Trends in Computational Social Choice

20

Cite as: Sylvain Bouveret. Social Choice on the Web. In Ulle Endriss (editor), *Trends in Computational Social Choice*, chapter 20, pages 387–402. AI Access, 2017.

http://www.illc.uva.nl/COST-IC1205/Book/

CHAPTER 20

Social Choice on the Web

Sylvain Bouveret

In this chapter, we will speak about the development and use of a web application dedicated to social choice, Whale.[1] The chapter is intended to be a compilation of lessons learned from the users of this web application, rather an extensive presentation of it. We will see through specific use cases what concrete problems have been encountered during the development of this web application.

20.1 Introduction

The history of social choice theory is paved with paradoxes. Of course, theoretical peculiarities like the Condorcet Paradox or Arrow's Theorem form the basis of the discipline. But social choice even seems to be inherently paradoxical at the epistemological level.

First example: Social Choice Theory has been being studied in its modern form for two centuries — if we date back the birth of modern social choice to the debates between Condorcet and Borda in the late 18$^{\text{th}}$ century. During the history of this discipline, a lot of theoretical knowledge about collective decision making and voting procedures has been accumulated. Yet, until recently, the occasions to really put this knowledge into practice in everyday life were rather infrequent and limited to political elections or particular situations (like scientific committees). As a consequence of this distortion between theory and practice of social choice, lessons and knowledge learned from the theoretical study of collective decision making did not seem to have much influence on the way people practice elections.[2] This can be well explained. In the rare occasions where people really practice voting, the stakes and the exogenous influences (political, cultural, ...) are usually so high that nobody would ever think of questioning the voting process itself.

Second example: although people only seldomly experiment with voting, practical situations of making collective decisions happen every day. A group of friends deciding which movie to watch, colleagues choosing a date for a meeting, a recruitment committee hiring a candidate for a job, ... All these situations that

[1] WHich ALternative is Elected. Current version accessible at: `https://whale.imag.fr/`

[2] The converse can also be observed to some extent, where the models at the heart of social choice are often based on very strong assumptions, the practical relevance of which can be questioned, or at least need to be confronted with experiment.

usually do not require a formal vote are nevertheless collective decision making situations. These everyday life situations would be perfect candidates for applying the formal knowledge brought by voting theorists. Usually the stakes remain quite low, but are still high enough to require an enlightened decision that voting theory can provide.

Bringing voting theory to these kinds of everyday situations was still hardly conceivable until recently. One major reason is that the logistical burden required to implement the least voting procedure (expressing a linear order in practice, tallying the votes, ...) is often too high to be reasonable. Another major reason is that most people are simply not aware of the fact that it is possible to vote other than by simply writing one name on the ballot sheet and choosing as the winner the candidate voted for most often.[3]

Two recent major scientific and technological breakthroughs that have nothing to do with social choice have dramatically changed things. The first one is the advent of computer networks and the World Wide Web. This has a major impact on the democratisation of social choice for several reasons. First, it makes the implementation of light remote voting systems possible. Human computer interfaces make possible the use of ways of expressing preferences that would have been unpractical otherwise, and enable the access to automatic tallying procedures. Second, it multiplies collective decision-making-like situations: social networks, recommender systems, and so on. Finally, computer networks have considerably changed the way knowledge disseminates. This also applies to voting: new citizens' initiatives to promote alternative voting procedures appear every day. They spread much faster and further than they used to, and as a consequence they disseminate in some way ideas from social choice theory.

The second major breakthrough is the advent of mobile devices. Ubiquity has a major impact on social choice. It now becomes possible to use computer-aided collective decision making in virtually every situation. More than that, it has brought connectivity and technology to a substantial part of the population that was before excluded from the technological sphere.

Social choice is facing a unique situation in its history. It is no longer limited to the political and scientific communities, but it is now ready to be widely applied. The goal of this chapter is to analyse the major difficulties that social choice theorists will probably have to face if they want to apply social choice in practice. The considerations proposed in this chapter are inspired by the active development of a web application dedicated to social choice. It is important to note that this is by no means intended to be the description of scientific results obtained by a rigorous experimental approach. It can be seen as nothing more than a compilation of lessons learned from the user experience with this platform. Moreover, this chapter voluntarily excludes the security issues posed by electronic voting. We make the assumption that the collective decision making situations analysed here are uncritical situations (further referred to as low- and middle-stake situations) where the voters can tolerate some lack of guarantees in terms of security, authentication and certification.

This chapter is organised as follows. First, we will quickly explain the main

---

[3]Except for the few countries — like Ireland — where people vote by ranking (or approving / disapproving) candidates for political elections.

features and objectives of the Whale web application in Section 20.2. Then, in Section 20.3, we will describe a particular use case and analyse the users' behaviour and impressions. Finally, in Section 20.4, we try to give more general lessons learned from the everyday use of the platform.

