International Journal of Information Security
https://doi.org/10.1007/s10207-019-00474-7

REGULAR CONTRIBUTION

Random dictatorship for privacy-preserving social choice

Vicenç Torra[1,2]

Abstract

Social choice provides methods for collective decisions. They include methods for voting and for aggregating rankings. These methods are used in multiagent systems for similar purposes when decisions are to be made by agents. Votes and rankings are sensitive information. Because of that, privacy mechanisms are needed to avoid the disclosure of sensitive information. Cryptographic techniques can be applied in centralized environments to avoid the disclosure of sensitive information. A trusted third party can then compute the outcome. In distributed environments, we can use a secure multiparty computation approach for implementing a collective decision method. Other privacy models exist. Differential privacy and k-anonymity are two of them. They provide privacy guarantees that are complementary to multiparty computation approaches, and solutions that can be combined with the cryptographic ones, thus providing additional privacy guarantees, e.g., a differentially private multiparty computation model. In this paper, we propose the use of probabilistic social choice methods to achieve differential privacy. We use the method called random dictatorship and prove that under some circumstances differential privacy is satisfied and propose a variation that is always compliant with this privacy model. Our approach can be implemented using a centralized approach and also a decentralized approach. We briefly discuss these implementations.

Keywords Privacy · Social choice · Probabilistic social choice · Differential privacy · Random dictatorship

1 Introduction

Social choice provides mechanisms for making collective decisions. They aggregate individual preferences and lead to a summary of these preferences. Procedures for building collective decisions have been defined and studied for a long time.

Examples of well-known social choice mechanisms include Condorcet and Borda methods, and quite a few electoral systems. Condorcet and Borda methods are functions of ranks that return a single winner. Electoral systems start from individual preferences expressed in a variety of ways and then return either a single winner or a set of winners. Methods to aggregate preferences also exist. See, e.g., [3] for some reference works in this area.

Arrow's impossibility theorem [2] showed that when designing a method for aggregating preferences, a very small set of requirements makes the problem overconstrained. That is, no method can be defined satisfying all requirements. Solutions exist when some of these requirements are dropped. For example, methods exist for restricted domains (e.g., single-peaked preferences) or, alternatively, methods do not satisfy the condition of independence of irrelevant alternatives. Gibbard–Satterthwaite theorem [10,21] is another impossibility theory for ordinal electoral systems with a single winner.

Impossibility results have fostered the development of solutions that try to minimize or overcome their negative effects. The study of single-peaked preferences is an example. When preferences are restricted to be of this type, there are social choice mechanisms that satisfy all the axioms in Arrow's theorem.

Probabilistic social choice is a sub-field of social choice in which the outcome of a mechanism is computed in terms of a probability distribution. Interesting mechanisms have been introduced. They can circumvent [5,7,19] the difficulties of the Gibbard–Satterthwaite theorem, and they satisfy some well-known properties of social choice. In this paper, we will use methods belonging to this area of probabilistic social choice. We will show that they can provide, under some conditions, differential privacy.

✉ Vicenç Torra
vtorra@ieee.org

[1] Hamilton Institute, Maynooth University, Maynooth, Ireland

[2] School of Informatics, University of Skövde, Skövde, Sweden

Published online: 16 October 2019

Social choice procedures are used for multiple purposes in multiagent systems. In this area, however, preferences and votes are sensitive information.

The goal of data privacy [25,28] is to provide mechanisms to avoid the disclosure of sensitive information. On the one hand, it provides a set of privacy models that provides computational definitions of privacy. On the other hand, it provides techniques so that systems and databases are compliant with the privacy models, and/or measures to evaluate different aspects related to privacy (e.g., disclosure risk and information loss measures).

In this paper, we study the use of probabilistic social choice as a way to implement private voting. As we discuss later, this permits to obtain a voting mechanism that satisfies differential privacy. The literature presents different alternatives with respect to privacy and voting. We review them in the next section and give more details on our goals.

1.1 Related work

The most straightforward way to implement any voting procedure is to follow a centralized approach using a trusted third party and use cryptography so that no information is leaked, and the only information learnt by the agents is the outcome of the function (i.e., the result of the voting procedure). Most secure e-voting protocols [15] follow this approach. Nevertheless, this approach has two main drawbacks.

- One is that it needs this trusted third party and it is, thus, a centralized approach not much appropriate in a multiagent system context. This problem can be solved using a distributed (decentralized) approach, i.e., using secure multiparty computation as our privacy model. However, in this latter case, computational costs are higher, and they can be prohibitive if the function to compute is complex and the number of parties is very large. Examples of distributed secure e-voting protocols are found in [9,20,29,30].
- Another drawback is that the outcome can give clues on the participation of certain agents. Differential privacy is an alternative privacy model that focuses on the result of a calculation (i.e., the function or query computed from a database). It is about this result being independent of the presence or absence of a single agent. This is to ensure the confidentiality on the participation of any agent. Note that secure e-voting does not avoid this type of disclosure.

Other privacy models exist as well. For example, k-anonymity, initially defined for databases, is about modifying (perturbing) a database so that for any record there are other $k-1$ records that are indistinguishable. Or, from another perspective, that for any record the anonymity set (the set of respondents that may have supplied the same information) is of size at least k. In the case of preferences of agents, this would correspond to having at least k agents with the same preferences.

