Journal of Information Security and Applications 36 (2017) 112–126

Contents lists available at ScienceDirect

Journal of Information Security and Applications

journal homepage: www.elsevier.com/locate/jisa

Privacy in elections: How *small* is “small”?

Annabelle McIver[a,*], Tahiry Rabehaja[a], Roland Wen[b], Carroll Morgan[b]

[a] Macquarie University, Sydney
[b] University of NSW, Sydney

ARTICLE INFO

Article history:

ABSTRACT

We study the impact on privacy in results reporting in elections. In the interests of transparency election commissions report details of aggregated vote counts; in the interests of privacy some of that information must be suppressed. We apply recent advances in Quantitative Information Flow (QiF) to describe several privacy properties in order to study the trade-off between transparency and privacy in results reporting.

We show that for some properties the impact on privacy in releasing detailed results data is minimal; on the other hand we identify some privacy properties that potentially reveal a great deal of information making results reporting in small batches problematic.

1. Introduction

Strong privacy is a core principle of secret ballot elections. Elections have many complex processes with the potential to leak information unintentionally, and consequently many countermeasures have been developed to help preserve privacy.

However in all elections, privacy is inevitably weakened to some extent during the counting process, where information is revealed intentionally in announcing the election result. Of course at minimum the identities of the winners must be published. But in practice a wealth of further information is frequently published for transparency, which is another core principle of elections.

This tension between privacy and transparency raises important questions over what trade-offs are appropriate and what are the risks to privacy. In this paper we consider privacy in the counting process and apply Quantitative Information Flow techniques to study privacy risks of information revealed as part of the election results.

The amount and type of information released during the counting varies widely by jurisdiction, according to electoral culture and practices. Typically elections publish intermediate results such as tallies for the candidates. Elections may also publish other information such as the number of spoiled ballots, and even the identities of all the voters who voted. Such information is necessary to assure integrity and build trust by enabling public scrutiny, and is also highly desirable to facilitate analysis by commentators, political scientists and political parties.

But what are the privacy risks of releasing such information?

A well-known risk is extreme scenarios such as when the tallies reveal if everyone voted the same way or nobody voted for particular candidates. But beyond such extreme scenarios, the risks remain not well-understood in general. Instead there is usually an assumption that large-scale elections have a sufficient number of voters to make such scenarios highly unlikely.

However this assumption is not always true. Indeed in large-scale elections similar extreme scenarios can still occur when small batches of votes are reported and the corresponding set of voters can be identified. For example a small polling place might receive only a small number of votes. All these votes could be for the same candidate, or none of the votes could be for minor candidates. Publishing the candidate tallies would then compromise privacy if it is also known which voters voted at that polling place. Some jurisdictions publish the list of who voted. But even if this is not the case, there may still be other means to deduce that a particular voter voted at that polling place, for instance through family or friends, social media posts or phone location tracking. (Note that regardless of what is published, it may be unavoidable for privacy to be compromised wrt election officials and scrutineers, who must have access to information on the tallies and who voted.)

Many other factors, including electoral practices, can increase the likelihood that small batches of votes can be distinguished. In Australia, the large area and small population means that certain polling places receive relatively few votes, especially in remote areas. Also Australian elections provide highly flexible voting arrangements. For example in national and state elections, a voter can vote at any polling place in their state. Many polling places are likely to receive few votes from each outside electorate. In addition there are numerous voting methods, including voting on election

[*] Corresponding author.
E-mail address: annabelle.mciver@mq.edu.au (A. McIver).

http://dx.doi.org/10.1016/j.jisa.2017.08.003
2214-2126/© 2017 Elsevier Ltd. All rights reserved.

day, early voting (in-person at a polling place), postal voting and mobile polling places (for hospitals and nursing homes). Some of these methods may be used by few voters. (I know how Grandma and all her nursing home buddies voted!)

Fine-grained reporting of such small batches of votes can increase privacy risks. In Australia tallies for first preferences and spoiled votes are reported but the granularity varies by jurisdiction and even by election. In some instances tallies are reported per polling place and per voting method. In other instances certain tallies are consolidated for reporting to reduce the privacy risk, albeit at the cost of reducing transparency.

An interesting question then is when is it appropriate to consolidate reports for privacy reasons?

Another privacy risk is the scenario where the number of possible voting options is relatively large compared to the number of voters. These elections are vulnerable to signature attacks, which compromise privacy by embedding unique signatures in votes. Australian elections are particularly vulnerable to signature attacks due to the use of preferential electoral systems with large numbers of candidates (often over 100). Hence there is an enormous number of possible voting options (exponential in the number of candidates), which increases the effectiveness of signature attacks in identifying vote signatures.

Australian elections publish detailed aggregate intermediate counting information about the preferential votes. This partial information weakens privacy as it can be exploited by sophisticated signature attacks. The published information again varies by jurisdiction and by election.

Another interesting question then is what impact does publishing different aggregate information have on privacy?

It is also worth mentioning that in fact when electronic counting is used (as is the case for most STV elections [1]), current practice in Australia is to publish all the (anonymised) individual votes so that anyone can verify the correctness of the counting. This trade-off in favour of transparency can cause a complete loss of privacy because it is trivial for signature attacks to identify vote signatures.

Cryptographic counting schemes have been proposed to count preferential votes without releasing individual votes. But these schemes still release aggregate information as a trade-off to provide verifiability and efficiency. Also different schemes release different partial information, and so it is difficult to objectively compare the levels of privacy afforded.

The above examples are some of the known scenarios where releasing information during the counting has privacy risks. Of course there are likely many other scenarios that are as-yet unknown. Also even when it is known that privacy risks exist, the extent of the risks is unclear and the effectiveness of possible countermeasures is difficult to measure. As a result trade-offs are likely to evolve ad hoc and may be guided by possibly flawed intuition because privacy is less concrete than other properties. For example in designing cryptographic counting schemes, verifiability is absolute (the scheme is either verifiable or not) and performance is fixed (the scheme must be able to compute the result for a given election size within a given time), whereas the level of privacy is currently hard to determine (what is the actual difference if a scheme reveals only the identities of the winners versus all the individual votes?).

What is desirable is a way to quantify the privacy risks. The purpose of this paper is to apply modern theories of Quantitative Information flow to study this issue. Our main contributions are listed below.

1. We show how general techniques from quantitative information flow can be applied reasonably in elections. The notion of a *gain function* allows us to formulate exactly the contents of privacy questions we are interested in (Section 3.1).
2. We formalise two general privacy principles postulating the expected "amount of privacy" in small and large electorates (Section 4).
3. We look at two relevant scenarios and study their privacy guarantees, and draw some general conclusions about privacy and the publication of election results (Section 5). In particular, we show that for some types of privacy questions, publishing more doesn't necessarily mean increasing the privacy risk (Theorem 5.2).
4. We also compute the information flow for small numbers of voters, giving some indication of how much privacy is compromised when data is published for small numbers of voters. These comparisons are summarised in graphical form throughout this paper.

We note that the quantities we compute are not particularly informative in isolation. What is compelling however is the way that our framework allows us to compare scenarios and to analyse qualitatively different ways of reporting results and their possible consequences to privacy. Most of our conclusions relate to this method of comparison.

2. Related work

In the literature on cryptographic election schemes, privacy of the voting process has been studied extensively, particularly the strong privacy notions of receipt-freeness and coercion resistance. For example a common method for defining privacy is indistinguishability with an ideal voting functionality, where the inputs are the votes and the output is the election result. Our work complements such approaches by studying privacy of the information released by the "ideal" election result.

Closely related to our work is the privacy measure of Bernhard et al. [4]. They develop a general, entropy-based privacy definition that measures privacy overall in both the cryptographic protocols and the election result. Their privacy measure is designed to be flexible and compatible with different notions of information-theoretic entropy. Our work also complements this approach since our privacy measure is based on information-theoretic entropy. An interesting avenue for future work would be to combine our notion of privacy with their definition.

In the context of statistical databases, Dwork introduced the notion of differential privacy [7]. A statistical query is differentially private if the probability of distinguishing two databases that differ at a single entry is negligible. Superficially, this idea seems to apply directly to the context of voting but the resulting definition of privacy is too discriminating. This situation was also observed elsewhere [4]. In particular, if the tallies are published then changing the vote of a single voter is always enough to distinguish the input ballot boxes. With respect to our definition, such a counting procedure do provide enough privacy, even though it is not differentially private, as long as the number of candidates is reasonably small.

For signature attacks, some ad hoc examples of the privacy implications have been considered previously [3,6,14]. A quantitative approach to measuring information leakage for signature attacks has also been suggested, along with early steps towards more formal analysis of the problem [15].

[1] Single Transferable Voting systems require each elector to rank candidates in his or her vote. For such systems, the counting algorithm is usually composed of multiple rounds. In the first round, each vote is allocated to the most preferred candidate and then proportionally transferred to other ranked candidates in each subsequent round.

3. Review of quantitative information flow

Informally, a secret is some value about which there is some uncertainty, and the greater the uncertainty the more difficult it is to know exactly what that value is. When some information about the secret becomes known to an observer (often referred to as an adversary) the uncertainty is reduced, and we say that information (about the secret) has leaked.

Quantitative Information Flow (QiF) makes this intuition mathematically precise: given a range of possible secret values of (finite) type $\mathcal{X}$, we model a secret as a probability distribution $\mathbb{D}\mathcal{X}$. Given $\pi : \mathbb{D}\mathcal{X}$ we write π_x for the probability that π assigns to $x{\in}\mathcal{X}$. Normally the uniform distribution over $\mathcal{X}$ models a secret which could equally take any one of the possible values drawn from its type, but there could be reasons for using some other distribution, for example if the secret was the height of an individual then a normal distribution might be more realistic. In any case, once we have a secret, we are interested in analysing whether an algorithm, or protocol, that uses it might leak some information about it.

The original QiF studies used Shannon Entropy to measure uncertainty in probability distributions modelling secrets, with the idea being that the "more secret" something is, the more uncertainty there is about its actual value from the perspective of an observer. More recent treatments of this idea have shown that Shannon entropy is not the best way to measure uncertainty in security contexts because it does not model scenarios where an observer is trying to guess the value of a secret. Moreover, there are some circumstances where a Shannon analysis actually gives a more favourable assessment of security than is actually warranted [13].

Recently Smith [2] proposed more general notions of how to measure uncertainty based on "gain functions". This is the notion we will use in this paper. A *gain function* models a very general scenario where a secret's uncertainty is measured directly according to the ability of an observer to find out information about the secret. We write $\mathcal{W}$ for a (usually finite) set of decisions available to an observer corresponding to an "attack scenario" where the adversary tries to guess something (e.g. some property) about the secret; for a given secret $x{\in}\mathcal{X}$ an adversary's choice of $w{\in}\mathcal{W}$ results in some gain which can vary depending on his choice and the value of the secret. The more effective is the adversary's choice in guessing the secret, the more he is able to gain from the attack.

Definition 3.1. Given a type $\mathcal{X}$ of secrets, a gain function $g: \mathcal{W}{\times}\mathcal{X} \to \mathbb{R}$ is a real-valued function such that $g(w, x)$ determines the gain to an attacker if he chooses w and the secret is x.

A simple example of a gain function is given by g_{id}, where $\mathcal{W} := \mathcal{X}$, and

$$g_{id}(x, x') := 1 \quad \textit{if} \quad x = x' \quad \textit{else} \quad 0 \ . \tag{1}$$

For this scenario, the attacker receives a gain of 1 if he correctly guesses the value of a secret. Elsewhere the utility and expressivity of gain functions for measuring information flow have been described [1,2]. Given a gain function we define the *vulnerability* of a secret in $\mathbb{D}\mathcal{X}$ relative to the scenario it describes: it is the maximum average gain to an attacker.

Definition 3.2. Let $g: \mathcal{W}{\times}\mathcal{X} \to \mathbb{R}$ be a gain function, and $\pi : \mathbb{D}\mathcal{X}$ be a secret. The *vulnerability* $V_g[\pi]$ of the secret wrt g is:

$$V_g[\pi] := \max_{w \in \mathcal{W}} \sum_{x \in \mathcal{X}} g(w, x) {\times} \pi_x \ .$$

For a secret $\pi : \mathbb{D}\mathcal{X}$, the vulnerability wrt g_{id} is $V_{g_{id}}[\pi] := \max_{x:\mathcal{X}} \pi_x$, i.e. the maximum probability assigned by π to possible values of x. The attacker's best strategy for optimising his gain would therefore be to choose the value x that achieves $V_{g_{id}}[\pi]$.

