for exactly one variable, and *fully irrational* if she is irrational for both. (The voter of the previous item is half-rational.)

3. Assume an agent has a semi-separable preference (w.l.o.g.) $TR \succ TL \succ BL \succ BR$. The value R of the non-separable variable Col is said to have a *large span*: if the voter is pivotal for this variable then it will lead to change the outcome either in favour of her best or of her worst alternative. Similarly, L is said to have a *small span*. The voter is said to be *optimistic* if she votes for the large-span value and *pessimistic* otherwise.[6]

4. Assume an agent has the fully non-separable preference (w.l.o.g.) $TR \succ BL \succ BR \succ TL$. Value T has large span and B has small span, therefore we say that such a voter is optimistic (resp. pessimistic) if she votes for T (resp. B). For Col, this is different: the value R leads to either the best or the second worst alternative in case the voter is pivotal for Col, while L leads to either the second best or the second worst alternative. We say that R is *stochastically dominant*. The voter is said to be *SD-rational* if she votes for the stochastically dominant variable, and SD-irrational otherwise.

We provide voters with utility matrices, as depicted on Figure 1.

	L (rat.)	R (irrat.)
T (rat.)	3	1
B (irrat.)	2	0

(a) Separable preference

	L (rat.)	R (irrat.)
T (opt.)	3	0
B (pess.)	2	1

(b) Semi-separable preference

	L (SD-rat.)	R (SD-irrat.)
T (opt.)	0	3
B (pess.)	2	1

(c) Fully non-separable preference

Figure 1: Visual explanation of voter behaviour with three types of preferences.

In a multiple referendum, *simultaneous voting* consists of asking each voter to cast a vote regarding each issue. For the two-issue case, each of the n voters has to vote twice: one for the column issue (L or R) and one for the row issue (T or B). When a voter casts a vote regarding an issue, she does not know the other votes regarding this issue, nor the votes regarding the other issue. For each issue, the collective decision is made via majority rule; from now on we assume n to be odd so that ties cannot occur.

Iterative voting for multiple referenda is defined as a series of simultaneous voting stages; at each stage, voters observe the distribution of votes regarding each issue at the previous stage. Thus, at stage 1, voters do not observe anything; at stage 2, they know the number of votes for L, R, T and B for stage 1; at stage 3, they know the number of votes for L, R, T and B for stages 1 and 2, and so on. The stopping criterion is defined exogeneously.

[6]See Section 5.3 for a discussion on the terminology *optimistic* vs. *pessimistic*.

3.1 Hypotheses

We formulate three main hypotheses on iterative voting in multiple referenda. First, for each preference type we formulate a prediction on the voters' behavior:

Hypothesis 1A *Voters avoid dominated choices.*

Hypothesis 1B *Voters choose SD-rationalizable actions.*

Hypothesis 1C *Voters choose optimistic actions more often than pessimistic actions.*

Hypotheses 1A and 1B are quite natural and do not need explanations. Verifying hypothesis 1A on the first game (containing only separable preferences) is also a good test to see if voters understood the rules. Hypothesis 1C is bolder. It has already been made by [13] and this is one of the most important hypotheses within our list.

Second, we test whether separability and pivotality have an effect on the observed changes in voters' ballots in iteration:

Hypothesis 2A *Voters with non-separable preferences change their choices more often than voters with separable preferences.*

Hypothesis 2B *Voters change their votes more often if they were pivotal in the previous round.*

Hypothesis 2A does not need any explanation. The reason to expect Hypothesis 2B is that changing one's vote has a cognitive cost and voters who foresee they can change the outcome at the next round, because they are pivotal in the current round, have more to gain in changing their vote than those who are not, and to pay the associated cognitive cost.

The last two hypotheses are classical in the literature on iterated voting:

Hypothesis 3A *Iterative voting increases average social welfare.*

Hypothesis 3B *Iterative voting increases Condorcet efficiency.*

4 Experimental Design

In our experimental design each participant took the role of a voter in nine elections. The first four of these elections were implemented as one-shot voting, in which subjects are told that they are in a group of 7 and will make choices for columns and rows based on their individual payoff matrices, without any further information. The collective choice for both the column and the row are reached with the majority rule to determine the final outcome.