Note that differential privacy and k-anonymity are not in contradiction with the privacy models that use cryptography. For example, we can have a system implementing both the secure multiparty computation model and the differential privacy model.

In this paper, we focus on privacy according to the differential privacy model. In Sect. 5.1, we discuss briefly its implementation in combination with the trusted third party and the multiparty computation model to have additional privacy guarantees.

We study mechanisms for ensuring privacy in the application of social choice methods for multiagent systems. The literature on this topic is limited. In [16–18], authors study privacy for mechanism design. They focus on the differential privacy model. In particular, [16] focuses on auctions, [17] on monopolist pricing and facility location, and [18] includes a differentially private poll. The approach in [18] is described as also following the secure multiparty computation model. Chen et al. [8] follow a similar approach, also in the context of multiagent systems. The paper proposes a mechanism (Mechanism 4.1) satisfying differential privacy for a two-candidate election.

In [24] (Chapter 5), k-anonymity is studied for preferences. The goal is to build sets of at least k preferences and build an aggregated preference for each set. To build these sets, two different types of distances between preferences are considered: discrete distance (where any pair of preferences is at a distance one) and swap distance (where a preference is at a distance one of another that only differs on a swap between consecutive options). The author considers two voting rules: the plurality rule and the Condorcet rule.

Our approach on privacy for social choice is different. We consider the differential privacy model in combination with probabilistic social choice. Probabilistic social choice is naturally based on randomness. The use of randomness in elections has an old history (see, e.g., [7]). Sortition or allotment as a way to elect officials was used already in ancient Athens and in medieval ages (e.g., North Italy and Catalonia). Probabilistic social choice models have been recently reconsidered. They have interesting properties and, as we have stated above, they are a way to circumvent the Gibbard–Satterthwaite theorem. From our perspective, probabilistic social choice mechanisms fit well with differential privacy, as methods compliant with differential privacy are typically implemented in terms of randomization. Although not all probabilistical social choice methods are differentially private, some satisfy this property under certain circumstances.

Random social choice is naturally not a valid procedure for all types of contexts in which decisions are made within multiagent systems. Nevertheless, there are situations in which this type of procedures is useful. They are cases where randomness does not have a dramatic effect and can even be beneficial. Consider the case of simulated annealing and other search procedures that include a randomization component. In our case, we are particularly interested in its application in the context of reinforcement learning and federated learning. This corresponds to situations in which the decision is not final but an iterative process. In this case, randomized social choice can be useful. Actual application in a real environment is out of the scope of this paper.

For example, consider a cohort of drones guiding a terrestrial vehicle. Drones vote continuously for landmarks that are used to guide the vehicle. The selected landmarks at any time are not so important as it is the overall set (a rough path) that will really affect the trajectory of the vehicle. Under this scenario, we would consider local differential privacy or differential privacy with a budget so that each drone can vote several times.

Similar scenarios can be considered in federated learning. In this case, mobile devices of users need to affect the learned model by means of their vote, but it is the long run what is of relevance, not the result of the voting at a particular time.

In this paper, we study random dictatorship in light of differential privacy, and then we introduce a variation of this method that satisfies differential privacy. Our definition follows a centralized approach that can be thus combined with a secure multiparty computation protocol to avoid the disclosure of the votes of the agents. Improving in this way the privacy is guaranteed by differential privacy.

1.2 Structure of the paper

The structure of this paper is as follows. In Sect. 2, we review briefly probabilistic social choice. In Sect. 3, we review the definition of differential privacy. In Sect. 4, we study random dictatorship in light of differential privacy and prove the main theorem of the paper. In Sect. 5, we introduce a differentially private method for random dictatorship. The paper finishes with some conclusions and lines for future work.

2 Probabilistic social choice

In this section, we review a few methods of probabilistic social choice. We will consider a set I of T agents and a set A of m alternatives. Each agent $i \in I = \{1, \ldots, T\}$ has associated a preference relation $\succeq_i$ on the set of alternatives.

All preference relations $\succeq_i$ are subsets of $A \times A$, and they are reflexive, complete and transitive. Recall that a relation $\succeq_i$ on A is reflexive if for all $a_r \in A$ it holds $a_r \succeq_i a_r$, it is complete if for any $a_r \in A$ and $a_s \in A$ we have either $a_r \succeq_i a_s$ or $a_s \succeq_i a_r$, and it is transitive if for any $a_r, a_s, a_t \in A$ it holds that $a_r \succeq_i a_s$ and $a_s \succeq_i a_t$ implies $a_r \succeq_i a_t$. As usual we use $\succ_i$ to denote the corresponding relation that is not reflexive (i.e., $a_r \succ_i a_s$ if and only if $a_r \succeq_i a_s$ but $a_r \neq a_s$).

Using the example in the introduction, I is the set of drones and T the number of drones. A corresponds to the set of possible landmarks. At any time, each drone i has a preference on the landmarks and this defines the preference relation $\succeq_i$ (i.e., $a_r \succeq_i a_s$ if the ith drone prefers the landmark r to the landmark s).