20.2 Whale: A Web Application for Social Choice

20.2.1 Goals and History

The genesis of Whale dates back to 2010 when we started to develop a small and lightweight web application dedicated to voting. At the beginning, this application was just intended to be an alternative to well-known and widely used poll applications like Doodle® or Framadate. The initial observation was that such applications were very efficient tools to help solving collective decision making problems that happen in our everyday life (like choosing a restaurant, a candidate for a position, a date for a meeting, ...), but were sometimes too simple to adequately represent the complexity of the individuals' preferences. For instance, most of these tools are limited to the expression of binary (Yes / No) preferences, but for some applications this is not enough. As an example, in situations where a group of people has to hire someone for a job, it is reasonable to assume that the participants have more in mind than a simple dichotomy between approved and disapproved candidates. When collectively choosing a date for a meeting, we are often faced with the situation where some date is not completely unavailable to us, but is not ideal either, whatever it may mean. In such situations, it could be useful to provide the participants with a way of expressing more subtle preferences than just approving / disapproving alternatives.[4]

Even if there is a clear method to choose the winning option when we deal with approval (Yes / No) ballots, namely the one which consists in choosing the option that is approved by a maximal number of participants, this is not so clear anymore once the participants can choose among strictly more than two possible evaluations for each option.[5] In this case, there exist several collective decision making procedures. All of them satisfy their own properties, but none is always better than the others, so the procedure has to be carefully chosen, depending on the context of the decision making problem.

Hence, as we can see, a lot of people use poll applications to make collective decisions nowadays, but these applications are sometimes too simple to adequately represent preferences and reflect the complexity of collective decision making. On the other hand, social choice theory provides all the mathematical background and models to make well-informed collective decisions. The main rationale of Whale is to try to bridge the gap between these two worlds.

As of January 2017, four versions of Whale have been released. The main reasons for this evolution are mainly technical: technologies tend to evolve more quickly in the context of web applications than in any other context. Developing

[4] The best evidence of this affirmation is that popular poll applications now propose to the participants an intermediate level "(Yes)" between Yes and No.

[5] Or as soon as we are in the multiwinner case, which we exclude from the scope of this chapter.

an upgradable, maintainable, reusable and durable code is a real challenge and it is often easier to restart everything from scratch than to make the code evolve.

The first version of Whale was issued in 2010. At this time, it was just a very simple web application, developed in PHP with a Postgres backend for the persistent data storage. The second version was developed by a team of students in 2012, using Java and the web framework Play (`https://playframework.com/`). Both versions had the same functionalities: registered users could create polls (date or classical) with several possible ballot types (approval, scores, rankings with or without ties), users could vote, and the results given by several voting rules where displayed using a simple table (as in Figure 20.1).

	April 2012					
	Monday, 2nd		Tuesday, 3rd		Wednesday, 4th	
	matin	après-midi	matin	après-midi	matin	après-midi
Hide the votes						
Sylvain B.	5	5	4	0	3	2
Herbert	0	0	4	5	4	2
Marvin	4	3	5	5	2	2
Alcybiade	5	5	3	3	5	5
Grogro	2	2	0	0	4	5
Borda	16	15	16	13	18	16
Majority	2	2	1	2	1	2
Veto	4	4	4	3	5	5
Bucklin	2	2	2	4	1	3
Leximin SWO	5	3	3	0	3	2
Leximax SWO	3	0	5	0	5	3
Nash SWO	540	432	600	144	1800	972
Condorcet Winner	(no winner)					
Copeland	3	1	-1	-4	2	-1
Maximin (Simpson)	1	0	1	1	2	1
Display Condorcet matrix						

Figure 20.1: Results of different voting procedures in the first version of Whale.

The third version of Whale was released in 2013, with the clear idea to focus on a better presentation of the results given by different voting rules. For that, we used data visualisation techniques, as we will see later. This version was developed in Scala / Java Servlet Pages and Postgres / JDBC for the backend, and the data visualisation was based on the D3.js library (`https://d3js.org/`). The last version of Whale (Whale 4) is in testing phase since September 2016. For this new version, the main focus was made on the simplicity of use and the responsiveness of the application (adaptation to mobile devices). It was developed in Python with the Django framework (`https://www.djangoproject.com/`), is based on the Bootstrap framework (`http://getbootstrap.com/`) for the responsiveness, and still on D3.js for data visualisation.

20.2.2 Features

Since the first version, the objectives of Whale have slightly evolved and this application now has three main goals:

- to be used by non-specialists as a poll and voting platform;
- to be used by researchers and teachers as a pedagogical resource to illustrate the concepts of voting;
- to be used by researchers as a source of real voting data for experiments.