In the remaining five elections, which were implemented as iterated voting, groups consisted of either 3, 5, or 7 voters. After each stage within an election, subjects are provided with the vote tally for each column and row and the final outcome in all previous stages. The final outcome of an iterated election was reached if an outcome is repeated for three stages. In case this is not achieved, the outcome at 12th stage is assigned as the final outcome. Subjects and groups are randomly assigned to voter profiles at each election.

There are notable differences between our experimental design and other related behavioural experiments in voting. First, about the information available to voters: we assume that voters initially ignore everything from other voters' preferences, while in Van der Straeten et al. [22], the preference profile is common knowledge. Second, voters change their vote simultaneously in each step, while in Meir et al. [19] (and in much of the literature on iterative voting) only a single voter can change their vote in each iteration step.

Voters 1 & 2:

	L	R
T	2	3
B	0	1

Voters 3 & 4:

	L	R
T	1	0
B	3	2

Voters 5 & 6:

	L	R
T	0	2
B	1	3

Voter 7:

	L	R
T	3	1
B	2	0

Table 1: Profile 1.

Profiles 1 to 4 are voted one-shot. In Profile 1, all preferences are fully separable. It allows us to look into the tendency to vote rationally, avoiding dominated choices.

In Profiles 2 to 4, no preference is separable in any issue. In all cases, BL is the Condorcet winner. Voters are placed in a dilemma between optimism and pessimism for one issue, while the other issue pertains to SD-rationality. Also, comparing the individual behaviour within these profiles allows us to investigate the behavioral effects of different utility scales.

Voters 1 & 2:

	L	R
T	b	d
B	c	a

Voters 3 & 4:

	L	R
T	b	c
B	d	a

Voters 5 & 6:

	L	R
T	c	a
B	b	d

Voter 7:

	L	R
T	c	b
B	a	d

Table 2: Profiles 2,3, and 4, where for Profile 2 the utility vector is $(a, b, c, d) = (0, 2, 3, 4)$, for Profile 3 it is $(a, b, c, d) = (0, 1, 2, 3)$, and for Profile 4 it is $(a, b, c, d) = (0, 1, 2, 4)$

Profiles 5 to 9 are voted in iteration. In Profile 5, the preference of voter 3 (resp. of voters 1 and 2) is separable (resp. semi-separable). Voters need to get away from a bad outcome by coordinating: if voter 3 is rational, and if 1 and 2 are SD-rational, and either both optimistic or both pessimistic, then the outcome is TR or BL, i.e., respectively the worst and the second worst outcomes for both.

Voter 1:

	L	R
T	1	0
B	1	2

Voter 2:

	L	R
T	2	0
B	1	1

Voter 3:

	L	R
T	1	3
B	0	2

Table 3: Profile 5.

In Profile 6, no voter has a fully separable preference. Voters 1 and 2 have a separable preference in Row, whereas voters 3, 4, and 7 have separable preferences in Column; TL (Condorcet winner) is a good compromise, but it may not be reached at the first stage. The profile is not 'pathological' so one would expect a rather rapid convergence to TL.

Voters 1 & 2:

	L	R
T	3	4
B	2	0

Voters 3 & 4:

	L	R
T	3	2
B	4	0

Voters 5 & 6:

	L	R
T	2	1
B	0	3

Voter 7:

	L	R
T	2	1
B	3	0

Table 4: Profile 6.

In Profile 7, only voter 7 has a (fully) separable preference; TR is the Condorcet winner. The outcome at the first stage is expected to be very bad: assuming that a large enough part of voters 1-6 vote SD-rationally, the outcome would occur to be BR, which gives utility 1 to all: for these six voters, BR is Pareto-dominated by BL and TR but they must in some way coordinate to elect one of them, which may take several steps.[7]

[7]Another interest of Profile 7 is that if the current outcome is BR, and if the margins of victory for both

Voters 1–3:

	L	R
T	0	3
B	4	1

Voters 4–6:

	L	R
T	0	4
B	3	1

Voter 7:

	L	R
T	0	2
B	1	3

Table 5: Profile 7.

In Profile 8, voters 1, 2, and 4 have fully separable preferences, whereas voter 5 has a separable preference only in the column issue; TL is the Condorcet winner. Profile 8 is different enough from the previous profiles because three voters out of five have separable preferences and are therefore expected not to change their votes, and voter 5 has a dominant action (L); therefore the outcome would be expected to be BL or TL.

