Building trust for sample voting

Abstract

This work explores how to build popular trust for voting systems that rely heavily on statistical tools, as they are generally counter-intuitive to the average citizen (and even to experts). By trying out the voting system in public and letting people tinker with it, a first level of familiarity can be achieved. Preliminary results from real-world experiments seem encouraging and point out the importance of psychological and sociological factors in election organization as well as the influence of user interface design. To go further, integration into a larger debating platform held by a national party could give first-hand experience to the majority of the people, and would progressively build trust as the political stakes grow higher. Finally, we look into how different e-democratic tools could interact in a mutually beneficial manner.

Keywords

Sortition, e-Democracy, Probabilistic Algorithms, End-to-End Verifiability, Citizen Trust, Rational Ignorance, Institution Design

Introduction: the rational voter model

In any system where participants' preferences are aggregated to produce a decision, such as the electoral systems of modern democracies, the degree of public participation is a crucial metric. Not only does it bring legitimacy to the system – thus making decisions more accepted by the population – it also protects minorities against some forms of abuse.

Despite the importance of public participation for the political health of the country, the direct impact on the voter is low, giving rise to the paradox of voter turnout. If the voter's interest lies only in their vote's impact, then large countries should have extremely high abstention, which isn't the case, with huge variability between countries and elections (Jackman, 1987, Blais et al. 2014).

This paradox has been extensively studied in the past five decades (Riker and Ordeshook 1968; Overbye, 1995; Kanazawa, 1998), and multiple factors have been shown to have a strong impact on public participation. Here are some of the best-established factors:

- Population size, as an increase in population comes with a decreased probability of casting the decisive vote for each voter, hence a lower turnout.
- Population stability (defined as the rate of demographic exchanges), as people vote more frequently when they feel they – and their neighbours – will stay in the constituency for a long time.
- Closeness of the election, as a perceived sure win for one party lowers the turnout.

- Whether the vote is mandatory or not, with a fine in case of abstention (although the fine is not enough to explain the magnitude of the effect, and social costs also factor in).

Many different models have been developed (Grofman, 1993; Bendor, Diermeier, and Ting, 2003) to explain the discrepancies between the original model and the observed variability of behaviour, by incorporating risk aversion, externalities such as social costs, or by proposing information asymmetry (Geys, 2006).

Moreover, the voters are often accused – legitimately – of being ignorant of the main issues at stake (Congleton, 2001), which can be explained using the same analyses, under the term of rational ignorance (Downs, 1957; Martinelli, 2006). Under this model, if the expected effect of each person's vote on an issue goes down, while the cost of acquiring expertise on an issue stays constant, the proportion of knowledgeable voters will go down. As a consequence, this imposes an externality on the knowledgeable voters as it further diminishes the effect of their own vote (Großer & Seebauer, 2016), which in turn lowers the interest in getting more information. In such cases, mandatory voting can in fact make matters worse.

Combining the different effects mentioned, we realize that large-scale democratic procedures are replete with adverse incentives, and leaving voters in direct control of policy has been decried as a suicidal move for any state for more than two centuries – Sieyès (1789), in the morrow of the French Revolution, stated that the average citizen's lack of education and time doesn't prevent him[1] from choosing a representative, but that he shouldn't concern himself with matters of state as it would lead to chaos. That was already an echo of Montesquieu (1748), for whom choosing the best suited for government was a task fit for the general population, unlike lawmaking which should be left to those in the government.

Thus, only two possibilities have been in the mainstream public discourse since the middle of the 18th century: direct democracies with uninformed voters, inefficient on a scale larger than a city-state, and "representative" democracies, where a small set of individuals – the best among their peers – are in power. This set of individuals was initially composed of the aristocracy, but the idea is still very present, with the system known now as technocracy, rule by the experts.

Despite the obvious problems with direct democracies, representative systems are increasingly criticized for putting the interests of the few ahead of those of the many (Gilens, 2014; Salamon, 1977), and calls for more direct forms of representation are heard in most western democracies (L'Observatoire de la Démocratie, 2016). Moreover, their proponents – such as the Five Star Movement in Italy (Five Stars, 2017) – are faring better and better in national elections. New institutions, like the Public Opinion Platform (Public Opinion Platform, 2017) are also being proposed and gaining visibility. To avoid the pitfalls of large-scale decision-making, some of those institutions are re-introducing the old method of random selection among the general population, or sampling. This method was prevalent in Antiquity, especially in the Greek city-states (Rackham, 1959), as well as during the middle ages in the Italian republics (Da

[1] Sieyès, unlike his contemporaries Condorcet and Olympe de Gouges, was opposed to granting women full citizen's rights.

