Secure MPC-Sortition: Consolidating Innovations in Democracy and Cryptography

Wouter Maas[1] , **Zekeriya Erkin**[1]
[1]Cyber Security Group
Department of Intelligent Systems
Faculty of Electrical Engineering, Mathematics & Computer Science
Delft University of Technology
Van Mourik Broekmanweg 6, 2628 XE Delft, The Netherlands
w.f.v.maas@student.tudelft.nl, z.erkin@tudelft.nl

Abstract

Globally, citizens' assemblies have been gaining momentum as a way to counter dissatisfaction in democracies. Central to the citizens' assembly is sortition, the process of randomly selecting political representatives given certain demographic criteria. In order to have an assembly representative of the population, personal data is necessary to perform the sortition, making participant privacy a matter of concern when guaranteeing fairness of the process. Secure multi-party computation (MPC) makes it possible to perform calculations on encrypted data from multiple sources without revealing it to the other processing parties or data contributors. This paper outlines two sortition designs that use MPC to guarantee participant privacy during the sortition process. Design 1 allows participants to hide their personal details to the citizens' assembly organisers. Design 2 makes it possible for involved local government to contribute citizen data to the sortition. Different domain experts were interviewed as a means to survey the needs of the sortition community and to provide feedback on the designs. We found that little research has yet been conducted on the effects of increased privacy on sign-up rates for deliberative events, creating a potential avenue for future research.

1 Introduction

Dissatisfaction with democracy is at an all-time high. The Centre for the Future of Democracy reports that in 2020 the share of individuals living in democracies who are dissatisfied with their democracy has risen by around 10%, from 47.9% to 57.5% since the mid-1990's [1]. With dissatisfaction rates this high, the call for democracies to innovate and to restore trust in the democratic system has become louder.

Deliberative democracy is a form of democracy that tries to restore this trust by making deliberation amongst citizens central to the political decision-making process. One of the most promising democratic innovations in the field of deliberative democracy is the *citizens' assembly*. In a citizens' assembly, a group of residents is randomly selected according to some demographic criteria such as gender, age, ethnicity, income level and/or political stance, such that the assembly constitutes a city or a country in miniature [2]. This process of selecting political representatives by lot is called *sortition*, which in the case of citizens' assemblies is achieved through *stratified sampling* [3].

The goal of a citizens' assembly is to perform an in-depth analysis of a given issue, deliberate over different solutions, weighing of the pros and cons, and then, making informed policy decisions [2]. In this way, citizens' assemblies have recently been used to come up with solutions for particularly difficult, emotionally loaded and controversial issues which previously caused cleavages in societies or communities.

Two recent and famous examples of citizens' assemblies are the Citizens' Convention for Climate in France [4] and the Irish Citizens' Assembly focused on abortion [5]. These examples also show the potential results citizens' assemblies can have, as in France all but three of the 149 recommendations by the assembly were accepted by French President Macron and the assembly in Ireland led to a referendum and the subsequent removal of the Eighth Amendment (article 40.3.3) from the Irish constitution (which previously banned abortion in almost all circumstances).

The collected data, necessary to do the sortition based on the aforementioned demographic criteria, is considered personal data under the definition of the General Data Protection Regulation (GDPR), as it is information that can be related to an identified or identifiable natural person [6]. Based on the interviews with domain experts done for this research, it is clear that all organisations in charge of organising citizens' assemblies take the collection and storage of this data very seriously, but also that their methods to do so vary wildly.

Organisers use either door-to-door registration to find participants or send a large number of mail invites in which people are asked to sign up online, for instance by collaborating with the local government. As indicated by the interviewed experts, the sensitive nature of the necessary data can make people unwilling to share their personal details when asked to sign up. On the other hand, due to privacy concerns, the involved local government might be unwilling to combine its tax records (which could be used when income is one of the

demographic selection criteria) with the sign-up information gathered by the organisers.

Secure multi-party computation (MPC) is a set of cryptographic techniques which has the potential to improve on these issues. MPC makes it possible to perform calculations on encrypted data from multiple sources without revealing it to the other computational parties or participating data providers [7]. For a long time MPC was considered too slow for many practical applications, but recent research, however, has focused on improving practical MPC implementations [7]. Additionally, usability has improved as well, with the recent development of protocols such as Web-MPC [8] it can be explored whether MPC can also be used for the sortition process fundamental to citizens' assemblies.

The aim of this research is, therefore: *how can MPC be used to make the sortition used in deliberative democracy fairer and more private?*

This paper is structured as follows: in Section 2 the methodology for this research is briefly outlined. In Section 3 a formal definition of sortition is given, alongside two use case scenarios for MPC and the requirements a secure MPC-sortition design should adhere to. Section 4 outlines related work. Subsequently, Section 5 briefly explains MPC and compares existing deployments which could be used. Section 6 uses these building blocks to outline two designs that could make sortition fairer and more private. In Section 7, the feedback from the interviewed experts is considered to discuss the advantages and disadvantages of these designs. Additionally, we present some pointers for future research. Section 8 briefly reflects on the ethical aspects concerning this research. Finally, Section 9 contains some concluding remarks.

2 Methodology

For this research, two methodologies were used. Firstly, a literature study was performed. Secondly, in-depth interviews with experts in the fields of deliberative democracy and sortition were conducted.