Voter 1:

	L	R
T	3	2
B	1	0

Voter 2:

	L	R
T	2	3
B	0	1

Voter 3:

	L	R
T	2	0
B	1	3

Voter 4:

	L	R
T	1	0
B	3	2

Voter 5:

	L	R
T	2	1
B	3	0

Table 6: Profile 8.

In Profile 9, only voter 2 has a fully separable preference. Voters 4 and 5 have separable preferences only in the column issue. There is no Condorcet winner.

Voter 1:

	L	R
T	1	3
B	2	0

Voter 2:

	L	R
T	2	0
B	3	1

Voter 3:

	L	R
T	2	1
B	0	3

Voter 4:

	L	R
T	0	3
B	1	2

Voter 5:

	L	R
T	3	0
B	2	1

Table 7: Profile 9.

Further details of our experimental procedure are described in our full paper [10], specifically in Section 4.2 and Appendix 2.

5 Rationalizability and optimistic behavior

5.1 Dominated choices

Table 8 gives the number of instances (voter and profile) in which subjects made a dominated choice on both issues (for separable preferences), a dominated choice on one issue with separable preferences, and a dominated choice on one issue with semi-separable preferences.

	Profile 1	Profiles $5-9$
Dominated choice in both issues	11(7.5%)	8 (4%)
Dominated choice in one issue (separable pref.)	26(17.7%)	36 (12.8%)
Dominated choice in one issue (semi-sep. pref.)	–	47 (16.4%)

Table 8: Number of dominated choices in one-shot elections and first stages of iterative elections. Percentages are calculated on the total number of potential dominated choices.

B and R are large, then voters 1-6 would feel that their probability of being pivotal at the next stage is low for each of both issues, and may be tempted to vote TL to try to be pivotal in either one them (the probability of being pivotal on both being negligible); this behaviour has been analyzed theoretically in [1].

Our findings show that violation of rationality is not prevalent in one-shot voting, supporting Hypothesis 1A. More details are in our full paper, Section 5.1.

5.2 SD-rationalizable choices

Recall that a voter with fully non-separable preferences necessarily has an issue for which one choice is SD-rationalizable. Here, we focus on the frequency of SD-rationalizable choices by subjects. As Figure 2 shows ("All" columns), both in the one-shot elections and in the first stages of iterated elections, most of the subjects with fully non-separable preferences made SD-rationalizable choices, in support of Hypothesis 1B.

To refine our analysis, we computed the same percentages among two populations of "irrational" voters: those who made at least once a dominated choice on two issues and those who made it once on one issue only. Our findings that are summarised in Figure 2 show that voters who make dominated choices are more likely to make SD-dominated choices (or less likely to make SD-rationalizable choices).

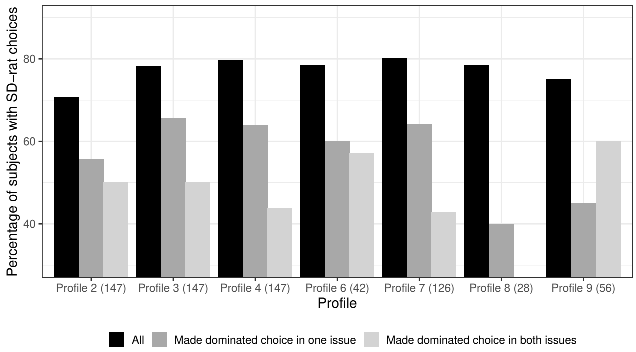

Figure 2: Percentage of subjects who made SD-rationalizable choices in one-shot elections and first stages of iterated elections including voters with fully non-separable preferences.

5.3 Optimistic/pessimistic choices

Recall that when a voter has a semi-separable or fully non-separable preference, she has necessarily an issue for which the two choices can be characterized exclusively as optimistic and pessimistic. Here we focus on the frequency of optimistic choices in these cases. Figure 3 shows the percentage of optimistic choices in one-shot elections and in the first stages of iterative elections. In both cases, we can conclude that a majority of subjects (between 63% and 82%) make optimistic choices, in support of Hypothesis 1C. This finding confirms our beliefs based on observations over many years.