Canal, 1275). Since then, it has progressively disappeared, and the only testament to its history is in citizen juries – which are also being phased out in many countries (Sparks and Butts, 2006).

Sampling combines the main advantage of direct democracy – that the interests of both the general population and the decision-makers are aligned – with the efficiency and expertise of technocracy. However, sampling naively – such as with a simple public lottery – would create more problems than it would solve. Thus, the new institutions appearing today need new tools to handle registration, discussion, drafting of legislation and voting when necessary, as well as to help organize public meetings.

Recent advances in cryptography are a first step in this direction, and allow for secure sampling, end-to-end verifiable voting, and even a combination of those in the case of Random-Sample Voting (Chaum 2017). However, all those advances rely both on randomness and on advanced cryptography, and the inner workings of the system, although public, are generally not considered accessible to the common citizen.

Although the security and representativeness of such systems are both proved mathematically, most people wouldn't be ready to trust it for important decisions, as its legitimacy is not established. The goal of this research is to see how to establish a popular appeal for voting systems that use randomness, and how to prove their legitimacy to voters. Moreover, we seek to study how an ecosystem combining those tools with new institutions like the Public Opinion Platform (POP, 2017) could appear and form mutually beneficial relationships between its actors.

Benefits of sampling

As stated earlier, sampling combines the positive features of direct democracy and technocracy. Its main advantages are as follows (in a context where it could be used to vote on legislative proposals, although there are other possibilities):

- **Expertise**. The set of decision-makers is reduced, meaning that each individual voice has a lot of importance, and all voters have a lot more power. This translates into better-informed voters who spend more time on issues and are less vulnerable to propaganda. Moreover, as independent samples can be taken for each different decision to make, the task is distributed among the general population, reducing individual loads. This is to contrast with professional parliamentarians, who sometimes admit that they rarely have the time to read legislative proposals before voting on them (Dosière, 2014).
- **Efficiency**. As the set of decision-makers is of a reduced size, the time needed to reach a conclusion can be shortened, while countering the accusations of skipping a public debate (as the decision-makers are representative members of the public). This compounds with the fact that a sample can be devoted to a single decision, instead of having one set of people that needs to take all the decisions.
- **Protection of minorities as well as the majority**. A problem that has appeared in multiple democracies is that of a vocal minority oppressing a smaller minority (as well as the uncaring majority, to a lesser extent) while the general population does not intervene.

This happens because certain groups take political decisions based on a single issue, and have extremely high participation rates when it comes to legislation going against their core beliefs. Politicians then have a strong incentive to tolerate – and even defend – those beliefs, under penalty of having a strong disadvantage in local elections, which means that there can be a rift between the opinions of the general population and those of the elected officials. For an example of this effect, one can look at devout catholics and abortion in both Northern Ireland and the Republic of Ireland – where abortions are illegal unless in case of deadly risk to the mother – where between 60% and 80% of the population has been in favor of reforms since the 1990s (Bouclin 2002; Millward Brown, 2016), whereas a strong majority in parliament had been opposed to reforms until very recently.

- **Alignment of interests**. The main selling point of these systems is that there is no technocratic class of people set apart from the rest of society. If the sample is representative of society in its entirety, its interests are entirely aligned with those of society. However, this is only true if the sample corresponds to a faithful snapshot of the entire society. This is the first point that we must prove before going further.

Sampling from a mathematical standpoint

To prove that the interests of a sample are indeed aligned with those of society at large, we must prove that the whole range of opinions held inside the sample corresponds (within a small bound) to that of society, with similar support for the same ideas. To do this, we first need to recall a standard theorem on probabilities:

Theorem 1 (Chernoff)

Let $\square = \frac{1}{\square}\sum_{\square=1}^{\square} \square_{\square}$ *where* $\square_{\square} = 1$ *with probability p and 0 otherwise, and where all the* $\square_{\square}$ *are independent. Then*

$$Pr(\square \geq (1+\square)\square) \leq \square^{\frac{-\square^2}{2+\square}\square\square} \text{ and}$$

$$Pr(\square \leq (1-\square)\square) \leq \square^{\frac{-\square^2}{2}\square\square}.$$

This means that the probability that a set of n independent experiments diverges from its expected value by a constant multiplicative factor is exponentially small in n. The proof can be found in Mitzenmacher and Upfal (2005).