2.1 Literature Study

The goal of the literature study was to deepen the understanding of MPC, deliberative democracy and in particular the sortition process supporting deliberative democracy. This was done by investigating citizens' assembly organiser manuals, reports and websites. Finally, academic literature was explored to find and compare MPC protocols and designs that could be used as a model for the design necessary to improve the current procedures.

2.2 Expert Interviews

Interviews with deliberative democracy experts and technical sortition experts were conducted at two stages in the research process. Early on, they were used to further the understanding of current problems for citizens' assembly organisers. Based on the preliminary findings of the literature study and initial interviews, two designs for secure sortition using MPC were outlined. In the second round of interviews, feedback on these designs was collected to evaluate their effectiveness.

Organisation Name	Interview Date
Extinction Rebellion Netherlands	15/05/2021 & 18/05/2021
De Nationale Denktank	15/05/2021
CitizenOS	18/05/2021
The Sortition Foundation	18/05/2021
The Dutch G1000	19/05/2021
De Transitie Motor / Dutch G1000	20/05/2021
The Centre for Deliberative Democracy at Stanford	26/05/2021
The Irish Citizens' Assembly / Amárach	03/06/2021
The Citizen Initiative Portal (rahvaalgatus.ee)	17/06/2021

Table 1: An overview of the interviewed organisations and their corresponding interview dates.

In total, 10 interview sessions were held with a total of 16 domain experts (some sessions included multiple experts). Table 1 presents an overview of the interviewed organisations.

3 Problem Description

In the following section sortition in its current form is summarised, after which two scenarios are outlined in which MPC could provide an improvement to sortition. Subsequently, a more formal problem description of sortition is presented alongside a set of requirements secure MPC-sortition would need to adhere to in order for it to be usable.

3.1 Sortition in Its Current Form

As previously mentioned, different organisers use different procedures for selecting the participants. Central to all of them is the process of *stratified sampling*. Briefly explained, stratified sampling works in the following way: 1. list several criteria. 2. make a *stratum* for each possible combination of criteria. E.g. gender, age, income: male, 20-40 years old, €10,000 to €20,000 annual income. 3. fill the individual *strata* through random sampling according to the relative size of that group. For a more exhaustive discussion of stratified sampling see [9].

Based on a wide set of sortition literature, including academic papers, citizens' assembly reports and practical sortition guides, a general outline of the steps involved in the sortition process for a citizens' assembly was summarised [10] [11] [12] [13] [14]:

1. An organiser, very often with the mandate of a local government, sets out to organise a citizens' assembly surrounding a certain societal issue.
2. The organiser determines the demographic criteria for the sortition proces. Standard criteria are: age, gender and geography (i.e. postal code). Situational criteria are ethnicity, nationality, socio-economic (e.g. income), disability, opinion on a certain issue (e.g. climate change).

   The interviewed experts indicated that using more than 8 criteria becomes infeasible as each added criteria dou-

bles the number of strata. Eventually, making them too small and too specific, and therefore making it unreasonably hard to select people, also see [9].

3. Based on the criteria and national or regional demographic statistics the strata for the stratified sampling are defined.
4. Subsequently, either participants are recruited door-to-door or they are asked to self-register online or via the telephone with the invites being sent by mail to a large number of random addresses throughout the region of interest.
5. People who indicated their willingness to participate administrate their relevant details corresponding to the selection criteria and consent to the use of this information for the stratified sampling.
6. Based on all the registrations an assembly is formed, as close as possible to the originally defined strata. A surplus of participants is selected, such that potential last-minute cancellations can be replaced.
7. The actual assembly begins, often meaning multiple weekends of deliberation and policy drafting.

3.2 Use Case Scenarios and Improvements to Sortition

This paper focuses on improvements MPC could provide regarding the personal data that is collected for the stratified sampling in steps 4 and 5, as described in Section 3.1. These improvements come in two flavours: the first one being improved privacy for the individual participants of the citizens' assembly, the second one being a privacy and security guarantee for local governments, opening up the possibility for organisers to make the sortition more accurate and advanced by using high quality governmental data.

Use Case Scenario 1. The first potential improvement stems from the fact that individuals currently have to be willing to provide their personal data to the organising party. Meaning, they have to trust this organiser with their personal details. As noted in step 2 in Section 3.1, these details can be rather intimate, and people could find it uncomfortable to elicit details such as political opinions to a recruiter or in an online form to the organiser, especially when the local government is a collaborator. Using MPC to guarantee that the organiser is not be able to inspect an individual participant's details could thus increase privacy and trust in the organiser.

This could even have a potentially positive effect on sign-up rates for deliberative democracy events, as current sign-up rates are typically around 3.7% [13]. A plurality of reasons can be named as the cause for this low statistic but generally speaking, anonymity correlates with higher levels of self-disclosure in computer-mediated communication [15]. For instance, research on online surveys indicates that there are measurable influences on the levels of self-disclosure based on the way participants' privacy is treated [16][17]. Additionally, anonymity fosters self-disclosure amongst individuals who feel stigmatised (e.g. because of illness [18])[1].

[1]This particular example bears importance considering the fact that the Irish Citizens' Assembly used disability to define one of its strata.