We used the adjectives "optimistic" and "pessimistic" because we had to give a name to these two observed behaviours, but one should be cautious about their interpretation. For instance, voters may be 'optimistic' simply because of simplicity: they identify their preferred cell out of the four, and project it on each dimension. In a sense, optimistic voters act as if they had forgotten that two interacting votes are taking place and thought instead that voting for, say, T and L, simply means voting for the TL cell. We do not say that all optimistic voters reason like this; but the fact that *some* do can explain why there are more optimists than pessimists. See Section 5.3 of our full paper for a detailed discussion.

We also tested whether the utility scale used in the matrix has an impact on Hypothesis 1C. Profiles 2, 3, and 4 are played one-shot and feature voters who have the same ordering

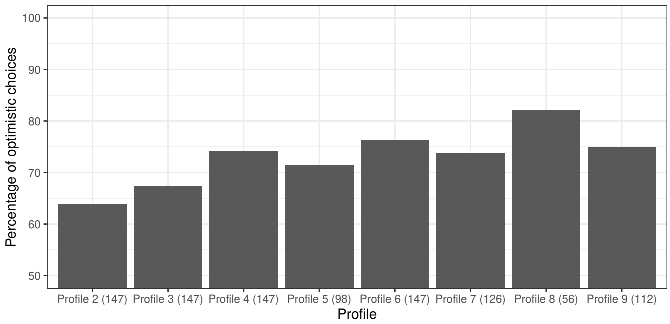

Figure 3: Percentage of optimistic choices by profile (first stage for the iterated elections of Profiles 5 to 9). The number of subjects are in parentheses.

Profile 3	Profile 6	Profile 8	Profile 9
67.3%	78.6%	78.6%	69.6%
(147)	(42)	(28)	(56)

Table 9: Comparison of the percentage of optimistic choices for voters with utility scale 3210 in one-shot elections (Profile 3) and first stages of iterative elections (Profiles 6, 8, and 9). The numbers of subjects are printed in parentheses.

over the outcomes but different utility scales. In Profile 2, the utility vector is $(0, 2, 3, 4)$, in Profile 3, it is $(0, 1, 2, 3)$, whereas in Profile 4, it is $(0, 1, 2, 4)$. One can argue that subjects would avoid the worst alternative more often in Profile 2 than in Profile 3 (hence tend to be more pessimistic), since this brings about a larger relative loss. Similarly, subjects would be attracted more by the best alternative in Profile 4 since it leads to greater relative gain (hence tend to be more optimistic, also as the small span action pays less). We observe that choices are significantly more optimistic going from Profile 2 to 3 and/or to Profile 4, supporting the intuition that more 'convex' utilities lead to more optimistic choices. This finding also gives evidence that a significant fraction of the voters are indeed "optimistic" rather than using a simple best-focused heuristic.

In our full paper (last two paragraphs of Section 5.3) we give results on whether the percentage of observed optimistic choices is correlated to the presence of iteration (the answer is negative), and on how obtaining the best outcomes relates to voting optimistically (there is a clear, but unsurprising, correlation).

6 Individual Voting Dynamics

We now focus on iterative elections and investigate how subjects' choices evolve in relation to the separability of their preferences and the pivotality of a voter in the previous round.

We first compare the number of subjects who change their votes at least once and those who never change, based on whether their preferences are separable or not. Results are presented in Table 10. According to Fisher's exact test, the difference is statistically significant hence we indeed observe a higher tendency to change when preferences are non-separable, in the line of Hypothesis 2A. Table 11 presents vote changing behavior when preferences are semi-separable. We see that subjects are more likely to change their votes for the non-separable issues, also in support of Hypothesis 2A. We give more details in our full paper, Section 6, on detecting the effect of separability on the vote changing behavior.

We now look at how subjects respond to the outcome of the previous round. We say

Profile	Separable			Non-separable		
	Voters	No Change	Any Change	Voters	No Change	Any Change
5	3	25	24	-	-	-
6	-	-	-	5,6	16	26
7	6	15	6	1-6	54	72
8	1,2,4	44	40	3	13	15
9	2	12	16	1,3	21	35
All		96	86		104	148

Table 10: Number of subjects with separable and non-separable preferences who change their votes at least once and who do not change at all.