In its current state, however, the theorem does not directly apply to the problem at hand, and we need a corollary that can handle populations instead of random variables. Let M be a set of m individuals among a population of size n. Let's now take a sample S of size k uniformly at random, and look at the k' individuals in $\square \cap \square$. Let $\delta < 1$.

Corollary A of Theorem 1

With probability at least $1 - \square^{\frac{-\square^2}{3}\square\frac{(\square-\square)}{\square}}$,

$$(1-\square)\frac{\square}{\square} \leq \frac{\square'}{\square} \leq (1+\square)\frac{\square}{\square}.$$

The proof is easy, and goes by simply adapting the original proof and bounding the individual probabilities of each member getting into the sample – varying from $\frac{\square-\square+1}{\square}$ to $\frac{\square}{\square-\square+1}$. This establishes the representativeness of any uniform sample of a population – as long as it is large enough, with 10 000 members being enough in practice.

This seems contradictory with common experience of political polling, where even with large sample sizes the error rates remain high (Desart & Holbrook, 2003), and where the average predictive ability of most electoral polls stagnates around 80%. This is the crucial difference between the accuracy of polling (where people are asked their opinions), and the representativeness of a sample. Because of nonresponse bias, even a strong representative sample can end up with a high error rate if the response rate differs greatly in a way correlated with the answers to the question – such as in the 1936 Literary Digest election poll which had a 19 point error with a sample of 2.6 million people, more than 5% of all voters. Despite improved quota methods, this effect has steadily increased in magnitude over time (Kohut et al., 2012; Zukin 2015), as response rates that were initially higher than 60% dropped progressively, first to 36% in the late 90s, and and then below 10% in most western countries nowadays.

Thus, despite the strong mathematical support for its representativeness, a sample will take the same decision as the general population only when abstention is low. According to the rational voter model, this should happen when each voter's choice has increased weight, which is exactly the kind of framework we are using. Altogether, this means that sampling, when used for direct control of legislative procedures – or at least in cases where each vote has a significant effect – could have all the advantages mentioned previously, including the alignment of interests. However, the very fact that each vote really counts gives rise to some complications.

Vulnerabilities of a sample

Sampling, when performed naively, has two main vulnerabilities that should be addressed. The two risks are the sample not being representative of the general population, and the sample being influenced by outside forces.

If a sample is taken uniformly at random, we have just proved that it is mathematically representative. The problem lies in proving the fact that it was taken uniformly at random. In the past, mechanical lotteries have been used for this purpose (such as the *kleroteria* in Ancient Athens), but modern technology makes it harder to trust such a machine, especially on a large scale. To guarantee the legitimacy of the decision taken, some procedure must be in place to prove that the sampling was fair, and that every citizen had an equal probability of being a member of the sample.

The second problem is directly linked to the superior power held by participants in a vote, as there are fewer of them. This raises the possibility of corruption, as with any reduced set of people holding a lot of power. Even in usual elections, votes are sometimes bought, for prices generally lower than 50 US$, mostly in South-East Asia and Latin America (Bouton et al., 2015) with the prevalence globally diminishing as GDP per capita increases. Lobbying and campaigning costs for major elections in OECD countries end up in the same order of magnitude

(around 54$ per US voter in the 2016 election). As a sample should generally take less than 1% of a population (and ideally, less than 0.1% in populous countries), buying votes becomes not only feasible but also cost-efficient. Moreover, even among an incorruptible population, coercion is still imaginable.

There are three different ways to prevent undue influences on the sample. The first one is by having strict and complete monitoring of its members during the whole voting period, which is increasingly difficult, not to mention that the surveillance apparatus itself could be used by its controllers to influence the members of the sample. The second one was used by the Athenians, and consists in isolating the members entirely from the moment they are chosen to the moment they take a decision, which was already logistically complex at the time, and impossible today – especially if we want them to spend more than a few hours thinking about the decision.

The last one is by enforcing the anonymity of the members of the sample. For example, one could draw the lots secretly and only inform the members, which would prevent lobbyists from contacting them. However, that just changes the onus of corruption, and members of the sample could contact lobbyists themselves to see how much they could gain by selling their vote. Therefore, we need a stronger version of anonymity where even members of the sample can't prove their own status.

Random-Sample Voting

To address both of those constraints, a voting protocol called Random-Sample Voting (RSV) was developed over the past few years by David Chaum and the RSV Project (Chaum, 2017). It features end-to-end verification and integrates voting with the sortition of a reduced voter pool from the general electorate. As such, the result of a vote cannot be modified through hacking – as there are public audits and a paper trail – and it is hard to force a denial of service as it handles long voting periods, which are too costly for polling place voting. Moreover, as with all sortition-based systems, it is possible to have very frequent votes – each on an independently generated sample – to guarantee a representative opinion while limiting voter fatigue.