Use Case Scenario 2. An important set of stigmas that is well documented and which are especially relevant when considering the criteria of income or wealth for sortition are those on living in poverty or in a low-income bracket. People with a low income overwhelmingly indicate that they believe "that other members of society tend to view them as a burden to society—as lazy, disregarding of opportunities, irresponsible, and opting for an easy life" [19, p. 1]. Confronting individuals with their socio-economic class by asking them to indicate their levels of wealth or income in a sign-up sheet might demotivate people to sign up in the first place. For the purpose of a citizens' assembly, a safe environment for participants to share their ideas is a prerequisite.

An alternative scenario would thus be if the question would never have to be asked in the first place, but that instead a participant could simply be informed that their income data would be used for the stratified sampling but in such a way it was never revealed to anyone in the organising organisation. Consent regarding the processing of this data (be it in an encrypted manner through MPC) as provided by the local tax office could still be asked (e.g. by ticking a consent box) but this would present a significantly lower threshold than actively confronting individuals with their socio-economic class through a questionnaire.

The second potential improvement could thus be seen as a more classical MPC issue, namely, the desire of different data providing parties to keep their databases separated and unrevealed to other parties. Currently, organisers of citizens' assemblies base the stratified sampling entirely on data they have collected themselves. Meaning, the only way to gather income data is by asking the participants directly, leaving the data unverified. Local governments might collaborate with organisers to send out the invitations for the sign-up phase, but they are unwilling to share citizen data such as home address or income information with the organiser directly. Organisers are thus limited by the honesty and willingness of the participants to share personal information. MPC could be used to overcome this by working directly with the tax office of local governments to use their tax data for the sortition.

3.3 Formally Defining Sortition

Formally, sortition can be defined through the *panel selection problem* as outlined by [20]. Since both [20] and this paper use this for the selection of citizens' assemblies, we refer to this as the *assembly selection problem* to avoid confusion in the context of citizens' assemblies.

We let $N = [n]$ be the *pool* of people who indicated to being willing to participate through the sign-up (i.e. the candidates). Let $F = \{f_t\}_t$ indicate a set of relevant *features* (i.e. the selection criteria defining the strata). Each feature function $f_t : N \to \Omega_t$ maps each candidate to their corresponding feature value, as such Ω_t is the set of possible values for f_t. For instance, feature $f_t =$ "income" could correspond to $\Omega_t = \{$"€0 to €10.000", "€10.001 to €20.000", "€20.001 to €35.000", "€35.001 to €80.000", "€80.001 and above"$\}$. Finally, candidate i's *feature vector* is defined

as $F(i) = (f_t(i))_t \in \prod_t \Omega_t$, meaning it's the vector encoding for all of candidate i's feature values in F.

Naturally, we impose the constrain that the chosen assembly A must be a subset of the candidate pool with size k. Additionally, we must impose that the assembly is representative of the population of the chosen community or region, as defined through the features in F. [20] impose this mathematically through *quotas*: for each feature f and corresponding value $v \in \Omega$, a lower quota $l_{f,v}$ and upper quota $u_{f,v}$ is defined. Together, these two quotas assure that the assembly must comprise of between $l_{f,v}$ and $u_{f,v}$ individuals with feature f.

Now, we can formally define an instance of the assembly selection problem as: given parameters (N, k, F, l, u)—a pool of candidates, assembly size, set of features (i.e. selection criteria), and sets of lower and upper quotas—randomly select a *feasible assembly*. A feasible assembly is any set of candidates C from the collection K, such that:

$$K := \{C \in \binom{N}{k} : l_{f,v} \leq |\{i \in C : f(i) = v\}| \leq u_{f,v} \text{ for all } f, v\}. \tag{1}$$

A *selection algorithm* is an algorithm that solves instances of the assembly selection problem [20]. Multiple implementations for this problem exist, optimising for different aspects of the selection problem. For instance, [20] focused on optimising transparency, meaning, that participants should be able to understand without much in-depth reasoning the probability each individual candidate has for being selected for the assembly.

One of the most important aspects, however, is how to optimise for fairness. Fairness can, for instance, be looked at from the individual candidate's perspective by looking at privacy, as is in this paper, or by making an attempt to equalise the chances for individual candidates to get selected given the demographic selection criteria [21]. Alternatively, fairness can be considered from the societal or organiser's perspective, in which case fairness is often defined by picking the most representative assembly given the limitations caused by the demographic characteristics of the actual registered candidates [22].[2]

For this paper, the selection algorithm can be thought of as a simple greedy implementation. A simple greedy implementation is not optimised for individual fairness and will only terminate when a sufficient number of candidates is available for each stratum (i.e. has signed up), something that cannot always be assumed. For any real-life sortition, more advanced implementations of the selection algorithm (such as [20] [21] [22]) should be considered, but for the secure MPC-sortition use case demonstration in this paper, a greedy solution suffices.

An outline of a greedy implementation is the following: after the strata have been defined by setting the lower and upper quotas for the different features, participants with the right characteristics are selected through simple random sampling until at least all features' lower quota are filled. Subsequently, the feature groups are filled by consecutively selecting a suitable participant for the feature group with the current lowest "fill / upper quota" ratio, stopping when k participants are selected. A full greedy implementation is available through the Sortition Foundation [23].

[2] A common issue is that mostly older people from the middle-income bracket register as they have the time and means available to free-up multiple weekends.