A mechanism is an abstract model of a protocol or algorithm that uses secrets. As the mechanism executes we assume that there are a number of observables that can depend on the actual value of the secret. We define $\mathcal{Y}$ to be the type for observables, and for each secret x each observable can be observed with some probability. Such observables could be sample timings in a timing analysis in cryptography. For our example in elections the observations could be the winners of the election, the margins by which they won, or distributions of preferences.

Definition 3.3. A mechanism is a *stochastic channel*[2] $C : \mathcal{X} \rightarrow\!\!\!\!\!\rightarrow \mathcal{Y}$. The value C_{xy} is the probability that if the secret is x then y is observed. Given a (prior) secret $\pi : \mathbb{D}\mathcal{X}$ we write $\pi{\triangleright}C$ for the joint distribution in $\mathcal{X}{\times}\mathcal{Y}$ defined

$$(\pi{\triangleright}C)_{xy} := \pi_x {\times} C_{xy} \ .$$

For each $y : \mathcal{Y}$, the probability that y is observed is $p_y := \sum_{x':\mathcal{X}} (\pi{\triangleright}C)_{x'y}$ and we define the posterior probability, i.e. that the secret is x given that y was observed, as the conditional $\pi|_y := (\pi{\triangleright}C)_{xy}/p_y$.

Intuitively, given a prior secret π, the entry $\pi_x \times C_{x,\,y}$ of the joint distribution $\pi{\triangleright}C$ is the probability that the actual secret value is x and the observation is y. This joint distribution contains two pieces of information: the probability p_y of observing y and the corresponding posterior which represents the attacker's updated view about the uncertainty of the secret's value. If the vulnerability of the posterior increases, then information about the secret has leaked and the attacker can use it to increase his gain. The attacker's average overall gain, taking the observations into account, is defined to be the average vulnerability wrt the posteriors:

$$V_g[\pi{\triangleright}C] := \sum_{y \in \mathcal{Y}} p_y {\times} V_g[\pi|_y] \ . \tag{2}$$

Leakage can now be defined: it is a measure of the increased gain to an attacker using the extra information to refine his strategy to discover facts about the secret.

Definition 3.4. Given a mechanism C, a gain function g and a secret π we can define the multiplicative leakage of C wrt π and g as:

$$\mathcal{L}_g[\pi, C] := V_g[\pi{\triangleright}C]/V_g[\pi] \ . \tag{3}$$

The multiplicative capacity maximises the leakage over all non-negative gain functions $g : \mathcal{W}{\times}\mathcal{X} \to \mathbb{R}_{\geq}$ and prior information about secrets $\pi : \mathbb{D}\mathcal{X}$.

$$\mathcal{L}_{\forall}[\forall, C] := \max_{g,\pi} V_g[\pi{\triangleright}C]/V_g[\pi] \ .$$

Intuitively, the leakage $\mathcal{L}_g[\pi, C]$ quantifies the amount of information leaked by the mechanism C for a fixed gain function g and distribution π, disregarding the prior information contained in π. The minimum value for $\mathcal{L}_g[\pi, C]$ is 1, i.e. $V_g[\pi{\triangleright}C] = V_g[\pi]$, and this means that the mechanism C leaks nothing. The higher the leakage value, the more information is leaked by the mechanism.

Similarly, the capacity $\mathcal{L}_{\forall}[\forall, C]$ measures the maximum amount of information leaked by the mechanism C, disregarding prior information. If C leaks nothing, then $\mathcal{L}_{\forall}[\forall, C] = 1$, i.e. $V_g[\pi{\triangleright}C] =$

[2] Stochastic means that the rows sum to 1.

$V_g[\pi]$ for all gain functions and priors, so that C has no effect on information flow. The uniform distribution has particular significance here because it gives a strong upper bound on capacity [2]:

$$\mathcal{L}_\forall[\forall, C] = \mathcal{L}_{g_{id}}[\upsilon, C] \ , \tag{4}$$

where υ is the uniform distribution over $\mathcal{X}$. This means that modelling a secret using the uniform prior, and taking the vulnerability wrt g_{id} gives a very robust bound on the vulnerability of the secret, wrt any scenario.[3]

Note that whilst vulnerability gives a measure of “risk” within a particular scenario, the leakage gives a measure of how much the information flowing from the channel has contributed to that risk. A secret could be vulnerable even before the channel has released any information because the adversary somehow possesses a great deal of prior knowledge, as modelled by a skewed probability distribution. In this case, it could be that the extra information released by the channel is not significantly adding to the risk that was already present. In this sense the leakage Definition 3.4 is about the channel’s contribution to that final posterior vulnerability, whereas the posterior vulnerability (2) is about the actual risk to the secret averaged over the possible observations.

We can compare mechanisms wrt their ability to preserve confidentiality of secrets: one channel is regarded as being more secure than another if and only if it never leaks more in any scenario.

Definition 3.5. Given mechanisms $C : \mathcal{X} \rightarrowtriangle \mathcal{Y}$ and $D : \mathcal{X} \rightarrowtriangle \mathcal{Y}'$ we say that D refines C, or $C \sqsubseteq D$ if, for all secrets $\pi : \mathbb{D}\mathcal{X}$ and all non-negative gain functions $g : \mathcal{W}{\times}\mathcal{X} \to \mathbb{R}_{\geq}$, we have:

$$\mathcal{L}_g[\pi, C] \geq \mathcal{L}_g[\pi, D] \ .$$

There is a different way to determine refinement [1,11]:

$C \sqsubseteq D$ if and only if there is some stochastic R such that

$$D = C{\cdot}R, \tag{5}$$

where $C \cdot R$ means post multiplication of C by R. It takes (possibly scaled) sums of columns of C to form columns of D, thereby obscuring the provenance of observations in C.

The main idea behind the post-multiplication is that the observations that are release by the mechanism C are processed by some other mechanism R. Thus, unless R is the identity matrix, that post-multiplication adds some extra noise to the information released by C to the point that they are less useful to the attacker.

Definition 3.5 summarises an important aspect of quantitative information flow, in that it promotes a comparative notion of information leakage. Whilst it can be difficult, or impossible, to completely secure a mechanism against leaks, Definition 3.5 provides the basis for comparing the information flow properties of ideal mechanisms with more realistic implementations of mechanisms. In general if D is an ideal mechanism and $C \sqsubseteq D$, if the adversary’s strategy to optimise the gain function g is the same in D as it is for C then the corresponding leakages will be the same for that gain function g. If this is the case then the extra information released in C cannot be used in the scenario modelled by g.

Most of the reporting channels that we will study below are *deterministic*, which means that for each row of the matrix, exactly one entry $C_{xy} = 1$; in this case we say that $C : \mathcal{X} \rightarrowtriangle \mathcal{Y}$ divides $\mathcal{X}$ into equivalence classes: $x \sim_C x'$ if and only if $C_{xy} = C_{x'y}$ for all $y : \mathcal{Y}$. Let $\#C$ be the number of equivalence classes in $\sim_C$, then $\mathcal{L}_\forall[\forall, C] = \#C$ [2].

For a given gain function g and reporting channel C, we have that

$$1 \leq \mathcal{L}_g[\pi, C] \leq \#C \ , \tag{6}$$

so for a particular privacy property modelled by g, depending on how close the leakage is to these extremal bounds indicates how much the adversary is able to use the information leaked.

3.1. Election announcements as reporting channels

In this section we consider a simple example to illustrate how some of the concepts described above can be applied to privacy in voting. Consider the scenario where there are two candidates, a and b, and exactly one of them is to be elected by simple majority. We assume that the voters are able to cast their ballots privately, so that the secret is how they voted, and we then assume that once the tallies have been computed, the results are announced.

Definition 3.6. We model the announcement process as a channel from $\mathcal{B} \rightarrowtriangle \mathcal{A}$, where $\mathcal{B}$ is the set of assignments $\mathcal{E} {\to} \mathcal{C}$ describing how each elector (in $\mathcal{E}$) cast their vote for candidates (in $\mathcal{C}$), and $\mathcal{A}$ is the set of possible announcements concerning the results after (or during) tallying.[4]

In Fig. 1 the matrices W and T describe channels for two different modes for announcing election results. Channel W simply releases the name of the winning candidate, whereas T announces the tallies for each candidate. The labels for each row represent a particular assignment for how the three voters cast their ballots. For example $\langle abb \rangle$ means that the first voter selected a, and the remaining two both selected b. We assume that the identities of each voter have been removed from the ballots, and when we refer to “the first voter” we simply mean the anonymous voter whose ballot happened to appear first in the ballot box. The columns in the matrices correspond to the different announcements. In W there are only two possible announcements: “a won” or “b won”. In T there are four: “a won (3, 0)”, “a won (2, 1)”, “b won (1, 2)” and “b won (0, 3)”. Thus T certainly releases more information than W, and for transparency it is desirable to use T, but we would like to know the impact on privacy.

Let $\upsilon : \mathbb{D}\mathcal{B}$ now be the secret i.e. the probability distribution over all possible assignments in $\mathcal{B}$; we will assume for the moment that it is uniform. When the winner is announced after counting via channel W, we would observe a as the winner with probability $\sum_{x:\mathcal{B}_a} \upsilon_x$, where $\mathcal{B}_a$ is the set of ballot assignments in which a has the majority. The posterior probability associated with a’s win is the uniform distribution over $\mathcal{B}_a$, so all that is known in the election is that a received the majority of the votes.

The information flow in T is a little different: by announcing the tallies there are now four possible observations, which in particular reveal whether or not one of the candidates did not receive any votes. If the secret is revealed to be definitely either $\langle aaa \rangle$ or $\langle bbb \rangle$ then the losing candidate knows for certain that no one voted for them, and indeed that the winner was chosen unanimously.

Observe that $T \sqsubseteq W$ because the first (last) two columns of T can be combined to give the first (last) column of W, so that T leaks strictly more information than does W in some scenarios defined by gain functions. Indeed since $\#W = 2$ and $\#T = 4$, (6) implies that the maximum leakage in T is twice as much as it is for C in some scenarios, but not all. In particular we can define a gain

[3] Furthermore, it can be shown [2] that $\log_2 \mathcal{L}_{g_{id}}[\upsilon{\triangleright}C]$ is an upper bound on the mutual information between $\mathcal{X}$ and $\mathcal{Y}$ defined by the joint distribution $\pi{\triangleright}C$, and so the leakage can also be understood in traditional terms of uncertainty measured by Shannon Entropy.

[4] Note that we assume that all identifying information has been removed from ballots, and the function in $\mathcal{E}{\to}\mathcal{C}$ simply means that the set of ballots is in one-to-one correspondence with the registered voters.

W	a	b
$\langle aaa\rangle$	1	0
$\langle aab\rangle$	1	0
$\langle aba\rangle$	1	0
$\langle baa\rangle$	1	0
$\langle bab\rangle$	0	1
$\langle abb\rangle$	0	1
$\langle bba\rangle$	0	1
$\langle bbb\rangle$	0	1

T	$(3,0)$	$(2,1)$	$(1,2)$	$(0,3)$
$\langle aaa\rangle$	1	0	0	0
$\langle aab\rangle$	0	1	0	0
$\langle aba\rangle$	0	1	0	0
$\langle baa\rangle$	0	1	0	0
$\langle bab\rangle$	0	0	1	0
$\langle abb\rangle$	0	0	1	0
$\langle bba\rangle$	0	0	1	0
$\langle bbb\rangle$	0	0	0	1

The options for the secret describes how the voters cast their ballots: $\langle baa\rangle$ means that the first voter chose b and the second and third chose a. Channel W simply announces the winning candidate but channel T also announces the tallies: $(2,1)$ labelling a column in T means that candidate a received 2 votes and candidate b only 1.

Fig. 1. Two reporting channels with three voters and two candidates.

function to model a privacy property that is relevant to the context of interest.