The user experience is relatively simple: using the complete list of eligible voters, public random bits are used to draw a reduced group of random citizens whose identities are kept secret. Every member of this reduced voting pool receives one ballot with a secret ID and four voting codes (two for YES, two for NO). The members of the voting pool then go to the voting website, input their ID and one of the two vote-codes corresponding to their choice. They can also give this ID and one code to a third party to securely vote for them (without knowing what they voted for), and check afterwards that this was performed accurately. This means that the voting part can be done over any medium: online, by phone, by mail, or by asking your neighbour, negating the usual problems with online systems.

Corruption and decoys

RSV overcomes the first vulnerability mentioned by publishing an encrypted list of all citizens in a given order first and drawing public random bits afterwards – corresponding to a set

of numbers, hence a set of citizens. The random bits come from a source that is impossible to manipulate (such as entropy from a stock exchange), with a simple algorithm that is visible and computable by everyone. This allows all citizens (and external observers) to check that the sample was drawn uniformly at random, with every citizen getting the same chance.

Although the voters are drawn publicly, the encrypted nature of the table means that only the organisers can initially find out who gets a ballot – the key being revealed after the vote. The anonymity is hence not perfect because:

- Some people know which citizens get a ballot, and
- Ballot owners can just go see lobbyists and sell their ballots.

To prevent the first one, an automated system can be used, and letters could be sent automatically to all citizens (or simply an added sheet in their yearly tax forms, informing them of whether or not they were selected in a sample). Moreover, control systems limiting administrative access to part of the database are also possible.

To prevent people from selling their ballots, *decoy ballots* can be added to the protocol. Those are identical in all features to a real one except for the fact that their votes are not counted in the final tally. By authorizing all citizens to ask for such a ballot and receive one at will, the vote-buying market is saturated by worthless decoys (and an equilibrium point with a negligible ballot price can be forced to make sure that it is always stays inundated). This means that any citizen who has asked for a decoy will know that their ballot is a decoy (as requested), whereas a citizen who receives a ballot without having asked for a decoy will know that theirs is authentic (this can be enforced with some cryptographic tools).

Finally, despite being able to prove that each vote was counted correctly, no-one can link a specific vote to a given voter. Thus, we can avoid both corruption and coercion.

Trust and voting

The protocol explained is not too complex and is strictly superior to most other voting systems (through provable security and anonymity, as well as secure proxy voting). However, it relies on three different problematic tools:

- Randomness to select the sample
- Cryptography to handle the voting
- Internet to handle the public part of the vote (although voting can be done without personal online access, by voting securely through a proxy).

Each of those represents a significant hurdle to overcome for the system to be perceived as trustworthy in the eyes of the public. For example, many random procedures are extremely counter-intuitive to humans, and convincing citizens of their correctness is an arduous task. Adding cryptography to the mix, which includes mathematics generally out of reach for most users, precludes a simple explanation. The simple fact that two different sets of codes are on each ballot and that voters are supposed to choose one and use it – which is an integral part of the security mechanism as it prevents election organisers from cheating – is by itself very confusing for voters. The use of internet is also problematic for a significant proportion of the population

who have doubts about security issues, whereas they generally do not care about those when voting with paper ballots (Riera & Cervelló, 2004).

However, as shown in by Oostveen and Van den Besselaar (2004), the actual features present in a voting system have little impact on voters' opinions and preferences, and their initial perception of the system is what matters most. Thus, particular attention has to be paid to voter psychology to maximize the probability that they trust the system.

To this end, a simple explanation – with proof – of the correctness and validity of the system are not enough. The first point that needs to be addressed is the representativeness of the sample. Most uninformed voters, when presented with systems that use sortition, are afraid of selecting a sample that entirely votes for one extreme party. To convince them of the practical impossibility of such a scenario, powerful images are needed to assuage the irrational fears that may be present.

For example, Corollary A is extremely powerful, and directly applicable, but makes no sense to the average citizen. We must then apply it to real data to make it intelligible. Moreover, we want the sample to be representative of more than a simple binary partition into those who are for or against a given piece of legislation. To this end, let's define a set P of partitions of the population into two sets M_i and $M_i{}^c$. For a sample S to be $\delta - representative$, we want that for each i,

$$(1-\delta)\frac{|M_i|}{n} \leq \frac{|S \cap M_i|}{k} \leq (1+\delta)\frac{|M_i|}{n}$$

This allows us to prove

Corollary B of Theorem 1

For any set of partitions P of size less than 5000, a sample S of size 10^6 *is 0.005-representative with probability at least* $1 - 10^{-21}$.