Profile	Voters	Separable issue		Non-separable issue	
		No Change	Any Change	No Change	Any Change
5	1-2	53	45	52	46
6	1-4,7	72	33	57	48
8	5	22	6	19	9
9	4,5	42	14	28	28
All		189	98	156	131

Table 11: Number of subjects with semi-separable preferences who change and who do not change their votes across profiles.

that a voter is pivotal when the outcome of the previous round has a margin of victory of one on any of the two issues. Thus, provided that all other voters keep the same vote, the voter can switch the outcome by unilateral deviation. Further, we call improvers those pivotal voters who can change the outcome in their favour, and unsatisfied those pivotal voters who can change the outcome for something they prefer less than the current outcome. We first tested whether the possibility of improving the outcome when pivotal implied more observed vote changes, and we answer in the positive. Out of 432 cases in which a voter was a pivotal improver we observed 200 vote changes. Out of 520 cases in which a voter was a pivotal unsatisfied we only observed 187 vote changes. The difference is statistically significant. Moreover, out of 760 cases in which a voter is not pivotal and did not get the best outcome we observed 267 vote changes. This ratio is significantly less than the vote change ratio of improvers shown above. In conclusion, we find support for Hypothesis 2B: subjects pay attention to pivotality when they can change the outcome in their favour.

7 Collective dynamics and quality of the final outcome

Now we focus on the iterated elections (Profiles 5 to 9). Recall that the termination rule puts an end to the iteration when the outcome does not change for three rounds. Many groups concluded their iterative voting process at the end of the third stage, *i.e.*, the earliest possible. Only 53, 29, 43, 36, and 46 % of groups remained in iteration beyond the third stage in profiles 5 to 9 respectively.[8] See Figure 5 in our full paper for details.

We measure the social welfare as the average utility of the election outcome within groups. Figure 4 shows the average social welfare in the first and last stages of profiles 5-8. The minimum and maximum attainable social welfare values are depicted by triangles. The confidence intervals are for 90% level. We observe that iterative voting increases average social welfare in all profiles, which is in line with the simulations in [4] and our Hypothesis 3A.

[8]Still, there are a significant number of subjects that change their votes, without however changing the result of the election.

In Profiles 5 to 7, voters in the (static) groups that terminate at the earliest possible stage (*i.e.*, the third stage) started with an election outcome that has a higher social welfare than the (dynamic) groups who go beyond the third stage. The latter eventually improve it to a comparable level in the last stage of iteration. Profile 9 has the same social welfare in the first and last stages, hence excluded.

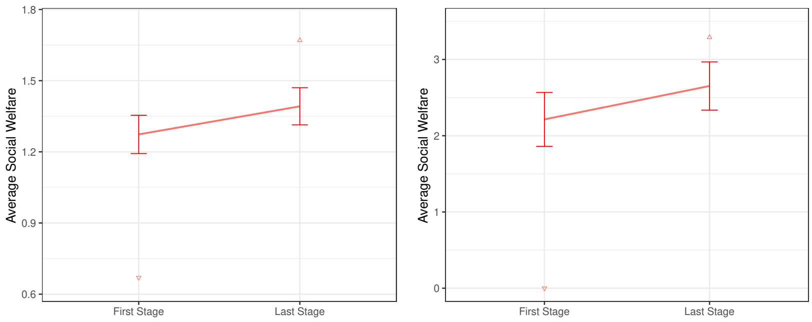

Figure 4: Average social welfare in first and last stage for profiles 5 and 7. Please refer to the full paper for the remaining figures.

Observe that our profiles were chosen so as to be pathological enough for the improvement along iterative voting to be most visible. If we had tested profiles with "no-dilemma" games, where players have no effect on the utility of their opponent, it is most likely that a socially satisfactory outcome would have been obtained without need for iteration, and that the game would have stopped after the minimal number of iterations. We did not test such profiles because they are not very interesting, so that we would have wasted subject time (and our money). Testing randomly generated profiles would have even be more time and money consuming. So our conclusion about this hypothesis should be tempered: *on profiles that are pathological enough so that if there is hope that iteration improves the quality of the outcome, it actually does.* This is in line with the main argument by [4]: we confirm by behavioural studies the result they obtained by simulations. Interestingly, these results contrast with the findings of [19] in online iterated elections, who do not find any statistically significant improvement in social welfare through iteration. This suggests that multi-issue elections are more likely to profit from iterative voting than classical plurality elections.