2.1 Random dictatorship

Given the set of agents I with the corresponding preference relations $\succeq_i$ for $i \in I$ on the alternatives A, uniform random dictatorship proceeds as follows.

Method 1 *It selects an agent i in I according to a uniform distribution on I, and then uses $\succeq_i$ to select the most preferred alternative by agent i as outcome.*

This method has been studied by several authors, and, for example, there is a theorem that characterizes it in terms of anonymity, strong SD-strategyproof and ex post-efficiency when preferences are strict. See [7,11]. The method as explained here is well defined for strict preferences. There are variations that allow for non-strictness. We do not consider them in this paper.

2.2 Other methods

There are other methods based on lotteries. See, e.g., [6,7] for details. There are methods based on Borda's rule, e.g., random selection among the alternatives with a maximal score, or random selection proportional to Borda scores. There are also methods based on maximal lotteries.

3 Differential privacy

Differential privacy is to ensure that the outputs of a function (or query) do not depend too much on a particular record or individual. In order to ensure this, the outputs of the function when applied to two databases that only differ in one record should be similar enough. The formalization assumes that the function is not deterministic but that we have a probability distribution on the space of outputs, and then it is expected that the two probability distributions obtained from the application of the function to the two databases are similar.

We give now its definition. In the definition, D_1 and D_2 are the two arbitrary databases that differ only in one record. Then, $A(D_1)$ is the application of the algorithm or function to D_1 and S is a subset of the space of outputs of the algorithm.

The definition depends on a parameter ϵ that permits us to control the level of privacy. A large ϵ permits that the function is changed significantly when a new record is added. In Sect. 4, we discuss briefly some values of ϵ used in the literature. Note that in general the selection of ϵ is an open problem.

Definition 1 An algorithm A gives ϵ-differential privacy if for all data sets D_1 and D_2 differing at most one element, and all $S \subseteq \text{Range}(A)$, it holds

$$\frac{\Pr(A(D_1) \in S)}{\Pr(A(D_2) \in S)} \leq \mathrm{e}^{\epsilon} \tag{1}$$

In this definition, ϵ is the parameter to control the privacy level. The smaller the ϵ, the greater the privacy.

When a set of agents I is considered, and we are interested in computing a function of agents opinions (i.e., $\succ_i$), we can consider two frameworks. The first one is as described in the definition of differential privacy above that D_1 and D_2 can differ in a record or agent. This implies that agents can opt out to take part in the collective decision-making process.

Alternatively, we may consider that all agents in our multiagent system are required to take part in the collective decision-making process. In this case, D_1 and D_2 will have the same size (i.e., the same set of agents) but there is one agent that differs in its contribution (i.e., $\succ_i$ has changed). This later approach is similar to the one in [17].

Consider the running example presented in the introduction. The first case is that in each voting process only some of the drones are available for voting. Here, *availability* can mean that they can communicate to each other, or that they have a preference on a landmark and want to communicate this preference for voting. It can also be the case that some drones decide not to participate because their privacy budget is exhausted, or because their internal budget management rules decide not to participate. In the second case, all drones vote in all voting processes. In the scenario for federated learning also mentioned in the introduction, both cases are also relevant. We may consider the case that only a subset of devices is used for computing partial decisions, and the case that all devices are required to update a machine learning model.

4 Differential privacy for probabilistic social choice functions

We study in this section conditions on the satisfaction of differential privacy. We begin considering random dictatorship (see Sect. 2.1) in the case that agents can decide whether they participate or not. The result assumes that the voting procedure takes place if at least T agents participate. Later, we consider the case where participation is mandatory.

The following theorem shows that although in general differential privacy is not satisfied, there are scenarios in which it holds. We will use these results later in Sect. 5 to define a method that always satisfies differential privacy.

Theorem 1 *Let us consider a set I of T agents, a set A with at least two alternatives, and preferences $\succ_i$ for $i \in I$. Then, random dictatorship satisfies the following properties with respect to differential privacy.*

- (i) *The procedure does not satisfy differential privacy in general.*
- (ii) *When there is at least one agent that supports each alternative, differential privacy is satisfied for any ϵ larger than*

  $$\log\left(\frac{2T}{T+1}\right) \tag{2}$$

  where $T \geq 3$ is the number of agents.
- (iii) *When there is at least an $\alpha\%$ of agents that support the least represented option, we have*
  - (iii.a) *that for a set of agents of size T, the method satisfies ϵ-differential privacy for ϵ s.t.*

    $$\epsilon \geq \log \frac{\alpha T + 100}{\alpha(T+1)} \tag{3}$$

  - (iii.b) *that ϵ-differential privacy is satisfied for any set of agents with a size larger than*

    $$size \geq \frac{100 - \alpha \mathrm{e}^{\epsilon}}{\alpha(\mathrm{e}^{\epsilon} - 1)},$$

  for any $\alpha \leq \frac{100T}{2T+1}$.

Proof To prove this theorem, we consider $A(D)$ as the algorithm for random dictatorship. Therefore, the range of $A(D)$ is the set of possible outcomes, i.e., the set of alternatives. We assume that there are at least two alternatives.