The proof goes by simply applying an union bound in conjunction with Corollary A. No explicit mention is made of the total population n in this result, which is a strong feature of the model: as population diverges towards infinity, the exact probability of error increases, but stays beneath the bound given (the variation is due to $\frac{m-k+1}{n}$getting closer to $\frac{m}{n}$ for a given k).

This is still not intuitive enough for the average person as large numbers and probabilities don't speak to us intuitively. This is because, as stated in (Riera & Cervelló, 2004), people are not as interested in the security of the systems they generally use, and tend to idealize them. Even in systems considered secure, more than 0.5% of ballots – and up to 2% – are generally badly counted (Goggin, Byrne, and Gilbert, 2012). In end-to-end verifiable voting, errors should not happen, and can be accurately estimated in case of an intrusion. Together with the previous result, we can finally show

Corollary C of Theorem 1

If we were to run an election by sampling 10^6 *people every second since the Big Bang, the probability that the margin of error – for any criterion in a set of 5000 – would be greater than for current elections is less than 0.1%.*

With this statement, we can break the general thought pattern that current systems are secure, and give a strong argument for sampling, not by showing that it is perfect, but by showing its superiority.

However, to truly convince a human being, logical and mathematical arguments are generally not enough, and first-hand experience is preferable. For example, in the case of the Monty Hall problem, one of the textbook counter-intuitive problems in probabilities, humans who repeat the game many times tend to adopt the dominant strategy, the proportion of players going from 25% to 55% in a few dozen games under certain conditions (Petrocelli and Harris, 2011). Although this does not convince everyone (Herbranson and Schroeder, (2010); Granberg, 1999), it is still an extremely powerful tool.

Real first-hand experience is, however, hard to attain for a system that by design wants to query only a minority of the electorate – such that most people only participate every few years. To this end, we created a graphical sampling simulator that allows users to set up a fake electorate and see how RSV would perform when compared to usual voting methods (on either user-created or historic elections). This can take the form of a game where they try to find errors (and can run hundreds of samples every second while recording the error rate). The hope is that by trying it out themselves, interested citizens could get an experimental conviction of the correctness of the system, which would be more efficient than purely intellectual arguments. A beta version can be found at (website removed for anonymity reasons). As an additional feature, to guarantee the authenticity of those experiments, the code is entirely visible, open-source, and runs on the user's machine.

In parallel, RSV has also been used in real-life trials to get public feedback. The first two were at international cryptography conferences where no-one opposed its legitimacy, but where most of the reactions concerned the security (no vulnerability was found). This legitimacy was questioned in the third test, which targeted a general audience composed of non-specialists in San Sebastián, as detailed below.

Experimental results

The Global Forum on Modern Direct Democracy (Global Forum, 2016) was held in San Sebastián between the 16th and the 19th of November 2016, gathering more than 200 participants from 27 countries, including political scientists, journalists, activists, professional politicians and people from many other backgrounds. The RSV Project was officially invited to present its technology by running a public demonstration of the voting system. The first problem arose with the absence of a list of participants (or even of their expected number), which meant that we had to simulate the random drawing of the sample. We decided to do that by holding two concurrent votes on different questions: "Should voting in national elections be compulsory?" and "Should negative campaigning be prohibited?", with participants getting to vote on one of the two randomly. The process went as follows:

- 146 ballots of each type were printed and put inside identical envelopes and then shuffled, no one knowing in the end what each envelope contained, thus simulating the random selection, and making sure that everyone had a chance to participate.
- During the introductory talk about RSV on the first day those envelopes were distributed to all the people in the room, ensuring that everyone got a ballot, but that no one knew who had what (around 130 ballots were distributed in the end because of low turnout at the conference that day).
- The polls opened at 00:00 on the night of the first talk and stayed open for nearly two days until 21:00 on the penultimate day.
- People could vote on their phones or computers and a public polling station was established for those who didn't have internet access, but they were also encouraged to vote through a third person (generally an RSV representative), to demonstrate that feature. This contributed greatly to alleviate concerns over online voting, especially when it comes to unequal access to Internet.

Two complications arose in the experiment: just before the beginning of the conference errors were found in the voting parameters – thankfully we had just enough time to change those before the daily publication of the random bits – forcing us to redraw and reprint all ballots. Another problem was that on some devices the voting interface would hide the last numbers of the input. We found a solution before the second day – which was as simple as holding one's phone horizontally – but many potential voters were confused.