For example, let *pd* be the "plausible deniability" gain function defined with two choices $\{u, \neg u\}$ to model a scenario where an adversary tries to determine whether the majority was unanimous or not:

$$\begin{aligned} pd(u, \langle hkl\rangle) &= 1 \quad iff \quad h{=}k{=}l \quad , \quad and \quad pd(\neg u, \langle hkl\rangle) \\ &= 1 \quad iff \quad \neg(h{=}k{=}l) \end{aligned} \tag{7}$$

We see that the prior vulnerability of the secret is $V_{pd}[\upsilon] = 3/4$ since the best strategy for the adversary is to choose $\neg u$ because the chance of a unanimous election is only 1/4. Moreover if only the winner is announced, we see that there is no change in gain: $V_{pd}[\upsilon, W] = 3/4$. However if the full tallies are released we find that $V_{pd}[\upsilon, T] = 1$ because it is always possible for the adversary to distinguish between unanimous and split majorities. Therefore $\mathcal{L}_{pd}[\upsilon, W] = 1$ whereas $\mathcal{L}_{pd}[\upsilon, T] = 4/3$, showing that if the tallies are not released then there is no way for the adversary to learn whether or not the winner achieved a majority with absolute certainty.

Consider now the gain function $\alpha_{\mathcal{E}}$ where the adversary tries to guess how some voter voted. This time $\mathcal{W} = \mathcal{C}\times\mathcal{E}$, so that the pair $(c, e){\in}\mathcal{C}\times\mathcal{B}$ models and the adversary guessing that voter e voted for candidate c:

$$\alpha_{\mathcal{E}}((c, e), \langle hkl\rangle){=}1 \quad if \quad c{=}\langle hkl\rangle_e \quad else \quad 0 \ , \tag{8}$$

where we write $\langle hkl\rangle_e$ for the e'th value in the list, so that $\langle hkl\rangle_1 = h$, for example. In this case $V_{\alpha_{\mathcal{E}}}[\upsilon] = 1/2$, because whichever pair (c, e) the adversary picks, only half of the set of ballots has voter e selecting candidate c. After the results are announced however, both W and T leak the same: $\mathcal{L}_{\alpha_{\mathcal{E}}}[\upsilon, W] = \mathcal{L}_{\alpha_{\mathcal{E}}}[\upsilon, T] = 3/2$ because, *in both cases*, the adversary improves his guess given the information released using the following reasoning.

If candidate a wins then the adversary *always* guesses (a, e) for any voter e, if b wins then the adversary always guesses (b, e). That is because in the post hoc situation the winner is the candidate that is most likely to have been selected by most voters. Surprisingly, although the maximal possible leakage for T is 4, the calculated leakage $\mathcal{L}_{\alpha_{\mathcal{E}}}[\upsilon, T] = 3/2$ means that much of the extra information given by releasing tallies is not useful to an adversary who is *only* interested in trying to guess how some voter voted, and indeed this particular attack cannot be mitigated in any way since at the very least the winner must be announced, and as we have seen $\mathcal{L}_{\alpha_{\mathcal{E}}}[\upsilon, W]$ is also 3/2. We discuss this phenomenon further in Section 5.1.

Next we set out three gain functions which capture three different aspects of privacy which we will use to evaluate privacy in our case studies below in Section 5.

Definition 3.7. The following gain functions can be used to investigate privacy in election announcements.

1. Identify how voters cast votes:

   $$\alpha_{\mathcal{E}}((c, e), b){:=}1 \quad if \quad b_e{=}c \quad else \quad 0 \ . \tag{9}$$

   This gain function generalises (8) for arbitrarily many voters. It models an adversary who tries to guess how a voter cast their ballot. The parameter b is the set of ballots $\mathcal{E}{\to}\mathcal{C}$; an adversary's guess (c, e) is of type $\mathcal{C}\times\mathcal{E}$. The adversary receives a gain of 1 if (s)he correctly guesses who voter e selected.

   Given $\pi{\in}\mathbb{D}\mathcal{B}$, the gain $V_{\alpha_{\mathcal{E}}}[\pi] := \max_{(c,e)\in\mathcal{C}\times\mathcal{E}} \sum_{b\in\mathcal{B}} \alpha_{\mathcal{E}}((c, e), b){\times}\pi_b$. It is the maximum probability that the adversary can guess correctly some voter's vote, assuming prior knowledge π.
2. The number of voters whose cast votes can be identified correctly:

   $$\#\alpha_{\mathcal{E}}(b', b){:=}\sum_{e\in\mathcal{E}} 1 \ if \ b_e{=}b'_e \ else \ 0 \quad . \tag{10}$$

   This gain function estimates the average number of ballots that are correctly guessed at (9) above, and uses both b and b' in $\mathcal{E}{\to}\mathcal{C}$, where b' summarises the adversary's guess for each voter. Here the gain to the attacker is the sum of correctly guessed ballots.

   Given $\pi{\in}\mathbb{D}\mathcal{B}$, the gain $V_{\#\alpha_{\mathcal{E}}}[\pi] := \max_{b'\in\mathcal{B}} \sum_{b\in\mathcal{B}} \#\alpha_{\mathcal{E}}(b', b){\times}\pi_b$. It is the maximum expected number of voters that the adversary can guess correctly, assuming prior knowledge π.
3. Identify how voters did not cast votes:

   $$\neg\alpha_{\mathcal{E}}((c, e), b){:=}1 \quad if \quad b_e{\neq}c \quad else \quad 0 \ . \tag{11}$$

   This gain function is the complement of (9), with the parameters defined in the same way. This gain function models an adversary who tries to guess which candidates for whom a voter did not choose. The gain is 1 if (s)he correctly guesses the candidate that the voter did not vote for.

We can use these gain functions to express privacy for election channels by analysing the behaviour of the derived leakages. For example the leakages corresponding to the gain functions (9) and (11) together express a strong form of privacy: if an adversary can *neither* guess a voter's selection *nor* who they did not select, then the voter is able to lie plausibly with some confidence.

A particular focus of our analysis however is to explore how privacy is affected by reporting results for small samples of voters. When the reporting channel C is a good preserver of privacy then $\mathcal{L}_{g_\mathcal{E}}[\upsilon_\mathcal{E}, C]$ will be close to 1, and the more that privacy could be compromised, for some voter, the larger (than 1) this leakage will be. If the leakage is comparable to the maximum possible leakage (6) then it means that the adversary is able to use much of the information leaked to find out how voters voted.

In the next section we formalise some general principles for privacy related to sample size.

4. Formalising privacy principles in elections

In this section we formalise some privacy principles in elections using the framework summarised in Section 3.

4.1. The law of large electorates

Our first principle is that the leakage in a reporting channel wrt a given privacy question must tend arbitrarily to zero the more voters in the election. This expresses the general principle that voter privacy is derived through reporting only the amalgamated results: if there are enough voters then each candidate has some chance of being picked.

Definition 4.1. Let $g_\mathcal{E}$ be a privacy property described in Definition 3.7, and let C be a reporting channel; we say that C and $g_\mathcal{E}$ satisfy the law of large electorates if the leakage $\mathcal{L}_{g_\mathcal{E}}[\upsilon_\mathcal{B}, C] \to 1$ as $|\mathcal{E}| \to \infty$ and $|\mathcal{C}|$ is fixed, where $\upsilon_\mathcal{B}$ is the uniform distribution over $\mathcal{B}$, and $|\mathcal{E}|$ is the number of voters.[5]

Definition 4.1 sets out a criterion which depends both on the method of reporting results. If a reporting channel does not satisfy the law of large electorates it means that it releases sufficient information in its reporting protocol that could compromise some voters' privacy.

For example if reporting channel A simply releases all information, i.e. the identities of voters as well as how they voted, then $\mathcal{L}_{\#\alpha_\mathcal{E}}[\upsilon_\mathcal{E}, A] = |\mathcal{B}|$ and so A, of course, does not satisfy the law of large electorates: privacy does not increase with increasing numbers of voters, since all information is leaked in this method of reporting.

We show in 5 that (9)–(11) satisfy Definition 4.1 for the typical results reporting of "first past the post" and "instant run off" elections.

4.2. The law of small samples

Our second principle deals with the question of privacy when reporting results for small samples of voters. This can happen for instance in Australia when voters do not vote at a polling place in their own constituency.

We assume that C and $g_\mathcal{E}$ satisfy the law of large electorates, so that privacy increases with the size of $\mathcal{E}$. We would like to analyse the degree to which privacy is lost in reporting results in small batches rather than aggregating them all first. Let $\mathcal{X} \subset \mathcal{E}$, and we assume that $\mathcal{L}_{g_\mathcal{X}}[\upsilon_\mathcal{X}, C] \geq \mathcal{L}_{g_\mathcal{E}}[\upsilon_\mathcal{E}, C]$. The greater the difference between $\mathcal{L}_{g_\mathcal{X}}[\upsilon_\mathcal{X}, C]$ and $\mathcal{L}_{g_\mathcal{E}}[\upsilon_\mathcal{E}, C]$, the greater the impact on privacy to voters in the sample $\mathcal{X}$. The next definition gives the proportional loss of privacy relative to the leakage $\mathcal{L}_{g_\mathcal{X}}[\upsilon_\mathcal{X}, C]$.

Definition 4.2. Let C be a reporting channel, and $g_\mathcal{E}$ a gain function defined in Definition 3.7. Let $\mathcal{X}$ be a subset of all ballots $\mathcal{E}$. The proportional loss of privacy to voters in $\mathcal{X}$ relative to the election consisting of voters in $\mathcal{E}$ is:

$$1 - \mathcal{L}_{g_\mathcal{E}}[\upsilon_\mathcal{E}, C] / \mathcal{L}_{g_\mathcal{X}}[\upsilon_\mathcal{X}, C] \ . \tag{12}$$

When the size of the electorate is so large that the reporting channel provides good privacy i.e. $\mathcal{L}_{g_\mathcal{E}}[\upsilon_\mathcal{E}, C] \approx 1$, we can approximate (12) to $1 - 1/\mathcal{L}_{g_\mathcal{X}}[\upsilon_\mathcal{X}, C]$. In any case, this provides an upper bound on the proportional loss of privacy to voters in sample $\mathcal{X}$, caused by the contribution of the information flow in the channel.

5. Case studies

In this section we investigate the privacy properties described in Definition 3.7 wrt two types of elections: first-past-the-post and instant runoff. We take primarily a comparative approach, recognising that no reporting channel can be absolutely risk free, but that some kinds of reporting leak more information than others, and that some of the information released could potentially impact privacy.

We make two kinds of comparison. In the first (for first past the post) we compare the scenario where only the winner(s) are announced to the scenario where the tallies for each candidate are announced, thus generalising our example at Fig. 1. This is an interesting comparison, because although it is not realistic to expect that only the winners (without the tallies) be announced, that scenario offers the greatest possible privacy to voters, and therefore allows us to identify what sort of privacy risks are related to releasing more information.

With respect to instant run-off we again define two types of reporting channels. One (called P) releases the aggregate tallies of the full preference list selected by voters, and the other releases the aggregate tallies of "first preferences" only (called F). Whereas information flow in F is closer to what is done in real elections, it is related to P via refinement. However information flow in P is feasible to calculate which then allows us to give bounds on the extent of information leakage in F.

In the second comparison, we consider the loss to privacy of releasing results in small samples. This sheds light on the increasing risks to privacy in making announcements in small batches.

In all our analyses we assume that the adversary knows nothing about the likely voting preferences, and throughout we use a uniform distribution over $\mathcal{B}$ to model the adversary's prior knowledge. We note however that the general method of QiF does not require this assumption.

5.1. First-past-the-post

First-past-the-post elections are used to elect a single candidate. Voters (in $\mathcal{E}$) pick exactly one candidate out of some given set $\mathcal{C}$, and the candidate who obtains the most votes is declared the winner. The next definition sets out a generalisation of Fig. 1's methods of reporting results.

Definition 5.1. Let $\mathcal{B} := \mathcal{E} \to \mathcal{C}$ model sets of ballots. Define $W{:}\mathcal{B} \rightarrowtriangle \mathcal{C}$ to be the reporting channel that announces the winner only:

$$W_{bc} := 1 \quad if \quad maj(b) = c \quad else \quad 0 \ , \tag{13}$$

where $maj(b)$ is the candidate with the maximum number of votes determined by b. If there is a draw then $maj(b) := c$ for the first candidate c with the highest tally.[6]

[5] Recall that when there is no information flow, then $\mathcal{L}_{g_\mathcal{E}}[\upsilon_\mathcal{B}, C] = 1$.