3.4 Requirements for Secure-MPC Sortition

Based on the interviews with the different stakeholders in the sortition process, a number of requirements were formulated which secure-MPC sortition should fulfil. These requirements were supplemented with the requirements elicited by [8] for user-friendly MPC designs.

Comprehensibility, meaning the ease with which the implementation of a design can be understood: For the involved local government and participants to trust a secure MPC-sortition design it must be simple to understand.

Auditability, meaning the ease with which the code of an implementation can be audited and inspected: For the involved local government, it would be best if the code of an implementation of a design is open-source to enable outside auditing. Additionally, the entire software should be able to run on servers of the organisers.

Accessibility, meaning the ease with which users can interact with a design: To minimise any hurdles that might discourage participation, secure MPC-sortition must require no set-up by the participants during the sign-up face, and no specialised software or hardware. The participants must be able to contribute data through a simple web interface. No technical expertise beyond this can be expected.

Asynchronicity, meaning the absence of a timing requirement for transmission: Participants should be able to contribute data at any point (before a certain registration deadline).

Idempotence, meaning the ease at which submission errors can corrected: Participants should have the opportunity to resubmit their data quickly after they uploaded if they discover they uploaded erroneous data.

Feedback, meaning the ability of a system to give feedback on erroneous input: Incorrectly formatted data from one contributor disturbs the entire final result. The interface should thus warn participant contributors proactively.

4 Related Work

Generally, few papers in the fields of computer science and statistics have been directly linked to sortition [24] [25]. This has started to change with the recent seminal work on sortition by Flanigan, Gölz and Procaccia [20] [21] [22], which has mostly focused on improving sortition based on different definitions of fairness. Additionally, open-source implementations of selection algorithms are available online [23] [26] [27].

To the best of our knowledge there has been no research into performing sortition or stratified sampling using MPC until now. Broadly speaking, however, secure MPC-sortition relates to the work done on secure MPC-questionnaires, as these require a user-friendly web-interface where a plurality

of clients without a technical background can input private data [28] [29] [30].

Notably, the protocol called 'Web-MPC', which was developed by [8] and [31], bears high similarity to the use cases in the scenarios as described in this paper. Web-MPC was developed for investigating the gender wage gap in the greater Boston area. The researchers put a strong emphasis on usability, giving the rationale that individuals and organisations without a technical background or large computing resources should be able to use the protocol. Many of its usability requirements can directly be translated into those for a secure MPC-sortition process.

However, one disadvantage of this protocol is that it only allows for a simple summation of numerical values. A slightly more advanced protocol for the underlying stratified sampling would thus be necessary as well. Something [28] demonstrated to be possible.

5 Building Blocks

The following section briefly explains how MPC works, followed by a discussion on existing MPC implementations which could be used to implement secure-MPC sortition.

5.1 Multi-Party Computation

In MPC, a set of n parties $P_1, ..., P_n$ desire to perform a joined calculation of function $y_i = f_i(x_1, ..., x_n)$ where $x_1, ..., x_n$ are the corresponding private inputs that need to stay private [32]. The security of such a system can be described through the *Real/Ideal Simulation Paradigm* [7]. In an *ideal world* the parties $P_1, ..., P_n$ could all submit their private inputs $x_1, ..., x_n$ individually to a trusted external party T. T would then compute the function $f_i(x_1, ..., x_n)$ and distribute result y_i to the input parties. However, in the *real world* no such trusted external party can be assumed to exist, meaning, the computing parties have to run a protocol amongst themselves emulating the ideal world. A protocol is considered secure if no more harm can be done by an adversary to a calculation in the real world compared to a calculation taking place in the ideal world.

An important aspect of this security definition is the number of parties that may be corrupted for the protocol to still perform the computation securely. Protocols exist for major honest majorities (2/3 of the parties is not corrupted), simple honest majorities and even honest minorities [7]. For this paper only the first is relevant as is explained in Section 6.2.

MPC protocols can be designed such that they are Turing complete, meaning that theoretically any calculation can be performed using MPC [7]. Damgård and Nielsen [33] put this more concretely by introducing and proving the idea that it is possible to translate any protocol into one using MPC using their idea of an arithmetic black box (ABB), which "can be thought of as a secure general-purpose computer. Every party can in private specify inputs to the ABB, and any majority of parties can ask it to perform any feasible computational task and make the result (and only the result) public." [33, p. 249].

MPC protocols with an honest majority often rely on a *secret sharing scheme* for their implementations [7]. Such a scheme is split into two algorithms, a randomised sharing algorithm and a recovery algorithm [29]. The sharing algorithm splits a secret value $s \in \mathbb{Z}_N$ into $s_1, ..., s_n$ shares which are send to the $P_1, ..., P_n$ computing parties. The shared value is recovered through a joint computation of the recovery algorithm by the computing parties. An example is the additive secret sharing scheme used in Sharemind [32], where a secret s is split into $s_1, ..., s_n \in \mathbb{Z}_{2^{32}}$ such that it can be reconstructed in the following way:

$$s_1 + s_2 + \cdots + s_n \equiv s \bmod 2^{32}. \quad (2)$$

5.2 Available MPC Implementations

The calculations necessary for stratified sampling are relatively simple (randomly picking names until the strata are filled), on relatively few data points (thousands rather than millions) and do not have to be calculated in real-time. This means that secure MPC-sortition does not need a specially optimised MPC protocol in order to work. Existing MPC frameworks and designs can thus be considered for its implementation.