Overall the process was a success as we got a lot of valuable data. Of the 130 ballots distributed, 34 were used, an abstention rate on par with previous tests. However, this time we were able to analyze the different reasons behind it by getting feedback:

- The voting site's web address written on the ballots was long and people didn't want to write a long string into the browser.
- People were mistaking 'g' and 'q' in that link, and couldn't access the voting interface.
- The secret IDs on the ballots suffered from the same problem, as they were 16 character long random strings.

Those problems, along with the interface bug, meant that many voters got frustrated and did not vote. Some were ready to try voting again after the fix was released (we recorded 6 ballots before the fix and 28 afterwards), but most had given up after the first try. Moreover, we could observe – but not measure – additional psycho-social factors: the concept of negative campaigning was foreign to people from countries where such practice is forbidden, resulting in low turnout. More importantly, the timeline of the vote was badly chosen: people did not vote on the first night as it was not urgent, and they often forgot their ballots in their rooms on the second day, meaning that a few dozen couldn't get to their codes before the deadline.

The simulator was also presented, but our preoccupation with fixing the voting system meant that little feedback was collected on it (the people who did use it generally said they were convinced, but we can’t quantify this effect). During the conference someone remarked that they got a ballot for one question and were upset that they didn't get the other. They thought it would

be a common feeling, and it is a real concern, although if RSV were implemented in practice the frequent votes would not let people feel they were chosen for one and not the other.

Multiple people in the crowd also voiced concerns about the popular reaction if the result of a vote differed from the popular expectation, which is unavoidable. We set up an anonymous feedback collection website accessible from the voting page, and found that people mostly trusted the accuracy and security of the system but were nearly all on the fence on its legitimacy for a mass election, meaning that we need to improve the popular appeal of the system.

New institutions

The public trial attracted the attention of one of the participants at the forum, Géza Tessényi, initiator of the Public Opinion Platform (POP, 2017). This new kind of institution for real-time democracy, created in 2015, aims to hack indirect democracy by acting as a political party whose members are elected normally but pledge to vote for the measures adopted by the population. Those measures are discussed on an online platform open to all – and not just to members of the party – and then voted upon. The platform is young, and multiple technological decisions need to be made to satisfy the following constraints:

- Everyone has to be able to easily access and contribute to the discussion and voting. This means including ways to counter the digital divide and promote public access.
- The online debate has to be productive, with concise arguments, and a system of expertise evaluation might be useful, as well as protection against lobbying.
- The whole system must be resistant to hacking and denial of service attacks. Moreover, there must be strong public trust in the results of the votes.
- Popular involvement should not be costly: both the time and the level of expertise needed to engage and contribute should be kept as low as possible.

The first constraint is far from easy to satisfy as any web-based platform will be inaccessible to a non-negligible proportion of the people – in 2013 at least 17% of the French population didn't use internet at all, and about 10% didn't have a mobile phone (Emmaüs Connect, 2017). However, it is crucial for the legitimacy of POP that its decision-making processes aren't discriminatory. To achieve this, it is necessary to distinguish between participation in the debates and participation in the voting process. It might be possible to partially forgot the first, but the second is essential.

As it happens, some of the technologies to solve the other constraints are already being used. For the second constraint, a potential system for public debate over legislation comes from République Numérique. It is also possible to go even further by condensing similar opinions with AI, raising the possibility of large scale debating. This kind of technology, developed by pol.is and vTaiwan (Pol.is 2017, vTaiwan 2017), is starting to show its usefulness and reliability, and we are currently collaborating with them to check potential adaptations. We need to test whether individual expertise evaluation is needed, potentially through a karma system. The last two constraints could be solved by using RSV, if we manage to make it seen as legitimate. This brings us to new ideas for institutions, such as the following:

- A decision is set to be discussed publicly (the originator of this decision can be the executive branch, a some parliamentary group or any citizen backed by a popular petition).
- A set of citizens is chosen randomly and given the power to decide on this question after a period of six months, and informed immediately.
- A general debate is open to the public, using best practices to make it as productive as possible (this would include AI-assistance and anonymous proposals to avoid extreme partisanship).
- After those six months, the sample, having been informed by the public debate, gets to vote on the issue.