[6] For simplicity we set $maj(b)$ to be a fixed candidate with the highest tally. The candidates can be ordered lexicographically and the first with a maximum tally is reported as the winner. We can also break draws randomly while keeping the re-

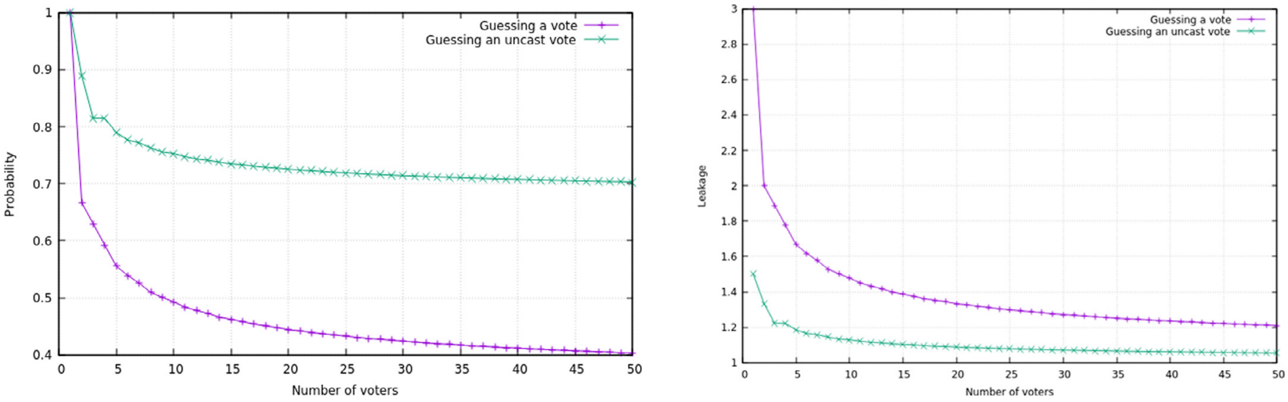

Fig. 2. (Left) Posterior vulnerabilities $V_{\alpha_{\mathcal{E}}}[\upsilon \triangleright W]$ (lower) and $V_{\neg\alpha_{\mathcal{E}}}[\upsilon \triangleright W]$ (upper). (Right) leakages $\mathcal{L}_{\alpha_{\mathcal{E}}}[\upsilon, W]$ (upper) and $\mathcal{L}_{\neg\alpha_{\mathcal{E}}}[\upsilon, W]$ (upper). Analysis run with 3 candidates.

Define $T{:}\mathcal{B} \rightarrowtriangle \mathbb{N}^{\mathcal{C}}$ to be the reporting channel that announces the set of tallies for each candidate:

$$T_{bf} := 1 \quad \mathit{if} \quad (\forall c : \mathcal{C} \cdot \#_c(b) = f(c)) \quad \mathit{else} \quad 0 \ , \tag{14}$$

where $\#_c(b)$ is the number of votes that candidate c received in the ballots b.

As described above, we have that $T \sqsubseteq W$ so that T releases much more information than does W. In fact, we can define a channel $R{:}\mathbb{N}^{\mathcal{C}} \rightarrowtriangle \mathcal{C}$ such that $R_{fc} = 1$ if and only if candidate c is reported as the winner when the tallies are given by f and c wins wrt the tie-breaking mechanism of W. It is then clear that $W = T{\cdot}R$ satisfies (5). Our interest in W is that it provides the most privacy to voters, and therefore it allows us to compare how much privacy is impacted by the additional announcements of T.

We show first that W and T satisfy the law of large electorates.

Theorem 5.1. *Both T and W satisfy the law of large electorates wrt* $\alpha_{\mathcal{E}}$, $\#\alpha_{\mathcal{E}}$ *and* $\neg\alpha_{\mathcal{E}}$.

Proof. In Appendix B. □

Theorem 5.1 shows that, as expected, privacy increases the more voters included in the reported results. This predicted trend can also be observed by Fig. 2 at right, where the leakages $\mathcal{L}g_{\mathcal{E}}[\upsilon_{\mathcal{E}}, C]$ are displayed for increasing sizes of electorate.

In the proof of Theorem 5.1, we have shown that

$$L_{\#\alpha}[\upsilon, T] = L_{\#\alpha}[\upsilon, W] \simeq 1 + \sqrt{\frac{2m \ln m}{n}} \ , \tag{15}$$

for n reasonably larger than m. This approximation is quite robust and can be used as an upper bound for any voting process that publishes at most the tally of each candidate. More specifically, we have shown that the quantity of votes that an attacker can guess using the best strategy is proportional to $\sqrt{n}$. In this case, the law of large electorate is equivalent to a decreasing relative amount of leaked votes, i.e., $\lim_{n\to\infty} \frac{\sqrt{n}}{n} = 0$.

Next we compare T and W in detail. Our first result shows that the extra information in the announcements of tallies cannot be used by the adversary neither to guess how a voter voted ($\alpha_{\mathcal{E}}$) nor to increase the expected number of correct guesses.

Theorem 5.2. *Let T and W be reporting channels defined at Definition 5.1 and* $\alpha_{\mathcal{E}}$, $\#\alpha_{\mathcal{E}}$ *be gain functions defined at Definition 3.7. Let* $\upsilon_{\mathcal{E}} \in \mathbb{D}\mathcal{B}$ *be the uniform prior. The following equalities hold:*

$$\mathcal{L}_{\alpha_{\mathcal{E}}}[\upsilon_{\mathcal{E}}, W] = \mathcal{L}_{\alpha_{\mathcal{E}}}[\upsilon_{\mathcal{E}}, T] \ \mathit{and} \ \mathcal{L}_{\#\alpha_{\mathcal{E}}}[\upsilon_{\mathcal{E}}, W] = \mathcal{L}_{\#\alpha_{\mathcal{E}}}[\upsilon_{\mathcal{E}}, T] \ .$$

Proof. (Informal sketch.) The prior gain $V_{\alpha_{\mathcal{E}}}[\upsilon_{\mathcal{E}}]$ is the same for both leakage calculations, therefore we only need show that $V_{\alpha_{\mathcal{E}}}[\upsilon_{\mathcal{E}} \triangleright W] = V_{\alpha_{\mathcal{E}}}[\upsilon_{\mathcal{E}} \triangleright T]$. Note that for any announcement of tallies for election with ballots given by b, it is still the case that most voters in $\mathcal{E}$ voted for the candidate with the majority, and so the adversary's optimum guessing strategy is to pick $(\mathit{maj}(b), e)$, which is exactly the same guessing strategy if only the winner is announced. Hence since the optimal guessing strategies are the same for both W and T, the leakage wrt $\alpha_{\mathcal{E}}$ must also be the same.

For $\#\alpha_{\mathcal{E}}$, a similar argument shows that the adversary's optimal strategy is to guess that all voters voted for the candidate who won.

We provide a detailed proof formalising this argument in Appendix A. □

On the other hand, compared to the most private mechanism W, releasing the tallies gives the adversary more scope to improve the effect of his guessing strategy for $\neg\alpha_{\mathcal{E}}$.

Theorem 5.3. *Let T and W be reporting channels defined at Definition 5.1 and* $\neg\alpha_{\mathcal{E}}$ *be the gain function defined at Definition 3.7. If* $|\mathcal{C}| > 2$ *then* $\mathcal{L}_{\neg\alpha_{\mathcal{E}}}[\upsilon_{\mathcal{E}}, W] < \mathcal{L}_{\neg\alpha_{\mathcal{E}}}[\upsilon_{\mathcal{E}}, T]$.

Proof. (Informal sketch.) When tallies are released, the adversary can increase his gain by guessing that the candidate who received the least number of votes is most likely not someone voters selected. This information is not available in W. □

5.1.1. Privacy risks in small samples

We computed the leakages and vulnerabilities to illustrate some aspects of privacy for first past the post channels W and T. We assumed a uniform distribution over the anonymised (set of) ballots and varied the number of voters in a "reporting batch" given by the size of the parameter $\mathcal{E}$, with three possible candidates to choose from. This set up is akin to an electoral commission releasing e.g. polling place tallies where only few voters cast their votes.

Computing the vulnerabilities, and thus the leakages, directly through (2) is computationally very expensive because that definition relies on summing over the set of all possible ballot boxes $\mathcal{B}$ which has $|\mathcal{C}|^{|\mathcal{E}|}$ elements. To compute leakages for reasonable electorate sizes, we had to transform these expensive sums into manageable ones and it turns out that the resulting expressions correspond to recent quantitative measures used in seemingly unrelated areas such as hash code analysis [8,12] and load balanc-

sults exactly the same. The important property to consider is that the tie breaking mechanism should be independent of the set of ballots b, otherwise it may reveal more information than expected.

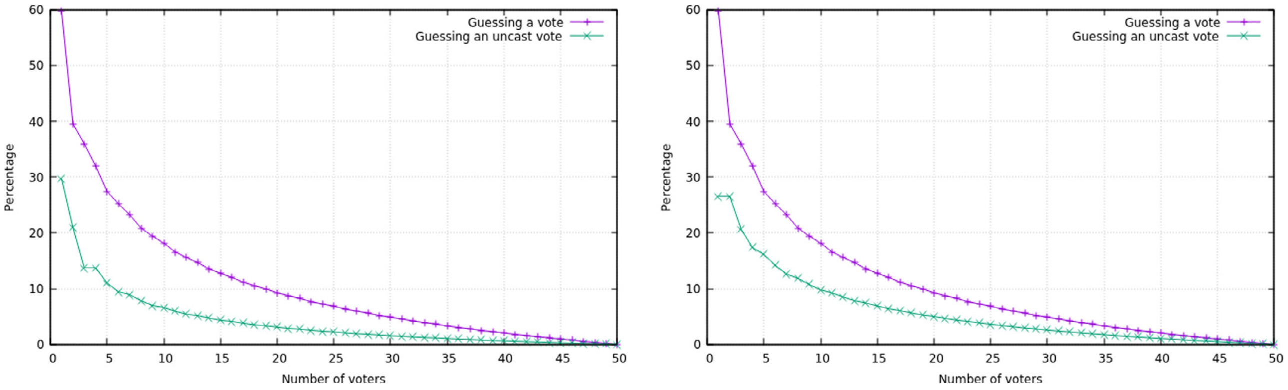

Fig. 3. Laws of Small Samples for W (left) (resp. T (right)) wrt $\alpha_\mathcal{E}$ (upper) and $\neg\alpha_\mathcal{E}$ (lower). Analysis run with 3 candidates.

ing [5] .[7] This unexpected relationship provides rich mathematical results that can be used to approximate the leakages for considerably large numbers of voters and candidates (see Appendix B and Appendix C for more details about the approximation). All the graphs in this paper were generated via these exact leakage formulas which can be found in the appendix.

In Fig. 2 at left we illustrate the average vulnerabilities for the privacy questions: "can the adversary guess which candidate a voter selected?" and "can the adversary guess which candidate a voter did not select?". The former corresponds to $V_{\alpha_\mathcal{E}}[\upsilon\triangleright W]$ and the latter $V_{\neg\alpha_\mathcal{E}}[\upsilon\triangleright W]$, where in both experiments we are using the channel W that releases only the winning candidate. Recall from Definition 3.2 and (2) that posterior vulnerabilities give the expected posterior vulnerability for a given privacy question — here this translates to computing the probabilities that an adversary can guess correctly the answer to the respective questions.

Since $T \sqsubseteq W$, we know that these are lower bounds on the more revealing reporting channel T which releases the tallies as well. For $V_{\alpha_\mathcal{E}}[\upsilon\triangleright W]$, we see that the vulnerability of guessing who voters voted for (i.e. the probability of correctly guessing the selected candidate on a given voter's ballot) rapidly approaches the prior vulnerability (of 0.33) (lower curve in Fig. 2, left graph), showing that privacy for this question behaves well in large reporting batches. However for the dual question $V_{\neg\alpha_\mathcal{E}}[\upsilon\triangleright W]$ the story is quite different: when the electorate consists of 50 voters there is still a 70% chance that an adversary can correctly guess who a voter did not vote for (upper curve in Fig. 2, left graph).

The graph on the right at Fig. 2 displays the leakages Definition 3.4 for the two privacy questions, illustrating the contribution of the channel W to the posterior vulnerabilities. Here we see that reporting the winner leaks more information about guessing who the voter voted for than who the voter did not vote for.