Multiple off-the-shelf MPC deployments exist (Sharemind [32], Viff [34], FairplayMP [35], SEPIA [36]). Sharemind and Viff were considered to be the most promising and are briefly compared.

Sharemind is a commercial MPC service provider which makes it easy for users to run secure calculations [32]. Sharemind's framework can run all the basic arithmetic and relational database operations in a secure manner. Clients can program the specific protocol in SecreC, a C based Sharemind specific language.

In particular, the implementation of Sharemind by [28] could make the fast deployment of secure-MPC sortition possible. [28] implemented a web interface for contributors to input data. This implementation creates the secret shares at the client side before sending them to the main severs running Sharemind.

A disadvantage of the use of Sharemind is that its partially closed-source design inhibits proper auditing by a third party. This could potentially be troublesome for the involved local government which might desire to know the exact details of the implementation before submitting data.

Viff can be programmed to do the same calculations as Sharemind, with the notable difference that Viff is completely open-source. However, as [8, p. 4] point out: "Viff unrealistically requires all contributors, service providers, and analyzers to run the Viff software on mutually available servers". Making it difficult to use in a scenario where there are hundreds of contributors with varied non-technical backgrounds who want to submit data asynchronously. Sharemind, therefore, seems to be the best off-the-shelf solution currently available.

6 Design

In this section an overview is given of the necessary elements for two secure MPC-sortition designs that could resolve the issues outlined in Section 3.2. A description of the different parties involved in the designs is given as well as a reflection on the assumed security model.

6.1 Roles in Secure MPC-Sortition

With respect to the two scenarios described in Section 3.2 the following roles can be identified with respect to the data analysis process in sortition:

- A large number of personal data *contributors*. In Scenario 1, these are solely the participants who want to register themselves to be part of the deliberation event. The number of contributors is unknown beforehand but can range from hundreds for small events to in the ten-thousands in the case of the French Climate Convention. In Scenario 2, one of these contributors would be the local tax office, which would provide significantly more data, which also needs to be matched on the participants' self-registered data.
- A small number of publicly accessible *service providers* that receive the encrypted data from the contributors in the form of distributed secret shares. The service providers should be online during the whole process such that the contributors can send their data at different times. Three is found to be a good number of service providers in terms of speed and security [32].

  For Scenario 1, it is suggested that the organiser, the involved local government and an independent third party form this group. For Scenario 2, it is suggested that the organiser, and two different departments of the involved local government form this group.
- A small number of *analysers*, possibly one: the organiser of the deliberative event, who receive(s) the results and might do part of the calculation.

  The result would in both scenarios be the list of names (and contact details) of the most accurate assembly possible given the selection criteria and applicants.

6.2 Security Assumptions and Trust Relations

The analysers need to put a small amount of trust in the contributors regarding their honesty to input valid data, there is, however, no clear incentive for a contributor (a possible participant) to be untruthful about their registered details.[3] Vice versa, the contributors need to trust that the service providers and analysers will not collude. For this, it's important that a range of processing parties is chosen which lack an incentive for this.

When the service providers are well picked we can assume a *static semi-honest attack model* to be sufficient in both scenarios. *Static* means that parties do not switch from being corrupted to not being corrupted or vice versa [37]. As will be seen, a reasonable assumption given the suggested service providers. *Semi-honest adversarial behaviour* means that even corrupted parties adhere to the protocol and, as long as they do so, can never learn more than what can be inferred from the resulting data [7]. The reason for assuming this attack model is threefold and is related to the suggested service providers.

[3]It would require a considerable amount of guessing with regards to which group will register in smaller numbers to figure out an optimal strategy to fill in "fake" details such that a candidate's chance of being selected is higher.

Firstly, the organiser lacks a clear incentive to perform adversarial behaviour. Its prime interest is to only receive an accurate result in terms of a list of names that fulfil the selection criteria. The underlying data to get to this result is mostly only of instrumental importance to achieve this goal. It has an additional interest to create a safe environment for the participants to deliberate in, hence being honest is vital to create such an environment.

Secondly, the local government which has given the mandate for the deliberative event has little incentive to collude as well. In Scenario 1, it has little to gain from colluding because most information that is being gathered is of little value to a government, as it already has access to this information (e.g. address, name, gender, age etc.). In Scenario 2, colluding is even more directly opposed to its interests as leaking citizens' tax data would be quite disastrous for a number of reasons.

In Scenario 1, the third service provider can be chosen to further increase trust for the contributors. It should, however, be another organisation lacking a clear incentive to perform adversarial behaviour. An example would be Sharemind whose reputation as an MPC service provider would be severely jeopardised if it would collude.

In Scenario 2, however, due to the extra sensitive nature of tax data we suggest that another governmental department fulfils this role. Different governmental departments are often by law restricted in which data they are allowed to share, whilst from the government's perspective, contributing two service providers implies an honest majority is always guaranteed and no tax data can be leaked.

6.3 Design Overview

A brief overview is given regarding how the necessary elements for the two designs for secure MPC-sortition fit together. Figure 1 illustrates a step-by-step representation of these designs.

In **Design 1** existing methods from the sortition process can be used to send an invite to sign up online to a large number of possible participants (e.g. via a physical mail invite containing a sign-up URL). Subsequently, interested participants contribute their personal data through an online form and this data is turned into secret shares at the client-side. Client-side encryption is vital for the private data of the candidates to be unreadable by the service providers.