Whale currently supports two kinds of decision problems, corresponding to two different situations:

- open-ballot polls, where anyone can participate, and everyone can see the preferences of the participants;
- sealed-ballot elections, which are accessible only at the invitation of the poll's creator. Each participant receives an individual 16-characters code that must be used to connect to the poll. A voter can modify her vote any number of times until the poll closes. When the poll closes, the ballots are publicly displayed, just hiding the names of the corresponding voters to guarantee anonymity.

Poll creators can currently choose between five ballot types: approval (Yes / No), scores (from 0 to 10), qualitative scale (--, -, 0, +, ++), rankings with ties, rankings without ties. Depending on the ballot type chosen, several methods are proposed to compute the collective preference and the election winner. Contrary to other voting platforms, the approach chosen was not to force the choice of a particular voting rule, but rather to provide decision makers with all the tools to understand the voters' preferences, discuss them, and make enlightened decisions. Hence, the users can see at any time the results suggested by several relevant voting procedures. In the first two versions of Whale, these results were given as tables of scores (see Figure 20.1). However, we observed that these tables were difficult to interpret, especially when the winners differed from one voting rule to another. In the subsequent versions of Whale, we have put a lot of effort into developing data visualisation modules to present these results in an easily understandable way. Four kinds of modules (illustrated in Figures 20.2, 20.3, 20.4 and 20.5 on the best poster election use case presented in Section 20.3) dedicated to four kinds of voting rules have been developed so far.

- Histogram-based visualisations for scoring rules (see Figure 20.2, left), and a specific view dedicated to approval voting (Figure 20.2, right).
- Representations of the majority graph for Condorcet-consistent methods. Here, we use two different representations of the majority graph: a node-link representation where the thickness of a link encodes the margin information (see Figure 20.3 left), and an incidence matrix representation, with lines and columns ordered according to a particular score chosen among Copeland 0, Copeland 1 and maximin (see Figure 20.3 right). We also use colour to

encode the information about the margin (for a case of the matrix) or the score of a candidate (for a node of the graph). This work on the majority graph has been extended to graph-compressed representations that seem to be quite efficient visualisation ways (Karanikolas et al., 2016), but this work has not been implemented in Whale yet.

- Run-off methods like plurality with run-off or Single Transferable Vote, where we simply display the list of candidates present in each round with the score they obtain in this round (see Figure 20.4).
- Randomized cups, displayed exactly like the board of a sport cup having candidates as participants (see Figure 20.5).

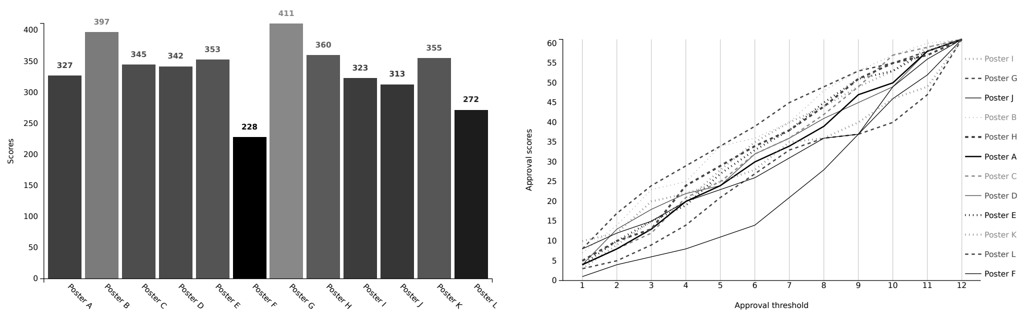

Figure 20.2: Graphical view for scoring functions in Whale 4 (for the best poster election). The left picture shows the Borda scores of the candidates. In the right picture, each curve corresponds to a candidate c and is a graphical representation of the function mapping a rank k to the k-approval score of c.

Beyond data visualisation techniques, the raw data of each poll can be accessed through an HTTP API, and the raw (and anonymous) data of all the polls of the database can be downloaded on the website. Currently three formats are supported: CSV, JSON and Preflib (see Mattei and Walsh, 2013, and Chapter 15 of this book). This data is provided for two main reasons. First, we hope that researchers will find valuable real-world data for their experiments. Second, providing an access to the data independently of any visualisation or aggregation technique gives the opportunity to the developers to extend the visualisation modules of Whale by developing their own representations of the voting profiles.