In Fig. 3 we see the law of small samples illustrated wrt W on the left and T on the right. In both graphs the upper curves correspond to the percentage loss in privacy for the question "can the adversary guess which candidate a voter selected?" and the lower curves for the question "can the adversary guess which candidate a voter did not select?" All curves are relative to the large sample size of 50 so that for the question "guess which candidate was selected" for both W and T, there is a greater than 30% loss in privacy in reporting batches consisting of fewer than 5 voters. By Theorem 5.2 these percentages are the same for this question.

However for the dual question "guess which candidate was not selected", as Theorem 5.3 suggests, there is a difference in the contribution to posterior vulnerability between the two channels W and T. For W (left-hand graph, lower curve) the contribution of W is approximately 10% for a reporting batch of size 5, whereas for T it is more than 15%. This increase means that the actual information contained in the tallies can be used effectively by the adversary to improve his ability to determine who voters did not vote for. In some election environments this could be problematic.

5.2. Instant run off elections

The second kind of reporting system we consider is instant run off. This is a system where voters are asked to rank candidates in order of preference. The counting then takes these preferences into account so that candidates are elected or eliminated after a number of rounds. In the first round the first preferences are tallied and any candidate who receives more than half of the votes is elected. If no candidate is elected in the first round then the candidate who received the least first preferences is eliminated, and their ballots are then re-distributed amongst the remaining candidates according to the second preference. This process is repeated until one of the candidates receives more than half the votes.

We define two reporting channels P and F associated with instant run-off. Whereas P publishes how many of each full preference ranking occurred, F publishes the aggregate first preferences only. We describe F because many elections in Australia do publish first preferences only and, although P is easier for us to analyse, we know that since F releases less information we must have $P \sqsubseteq F$. Therefore our privacy results for P give upper bounds for those of F.

Definition 5.2. Let $\mathcal{B}\in\mathcal{E}\to\mathcal{R}$ model sets of ballots; each ballot in $\mathcal{R}$ now gives a voter's ranked preference of candidates. Define $P\in\mathcal{B}\rightarrow\!\!\triangleright\,\mathbb{N}^{\mathcal{R}}$ to be the reporting channel that announces tallies for each possible ranking.

$$P_{bk}:=1 \quad \mathit{if} \quad (\forall r:\mathcal{R}\cdot\#_r(b)=k(r)) \quad \mathit{else} \quad 0 \;, \tag{16}$$

where $\#_r(b)$ is the total number of ballots that listed r as ranking in the instant runoff using ballots b.

Let $F\in\mathcal{B}\rightarrow\!\!\triangleright\,\mathbb{N}^{\mathcal{C}}$ model the reporting channel that publishes tallies for first preferences only.

$$F_{bf}:=1 \quad \mathit{if} \quad (\forall c:\mathcal{C}\cdot\#_c(b)=f(c)) \quad \mathit{else} \quad 0 \;, \tag{17}$$

[7] The intuition here is that, if voters are randomly choosing candidates then the election process is equivalent to throwing each of the n voters into a "bin" labelled by the selected candidate. This is analogous to storing n random strings into a table where two strings are inserted in the same column iff they have the same hash. We show in Appendix B that the posterior vulnerability is equal to the average maximal length of the columns of that table.

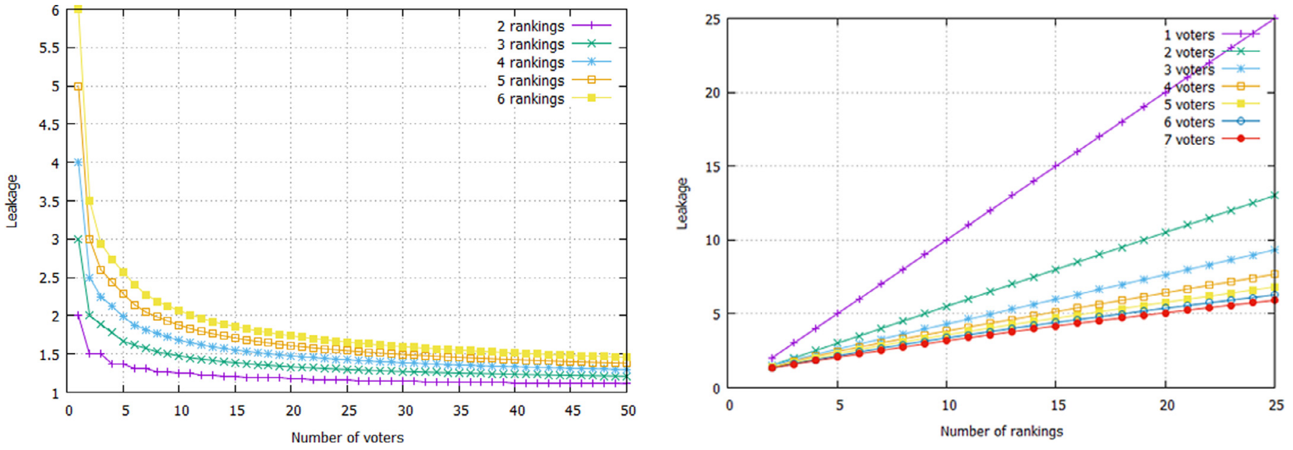

Fig. 4. (Left) Graph of $\mathcal{L}_{\alpha_\mathcal{E}}[\upsilon, P]$ with 2 to 6 possible rankings (2 to 3 candidates). (Right) Graph of $\mathcal{L}_{\alpha_\mathcal{E}}[\upsilon, P]$ with 1 to 7 voters and number of candidates ranging from 1 to 4 (i.e. 1 to 24 possible rankings).

where $\#_c(b)$ is the total number of first preferences that candidate c received in the instant runoff using ballots b.

A QiF analysis of P (and therefore also F) shows that it satisfies the law of large electorates for the privacy properties adapted to the type of secret. Now an adversary tries to guess the full preference list.

Theorem 5.4. *Reporting channels P and F satisfy the law of large electorates wrt $\alpha_\mathcal{E}$, $\#\alpha_\mathcal{E}$ and $\neg\alpha_\mathcal{E}$. Here the type of secret is now the whole preference list.*

Proof. For P this is a corollary of Theorem 5.1 adapted to the new types: Theorem 5.1 relies only on the fact that the number of candidates remains constant as the number of voters increased, and the same is true here.

Since $P \sqsubseteq F$ the leakages for P are greater than for F and therefore F also satisfies the law of large electorates. □

Although Theorem 5.4 shows that when the voting population is large compared to the observations in the channel, our experimental analysis shows that privacy is acutely sensitive to the number of observations. When tallies of full preference lists are announced, there are factorial $|\mathcal{C}|$ possible observations and, as pointed out elsewhere [15], this leaves open the possibility of using the information released to perform a signature attack.

5.2.1. Privacy risks in small samples

In Fig. 4 we illustrate some results on privacy for the instant run off election using the channel P. As before we assume that the votes are anonymised and that the aggregate rankings are reported, as defined at Definition 5.2. However the number of observations in the channel is equal to $|\mathcal{C}|!$ and even with only 4 candidates there are 24 possible ways to rank them.

In Fig. 4 (left) we display the leakages associated with privacy question Definition 3.7(1) trying to guess how a voter ranked the candidates. Recall that the leakages are a measure of how the channel P's information leak helps the adversary improve his guess.

The lowest curve corresponds to 2 candidates (therefore only two rankings) and the top-most curve corresponds to 3 candidates (therefore 6 possible rankings), with the curves in between illustrating leakages corresponding to numbers of observations between 2 and 6. As for first past the post we see that the leakages decrease as the number of voters increases: when there are 50 voters the leakages for all curves are below 1.5. However for about 5 voters the leakage is three times as much in the case that there are 3 candidates (leakage 2.5) compared to 2 candidates (0.8).

In Fig. 4 (right) we examine the relationship between numbers of voters versus numbers of candidates. The horizontal axis corresponds to the numbers of observations (related to $|\mathcal{C}|!$) and the vertical axis corresponds to the leakage. The top-most line is the leakage when there is 1 voter (thus his vote is revealed entirely by P) and the lowest line corresponds to 7 voters. The lines in between correspond to numbers of voters in between 1 and 7. These lines show that for small samples of voters and relatively small numbers of candidates, a great deal of privacy is potentially at risk: with 7 voters and 4 candidates the leakage is more than three times that for 50 voters and 3 candidates.

Finally in Fig. 5 (left) we illustrate similarly the leakages associated with privacy question Definition 3.7(3) (dual to those of Definition 3.7(1)) where the adversary tries to guess a ranking that a voter did not select. This time the lowest curve corresponds to leakage when there are 3 candidates and the upper when there are only 2, with the in between curves illustrating leakages corresponding to numbers of observations between 2 and 6. Dual to Fig. 4 (left), the leakage (and therefore contribution of P) decreases with increasing numbers of candidates. Moreover in Fig. 5 (right) where each curve (as for Fig. 4 (right)) corresponds to a fixed number of voters, ranging from 1 to 7, and the horizontal axis corresponds to numbers of observations. We can see clearly now that the contribution of P to this question is much less for this question than for Definition 3.7(1).

6. Discussion and conclusions

In this paper we have studied how voters' privacy is affected through reporting results in elections. Information leakage is an inevitable part of the election process, and whilst there is a common awareness that the results reporting can impact privacy, there has been limited formal study of the extent to which privacy is impacted.

Our goal has been to try to formalise notions of privacy using novel techniques in quantitative information flow which supports adversary-focussed modelling of scenarios through gain functions and their associated uncertainty measures. A significant feature of our approach is that it enables us to compare rigorously different

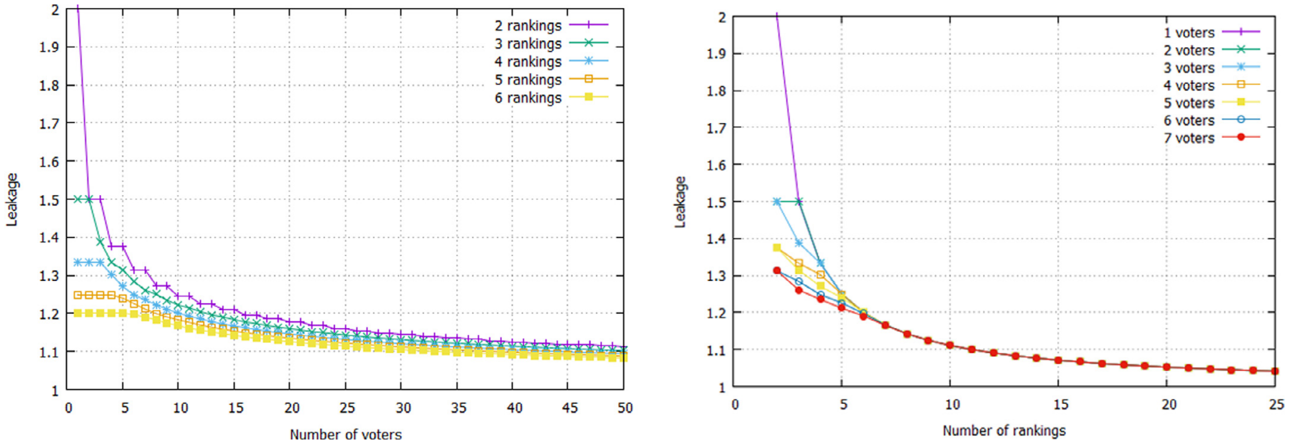

Fig. 5. (Left) Graph of $\mathcal{L}_{\neg\alpha_\varepsilon}[\upsilon, P]$ with 2 to 6 possible rankings (2 to 3 candidates). (Right) Graph of $\mathcal{L}_{\neg\alpha_\varepsilon}[\upsilon, P]$ with 1 to 7 voters and number of candidates ranging from 1 to 4 (i.e. 1 to 24 possible rankings).

rules for releasing information by controlling the observations in the definitions of channels.

As noted by other researchers [2] there are different kinds of privacy, all expressing different concerns. We concentrated on three questions, showed how they relate to each other and computed their relative values for differing scenarios of electorate and candidate sizes for two standard election styles: first past the post and instant run-off.