We then investigated the impact of iterative voting on Condorcet efficiency, *i.e.*, the probability of electing a Condorcet winner (CW) when there is one. In all our profiles but the last one (9) there is a CW. Figure 5 shows the frequency of Condorcet winners being chosen in one-shot votes (Profiles 1–4) and in the first stage of iterated elections (Profiles 5–8). The observed level of Condorcet efficiency is quite low.[9] As for the effects of iteration on Condorcet efficiency, there is no clear indication of increase or decrease. In all profiles played in iteration, the CW is also the social welfare optimum; this is not a guarantee of this outcome to be elected: in Profile 8, TL is elected in less than 50% of the cases. Given our findings, the observed increase in social welfare is due more to voters avoiding bad outcomes than improving good ones. This is corroborated by our results on pivotality showing that voters tend to change their vote only if they get a particularly bad outcome.

[9]No definitive conclusion can be drawn from these numbers, as most of our profiles were chosen so as to be at least somewhat pathological in the first place and not fit for this particular question.

Profile	Frequency of chosen CWs
1	24 %
2	19 %
3	24 %
4	5 %

Profile	Frequency of chosen CWs First iter.	Last iter.
5	35 %	49 %
6	24 %	24 %
7	52 %	29 %
8	29 %	18 %

Figure 5: Frequency of chosen Condorcet winners by profile in one-shot votes and in iterated profiles in the first and last iteration.

8 Conclusions

We presented results of lab experiments on iterative voting in multiple referenda. Voters faced two binary referenda on the choice of the column and the row of a matrix, with each voter attaching a different utility to each cell. Our findings can be summarised as follows. In presence of separability, subjects tend to avoid dominated choices, while in presence of non-separability, subjects tend to play a stochastically dominant action when they have one, and otherwise tend to vote optimistically more often than pessimistically; the frequency of optimism increases with the convexity of payoffs (but does not increase with time in iterated voting). We did not observe any significant difference in voters' behaviour in one-shot elections and the first stages of iterated elections. In iterated elections subjects with non-separable preferences tend to change their choices more often than voters with separable preferences, as do voters that are pivotal in the previous round, provided they can change the result to an outcome they prefer. Finally, iterative voting increases social welfare, which confirms the findings by [4], obtained from simulations. On the other hand, our experiments do not allow us to draw any conclusion about whether iterative voting increases the Condorcet-efficiency.

An important caveat of our results is the restriction to two (binary) issues. This choice was clearly motivated by feasibility: implementing an experiment with three issues or more makes it difficult to write the payoff matrices in a user-friendly way, the experiment would have been more cognitively complex and the risk would then have been that a significant number of voters do not understand it well or that they find it too demanding and end up playing randomly. As a consequence, it makes it difficult to compare directly with the simulation results in [4]. This leads us to the following important question: which of our conclusions would still hold for more than two issues? For some of them it can be reasonable to assume that they would continue to hold (even if we did not make further experiments): hypotheses 1, 2, 6 and 8; this is also the case for hypotheses 3-5 but it is more difficult to define optimism and pessimism because these notions become gradual as soon as we have more than two variables. It is more difficult to make conjectures for hypotheses 7 and 9, although we do not see a clear reason why they would not hold for more than two issues.

In any case, even if our results held for only two issues, they would still be meaningful: in most practical cases, voters' preferential dependencies between issues in multiple referenda can be restricted to sets of issues of size two (although situations where more than two issues interact do occur, especially for funding public projects with budget constraints).

The main practical impact of our study is a confirmation that iterative voting is a viable way of handling nonseparability in multiple referenda, confirming the simulation-based results of Bowman et al. [4]. To push further the inclusion of deliberative tools in the iterative voting protocol, in future work we will allow "cheap talk" among voters in between the iterations on each vote, testing whether this leads to a decrease in the number of iterations or if it modifies their individual voting behaviour.

References

[1] David S. Ahn and Santiago Oliveros. Combinatorial Voting. *Econometrica*, 80(1): 89–141, 2012.

[2] Stéphane Airiau, Umberto Grandi, and Filipo Studzinski Perotto. Learning agents for iterative voting. In *Proceedings of the 5th International Conference on Algorithmic Decision Theory (ADT)*, 2017.

[3] Dorothea Baumeister, Ann-Kathrin Selker, and Anaëlle Wilczynski. Manipulation of opinion polls to influence iterative elections. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2020.

[4] Clark Bowman, Jonathan K. Hodge, and Ada Yu. The potential of iterative voting to solve the separability problem in referendum elections. *Theory and Decision*, 77(1): 111–124, 2014.