Other Existing Voting Platforms. Whale is certainly not the only online voting platform project. If we exclude popular aforementioned poll applications, we can cite two interesting and successful projects from academia: Pnyx (`https://pnyx.dss.in.tum.de/`) and Robovote (`http://robovote.org/`). Both also aim at bringing social choice theory to ordinary people by providing user-friendly collective decision making interfaces. Pnyx (Brandt et al., 2015) proposes several ways of expressing preferences (first-past-the-post, approval, ranking with ties, rankings without ties, pairwise comparisons), and several outputs are possible: unique winner, lottery, ranking without ties. Depending on the input and output

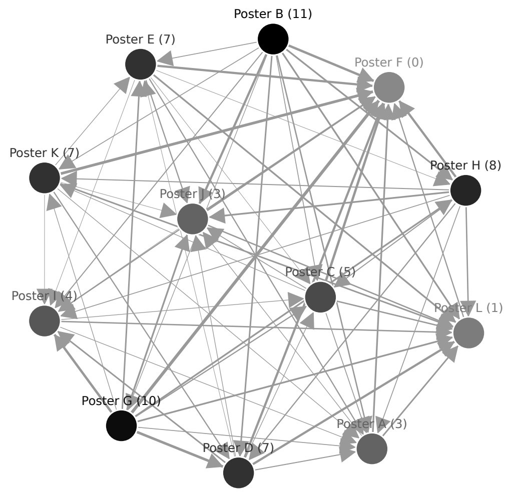

	Poster B	Poster G	Poster H	Poster D	Poster E	Poster K	Poster C	Poster I	Poster A	Poster J	Poster L	Poster F	Scores
Poster B	0	32	38	37	33	33	33	34	36	41	38	42	11
Poster G	29	0	38	43	37	32	35	42	33	38	38	46	10
Poster H	23	23	0	35	30	34	34	33	33	38	36	41	8
Poster D	24	18	26	0	31	26	32	37	34	32	41	41	7
Poster E	28	24	31	30	0	29	31	31	35	35	38	41	7
Poster K	28	29	27	35	32	0	29	31	32	31	37	44	7
Poster C	28	26	27	29	30	32	0	29	33	33	36	42	5
Poster I	27	19	28	24	30	30	32	0	32	30	36	35	4
Poster A	25	28	28	27	26	29	28	29	0	32	37	38	3
Poster J	20	23	23	29	26	30	28	31	29	0	36	38	3
Poster L	23	23	25	20	23	24	25	25	24	25	0	35	1
Poster F	19	15	20	20	20	17	19	26	23	23	26	0	0

Figure 20.3: The majority graph and its adjacency matrix for the best poster election. In the matrix, the rows and the columns are ordered by decreasing Copeland score.

Figure 20.4: Graphical view for run-off methods in Whale 4 (for the best poster election). The picture shows the candidate ordering in the different rounds of the Single Transferable Vote method.

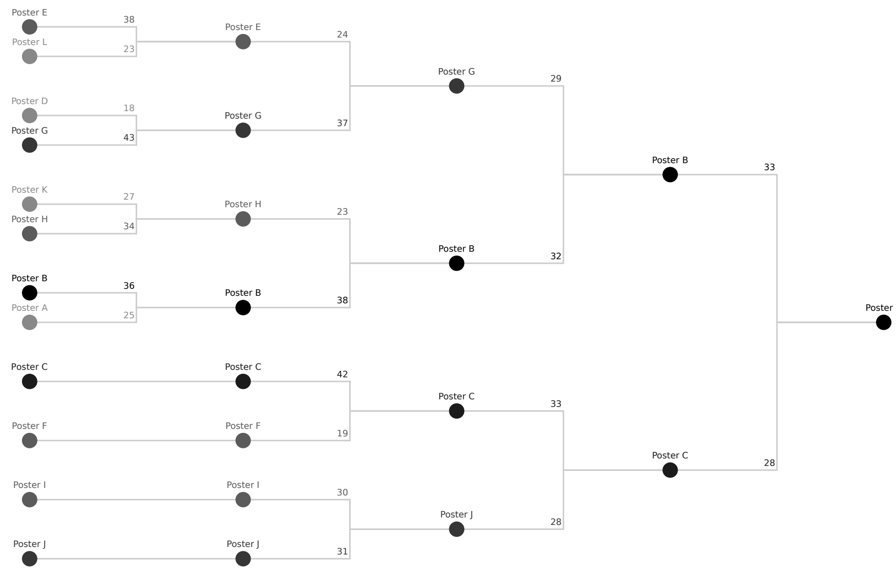

Figure 20.5: Graphical view of randomized cups in Whale 4 (for the best poster election).

types, a different aggregation rule is used. Robovote (Caragiannis et al., 2017) adopts a different point of view. Preferences are always given as rankings, but a different aggregation rule is used according to whether these preferences are supposed to be estimators of a ground truth ("objective" preferences) or totally subjective. These two platforms have clearly a different objective from Whale, as they impose a voting rule *a priori*, that the decision maker has to trust. In Whale, voting rules are more seen as voting profile exploration tools that give several points of view on the voters' preferences.