Our approach can be used to help inform decisions on privacy in reported election results. This principle can furthermore be improved if we have access to approximate leakage computations that generalise (15). [8] In Australia it is common for small batches of votes to be revealed because results are usually reported at a fine-grained level for transparency reasons. It is currently unclear what are the privacy risks.

An example is the 2015 results for the New South Wales Legislative Assembly,[9] which used instant run-off voting and included reports for first preferences by polling place and voting method.

Several voting methods were used by relatively few voters and thus were likely to have small batches of votes reported. One such method is mobile polling (also known as declared institution voting), which had around 14,000 votes spread across 93 electorates. These votes are reported per electorate, and so the expected batch size is 150 votes if the votes are spread uniformly. But in reality some batches were much smaller. For example the electorate of Canterbury had only 6 mobile polling votes: 5 votes for the Labour candidate, 1 vote for the Liberal candidate and no votes for the three other candidates.

Consolidated reporting is a countermeasure that is sometimes used in an attempt to mitigate the privacy risks of reporting these small batches. For example the 2011 results[10] consolidated Internet voting (iVote) with postal voting for each electorate; Internet voting was used for the first time and had only around 46,000 votes. Then the 2015 results reported Internet voting separately; this time around 280,000 were cast with this voting method.

At present such trade-offs between privacy and transparency are made based on intuition. This can result in inappropriate trade-offs because the impact on privacy can be counterintuitive. For example Theorem 5.2 shows that there are unexpected cases where reducing the information released would not reduce certain privacy risks. So perhaps the reduction in transparency from consolidated reporting of Internet voting and postal voting in 2011 was unnecessary, especially considering that other voting methods with fewer votes were reported separately. Or perhaps the separate reporting of Internet voting in 2015 still had substantial privacy risks despite the larger number of votes.

Applying our techniques would help to understand these privacy risks and enable evidence-based decisions by quantitatively assessing how much the risks would be reduced by particular countermeasures such as consolidated reporting.

Note however that consolidated reporting does not reduce the privacy risks wrt election officials and scrutineers, who have access to additional raw, finer-grained information. So analysing the reported election results provides a lower bound on the privacy risks.

Also we note that the calculations in this paper assume that the ballots are anonymised but distinct, as in there is a 1-1 correspondence between the voters on an electoral role and the ballots. But again election officials and scrutineers could have access to (partial) information about that correspondence.

Our framework could be adapted to modelling these “insider” adversaries by including additional observations in the reporting channel, and the contribution to privacy risks of that extra information could then be compared with the reported results.

Acknowledgement

We acknowledge support from the Australian Research Council DP140101119.

[8] Such a principle can be developed in the style of risk-limiting tally-audition techniques [10].

[9] http://pastvtr.elections.nsw.gov.au/SGE2015/la-home.htm.

[10] http://pastvtr.elections.nsw.gov.au/SGE2011/la_landing-fc.htm.

Detailed proofs

In the following appendices, we simply write $\upsilon{:}\mathbb{D}\mathcal{B}$ for the uniform prior over a fixed set of ballot boxes $\mathcal{B} = \mathcal{E} \to \mathcal{C}$ where $\mathcal{C}$ is the set of candidates and $\mathcal{E}$ is the set of electors or ballot identifiers.

Appendix A. Proof of Theorem 5.3

Theorem 5.3 Let T and W be reporting channels defined at Definition 5.1 and $\alpha_{\mathcal{E}}$, $\#\alpha_{\mathcal{E}}$ be gain functions defined at Definition 3.7. Let $\upsilon{:}\mathbb{D}\mathcal{B}$ be the uniform prior. The following equalities hold:

$\mathcal{L}_{\alpha_{\mathcal{E}}}[\upsilon, W] = \mathcal{L}_{\alpha_{\mathcal{E}}}[\upsilon, T]$ *and* $\mathcal{L}_{\#\alpha_{\mathcal{E}}}[\upsilon, W] = \mathcal{L}_{\#\alpha_{\mathcal{E}}}[\upsilon, T]$.

This theorem will follow from a series of intermediate results. In particular, it is a straightforward application Definition 3.4, Corollary A.6 and Lemma A.7.

Firstly, we assume that $\mathcal{E} = [0, 1, \ldots, n-1]$ and $\mathcal{C} = [\mathtt{0}, \mathtt{1}, \ldots, \mathtt{m}-1]$ so that there are exactly n voters and m candidates. Voters and candidates are thus numbered by integers and the ballot box where the votes are unanimously $\mathtt{0}$ (resp. $\mathtt{1}$) is denoted $\mathbf{0}$ (resp. $\mathbf{1}$). Secondly, notice that

$$\#\alpha_{\mathcal{E}}(b', b) = n - \mathbf{h}(b', b) \tag{18}$$

where $\mathbf{h}(b', b) = \sum_{e\in\mathcal{E}} 1$ *if* $b_e \neq b'$ *else* 0 is the Hamming distance between the ballot boxes b and b' when viewed as lists or strings. If b' is the guess of the attacker and the actual ballot box is b then $\mathbf{h}(b, b')$ is the number of votes that the attacker got wrong. This observation is particularly useful because $\mathbf{h}(,)$ is a proper metric.

The proof of Theorem 5.2 is achieved by explicitly computing the prior and posterior vulnerabilities and then deducing the ratio $\frac{\text{posterior vulnerability}}{\text{prior vulnerability}}$ which gives the leakage.

Computing prior vulnerabilities

Lemma A.1. *Let υ be the uniform prior on $\mathcal{B}$. then* $V_{\#\alpha_{\mathcal{E}}}[\upsilon] = \frac{n}{m}$.

Proof. We have

$$\begin{array}{lll}
 & V_{\#\alpha_{\mathcal{E}}}[\upsilon] & \\
= & \max_{b'\in\mathcal{B}} \sum_{b\in\mathcal{B}} \#\alpha_{\mathcal{E}}(b', b)\times\upsilon_b & \text{“Def. } V_{\#\alpha_{\mathcal{E}}}\text{”} \\
= & \frac{\max_{b'\in\mathcal{B}} \sum_{b\in\mathcal{B}}(n-\mathbf{h}(b',b))}{m^n} & \text{“(18) and } \upsilon_b = \frac{1}{m^n}\text{”} \\
= & \frac{\sum_{b\in\mathcal{B}}(n-\mathbf{h}(\mathbf{0},b))}{m^n} & \text{“See (†) below”} \\
= & \frac{\sum_{k=0}^{n}(n-k)\times|\{b \mid \mathbf{h}(\mathbf{0},b)=k\}|}{m^n} & \text{“}\bigcup_k\{b \mid \mathbf{h}(\mathbf{0}, b) = k\}\text{ is a partition of }\mathcal{B}\text{”} \\
= & \frac{\sum_{k=0}^{n}(n-k)\times\binom{n}{n-k}\times|\{b \mid b \text{ has } k \text{ non-zero entries}\}|}{m^n} & \text{“}\mathbf{h}(\mathbf{0}, b) = k\text{ means }b\text{ has }n-k\text{ zeros”} \\
= & \frac{\sum_{k=0}^{n}(n-k)\times\binom{n}{n-k}\times\sum_{y_1+\cdots+y_{\mathtt{m}-1}=k}\binom{k}{y_1\ldots y_{\mathtt{m}-1}}}{m^n} & \text{“Computing size of }\{b \mid b\text{ has }k\text{ non-zero entries}\}\text{”} \\
= & \frac{\sum_{y_0+\cdots+y_{\mathtt{m}-1}=n} y_0\times\binom{n}{y_0\ldots y_{\mathtt{m}-1}}}{m^n} & \text{“Arith.”} \\
= & \frac{\sum_{y_0-1+\cdots+y_{\mathtt{m}-1}=n} n\times\binom{n-1}{y_0-1\ldots y_{\mathtt{m}-1}}}{m^n} & \text{“Arith.”} \\
= & \frac{nm^{n-1}}{m^n} & \text{“Multinomial Theorem”} \\
= & \frac{n}{m} & \text{“Arith.”}
\end{array}$$

(†) This invariance wrt the guess follows from the fact that $\sum_x \mathbf{h}(w_1, x) = \sum_x \mathbf{h}(w_2, x)$ when w_1 and w_2 only differ at on place. By transitivity, we have that $\sum_x \mathbf{h}(w_1, x) = \sum_x \mathbf{h}(\mathbf{0}, x)$ for every guess w. □

Similarly, for the gain function $\alpha_{\mathcal{E}}$, we have

Lemma A.2. *Let υ be the uniform prior on $\mathcal{B}$. then* $V_{\alpha_{\mathcal{E}}}[\upsilon] = \frac{1}{m}$.

Proof. $V_{\alpha_{\mathcal{E}}}[\upsilon]{=}\max_{e,c} \sum_{b\in\mathcal{B}}(1 \text{ if } b_e = c \text{ else } 0)\times\upsilon_b{=}\frac{\max_c \sum_{b\in\mathcal{B}}(1 \text{ if } b_0=c \text{ else } 0)}{|\mathcal{E}|}{=}\frac{m^{n-1}}{m^n}{=}\frac{1}{m}$ □

Computing posterior vulnerabilities

Computing the posterior vulnerabilities is more involved. We start with the gain function $\#\alpha_{\mathcal{E}}$. Firstly, we show that any majority gives the optimal guess in the computation of $V_{\#\alpha_{\mathcal{E}}}[\upsilon{\triangleright}T]$. We focus on the reporting channel T that publishes the tallies in the following two lemmas.

Lemma A.3. *Let $y = (y_0, \ldots, y_{\mathtt{m}-1})$ be an observed tally where $y_{\mathtt{i}}$ is the tally for candidate $\mathtt{i}$ and assume wlog that* $\max(y) = y_0$[11]. *Then*

$$\max_{b'\in\mathcal{B}} \sum_{b\in\mathcal{B}} \#\alpha_{\mathcal{E}}(b', b)\times T_{by} = \sum_{b\in\mathcal{B}} \#\alpha_{\mathcal{E}}(\mathbf{0}, b)\times T_{by} \quad . \tag{19}$$

[11] This assumption says that candidate $\mathtt{0}$ has a majority, which can be possibly tied with other candidates.

Proof. Recall that $\#\alpha_{\mathcal{E}}(b', b) = n - \mathbf{h}(b', b)$, thus

$$\max_{b' \in \mathcal{B}} \sum_{b \in \mathcal{B}} \#\alpha_{\mathcal{E}}(b', b) \times T_{by} = \sum_{b \in \mathcal{B}} nT_{by} - \min_{b' \in \mathcal{B}} \sum_{b \in \mathcal{B}} \mathbf{h}(b', b) \times T_{by}$$

which is equivalent to $\mathbf{0}$ minimizing the negative term. It is thus enough to show that

$$\sum_{b \in \{x | T_{xy}=1\}} \mathbf{h}(b', b) \geq \sum_{b \in \{x | T_{xy}=1\}} \mathbf{h}(\mathbf{0}, b)$$

for every guess $b' \in \mathcal{B}$. We reason by induction on the number of voters n.