For the actual secret share creation, the protocol as outlined in [28] is a good off-the-shelf solution as explained in Section 4. Subsequently, the secret shares are sent to the service providers, which perform the stratified sampling calculation. The created shares can either be destroyed on reception or stored on the service provider servers for later processing.

Finally, the result of the calculation (i.e. a list of names) will be sent to the analysers (i.e. the organiser of the event) such that they can further process the results.

Design 2 differs only in the fact that one of the contributors is the involved local government. This party would need to create secret shares of the tax records of all the citizens that got a physical mail invite. Full name in combination with full address can be used as unique identifiers for citizens to match the data submitted by the government and the participants.

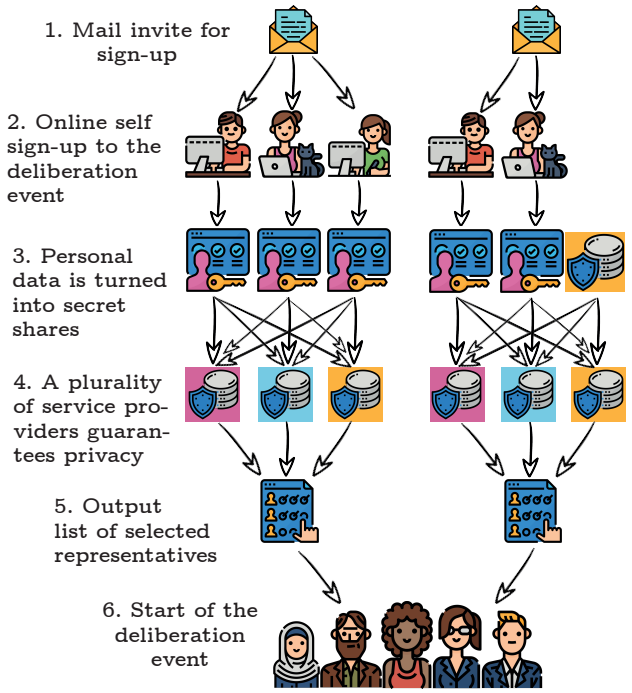

Figure 1: A step-by-step illustration of the two outlined designs. The figure on the left represents Design 1, the case where all information is gathered through the self sign-up. The figure on the right represents Design 2, the case in which the involved local government supplies some of the data.

The designs can be used in combination with each other as is illustrated in the steps on the right in Figure 1. Design 2, however, could also be used in combination with a classic sign-up phase, this would simply entail that all participant information is first gathered to a central server owned by the organiser. Effectively reducing the number of data contributors in the MPC phase to two.

7 Discussion and Future Research

In this research, a bridge was attempted to be made between innovations in cryptography and democracy. The secure MPC-sortition designs as outlined in the previous sections demonstrate an unconventional use case for MPC and contribute to the growing body of research on MPC in which a plurality of parties contributes data.

Design 1, as outlined in this paper, provides a way to improve privacy levels for the participants who register for deliberative events. Instead of having to put their trust in one organiser, participants instead put their trust in a set of parties with mutually incompatible reasons to collude. As such, the participants can be presented with a lightweight and user-friendly sign-up application in the form of a web interface whilst not having to run specialised hardware or software to have a stronger guarantee that their private data remains private. In this way, the presented designs make the sortition process more private.

Additionally, Design 2 allows the quality of the sortition itself to be improved by allowing governments to contribute tax data as a data source for the stratified sampling, making the process less dependent on participant supplied data and therefore more accurate. This also allows for the potential inclusion of participants who would otherwise have considered not to take part for the reasons mentioned in Section 3.2. In this sense, Design 2 makes the sortition process fairer.

An integral part of this research was to interview domain experts to investigate their ideas on sortition and the potential for MPC. These interviews also provided a way to gather feedback on the proposed designs. The most important feedback can be grouped into two themes: *the necessity of the improvements* and *the verifiability of the results*. It is immediately interesting to note the significant differences between the feedback given by different organisations. This can be explained by the diverse background of the interviewed organisations. Roughly speaking, four kinds of organisations were interviewed 1) *deliberative democracy event organisers*, involved in running a citizens' assembly or similar event itself, 2) *sortition organisations* (e.g. poling companies hired by a citizens' assembly organiser to do the sortition), 3) *activist organisations*, existing to promote the idea of citizens' assemblies or deliberative democracy in general, 4) *research institutes*.

Starting with the necessity of the improvements, one citizens' assembly organiser questioned the relevance of using MPC for sortition, stating that participants are happy with the level of privacy they are currently given through anonymisation and informed consent. Additionally, the absolute anonymity MPC could provide through the use of secret shares would also make further analysis of the gathered data almost impossible, something this organiser deemed unacceptable. Organisations focused on awareness and activism disagreed with this point of view. Stating that improved anonymity for participants outweighs the potential benefits of being able to process the gathered data for other purposes after the sortition is done. The sortition organisations and research institutes were more neutral towards this point, possibly because they have a less vested interest in the specific sortition data itself.