The first consideration is that we do not need to deal with all sets S, but only to consider that it holds for alternatives a_0 in A. That is, we have ϵ-differential privacy when

$$\frac{\Pr(A(D_1) = a_0)}{\Pr(A(D_2) = a_0)} \leq \mathrm{e}^{\epsilon}$$

for all $a_0 \in A$. This is so because if this equation holds for a_0 and a_1, then it holds for $S = \{a_0, a_1\}$ and by induction for all $S \subseteq A$.

The definition of differential privacy requires that Eq. 1 is satisfied for all set of agents D_1 and D_2 that differ in one agent. We will consider two cases, depending on whether D_1 is the larger or the smaller set. Using $D \cup \{a\}$ to denote the set of agents D enlarged with one additional agent, these two cases correspond to the following equations

- *Case 1*

$$\frac{\Pr(A(D \cup \{a\}) = a_0)}{\Pr(A(D) = a_0)} \leq \mathrm{e}^{\epsilon}.$$

- *Case 2*

$$\frac{\Pr(A(D) = a_0)}{\Pr(A(D \cup \{a\}) = a_0)} \leq \mathrm{e}^{\epsilon}.$$

Now, taking into account the algorithm for random dictatorship, $\Pr(A(D) = a_0)$ corresponds to the proportion of agents that have a_0 as their first option. Let N_{a_0} be the number of agents with such first option. Then, $\Pr(A(D) = a_0) = N_{a_0}/T$ where T is the number of agents in D. That is, $T = |D|$. We can obtain the following expressions for cases 1 and 2 when the added agent does not select a_0.

- *Case 1a*

$$\frac{N_{a_0}/(T+1)}{N_{a_0}/T} = \frac{N_{a_0}T}{N_{a_0}(T+1)} \leq \mathrm{e}^{\epsilon}.$$

- *Case 2a*

$$\frac{N_{a_0}/T}{N_{a_0}/(T+1)} = \frac{N_{a_0}(T+1)}{N_{a_0}T} \leq \mathrm{e}^{\epsilon}.$$

We can summarize these two inequalities as a condition comparing the size of the set of agents and ϵ. The equation is as follows.

$$\frac{1}{\mathrm{e}^{\epsilon}} \leq \frac{T}{T+1} \leq \mathrm{e}^{\epsilon}. \tag{4}$$

In contrast, when the added agent selects a_0 we can deduce the following inequalities:

- *Case 1b*

$$\frac{(1+N_{a_0})/(T+1)}{N_{a_0}/T} = \frac{(1+N_{a_0})T}{N_{a_0}(T+1)} \leq \mathrm{e}^{\epsilon}.$$

- *Case 2b*

$$\frac{N_{a_0}/T}{(N_{a_0}+1)/(T+1)} = \frac{N_{a_0}(T+1)}{(N_{a_0}+1)T} \leq \mathrm{e}^{\epsilon}.$$

It is easy to see that the most sensible case is when N_{a_0} is equal to zero and, thus, it corresponds to the case that we add an agent that selects an alternative that has not been considered before by any other agent. The probability of this alternative being selected will change from zero to $1/(T+1)$, and, thus, the equation of case 1b does not hold for any ϵ and, therefore, no ϵ satisfies differential privacy. This proves the first result [case (i)].

In order to prove the second result [case (ii)], we consider the most extreme cases that are when $N_{a_0} = 1$ and when $N_{a_0} = T - 1$. It is easy to see that case 1b when $N_{a_0} = 1$ corresponds to:

$$\frac{2T}{T+1} \leq \mathrm{e}^{\epsilon} \tag{5}$$

and, thus, it is implied by Eq. 2.

In this case, we do not need to consider case 2b because (2b) with $N_{a_0} = 1$ corresponds to $(T+1)/(2T) \leq \mathrm{e}^{\epsilon}$ which for $T \geq 1$ is true for any $\epsilon \geq 0$.

When $N_{a_0} = T - 1$, case 1b leads to $T^2/((T-1)(T+1)) \leq \mathrm{e}^{\epsilon}$. This function is decreasing and with $T = 3$ equals 0.117. In contrast, $2T/(T+1)$ is increasing and with $T = 3$ is 0.405. Therefore, $T \geq 3$ satisfying Eq. 2 will also satisfy case 1b. Case 2b is always true: $(T-1)(T+1)/T^2 = (T^2 - 1)/T^2 < 1 \leq \mathrm{e}^{\epsilon}$ for $\epsilon \geq 0$.

Let us now consider the proof of Eq. 4 (i.e., cases 1a and 1b). Observe that now the value of N_{a_0} does not play a role. First, we have that Eq. 2 $\mathrm{e}^{\epsilon} > 2T/(T+1)$ implies the right-hand side of Eq. 4 because

$$\frac{T}{T+1} \leq \frac{2T}{T+1} < \mathrm{e}^{\epsilon}. \tag{6}$$

Equation 2 is equivalent to $1/\mathrm{e}^{\epsilon} < (T+1)/2T$. Naturally if $(T+1)/2T < T/(T+1)$, then the left-hand side of Eq. 4 is proven. To prove that Eq. 2 implies $(T+1)/2T < T/(T+1)$ for $T \geq 3$, just rewrite it as $0 < 2T^2 - (T+1)(T+1) = 2T^2 - T^2 - 2T - 1 = T^2 - 2T - 1$. So, the equality holds for $T = 1 \pm \sqrt{2}$ and it naturally holds for $T \geq 3$. This completes the proof of the second result of the theorem [i.e., (ii) above].