$n = 1$: there is a single voter which must have voted for candidate $\mathtt{0}$. In this case, the best guess is $\mathtt{0}$.
$n > 1$: let us assume that at least one voter has NOT voted $\mathtt{0}$ (otherwise, we would see in the published tally that $\mathtt{0}$ is elected unanimously and the best strategy b' is trivially $\mathbf{0}$). *wlog*, we can assume that the first entry of b' is $\mathtt{1}$. The sum $\sum_{b \in \{x | T_{xy}=1\}} \mathbf{h}(w, x)$ is split into three disjoint sub-sums according to whether $b_0 = b'_0$, or $b_0 \neq b'_0$ and $b_0 = \mathtt{0}$, or $b_0 \neq b'_0$ and $b_0 \neq \mathtt{0}$:

1. $b_0 = b'_0 = \mathtt{1}$ and thus $\mathbf{h}(b', b) = \mathbf{h}(b'^*, b^*)$ where $b' = \mathtt{1}{:}b'^*$ and $b = \mathtt{1}{:}b^*$. Let $y^* = (y_0, y_1 - 1, \ldots, y_{\mathtt{m}})$ so that $T_{b^* y^*} = 1$ and $\max(y^*) = y_0$. On the one hand, by the Induction Hypothesis, we have

   $$\sum_{b^* \in \{x | T_{xy^*}=1\}} \mathbf{h}(b'^*, b^*) \geq \sum_{b^* \in \{x | T_{xy^*}=1\}} \mathbf{h}(\mathbf{0}, b^*)$$

   for every b'^*. On the other hand, $\mathbf{h}(\mathbf{0}, b) = \mathbf{h}(\mathbf{0}, \mathtt{1}{:}b^*) = 1 + \mathbf{h}(\mathbf{0}, b^*)$. Thus

   $$\sum_{b \in \{x | T_{xy}=1 \wedge x = \mathtt{1}:x^*\}} \mathbf{h}(b', b) \geq \sum_{b^* \in \{x | T_{xy^*}=1\}} \mathbf{h}(\mathbf{0}, b^*) = \sum_{b \in \{x | T_{xy}=1 \wedge x = \mathtt{1}:x^*\}} (\mathbf{h}(\mathbf{0}, b) - 1)$$

2. $b_0 \neq b'_0$ and $b_0 = \mathtt{0}$ and thus $\mathbf{h}(b', b) = 1 + \mathbf{h}(b'^*, b^*)$ where $b' = \mathtt{1}{:}b'^*$ and $b = \mathtt{0}{:}b^*$. This case is furthermore split into two subcases according to the tally $y^* = (y_0 - 1, y_1, \ldots, y_{\mathtt{m}})$ corresponding to this case:
   (a) $\max(y^*) = y^*_0 = y_0 - 1$: in this case, we can apply the Induction Hypothesis and obtain

   $$\sum_{b^* \in \{x | T_{xy^*}=1\}} \mathbf{h}(b'^*, b^*) \geq \sum_{b^* \in \{x | T_{xy^*}=1\}} \mathbf{h}(\mathbf{0}, b^*) \quad ,$$

   for every b'^*. Since $b = \mathtt{0}{:}b^*$, we have $\mathbf{h}(\mathbf{0}, b^*) = \mathbf{h}(\mathbf{0}, b)$ and therefore

   $$\sum_{b \in \{x | T_{xy}=1 \wedge x = \mathtt{0}:x^*\}} \mathbf{h}(b', b) \geq \sum_{b^* \in \{x | T_{xy^*}=1\}} (1 + \mathbf{h}(\mathbf{0}, b^*)) = \sum_{b \in \{x | T_{xy}=1 \wedge x = \mathtt{0}:x^*\}} (1 + \mathbf{h}(\mathbf{0}, b))$$

   (b) $\max(y^*) = y_1 > y_0 - 1$: in this case, the Induction Hypothesis corresponds to choosing all $\mathtt{1}$, i.e.,

   $$\sum_{b^* \in \{x | T_{xy^*}=1\}} \mathbf{h}(b'^*, b^*) \geq \sum_{b^* \in \{x | T_{xy^*}=1\}} \mathbf{h}(\mathbf{1}, b^*) \quad ,$$

   for all b'^*. Moreover, the maximality of y_0 in y implies that $y_0 = y_1$. Therefore, $\mathbf{h}(\mathbf{0}, b) = \mathbf{h}(\mathbf{1}, b) = \mathbf{h}(\mathbf{1}, b^*)$ for all ballot box b satisfying this case number 2. Hence, as in the case 2.(a), we obtain

   $$\sum_{b \in \{x | T_{xy}=1 \wedge x = \mathtt{0}:x^*\}} \mathbf{h}(b', b) \geq \sum_{b \in \{x | T_{xy}=1 \wedge x = \mathtt{0}:x^*\}} (1 + \mathbf{h}(\mathbf{0}, b))$$

3. $b_0 \neq b'_0$ and $b_0 \neq \mathtt{0}$ and thus $\mathbf{h}(b', b) = \mathbf{h}(b'^*, b^*)$ where $b' = \mathtt{1}{:}b'^*$ and $b = \mathtt{2}{:}b^*$ or $b = \mathtt{3}{:}b^*$ or...or $b = \mathtt{m}{:}b^*$. Since y_0 remains the maximum of the y's associated to this case, we deduce from multiple application of the Induction Hypothesis and some arithmetic that, for every guess b',

   $$\begin{array}{lll} & \sum_{b \in \{x | T_{xy}=1 \wedge x_0 \neq \mathtt{0} \wedge x_0 \neq \mathtt{1}\}} \mathbf{h}(b', b) & \\ = & \sum_{\mathtt{c} \neq \mathtt{0} \wedge \mathtt{c} \neq \mathtt{1}} \sum_{b \in \{x | T_{xy}=1 \wedge x = \mathtt{c}:x^*\}} \mathbf{h}(b'^*, b^*) & \\ \geq & \sum_{\mathtt{c} \neq \mathtt{0} \wedge \mathtt{c} \neq \mathtt{1}} \sum_{b \in \{x | T_{xy}=1 \wedge x = \mathtt{c}:x^*\}} \mathbf{h}(\mathbf{0}, b^*) & \text{“Induction Hypothesis”} \\ = & \sum_{b \in \{x | T_{xy}=1 \wedge x_0 \neq \mathtt{0} \wedge x_0 \neq \mathtt{1}\}} \mathbf{h}(\mathbf{0}, b) & \end{array}$$

Now, using some counting technique, notice that case 1 contributes $-\binom{n-1}{y_0, y_1 - 1, \ldots, y_{\mathtt{m}}}$ to the aggregated sum while case 2 contributes $+\binom{n-1}{y_0 - 1, y_1, \ldots, y_{\mathtt{m}}}$ to it. The case 3 does not have any extra contribution. By hypothesis, $y_1 \leq y_0$ which ensures that $\binom{n-1}{y_0 - 1, y_1, \ldots, y_{\mathtt{m}}} - \binom{n-1}{y_0, y_1 - 1, \ldots, y_{\mathtt{m}}} \geq 0$, i.e., the aggregated extra contribution from these two cases is positive. By summing up the above three disjoint case, we obtain the desired result:

$$\sum_{b \in \{x | T_{xy}=1\}} \mathbf{h}(b', b) \geq \sum_{b \in \{x | T_{xy}=1\}} \mathbf{h}(\mathbf{0}, b)$$

for every guess b'.

To prove inequality (19), we can choose any tally y such that $R_{y\mathtt{0}} = 1$ where R is the channel satisfying $W = T{\cdot}R$. In this case, The previous reasoning can be applied to channel T and the fixed tally y such that $\max(y) = y_0$ which again makes $\mathbf{0}$ an optimal guess □

Lemma A.4. *Let υ be the uniform prior on $\mathcal{B}$. We have*

$$V_{\#\alpha_{\mathcal{E}}}[\upsilon \rhd T] = \frac{1}{m^n} \sum_{y \in \mathbb{N}^{\mathcal{C}}} \max(y) \binom{n}{y} \tag{20}$$

where $\binom{n}{y} = \binom{n}{y_0,\ldots,y_{\mathtt{m}-1}}$ *is the multinomial coefficient and we assume that* $\binom{n}{y} = 0$ *if* $\Sigma_i y_i \neq n$*.*

Proof. We reason as follows

$$\begin{array}{lll}
 & V_{\#\alpha_{\mathcal{E}}}[\upsilon \rhd T] & \\
= & \sum_{y\in\mathbb{N}^{\mathcal{C}}} \max_{b'\in\mathcal{B}} \sum_{b\in\mathcal{B}} \#\alpha_{\mathcal{E}}(b',b)\times T_{by}\times\upsilon_b & \text{“Def. } V_{\#\alpha_{\mathcal{E}}}\text{”} \\
= & \frac{1}{m^n} \sum_{y\in\mathbb{N}^{\mathcal{C}}} \max_{b'\in\mathcal{B}} \sum_{b\in\mathcal{B}} (n-\mathbf{h}(b',b))\times T_{by} & \text{“}\upsilon\text{ is uniform and } \#\alpha_{\mathcal{E}}(b',b) = n-\mathbf{h}(b',b)\text{”} \\
= & \frac{1}{m^n} \sum_{y\in\mathbb{N}^{\mathcal{C}}} \sum_{b\in\mathcal{B}} (n-\mathbf{h}(\mathbf{i}_{\mathbf{y}},b))\times T_{by} & \text{“By Lem. A.3. Let } \mathbf{i}_y = (i_y,\ldots,i_y) \text{ such that } \max(y) = y_{i_y}\ (\dagger)\text{”} \\
= & \frac{1}{m^n} \sum_{y\in\mathbb{N}^{\mathcal{C}}} \sum_{b\in\mathcal{B}} \max(y)\times T_{by} & \text{“if } T_{by} > 0 \text{ then } n-\mathbf{h}(\mathbf{i}_y,b) = y_{\mathbf{i}_y} = \max(y)\text{”} \\
= & \frac{1}{m^n} \sum_{y\in\mathbb{N}^{\mathcal{C}}} \max(y) \sum_{b\in\mathcal{B}} T_{by} & \text{“Arith.”} \\
= & \frac{1}{m^n} \sum_{y\in\mathbb{N}^{\mathcal{C}}} \max(y) \binom{n}{y} & \text{“There are exactly } \binom{n}{y} = \binom{n}{y_0\ldots y_{\mathtt{m}-1}} \text{ different } b\text{'s for which } T_{by} = 1\text{ .”}
\end{array}$$

In the above reasoning, $y = (y_0,\ldots,y_{\mathtt{m}-1})$ is the observed tally. □

Now, we state the result for the reporting channel W: (20) gives a lower bound for the vulnerability of W wrt $\#\alpha_{\mathcal{E}}$.

Lemma A.5. *Let υ be the uniform prior on $\mathcal{B}$. We have*

$V_{\#\alpha_{\mathcal{E}}}[\upsilon \rhd T] \leq V_{\#\alpha_{\mathcal{E}}}[\upsilon \rhd W]$.

Proof. It suffices to show that the lower bound is achieved by guessing the published winner.

$$\begin{array}{lll}
 & V_{\#\alpha_{\mathcal{E}}}[\upsilon \rhd W] & \\
= & \sum_{c\in\mathcal{C}} \max_{b'\in\mathcal{B}} \sum_{b\in\mathcal{B}} \#\alpha_{\mathcal{E}}(b',b)\times W_{bc}\times\upsilon_b & \text{“Def. (2)”} \\
\geq & \sum_{c\in\mathcal{C}} \sum_{b\in\mathcal{B}} \#\alpha_{\mathcal{E}}(\mathbf{c},b)\times W_{bc}\times\upsilon_b & \text{“Let } b' = \underbrace{(c,\ldots,c)}_{n} = \mathbf{c}\text{”} \\
= & \frac{1}{m^n} \sum_{c\in\mathcal{C}} \sum_{b\in\mathcal{B}} (n-\mathbf{h}(\mathbf{c},b))\times W_{bc} & \text{“}\upsilon_b = \tfrac{1}{m^n} \text{ and } \#\alpha_{\mathcal{E}}(b',b) = n-\mathbf{h}(b',b)\text{”} \\
= & \frac{1}{m^n} \sum_{c\in\mathcal{C}} \sum_{b\in\mathcal{B}} (n-\mathbf{h}(\mathbf{c},b))\times \sum_{y\in\mathbb{N}^{\mathcal{C}}} T_{by}\times R_{yc} & \text{“}W = T{\cdot}R\text{”} \\
= & \frac{1}{m^n} \sum_{y\in\mathbb{N}^{\mathcal{C}}} \sum_{c\in\mathcal{C}} R_{yc}\times \sum_{b\in\mathcal{B}} (n-\mathbf{h}(\mathbf{c},b))\times T_{by} & \text{“Arith.”} \\
= & \frac{1}{m^n} \sum_{y\in\mathbb{N}^{\mathcal{C}}} \sum_{b\in\mathcal{B}} (n-\mathbf{h}(\mathbf{c}_y,b))\times T_{by} & \text{“Let } \mathbf{c} = \mathbf{c}_y \text{ such that } R_{yc} = 1\text{”} \\
= & V_{\#\alpha_{\mathcal{E}}}[\upsilon \rhd T] & \text{“This is } (\dagger) \text{ in the proof of Lem. A.4”}
\end{array}$$

□

Corollary A.6. *Let υ be the uniform prior on $\mathcal{B}$. We have*

$V_{\#\alpha_{\mathcal{E}}}[\upsilon \rhd T] = V_{\#\alpha_{\mathcal{E}}}[\upsilon \rhd W]$.