20.3 The Best Poster Award Use Case

After having introduced the main goals and features of the Whale web application, we will present in this section a situation where it has been used as a voting platform in a real context. This will be an occasion of discussing the main strengths and weaknesses of the application by analysing the results and feedback received from the users after the election.

20.3.1 Organisation

The situation we will speak about in this section is the election of the best poster in the main French-speaking conference in Geomatics, SAGEO'14 (`https://sageo2014.sciencesconf.org/`). The conference involved about 120 participants, and 12 posters — referred to as Poster A to Poster L — were competing for the best poster award. It is important to note that most of the participants (*i.e.*, the pool of voters) were from the field of geography, mathematics or computer science, but to the best of our knowledge none of them had a background in collective decision making.

From a practical point of view, each participant was given — together with the conference kit distributed at the registration desk — a voting sheet containing the following information:

- the direct URL to the voting page;
- a personal 16-characters voting code (randomly generated by Whale);
- the list of posters with labels (A...L), authors, titles.

The vote was opened during the first two days of the conference and was closed just a few minutes before the session dedicated to the best poster award (which shows the advantages of automatic tallying). People were free to vote using their own laptops or smartphones, but could also use a computer which was provided by the organisers.

The voting process was presented during the opening session of the conference. During this session, we presented the technical details concerning the connection to the voting page and the voting procedure that would be used to elect the winner. The choice we made was to ask participants to give complete rankings of the posters and to give the best poster award to the Condorcet winner if there is one, and to the Borda winner otherwise. This choice was made both for pedagogical and for simplicity purposes: the goal was mainly to promote alternative voting procedures which are easy to explain, to implement and to understand, and have well-known and good properties. As we will see later, this choice is debatable.

20.3.2 Feedback and Results

In the end, 61 persons took part to the vote, in other words, slightly more than half of the attendees, which is a reasonable score. We collected some feedback on the voting process by directly and informally discussing with the voters. Three kinds of remarks were made by several participants.

- Most people had never voted electronically before for this kind of elections, and had never used rankings to vote for an election. Most of them had not heard at all of the voting procedures we used before the presentation during the opening session. They seemed to be very happy to discover these kinds of methods, and had a real interest in alternative voting procedures. From this point of view, the data visualisation, may they be simple and imperfect, were helpful tools to better understand how these procedures work.
- On the negative side, rather surprisingly, a common remark is that having to type a 16-characters code to access the personal voting page was quite burdensome. This remark speaks for the use of light authentication processes for non-critical settings, at the price of lower security guarantees (see Section 20.4.1).

- Concerning the voting procedure itself, a very important remark is that asking for complete rankings was completely inappropriate in an election setting with 12 candidates where people usually have clear preferences over not more than 4 or 5 candidates.

When the poll closed, we used Whale to tally the votes and determine the winner of the best poster award. The majority graph of the election is shown in Figure 20.3.

It turns out that there was a Condorcet winner — Poster B — that should have been elected according to the voting rule chosen. However, if we have a closer look at the majority margin matrix shown in Figure 20.3, we notice that Poster G is very close to being a Condorcet winner but is not because it loses against Poster B with a margin of 3 votes (32 voters in favour of Poster B, 29 in favour of Poster G).

Let us have a look at the winner given by the most classical scoring voting rules, which are not Condorcet-consistent. Poster B is the (co-)winner only for $\{5, 9, 10, 11\}$-approval (hence Veto), whereas Poster G wins for $\{2, 3, 4, 5, 6, 7, 8, 9\}$-approval and for Borda, and defeats Poster B for the plurality rule.

The settings where different voting rules yield different winners are not uncommon, and are simply due to the fact that they are based on different normative definitions of a democratic consensus and on different interpretations of the voters' preferences. However, a basic analysis of the profile reveals that the majority in favour of Poster B might not be only due to a natural interpretation of the voters' rankings.

Namely, if we look at the voting profile, we observe a clear bias in favour of candidates appearing earlier in alphabetical order. This bias seems to increase as candidates approach from the end of the rankings. This impression is confirmed by a numerical analysis of the profile.[6] Table 20.1 shows the normalized Kendall-Tau distance (that is, the number of pairwise disagreements divided by the total number of pairs) between alphabetical orders and voters' rankings for the k worst and the k best candidates in the rankings.

k	2	3	4	5	6
δ (%) for the k worst	34.4	42.6	38.3	40.2	44.3
δ (%) for the k best	57.4	51.4	49.7	49.8	49.5

Table 20.1: Normalized Kendall-Tau distance between alphabetical orders and voters' rankings for the k worst candidates and the k best candidates in the rankings.

This table clearly shows that near the bottom of the ranking, the voters tend to agree more with alphabetical order than near the top. How does it apply to Poster B and Poster G? Table 20.2 shows the number of voters for which B and G both appear amongst the last k candidates of the ranking and preferring Poster B (resp. G) to Poster G (resp. B).