To prove the third result [i.e., (iii) above], we consider again the two cases (a) and (b) taking N_{a_0} as $\alpha T/100$. For case (a), we need to see (case 2a) that $(T+1)/T \leq \mathrm{e}^{\epsilon}$ holds with $\epsilon \geq (\alpha T + 100)/(\alpha(T+1))$. Then, if

$$\frac{T+1}{T} \leq \frac{\alpha T + 100}{\alpha(T+1)} \leq \mathrm{e}^{\epsilon}$$

the case is proven. As

$$\frac{T+1}{T} \leq \frac{\alpha T + 100}{\alpha(T+1)}$$

can be rewritten as $\alpha \leq 100T/(2T+1)$ and this is precisely the condition for α in (iii), the case is proven.

The other condition (i.e., case 1a) $T/(T+1) \leq \mathrm{e}^{\epsilon}$ holds as $T/(T+1) \leq 1 \leq \mathrm{e}^{\epsilon}$ for $\epsilon \geq 0$.

For case b, case (2b) is trivial because

$$\frac{\alpha(T+1)}{\alpha T + 100} = \frac{\alpha T + \alpha}{\alpha T + 100} \leq 1 \leq \mathrm{e}^{\epsilon}$$

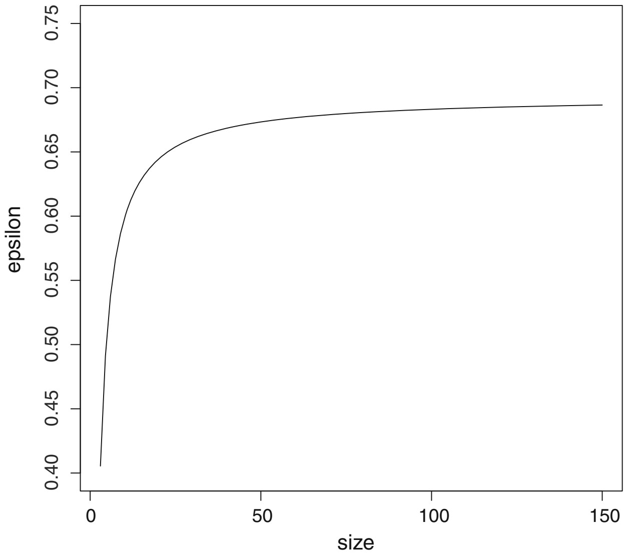

Fig. 1 Values of ϵ according to Eq. 2 (i.e., $\epsilon = log(2 * t/(t + 1))$

because $\alpha \leq 50 < 100$.

Case (1b) is just the condition of the third result ($(\alpha T + 100)/(\alpha(T + 1)) \leq e^{\epsilon}$.

This completes the proof of (iii.a). Condition (iii.b) about the size is obtained rewriting the one for ϵ and observing that $\alpha(e^{\epsilon} - 1)$ is always positive for $\epsilon \geq 0$. □

We illustrate now this theorem with some examples.

We need to consider the problem of setting the parameter ϵ. This is an open problem in differential privacy. It is often used an ϵ below one for an effective privacy level. We recall that in [14], values for ϵ were recommended when computing mean and median. These values were $\epsilon = 0.3829$ and $\epsilon = 2.776$, respectively. In addition, according to [4], Apple uses in some applications an $\epsilon = 4$, and according to [12] uses $\epsilon = 6$ and $\epsilon = 14$. Recall that the larger the ϵ, the lesser the protection.

The second result of this theorem implies that in a multiagent system, a small ϵ can be obtained (under the assumption that there is at least one agent supporting each alternative). Note that when the number of voters is 10, we have differential privacy with $\epsilon = 0.5978$. For large number of voters, ϵ tends to be $\log(2) = 0.6931$. Figure 1 shows different ϵ for a set of voters between 3 and 150. This result can also be seen in the other way round, i.e., the appropriate size for a given ϵ.

In relation to the third result of this theorem, note that the condition of $\alpha \leq 100T/(2T + 1)$ (this function is displayed in Fig. 2) is a necessary condition when considered together with Eq. 3, but this is not a sufficient condition. Any ϵ larger than $\log[(T + 1)/T]$ will be valid for all α. In this result, we assume that there is always a minimum percentage for the alternative with the minimum support. This implies that