An interesting remark is that all of the interviewed experts indicated that no research has been done on the influence of anonymity in the sortition process on sign-up rates. Most indicated travel time to have the biggest influence on this statistic but one of the citizens' assembly organisers mentioned that at least some people had made inquires regarding the processing and storage of their personal data, with a could-have-been participant citing it as a reason not to participate, something which the organiser saw as an indication that there is a real concern about this from at least some of the potential participants. In Section 3.2 similar research on online surveys was mentioned. Future research could use this as a base and explore whether anonymity as guaranteed through MPC can indeed influence sign-up rates.

The second feedback theme, regarding verifiability, was mentioned by all but the interviewee with the most technical background. None of the interviewees had heard of MPC before and it took some time to explain the basic underlying cryptographic concepts guaranteeing the security of MPC.

Still, there were some concerns that the results of the joined calculation would not be verifiable or for that matter explainable towards participants or clients.

In practice, MPC implementations get their robustness from the mathematical proofs underlying them, as became clear from the interviews this is not obvious or intuitive for everyone. Therefore, if actual MPC-sortition implementations are used, great care needs to be taken to explain this robustness to participants and clients in layman's terms. This is especially important when local governments are involved, as political officials with a non-technical background will want to be able to understand how security and privacy regarding tax data are guaranteed.

A final comment which also bears relevance for future research is that the members of the polling companies were especially interested in Design 2. They saw a lot of potential for the combination of data from different databases holding more sensitive data since getting access to these databases is normally a difficult affair for them. Future research could focus on exploring how MPC can be used to aid poling companies in this effort.

8 Responsible Research

The processing of private data always introduces ethical considerations, as the leakage of such data has great implications for the contributor. The first design outlined in this paper has the potential to reduce the chances of such a leakage happening, making the sortition process more private for participants and therefore arguably more ethical.

Design 2 presents a different ethical consideration. Until now, governments had no incentive to participate by contributing their tax information for sortition as there was no secure way of doing so. MPC changes this but therefore also introduces new risks. Human error or malicious behaviour during the preprocessing of the data is a potential risk that MPC cannot solve but does introduce in this new scenario. The answer for how to solve this is not immediately clear but it is something that needs to be considered when switching to an MPC based sortition process.

The reproducibility of this study is relatively high. The literature used as a background for the described scenarios and the presented MPC designs is well referenced and accessible. The aspect of this study that is the most difficult to reproduce is the interviews. The names of the interviewees were not included due to privacy concerns. However, the names of the interviewed organisations were mentioned in the methodology section. These organisations all have general inquiry mail addresses which can be used as a start for reproducing the results.

9 Conclusion

This paper explored whether MPC can be used to make sortition more secure and fair by outlining two secure MPC-sortition designs and gathering feedback from domain experts on these designs. Design 1 makes it possible for citizens' assembly participants to contribute their data in a more private way, whilst Design 2 allows involved governments to contribute their tax data to allow a more accurate stratified sample to be drawn.

The opinion of experts regarding these improvements differed mostly depending on the role these experts had in the sortition process. Some citizens' assembly organisers found the current methods of anonymisation used for participant sign-up to be sufficient, whilst more activist organisations were interested in seeing an increase of privacy for participants. On the other hand, sortition organisations (such as poling companies) were particularly interested in the second design presented in this paper, seeing the potential for combining data from different databases in a private and secure fashion.

Future research could focus on exploring these two findings further. Firstly, by looking into the effects of anonymity on sign-up rates for deliberative events. Secondly, by exploring which other use cases could be of benefit to sortition organisations and poling companies with regards to the combination of different sensitive data sources using MPC.

References

[1] Roberto Foa, Andrew Klassen, Mike Slade, Alex Rand, and Kirsten Collins. The global satisfaction with democracy report 2020. 2020.

[2] Marcin Gerwin. Citizens' Assemblies – Democracy that works, 2021. *URL https://www.citizensassemblies.org Accessed June 2021.*

[3] John Gastil and Erik Olin Wright. Legislature by lot: envisioning sortition within a bicameral system. *Politics & Society*, 46(3):303–330, 2018.

[4] Peter Yeung. 'It gave me hope in democracy': how French citizens are embracing people power, 2021. *The Guardian, URL https://www.theguardian.com/world/2020/nov/20/it-gave-me-hope-in-democracy-how-french-citizens-are-embracing-people-power Accessed June 2021.*

[5] Irish abortion law: citizens' assembly recommends unrestricted access to terminations, 2017. *BBC News, URL https://www.bbc.com/news/world-europe-39687584 Accessed June 2021.*

[6] European Commission. 2018 Reform of EU Data Protection Rules, 2018. *URL https://gdpr-info.eu/ Accessed June 2021.*

[7] Yehuda Lindell. Secure Multiparty Computation (MPC). *IACR Cryptol. ePrint Arch.*, 2020:300, 2020.

[8] Andrei Lapets, Nikolaj Volgushev, Azer Bestavros, Frederick Jansen, and Mayank Varia. Secure MPC for analytics as a web application. In *2016 IEEE Cybersecurity Development (SecDev)*, pages 73–74. IEEE, 2016.

[9] Jan E Trost. Statistically nonrepresentative stratified sampling: A sampling technique for qualitative studies. *Qualitative sociology*, 9(1):54–57, 1986.

[10] Marcin Gerwin. Citizens' Assemblies - Guide to democracy that works, 2018. *URL https:*

//www.peoplepowered.org/resources-content/citizens-assemblies-guide-to-democracy-that-works Accessed June 2021.