[6] This analysis just shows some informal trends, as a rigorous analysis would have required to start from a hypothesis on the probability distribution of the votes and evaluate it as regards to the data obtained during the election.

k	2	3	4	5	6	7	8	9	10	11	12
$\#(B \succ G)$	1	1	2	6	8	8	11	12	20	25	32
$\#(G \succ B)$	0	0	0	0	1	3	6	10	14	21	29
Ratio (%)	0	0	0	0	11.1	27.3	35.3	45.5	41.2	45.7	47.5

Table 20.2: Number of voters preferring Poster B (resp. G) to Poster G (resp. B) when B and G both appear amongst the last k candidates in the rankings.

Since about 47.5% of the voters prefer Poster G to B in general, there is no special reason that this ratio changes a lot when the preferences are limited to the bottom k ranks. Table 20.2 shows exactly the contrary. When B and G are both at the bottom of the ranking, there is a much higher probability that G is ranked after B (hence respecting alphabetical order) than the other way around.

The reasons of this obvious bias are pretty clear and are to be related to the aforementioned feedback on the voting procedure: most people have clear preferences on the first 4 or 5 candidates at most, but tend to order the rest of the candidates randomly.[7] In Whale, the interaction effort required by adding candidates at the end of the ranking is significantly lower if they are added in the same order as they appear on the voting page (in alphabetical order in the case of the best poster election). That clearly explains the bias we observe at the bottom of the rankings. This phenomenon could have been anticipated, since this bias has already been clearly observed and measured by groups of researchers in the context of real political elections (see *e.g.* Krosnick, 2006; Ho and Imai, 2008).

We can learn at least two lessons from this analysis. First, where the context permits it, the voting interface should be designed so as to break any *a priori* order in the presentation of the candidates. A solution can be for instance to shuffle the candidate list every time a new voter accesses to the voting page. Secondly, when the number of candidates exceeds 4 or 5, one should absolutely avoid asking for complete rankings, except in very specific contexts where each voter absolutely knows every candidate and needs to give a clear opinion on each of them.

Epilogue. In the end, it was quite clear from the analysis of the voting profile that even if Poster B was the Condorcet winner, it probably was for exogenous reasons that had nothing to do with the quality of the work perceived by the voters, whereas there seemed to be strong arguments justifying the fact that Poster G is actually a better candidate than Poster B. We hence decided to manipulate the election by changing the rules and giving the best poster award to both Poster B and G. We carefully explained the reasons of this change to the voters. This change was possible because we were in a rather small community of well-intentioned candidates and voters. In a higher-stake context, things would have been very different and it would have been all the more important to carefully choose the procedure and test the interface before proceeding to the election.

[7]This is not always the case. For political elections, for instance, we can reasonably assume that people might have a clear idea about both the top and the bottom of the ranking, with a pool of candidates with unclear order in-between.

20.4 Lessons Learned from Ordinary Users

We have described in the previous section the feedback received from the participants to a real election having used Whale as a voting platform. In this section we will now give some informal feedback received from people using Whale on a regular basis for everyday collective decision making. This feedback provides very useful insights on how to improve the user experience on a poll / voting platform. More importantly, it also gives directions and priorities for future research on social choice, by pointing to some practical problems that still need appropriate theoretical models and solutions.

20.4.1 Lighter is (Often) Better

The first important lesson learned from personal experience is that the success of a collective decision making platform like Whale largely depends on a careful analysis of its use cases and of its intended users. At the heart of this analysis, it is especially important to understand the stakes of the typical voting situations that the application will have to process. In the context of voting, we can typically distinguish three kinds of situations:

- high-stake elections like political elections, for which the requirements in terms of confidentiality, verifiability, transparency and availability are so high (and attackers potentially so equipped) that they require specific tools with mathematically provable properties[8] — like Belenios for instance (see Cortier et al., 2013);
- middle-stake elections like local political elections or recruitment for a job, where some mild guarantees must be given on the transparency of the process and the anonymity of the ballots, but where it is acceptable for the users to trust a third-party application and nobody has a high incentive to attack the voting server to manipulate the election;
- low-stake elections where the participants trust each other and the voting platform is just there to help people express their preferences about a set of options and discuss about them.

We can observe that most collective decision making we are faced with in our everyday life falls in the third category. In this case, it is crucial for the process to be as easy, light, and permissive as possible. These features are often contradictory with the standard point of view in voting that usually assumes that: (i) the set of candidates is fixed beforehand and does not evolve in time, (ii) votes are personal, and should not be altered or removed by anyone other than the voter herself, and (iii) it is forbidden for the same person to vote several times. In traditional voting theory, any action contradicting these three assumptions is seen as a (malicious) manipulation, either by the voters or by the chair herself. Experience shows that in many situations, it is perfectly fine for the system to allow these actions, mostly because we are the context of small communities

[8] It is even doubtful that any form of electronic or online voting can guarantee these properties.

where participants all know and trust each other and have a strong incentive to find a consensus that is beneficial to everyone. If the system is too restrictive, it can be observed that participants will try to tweak it to do what they want to do (for instance vote several times if they cannot modify votes, create new polls if they cannot add / remove candidates...), or simply not use it and switch to a more permissive system.