Even this quick sketch of a political decision system has the following properties:

- Efficient for legislative purposes (maximum 6 months between proposal and decision)
- Well-informed, as the sample has a strong rational interest in documenting themselves as seen earlier.
- Fair, as it takes according to the interests of the general population instead of a single class of individuals.
- Resistant to lobbying, coercion and corruption.
- Politically engaging, as it allows every citizen to have their voice heard (as the debating platform guides the reflexion of the sample members), while preventing voter ignorance.

As new tools appear, new institutions could likewise benefit from those tools and make the public not only aware of them, but also proficient in their use through regular practice. Just like Greek city-states, all had different kinds of democratic institutions taking inspiration from each other, multiplying the debating platforms, voting tools, and innovative democratic institutions would benefit all actors in the system. Over time, those could converge towards a better solution that could be reached by only supporting self-sufficient initiatives.

Future work

This work is currently being extended in four directions. The first part concentrates on how to make RSV more palatable to the public and address the issues raised in San Sebastián. The interface changes have been made, with ballots much more user-friendly. A simple website providing a way to access all ongoing votes is also being developed (with shortened URLs and maybe QR codes). Two other experiments have been organized recently to study how people type in codes to avoid the q/p confusion noticed in San Sebastián, and a new code structure has been designed which is more secure as well as easier and faster to type in – with a much lower error rate. Another demo experiment is set to take place in November 2017, to vote on the official recommendations of the World Forum for Democracy. This will happen in parallel with another voting system – evaluation voting – and will allow us to compare turnout and see how concentrating on one issue affects voters' perceptions. The RSV simulator is also being improved and will soon be advertised to check its effectiveness. Finally, a tool is being developed to work jointly with the simulator, whose objective is to let people create their own RSV votes on a limited scale. A psychological study of the effects of such tools on people's opinion might be a good idea once they are already well tested.

The second part lies in restarting the debate on sortition and sampling itself. This is a much harder problem as some deeper issues are present. In some countries, such as France, the strongest proponents of sortition (and sometimes the only ones) tend to be in fringe parties, which changes the focus of the debate and could damage the idea by association. To change this, a book is currently being published, which will hopefully create a more balanced perspective on the subject and move the debate closer to the general public's interest. To contribute to this debate, a better understanding of the history of the idea would be very useful. In particular, there is a need for more research on the political mechanisms that led to its consideration for potential adoption in the mid 18th century before it disappeared entirely from the public discourse – which led to a complete redefinition of the term "democracy" itself between 1770 and 1860.

The third part corresponds to current work with POP, focusing on a technological watch and cooperation with major actors in debating and AI software to settle on an architecture and launch its first public interface. This is leading to more exchanges between both projects, as the use of RSV by POP would make the voting system more familiar and, hopefully, allow it to be tested in real-world situations of increasing importance as the influence of the platform-party grows. To improve the accessibility for people with limited online access, we are also considering authentication alternatives, such as sending code books by physical mail and letting people vote by phone, and whether a multi-tiered authentication system could work. This could introduce some vulnerabilities in the system but a smart design might limit those.

Finally, more theoretical work is underway, focusing on institution and infrastructure design to take advantage of new tools. One innovative design for a global governance system was developed in the context of the Global Challenges Foundation's New Shape Prize, and research is ongoing to see how sampling could be used efficiently in different political systems.

References

Bendor, J., Diermeier, D., & Ting, M. (2003). A behavioral model of turnout. American Political Science Review, 97(2), 261-280.

Blais, A., Pilet, J. B., Van der Straeten, K., Laslier, J. F., & Héroux-Legault, M. (2014). To vote or to abstain? An experimental test of rational calculus in first past the post and PR elections. Electoral studies, 36, 39-50.

Bouclin, S. (2002), Abortion in Post-X Ireland. Windsor Review of Social and Legal Issues. Available at SSRN: https://ssrn.com/abstract=1395324

Bouton L, Conconi P, Pino F, Zanardi M, Facchini G, Steinhardt M, Giavazzi F, McMahon M, Mian A, Sufi A, Trebbi F (2015). Detecting vote buying in monetary aggregates: New evidence on monetary political cycles. Vox, retrieved from voxeu.org

Chaum, D. (2017) Random-Sample Voting, White paper. Retrieved from ww.rsvoting.org

Congleton, R.D. (2001). Rational Ignorance, Rational Voter Expectations, and Public Policy: A Discrete Informational Foundation for Fiscal Illusion, Public Choice, 107-35. https://doi.org/10.1023/A:1010337412291

Da Canal, M. (1275), Les Estoires de Venise; Cronique des Veniciens.

DeSart, J., & Holbrook, T. (2003). Campaigns, polls, and the states: Assessing the accuracy of statewide presidential trial-heat polls. Political Research Quarterly, 56(4), 431-439.