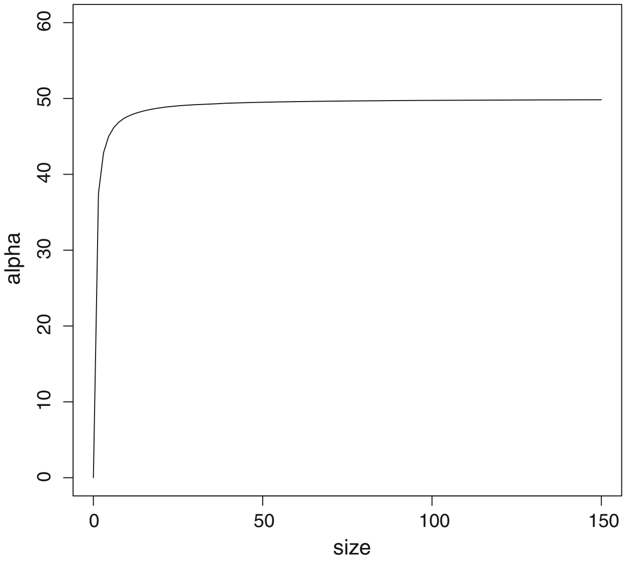

Fig. 2 Upper bound for valid values of α according to case (iii) in Theorem 1 [i.e., $\alpha \leq 100T/(2T + 1)$]

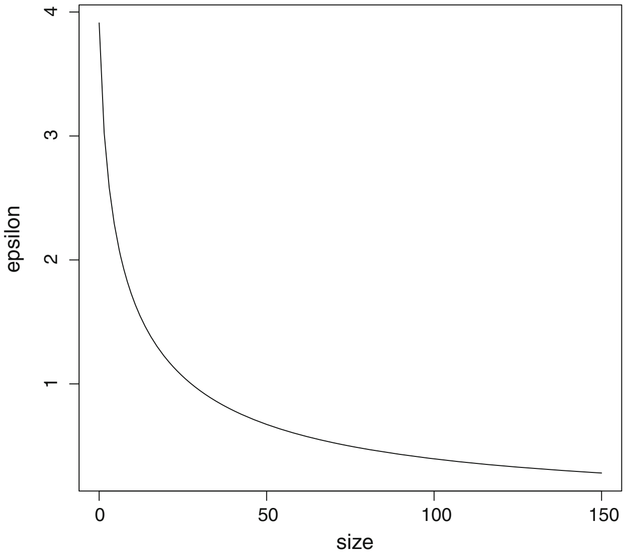

Fig. 3 Values of ϵ according to Eq. 3 for $\alpha = 2\%$ [i.e., $log((2 * t + 100)/(2 * (t + 1)))$]

the shape of the function is different from the previous result where the percentage of supporting agents is decreasing when the number of agents increases. Figure 3 shows the values of ϵ for $\alpha = 2\%$ (and for different number of agents). The same applies when a larger number of votes is obtained by the option with less support.

These results show that while differential privacy is not always satisfied, under the assumption that there is at least one vote per alternative, differential privacy is guaranteed with a reasonable ϵ.

We have discussed in Sect. 3 that we can consider two frameworks in the case of agents voting in a multiagent system. One is that the agents can opt out in the voting process and another in which this is not possible and all are required to vote. Theorem 1 considers the case that participation was not mandatory. Let us consider now the case that all agents are required to vote.

Theorem 2 *Let us consider the same assumptions as in Theorem 1. Then, when agents cannot opt out, the following holds for the random dictatorship.*

- *The procedure does not satisfy differential privacy in general.*
- *When there is at least one agent that supports each alternative, differential privacy is satisfied for any ϵ larger than $\log(2) = 0.6931$.*
- *When there is at least an $\alpha\%$ of agents that support the least represented option, we have*
  - *that for a set of agents of size T, the method satisfies ϵ-differential privacy for ϵ s.t.*

$$\epsilon \geq \log \frac{\alpha T + 100}{\alpha T} \tag{7}$$

  - *that ϵ-differential privacy is satisfied for any set of agents with a size larger than*

$$size \geq \frac{100}{\alpha(e^{\epsilon} - 1)}. \tag{8}$$

Proof In order to prove this result, we can reconsider the proof of the previous theorem. We can use the example used in the proof of the first statement in the previous theorem to prove the first statement here (i.e., $N_{a_0} = 0$).

The proof of the second statement uses the expressions for cases 1 and 2. Cases 1a and 2a are always satisfied as they correspond to the statement $1 \leq e^{\epsilon}$. Then, case 1b reduces to

$$\frac{1 + N_{a_0}}{N_{a_0}} \leq e^{\epsilon}.$$

The most extreme case for this expression is when $N_{a_0} = 1$ and $N_{a_0} = T - 2$, and the maximal value for ϵ will be at $N_{a_0} = 1$ where it results $\log(2) \leq \epsilon$. Values for large N_{a_0} lead to smaller ϵ. In contrast, Case 2b is always true. Therefore, the second statement of the theorem is proven.

The proof of the third statement follows the same structure considered in the previous theorem. For cases 1b and 2b, we obtain the following expressions:

- Case 1b

$$\frac{\alpha T + 100}{\alpha T} \leq e^{\epsilon}.$$

- Case 2b

$$\frac{\alpha T}{\alpha T + 100} \leq e^{\epsilon}.$$

Then, as $\alpha T \leq \alpha T + 100$, Case 2b is true for any $\epsilon \geq 0$. So, we only need to consider the expression in Case 1b, and from it, we get Expression 7. Similarly, we obtain Expression 8. □

This theorem shows that this framework permits to achieve differential privacy with a low value for ϵ but that this value does not change with the number of agents. That is, $\epsilon = \log(2) = 0.6931$ when there is at least one voter for each alternative. In contrast, Eq. 7 is similar to the one obtained in the previous theorem.