These observations thus argue for having two different sets of parameters: one dedicated to middle-stake elections that gives some mild guarantees on the verification of the usual assumptions (fixed set of candidates and voters, one vote for each voter, a vote cannot be altered...), and one dedicated to everyday situations and permitting every action. This last situation has not been studied so much by social choice theorists to the best of our knowledge, which shows a first divergence between theory and what people really do in practice.

20.4.2 Preference Representation

The second divergence concerns preference representation. In practice, voting procedures usually ask for one name on a ballot. There are many reasons for that: it is cognitively simple, easily understandable, and simplifies tallying.

At the other end of the scale, social choice theory is usually based on the assumption that a voting procedure takes a profile of linear orders (complete rankings) as input, and outputs either a collective ranking, a winner, or a set of co-winners. Hence, the traditional assumption is that each voter is able to rank all the candidates. As we have seen in Section 20.3, this assumption is just completely unrealistic as soon as there are more than four or five candidates to rank. As a result, the order given for the candidates at the bottom of the scale (or around the middle, depending on the context) is meaningless, or, worse, can be strongly biased.

There are many intermediate ways of expressing preferences on a ballot. The right level for most applications is probably to ask people to give scores to candidates, scores being taken from a small qualitative or numerical scale (*e.g.*, $\{\text{Yes}, \text{No}\}$, $\{-1, 0, 1\}$, $\{0, \ldots, 5\}$, $\{--, -, 0, +, ++\}$, ...). Even if a setting like approval voting is well understood and has been widely studied, this does not seem to be the case for other score-based ballot settings,[9] in spite of the obvious cognitive interest in this kind of ballots.

20.4.3 Incomplete Preferences

The third divergence that we can observe between theory and practice concerns potentially incomplete preferences. Whereas most work in voting theory assumes that the voters express complete preferences, in the sense that they express a clear opinion (rank, score, ...) on each candidate, we cannot rule out the possibility that a voter is unable to do so in practice, particularly when the number of candidates is high as in the example we have discussed in Section 20.3. Moreover, as soon as we allow the set of candidates to be dynamic and candidates to

[9] Apart from work in experimental voting that analyses the behavioural differences induced by different scoring scales.

be added during the election, as we have seen in Section 20.4.1, we need to be able to deal with incomplete ballots, since at the time we add new candidates, we do not know yet the opinion on these candidates of the voters that have already voted.

There are two usual ways of dealing with incomplete preferences. The first possibility is to ignore them and force people to give an opinion for every single candidate (*e.g.*, a complete ranking).

The second way of dealing with incomplete preferences is to use a kind of ballot with possible indifferences between candidates (*e.g.*, rankings with ties, truncated rankings, approval ballots, ...) and use a default value for each candidate not evaluated by the voter (for instance “No” for approval, not-ranked for truncated rankings, 0 for the numerical scale $\{0, \ldots, 5\}$).

None of these methods is satisfactory. We have seen earlier how bad the first method can be. The second method can be seen as a workaround, but turns out to be really unsatisfactory in some contexts, because giving a default value to a given candidate is different from the absence of information concerning this candidate. Let us take for instance the voting situation considered by Karanikolas et al. (2016) about the election of the best movie among the ten most popular movies of the last decades according to IMDB registered users. Most voters only had a partial opinion about the movies, just because they had not watched them all. If we ask their opinion on a $\{0, \ldots, 5\}$ numerical scale, assuming that the default value is 0, we will simply not be able to distinguish between a movie they have not watched and a movie they have not liked at all.

These examples show how crucial it is to provide voters with a way not to give any opinion on some candidates. Otherwise, the process of voting will be painful to them, and the result given by the voting procedure might be biased or irrelevant. However, the ways to deal with incomplete ballots is still not clear. Collective decision making in this context probably comes down to the following question: *between a well-known candidate which is just moderately appreciated by everyone and a candidate that just a few voters know but about which all of them are very enthusiastic, which one should be elected?* The answer to this question probably depends on the context, but incomplete preferences would probably deserve a careful attention from the scientific community.