[11] The Irish Citizens' Assembly. Report of the Citizens' Assembly on Gender Equality, 2021. *URL https://www.citizensassembly.ie/en/about-the-citizens-assembly/meetings/the-citizens-assembly-publishes-final-report-on-gender-equality/ Accessed June 2021.*

[12] Mark Diffley. The citizens' assembly of scotland: Recruitment of members, 2020. *URL https://www.citizensassembly.scot/about/assembly-members Accessed June 2021.*

[13] The Sortition Foundation. Scotland's Climate Assembly Recruitment Report, 2020. *URL https://www.sortitionfoundation.org/scotlands_climate_assembly_recruitment_report Accessed June 2021.*

[14] Extinction Rebellion Citizens' Assembly Working Group. The extinction rebellion guide to citizens' assemblies, 2019. *URL https://www.peoplepowered.org/resources-content/climate-assemblies-and-juries-yjy98 Accessed June 2021.*

[15] Cathlin V Clark-Gordon, Nicholas D Bowman, Alan K Goodboy, and Alyssa Wright. Anonymity and online self-disclosure: A meta-analysis. *Communication Reports*, 32(2):98–111, 2019.

[16] Adam N Joinson and Ulf-Dietrich Reips. Personalized salutation, power of sender and response rates to web-based surveys. *Computers in Human Behavior*, 23(3):1372–1383, 2007.

[17] Adam N Joinson, Alan Woodley, and Ulf-Dietrich Reips. Personalization, authentication and self-disclosure in self-administered internet surveys. *Computers in Human Behavior*, 23(1):275–285, 2007.

[18] Stephen A Rains. The implications of stigma and anonymity for self-disclosure in health blogs. *Health communication*, 29(1):23–31, 2014.

[19] Linda I Reutter, Miriam J Stewart, Gerry Veenstra, Rhonda Love, Dennis Raphael, and Edward Makwarimba. "Who do they think we are, anyway?": Perceptions of and responses to poverty stigma. *Qualitative Health Research*, 19(3):297–311, 2009.

[20] Bailey Flanigan, Gregory Kehne, and Ariel D Procaccia. Fair Sortition Made Transparent, 2021. *URL http://procaccia.info/publications/ Accessed June 2021.*

[21] Bailey Flanigan, Paul Gölz, Anupam Gupta, and Ariel Procaccia. Neutralizing self-selection bias in sampling for sortition. *arXiv preprint arXiv:2006.10498*, 2020.

[22] Gerdus Benadè, Paul Gölz, and Ariel D Procaccia. No stratification without representation. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 281–314, 2019.

[23] Paul Gölz and Brett Hennig. Sortition Foundation, 2019. *GitHub repository, URL https://github.com/sortitionfoundation/stratification-app Accessed June 2021.*

[24] Rene Saran and Norovsambuu Tumennasan. Whose opinion counts? implementation by sortition. *Games and Economic Behavior*, 78:72–84, 2013.

[25] Toby Walsh and Lirong Xia. Lot-based voting rules. In *AAMAS*, volume 12, pages 603–610. Citeseer, 2012.

[26] Nicolas Jacobeus and Dominique Lebrun. Belighted, 2020. *GitHub repository, URL https://github.com/belighted/decidim-module-castings Accessed June 2021.*

[27] Paul Gölz and Gili Rusak. Panelot, 2020. *URL https://panelot.org Accessed June 2021.*

[28] Riivo Talviste. Deploying secure multiparty computation for joint data analysis—a case study. Master's thesis, Institute of Computer Science, University of Tartu, 2011.

[29] Sven Laur, Jan Willemson, and Bingsheng Zhang. Round-efficient oblivious database manipulation. In *International Conference on Information Security*, pages 262–277. Springer, 2011.

[30] Meril Vaht. The analysis and design of a privacy-preserving survey system. Master's thesis, Institute of Computer Science, University of Tartu, 2015.

[31] Lucy Qin, Andrei Lapets, Frederick Jansen, Peter Flockhart, Kinan Dak Albab, Ira Globus-Harris, Shannon Roberts, and Mayank Varia. From usability to secure computing and back again. In *Fifteenth Symposium on Usable Privacy and Security ({SOUPS} 2019)*, 2019.

[32] Dan Bogdanov, Sven Laur, and Jan Willemson. Sharemind: A framework for fast privacy-preserving computations. In *European Symposium on Research in Computer Security*, pages 192–206. Springer, 2008.

[33] Ivan Damgård and Jesper Buus Nielsen. Universally composable efficient multiparty computation from threshold homomorphic encryption. In *Annual International Cryptology Conference*, pages 247–264. Springer, 2003.

[34] VIFF Developement Team. Viff, the virtual ideal functionality framework, 2009. *URL http://viff.dk/. Accessed June 2021.*

[35] Assaf Ben-David, Noam Nisan, and Benny Pinkas. Fairplaymp: a system for secure multi-party computation. In *Proceedings of the 15th ACM conference on Computer and communications security*, pages 257–266, 2008.

[36] Martin Burkhart, Mario Strasser, Dilip Many, and Xenofontas Dimitropoulos. SEPIA: Privacy-preserving aggregation of multi-domain network events and statistics. *Network*, 1(101101), 2010.

[37] Claudio Orlandi. Is multiparty computation any good in practice? In *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5848–5851. IEEE, 2011.