5 Differentially private random dictatorship

The results obtained in the previous section analyzing random dictatorship in light of differential privacy permit us to introduce a simple variation of the method that satisfies differential privacy.

The procedure is as follows. We call this method differentially private random dictatorship.

Method 2 *Let $A = \{a_1, \ldots, a_m\}$ be the set of alternatives. Then, given the set of agents I with the corresponding preference relations $\succeq_i$ for $i \in I$ on the alternatives A, enlarge I with a set of agents $I_0 = \{e_1, \ldots, e_m\}$ such that $\succ_i$ for $i \in I_0$ has as its preferred alternative the ith alternative in A.*

Then, apply uniform random dictatorship on $I \cup I_0$.

We have the following result for this method.

Lemma 1 *Differentially private random dictatorship satisfies differential privacy for any $\epsilon \geq \log \frac{2|I \cup I_0|}{|I \cup I_0| + 1}$.*

The proof of this lemma is a consequence of the second condition in the results of the previous section (Expression 2).

The procedure we have introduced here can also be applied when agents' participation to the voting is compulsory. This is defined as follows. We call this method differentially private compulsory random dictatorship.

Method 3 *Define one agent per alternative as in Method 2, and then select one among the previously existing ones union the ones created, and then use their decision.*

The following lemma is a consequence of the corresponding theorem proven in the previous section.

Lemma 2 *Differentially private compulsory random dictatorship satisfies differential privacy for any $\epsilon \geq 0.6931$.*

5.1 Implementation

The implementation of differentially private random dictatorship using a trusted third party can be done using cryptography and secure communication techniques. This will result into a method that will protect both the privacy of the votes and the participation or not of an agent (i.e., satisfy differential privacy). Nevertheless, this approach will be centralized.

In order to implement these methods in a distributed way so that they satisfy not only differential privacy but also the secure multiparty model, we can use some methods that have already been defined in the literature.

A detailed implementation and a security proof is out of the scope of this paper. Some hints for its implementation follow. On the one hand, we can first build a random order of the agents. See, e.g., [1,22] that consider a secret ordering of the agents. For a non-secret ordering, see, e.g., [23]. Then, for a generation of a random number, see, e.g., [13]. Finally, we will need to get privately the answer of the selected agent (a private selection protocol). Private information retrieval methods (see, e.g., [25]) could also be used to provide a partially distributed approach.

6 Summary and conclusions

In this paper, we have considered a method for collective decision that belongs to the area of probabilistic social choice in light of differential privacy. That is, we want to avoid disclosing the participation of a certain agent. At the same time, we want that the voting procedure has a low information loss / large utility, in the sense that the preferences of the agents are taken into account. We have proven that under some conditions, a probabilistic social choice method satisfies differential privacy for appropriate values of ϵ. We have introduced Methods 2 and 3 that provide differential privacy (see Lemmas 1 and 2, respectively).

Then, taking advantage of the results obtained, we have defined a variant that satisfies differential privacy. We have proven that some methods of probabilistic social choice integrate well with differential privacy. Not all methods within probabilistic social choice are suitable for differential privacy. For example, Borda$_{\max}$ (see [7]) that selects alternatives randomly among the ones with a maximal Borda count seems to be unsuitable for achieving differential privacy (or may require an important modification that can invalidate known theoretical results).

Our goal is to provide differentially private solutions that are also useful and sound from a social choice perspective. For this purpose, we can consider two approaches. One is to adapt existing classical social choice methods, which may result into new mechanisms that in the adaptation process will probably lose the nice properties that the original mechanisms satisfy. The other is to start with probabilistic social choice methods. If minimal or no transformation is needed to satisfy differential privacy, the methods will still keep the properties of the original method. We consider that the second option is preferable, and it is the one considered in this paper.

Solutions for data privacy are often compared in terms of the trade-off between utility and privacy. In differential privacy, once the ϵ parameter is settled, competing methods are evaluated in terms of utility. In our context, we consider that utility will be strongly dependent on the application, and, thus, difficult to evaluate. It is relevant to observe that in voting, it is often assumed that voters possess ordinal utility functions. That is, it is the order between alternatives that is relevant, and when *numbers* are used for ranking, they can only be compared with respect to order but neither added nor subtracted. Note that an arbitrary assignment of numerical utilities consistent with preferences, majority voting does not necessarily maximize the sum of the utilities (i.e., an alternative considered by all as second can accumulate more utility than the first alternative of the majority). If cardinal scales are used for utilities, aggregation and selection (as in multicriteria decision making) are preferable to voting [26,27]. Experimentation in different scenarios as the ones mentioned in the introduction is left for future work.

As future work, we will also consider other methods in the area, e.g., applying our approach to the method that selects alternatives according to a probability distribution proportional to Borda scores. We also consider the case of coalitions of agents to infer opinions from others.

Acknowledgements Open access funding provided by University of Skövde. Partial support of the project Swedish Research Council (Vetenskapsrådet) (Grant No. VR 2016-03346) is acknowledged. Discussion with Guillermo Navarro-Arribas is gratefully acknowledged.