20.4.4 Multiple Voting Procedures

The final issue we will quickly discuss is a well-known difficulty related to the nature of social choice itself, and that is as old as the scientific discussion between Jean-Charles de Borda and Nicolas de Condorcet about the respective merits of different voting systems. As soon as there are more than two candidates to choose from and we ask the voters for more than approval ballots, there are several ways of electing the winner of the election, all corresponding to a different notion of consensus. Worse, there are extreme cases where each candidate is declared winner of the election by a different voting rule. The problem of choosing the *right* voting rule, whatever it means, is not simple in this case.

There are several possible approaches to this problem. First, we can impose a voting rule *a priori* and choose the winner according to this voting rule. This is the approach used by several voting platforms.[10] We could also think of letting election chairs (or even the voters themselves, see Laslier, 2012) choose amongst several proposed voting rules before the election begins. The real question here is to provide enough explanation to ensure that stakeholders make an informed decision about which voting rule they choose to use. Several recent works in social choice have started to study this question of making people choose the voting rule that most fits their needs (see, *e.g.*, Cailloux and Endriss, 2016).

In the aforementioned context of low-stake elections, another point of view is possible. Since in this context the voting platform is mainly used to elicit preferences and discuss about them, voting rules can be seen as many ways of aggregating the preference information contained in the profile and as many ways of exploring this information. This is the approach we have followed in the first two versions of Whale, where the raw results given by several voting rules were displayed to the user, as shown in Figure 20.1. However, as we might expect, simply displaying the potentially contradictory information given by the different voting rules does not help much making a decision, and people were often lost especially when different voting rules yield different winners.

For the last two versions of Whale, we have tried a different approach based on a literal interpretation of "exploring the voting profile". Namely, we have tried to use information visualisation techniques to provide interactive visualisation of the results given by voting procedures. This has lead to the set of basic graphical views that are proposed in Whale 3 and 4 (see Figures 20.2, 20.3, 20.4 and 20.5). This work is just at the beginning, but we believe that it is a promising approach, both to better understand and analyze the voting profiles, and as pedagogical tools to explain to laypersons how the voting procedures work.

20.5 Conclusion

In this chapter, we have presented a web application dedicated to collective decision making. The aim was not exactly to advertise a particular platform, but rather to put into light some practical problems that have not received a lot of attention by the scientific community yet, or, at least, would deserve more attention. Amongst these topics, the design and analysis of voting rules that are usable in practice for everyday collective decision making seems to be of the utmost importance. That concerns the analysis of voting rules that work on small scales of scores rather than on linear orders and also voting rules that are robust to the absence of information concerning some candidates. It would be also crucial to put some efforts into making voting rules be more than only black boxes or oracles that elect a winner for a given set of ballots. It could mean developing ways of explaining these voting rules and providing people with concrete means to argue to choose the right one, or use them as voting profile exploration tools, through graphical visualisation for instance. It is a safe bet that some of the problems listed above will be amongst the future trends of the discipline.

[10]And it is also, to some extent, used by by poll platforms like Doodle® or Framadate, which recommend to choose amongst the set of co-winners that lexicographically maximise the number of "yes" and the number of "(yes)" choices.

Acknowledgments

Most work about information visualization presented in this chapter has been made in collaboration with Renaud Blanch. Several people have contributed to Whale, especially Marie-Jeanne Natete, who was the main developer of Whale 4. Whale is supported by ANR project CoCoRICo-CoDec.

Bibliography

F. Brandt, G. Chabin, and C. Geist. Pnyx: A powerful and user-friendly tool for preference aggregation. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1915–1916, 2015.

O. Cailloux and U. Endriss. Arguing about voting rules. In *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 287–295, 2016.

I. Caragiannis, S. Nath, A. D. Procaccia, and N. Shah. Subset selection via implicit utilitarian voting. *Journal of Artificial Intelligence Research*, 58:123–152, 2017.

V. Cortier, D. Galindo, S. Glondu, M. Izabachene, et al. A generic construction for voting correctness at minimum cost-application to helios. *IACR Cryptology EPrint Archive*, 2013:177, 2013.

D. E. Ho and K. Imai. Estimating causal effects of ballot order from a randomized natural experiment the California alphabet lottery, 1978–2002. *Public Opinion Quarterly*, 72(2):216–240, 2008.

N. Karanikolas, R. Blanch, and S. Bouveret. Edge-compressed majority graph: Where social choice meets information visualization. In *Proceedings of the 6th International Workshop on Computational Social Choice (COMSOC)*, 2016.

J. A. Krosnick. In the voting booth, bias starts at the top. *The New York Times*, November 2006.

J.-F. Laslier. And the loser is... plurality voting. In D. S. Felsenthal and M. Machover, editors, *Electoral Systems: Paradoxes, Assumptions, and Procedures*, pages 327–351. Springer, 2012.

N. Mattei and T. Walsh. PrefLib: A library for preferences. `http://www.preflib.org`. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT)*, Lecture Notes in Artificial Intelligence. Springer, 2013.