Dosière, R (2014), Le Métier d'élu local, Le Seuil.

Downs, A. (1957). An Economic Theory of Democracy, HarperCollins Publishers.

Geys, B. (2006). Explaining voter turnout: A review of aggregate-level research. Electoral studies, 25(4), 637-663.

Gilens, M., & Page, B. (2014). Testing Theories of American Politics: Elites, Interest Groups, and Average Citizens. Perspectives on Politics, 12(3), 564-581. doi:10.1017/S1537592714001595

Granberg, D. (1999). Cross-cultural comparison of responses to the Monty Hall dilemma. Social Behavior and Personality, 27(4), 431– 438

Grofman, B. (1993). Is turnout the paradox that ate rational choice theory. Information, participation, and choice: An economic theory of democracy in perspective, 93-103.

Großer, J. & Seebauer, M. (2016). The curse of uninformed voting: An experimental study, Games and Economic Behavior, Volume 97, Pages 205-226, https://doi.org/10.1016/j.geb.2016.04.009.

Herbranson, W. T., & Schroeder, J. (2010). Are birds smarter than mathematicians? Pigeons (Columba livia) perform optimally on a version of the Monty Hall Dilemma. Journal of Comparative Psychology, 124(1), 1.

Jackman, R. (1987). Political Institutions and Voter Turnout in the Industrial Democracies. *American Political Science Review*, 81(2), 405-423. doi:10.2307/1961959.

Kanazawa, S. (1998). A possible solution to the paradox of voter turnout. The Journal of Politics, 60(4), 974-995.

Kohut, A., Keeter, S., Doherty, C., Dimock, M., & Christian, L. (2012). Assessing the representativeness of public opinion surveys. Pew Research Center, Washington, DC.

Martinelli, C. (2006). Would rational voters acquire costly information?, Journal of Economic Theory, Volume 129, Issue 1, Pages 225-251, https://doi.org/10.1016/j.jet.2005.02.005.

Mitzenmacher, M. & Upfal, E. (2005). Probability and Computing: Randomized Algorithms and Probabilistic Analysis. Cambridge University Press, New York, NY, USA.

Montesquieu (1749), L'Esprit des lois.

Oostveen, A. M., & Van den Besselaar, P. (2004). Security as belief: user's perceptions on the security of electronic voting systems. Electronic voting in Europe: Technology, law, politics and society, 47, 73-82.

Overbye, E. (1995). Making a case for the rational, self- regarding,'ethical' voter… and solving the 'Paradox of not voting' in the process. European Journal of Political Research, 27(3), 369-396.

Petrocelli, J. V., & Harris, A. K. (2011). Learning inhibition in the Monty Hall Problem: The role of dysfunctional counterfactual prescriptions. Personality and Social Psychology Bulletin, 37(10), 1297-1311.

Rackham, H. (1959). Aristotle: politics. Vol. 264. Heinemann.

Riera, A., & Cervelló, G. (2004). Experimentation on secure internet voting in spain. Electronic Voting in Europe Technology, Law, Politics and Society.

Riker, W. H., & Ordeshook, P.C. (1968) "A Theory of the Calculus of Voting." The American Political Science Review, vol. 62, no. 1, 1968, pp. 25–42. JSTOR

Salamon, L., & Siegfried, J. (1977). Economic Power and Political Influence: The Impact of Industry Structure on Public Policy. American Political Science Review, 71(3), 1026-1043.

Sparks, S., & Butts, G. (2006). Disappearing Juries and Jury Verdicts. Tex. Tech L. Rev., 39, 289.

Sieyès, E.-J. (1789), Dire de l'abbé Sieyès, sur la question du Veto royal, à la séance du 7 septembre 1789.

Zukin, C. (2015), What's the Matter With Polling?, retrieved from www.nytimes.com

Emmaüs Connect (2012). Les connexions solidaires d'Emmaüs Connect : de l'enquête anthropologique aux interfaces inclusives. Presses de l'enssib.

Five Stars Movement (2017). Retrieved from http://www.movimento5stelle.it/

Global Forum on Modern Direct Democracy 2016. Retrieved from www.2016globalforum.com/

L'Observatoire de la Démocratie (2016). retrieved from www.institut-viavoice.com

Millward Brown Omnibus Survey (2016), retrieved from www.amesty.org.uk

Pol.is (2017), retrieved, from www.pol.is

Public Opinion Platform (2017), retrieved from www.poplatform.org

vTaiwan (2017), retrieved from www.vtaiwan.tw