Funding This study was funded by the Swedish Research Council (Vetenskapsrådet) (Grant No. VR 2016-03346).

Compliance with ethical standards

Ethical approval This article does not contain any studies with human participants or animals performed by any of the authors.

References

1. Adida, B., Wikström, D.: How to shuffle in public. Proceedings of TCC 2007. LNCS 4392, pp. 555–574 (2007)

2. Arrow, K.J.: A difficulty in the concept of social welfare. J. Political Econ. **58**(4), 328–346 (1950)
3. Arrow, K.J., Sen, A., Suzumura, K.: Handbook of Social Choice and Welfare, Volume 1, Volume 2 (2011). North Holland, Amsterdam (2002)
4. Apple. Apple Differential Privacy Technical Overview. (2017). https://images.apple.com/au/privacy/docs/Differential_Privacy_Overview.pdf. Accessed 14 Oct 2019
5. Barberà, S.: Majority and positional voting in a probabilistic framework. Rev. Econo. Stud. **42**(2), 379–389 (1979)
6. Brandl, F., Brandt, F., Seedig, H.G.: Consistent probabilistic social choice. Econometrica **84**, 1839–1880 (2016)
7. Brandt, F.: Rolling the dice: recent results in probabilistic social choice. In: Endriss, U. (ed.) Trends in Computational Choice, AI Access, pp. 3–26 (2017). Amsterdam, NL
8. Chen, Y., Chong, S., Kash, I.A., Moran, T., Vadhan, S.: Truthful mechanisms for agents that value privacy. ACM Trans. Econ. Comput. 4:3, article 13 (2016)
9. Dini, G.: Electronic voting in a large-scale distributed system. Networks **38**(1), 22–32 (2001)
10. Gibbard, A.: Manipulation of voting schemes: a general result. Econometrica **41**(4), 587–601 (1973)
11. Gibbard, A.: Manipulation of schemes that mix voting with chance. Econometrica **45**(3), 665–681 (1977)
12. Greenberg, A.: How One of Apple's Key Privacy Safeguards Falls Short, Wired. (2017) https://www.wired.com/story/apple-differential-privacy-shortcomings/. Accessed 14 Oct 2019
13. https://stackoverflow.com/questions/224058/distributed-random-number-generation. Accessed 14 Oct 2019
14. Lee, J., Clifton, C.: How Much Is Enough? Choosing ϵ for Differential Privacy. In: Proceedings of ISC 2011. LNCS 7001, pp. 325–340 (2011)
15. Liaw, H.-T.: A secure electronic voting protocol for general elections. Comput. Secur. **23**(2), 107–119 (2004)
16. McSherry, F., Talwar, K.: Mechanism design via differential privacy. In: Proceedings of FOCS, IEEE Computer Society, 94–103 (2007)
17. Nissim, K., Smorodinsky, R., Tennenholtz, M.: Approximately optimal mechanism design via differential privacy. In: Innovations in Theoretical Computer Science, pp. 203–213 (2012)
18. Nissim, K., Orlandi, C., Smorodinsky, R.: Privacy-aware mechanism design. In: Proceedings of the 13th ACM Conference on Electronic Commerce (EC'12), pp. 774–789 (2012)
19. Procaccia, A. D.: Can approximation circumvent Gibbard–Satterthwaite? In: Proceedings of 24th AAAI Conference on Artificial Intelligence (AAAI), pp. 836–841 (2010)
20. Riemann, R.: Towards Trustworthy Online Voting : Distributed Aggregation of Confidential Data, Ph.D. Dissertation, Université de Lyon (2017)
21. Satterthwaite, M.A.: Strategy-proofness and Arrow's conditions: existence and correspondence theorems for voting procedures and social welfare functions. J. Econ. Theory **10**(2), 187–217 (1975)
22. Studholme, C., Blake, I.: Multiparty computation to generate secret permutations. (2007) https://eprint.iacr.org/2007/353.pdf. Accessed 14 Oct 2019
23. Sweeney, L., Shamos, M.: Multiparty computation for randomly ordering players and making random selections. Technical Report CMU-iSRI-04-126 (2004)
24. Talmon, N.: Algorithmic aspects of manipulation and anonymization in social choice and social networks. Universitätsverlag der TU Berlin, Foundations of Computing #4 (2016)
25. Torra, V.: Data Privacy: Foundations, New Developments and the Big Data Challenge. Springer, Berlin (2017)
26. Torra, V., Narukawa, Y.: Modeling Decisions: Information Fusion and Aggregation Operators. Springer, Berlin (2007)
27. Torra, V.: When mathematics goes to the polls: decision processes, RBA, 2014 (2017)
28. Vaidya, J., Clifton, C., Zhu, M.: Privacy Preserving Data Mining. Springer, Berlin (2006)
29. Yi, H.: Securing e-voting based on blockchain in P2P network. EURASIP J. Wirel. Commun. Netw. **2019**, 137 (2019)
30. Zwierko, A., Kotulski, Z.: A light-weight e-voting system with distributed trust. Electron. Notes Theor. Comput. Sci. **168**, 109–126 (2007)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.