[5] Steven J Brams, D Marc Kilgour, and William S Zwicker. Voting on referenda: the separability problem and possible solutions. *Electoral Studies*, 16(3):359 – 377, 1997.

[6] Steven J. Brams, D. Marc Kilgour, and William S Zwicker. The paradox of multiple elections. *Social Choice and Welfare*, 15(2):211–236, 1998.

[7] Simina Brânzei, Ioannis Caragiannis, Jamie Morgenstern, and Ariel D Procaccia. How bad is selfish voting? In *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI)*, 2013.

[8] Roy Fairstein, Adam Lauz, Reshef Meir, and Kobi Gal. Modeling people's voting behavior with poll information. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, 2019.

[9] Umberto Grandi, Andrea Loreggia, Francesca Rossi, Kristen Brent Venable, and Toby Walsh. Restricted manipulation in iterative voting: Condorcet efficiency and Borda score. In *Proceeding of the 3rd International Conference on Algorithmic Decision Theory (ADT)*, 2013.

[10] Umberto Grandi, Jérôme Lang, Ali I. Ozkes, and Stéphane Airiau. Voting behavior in one-shot and iterative multiple referenda. *Social Choice and Welfare*, 2022. doi: https://doi.org/10.1007/s00355-022-01436-0.

[11] Jonathan K. Hodge. Permutations of separable preference orders. *Discrete Applied Mathematics*, 154(10):1478–1499, 2006.

[12] Joshua Kavner and Lirong Xia. Strategic behavior is bliss: Iterative voting improves social welfare. In *Advances in Neural Information Processing Systems 34 (NeurIPS)*, 2021.

[13] Dean Lacy and Emerson M. S. Niou. A problem with referendums. *Journal of Theoretical Politics*, 12(1):5–31, 2000.

[14] Jérôme Lang and Lirong Xia. Voting in combinatorial domains. In *Handbook of Computational Social Choice*, pages 197–222. Cambridge University Press, 2016.

[15] Omer Lev and Jeffrey S. Rosenschein. Convergence of iterative scoring rules. *Journal of Artificial Intelligence Research*, 57:573–591, 2016.

[16] Reshef Meir. Iterative voting. In Ulle Endriss, editor, *Trends in Computational Social Choice*, pages 69–86. AI Access, 2017.

[17] Reshef Meir, Maria Polukarov, Jeffrey S. Rosenschein, and Nicholas R. Jennings. Convergence to equilibria in plurality voting. In *Proceedings of the Twenty-fourth conference on Artificial Intelligence (AAAI)*, 2010.

[18] Reshef Meir, Maria Polukarov, Jeffrey S. Rosenschein, and Nicholas R. Jennings. Iterative voting and acyclic games. *Artificial Intelligence*, 252:100–122, 2017.

[19] Reshef Meir, Kobi Gal, and Maor Tal. Strategic voting in the lab: compromise and leader bias behavior. *Autonomous Agents and Multi Agent Systems*, 34(1):31, 2020.

[20] Svetlana Obraztsova, Evangelos Markakis, Maria Polukarov, Zinovi Rabinovich, and Nicholas R. Jennings. On the convergence of iterative voting: How restrictive should restricted dynamics be? In *Proceedings of the 19th AAAI Conference on Artificial Intelligence (AAAI)*, 2015.

[21] Annemieke Reijngoud and Ulle Endriss. Voter response to iterated poll information. In *Proceedings of the 11th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, June 2012.

[22] Karine Van der Straeten, Jean-François Laslier, Nicolas Sauger, and André Blais. Strategic, sincere, and heuristic voting under four election rules: an experimental study. *Soc. Choice Welf.*, 35(3):435–472, 2010.

[23] Anaëlle Wilczynski. Poll-confident voters in iterative voting. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, 2019.

Umberto Grandi
IRIT, Université Toulouse Capitole
Toulouse, France
Email: `umberto.grandi@irit.fr`

Ali I. Ozkes
SKEMA Business School, Université Côte d'Azur (GREDEG)
Université Paris Dauphine (LAMSADE) Paris, France
Email: `ali.ozkes@skema.edu`

Stéphane Airiau, Jérôme Lang
LAMSADE, PSL, CNRS
Paris, France
Email: `stephane.airiau@dauphine.fr,jerome.lang@lamsade.dauphine.fr`