Proof. Lemma A.5 gives one inequality while $T \sqsubseteq W$ and Definition 3.5 gives the other. □

Now, let us prove the other part of Theorem 5.2. We are going to prove a much more interesting result giving an direct correspondence between $\#\alpha_{\mathcal{E}}$ and $\alpha_{\mathcal{E}}$.

Lemma A.7. *Let υ be the uniform prior on $\mathcal{B}$. We have $V_{\#\alpha_{\mathcal{E}}}[\upsilon \rhd T] = n\times V_{\alpha_{\mathcal{E}}}[\upsilon \rhd T]$ and $V_{\#\alpha_{\mathcal{E}}}[\upsilon \rhd W] = n\times V_{\alpha_{\mathcal{E}}}[\upsilon \rhd W]$.*

Proof. Let us work with the reporting channel T first. We reason as follows

$$\begin{array}{lll}
 & V_{\alpha_{\mathcal{E}}}[\upsilon \rhd T] & \\
= & \sum_{y\in\mathbb{N}^{\mathcal{C}}} \max_{e\in\mathcal{E},c\in\mathcal{C}} \sum_{b\in\mathcal{B}} \alpha_{\mathcal{E}}((c,e),b)\times T_{by}\times\upsilon_b & \text{“Def. (2)”} \\
= & \frac{1}{m^n} \sum_{y\in\mathbb{N}^{\mathcal{C}}} \max_{e\in\mathcal{E},c\in\mathcal{C}} \sum_{b\in\mathcal{B}} (1 \text{ if } b_e = c \text{ else } 0)\times T_{by} & \text{“}\upsilon_b = \tfrac{1}{m^n}\text{, Def. } \alpha_{\mathcal{E}}\text{”} \\
= & \frac{1}{m^n} \sum_{y\in\mathbb{N}^{\mathcal{C}}} \max_{c\in\mathcal{C}} \sum_{b\in\mathcal{B}} (1 \text{ if } b_0 = c \text{ else } 0)\times T_{by} & \text{“Symmetry wrt. optimisation variable } e\text{”} \\
= & \frac{1}{m^n} \sum_{y\in\mathbb{N}^{\mathcal{C}}} \sum_{b\in\mathcal{B}} (1 \text{ if } b_0 = i_y \text{ else } 0)\times T_{by} & \text{“Let } i_y \text{ such that } \max(y) = y_{i_y}\text{”} \\
= & \frac{1}{m^n} \sum_{y\in\mathbb{N}^{\mathcal{C}}} |\{b | b_0 = i_y \text{ and } T_{by} = 1\}| & \text{“Rewriting } \sum_{b\in\mathcal{B}} (1 \text{ if } b_0 = i_y \text{ else } 0)\times T_{by}\text{”} \\
= & \frac{1}{m^n} \sum_{y\in\mathbb{N}^{\mathcal{C}}} \binom{n-1}{y_0\ldots y_i-1\ldots y_{\mathtt{m}-1}} & \text{“Counting”} \\
= & \frac{1}{m^n} \sum_{y\in\mathbb{N}^{\mathcal{C}}} \frac{\max(y)}{n} \binom{n}{y_0\ldots y_{\mathtt{m}-1}} & \text{“for } y_{i_y} = \max(y) \text{ we have } \binom{n-1}{y_0\ldots y_{i_y}-1\ldots y_{\mathtt{m}}} = \frac{\max(y)}{n}\binom{n}{y_0\ldots y_{\mathtt{m}-1}}\text{”} \\
= & \frac{V_{\#\alpha_{\mathcal{E}}}[\upsilon \rhd T]}{n} & \text{“Lem. A.4”}
\end{array}$$

The proof of the same result for W is similar. □

Appendix B. Proof of Theorem 5.1

Theorem 5.1 Both T and W satisfy the law of large electorates wrt $\alpha_{\mathcal{E}}$, $\#\alpha_{\mathcal{E}}$ and $\neg\alpha_{\mathcal{E}}$.

Proof. Let us start by proving that T satisfies the law of large electorates wrt $\#\alpha_{\mathcal{E}}$. Lemma A.4 tells us that

$$V_{\#\alpha_{\mathcal{E}}}[\upsilon \triangleright T] = \sum_{y_0+\cdots+y_{m-1}=n} \max(y) \frac{\binom{n}{y_0 \ldots y_{m-1}}}{m^n}$$

which is the expected value of the random variable $\max : \mathbb{N}^{\mathcal{C}} \rightarrow \mathbb{R}$ wrt the multinomial distribution. This expected value appears quite often in “balls into bins” combinatoric games and asymptotic behaviours have been given [8,9,12]. In particular, we use Theorem 1 in [12] [12] which implies that when the number of candidates $|\mathcal{C}| = m \geq 3$ [13] is fixed and the number of voters $|\mathcal{E}| = n$ is large enough $(n \gg m)$ [14], we have

$$V_{\#\alpha_{\mathcal{E}}}[\upsilon \triangleright T] = \frac{n}{m} + \sqrt{\frac{2n \ln m}{m}\left(1 - O\left(\frac{1}{\ln n}\right)\right)} \quad . \tag{21}$$

Thus, since $\mathcal{L}_{\#\alpha}[\upsilon, T] = \frac{V_{\#\alpha_{\mathcal{E}}}[\upsilon \triangleright T]}{V_{\#\alpha_{\mathcal{E}}}[\upsilon]}$, we have

$$\lim_{|\mathcal{E}| \rightarrow \infty} \frac{V_{\#\alpha_{\mathcal{E}}}[\upsilon \triangleright T]}{V_{\#\alpha_{\mathcal{E}}}[\upsilon]} = \lim_{n \rightarrow \infty} \frac{\frac{n}{m} + \sqrt{\frac{2n \ln m}{m}\left(1 - O\left(\frac{1}{\ln n}\right)\right)}}{\frac{n}{m}} = 1 + \lim_{n \rightarrow \infty} \sqrt{\frac{2m \ln m}{n}} = 1 \tag{22}$$

It follows from (22), Lemma A.7 and Theorem 5.2 that T and W satisfy the law of large electorates wrt $\#\alpha_{\mathcal{E}}$ and $\alpha_{\mathcal{E}}$. In fact, the term $\sqrt{\frac{2m \ln m}{n}}$ provides an adequate approximation of the residual leakage when $n \gg m$.

Now, let us prove that T satisfy the same law but wrt the privacy question given by $\neg\alpha_{\mathcal{E}}$. We start by computing the prior vulnerability

$$V_{\neg\alpha_{\mathcal{E}}}[\upsilon] = \max_{e \in \mathcal{E}, c \in \mathcal{C}} \sum_{b \in \mathcal{B}} (1 \text{ if } b_e \neq c \text{ else } 0) \times \upsilon_b = \frac{\sum_{b \in \mathcal{B}} (1 \text{ if } b_0 \neq 0 \text{ else } 0)}{m^n} = \frac{m-1}{m}$$

Next, we compute the posterior vulnerability.

$$\begin{array}{rlr}
 & V_{\neg\alpha_{\mathcal{E}}}[\upsilon \triangleright T] & \\
= & \sum_{y \in \mathbb{N}^{\mathcal{C}}} \max_{e \in \mathcal{E}, c \in \mathcal{C}} \sum_{b \in \mathcal{B}} \neg\alpha_{\mathcal{E}}((c, e), b) \times T_{by} \times \upsilon_b & \text{“Def. } V_{\neg\alpha_{\mathcal{E}}}[\upsilon \triangleright T]\text{”} \\
= & \frac{1}{m^n} \sum_{y \in \mathbb{N}^{\mathcal{C}}} \max_{c \in \mathcal{C}} \sum_{b \in \mathcal{B}} (1 - \alpha_{\mathcal{E}}((c, 0), b)) \times T_{by} & \text{“}\upsilon \text{ is uniform, symmetry wrt. } e \text{ and } \neg\alpha_{\mathcal{E}} = 1 - \alpha_{\mathcal{E}}\text{”} \\
= & 1 - \frac{1}{m^n} \sum_{y \in \mathbb{N}^{\mathcal{C}}} \min_{c \in \mathcal{C}} \sum_{b \in \mathcal{B}} \alpha_{\mathcal{E}}((c, 0), b) \times T_{by} & \text{“Arith.”} \\
= & 1 - \frac{1}{m^n} \sum_{y \in \mathbb{N}^{\mathcal{C}}} \sum_{b \in \mathcal{B}} (1 \text{ if } b_0 = i_y \text{ else } 0) \times T_{by} & \text{“Let } i_y \text{ such that } \min(y) = y_{i_y}. \text{ Def. } \alpha_{\mathcal{E}} \ (\dagger)\text{”} \\
= & 1 - \frac{1}{m^n} \sum_{y \in \mathbb{N}^{\mathcal{C}}} |\{b | b_0 = i_y \text{ and } T_{by} = 1\}| & \text{“Rewrite } \sum_{b \in \mathcal{B}} (1 \text{ if } b_0 = i_y \text{ else } 0) \times T_{by}\text{”} \\
= & 1 - \frac{1}{m^n} \sum_{y \in \mathbb{N}^{\mathcal{C}}} \frac{\min(y)}{n} \binom{n}{y_0 \ldots y_{m-1}} & \text{“Counting and } \min(y) = y_{i_y}\text{”} \\
= & 1 - \sum_{y \in \mathbb{N}^{\mathcal{C}}} \frac{\min(y)}{n} \frac{\binom{n}{y_0 \ldots y_{m-1}}}{m^n} & \text{“Arith.”}
\end{array}$$

(†) This step says that the best guess is to choose one candidate with the least tally which is available since T publishes the tallies [15]. A rigorous and detailed proof can be achieved as in the proof of Lemma 19.

In the last expression, we again have to compute the expected value of the random variable $\min : \mathbb{N}^{\mathcal{C}} \rightarrow \mathbb{R}$ wrt the multinomial distribution, i.e.

$$\sum_{y \in \mathbb{N}^{\mathcal{C}}} \min(y) \frac{\binom{n}{y_0 \ldots y_{m-1}}}{m^n}$$

By symmetry, this expected value must converge to $\frac{n}{m}$ as well since all candidates should have the same number of votes in average (due to uniformity). This implies that

$$\lim_{n \rightarrow \infty} 1 - \sum_{y \in \mathbb{N}^{\mathcal{C}}} \frac{\min(y)}{n} \frac{\binom{n}{y_0 \ldots y_{m-1}}}{m^n} = 1 - \frac{1}{m} = V_{\neg\alpha_{\mathcal{E}}[\upsilon]}$$

Thus, T satisfies the law of large electorates wrt $\neg\alpha_{\mathcal{E}}$. Since $T \sqsubseteq W$, it follows that W also satisfies that law wrt $\neg\alpha_{\mathcal{E}}$. □

Appendix C. Computing leakages efficiently

For really large values of m (number of candidates) and n (number of voters), computing the various exact leakages is computationally very expensive. For large m and n, approximations such as (22) provide accurate results. For smaller but non-trivial values of n and m, the exact leakage can be calculated by finding a way to remove any term involving a sum over $\mathcal{B}$ which is of the order of m^n. For instance, Lemma A.4 involves a sum over $\{y | \mathbb{N}^{\mathcal{C}}$ and $\sum_i y_i = n\}$ which has $\binom{n+m-1}{m-1}$ elements instead.

- To compute $\mathcal{L}_{\alpha_{\mathcal{E}}}[\upsilon, C]$, for $C \in \{T, W\}$, we use Lemma A.4 and Theorem 5.1.
- To compute $\mathcal{L}_{\neg\alpha_{\mathcal{E}}}[\upsilon, T]$, we use the derivation in the proof of Theorem 5.1.

[12] Raab and Steger use m as our n and n as our m.

[13] For $m = 2$, an application of Sterling's approximation gives $V_{\#\alpha_{\mathcal{E}}}[\upsilon \triangleright T] \approx \sqrt{\frac{n}{2\Pi}}$ for n large and where $\Pi \approx 3.14$.

[14] The exact condition is that $n \gg m(\ln m)^3$ but it will be satisfied when n goes to infinity and m remains fixed.

[15] For W, this strategy doesn't work anymore since all that is published is the winner. Thus the best the adversary can do is to “not choose the winner”.

- To compute $\mathcal{L}_{\neg\alpha_{\mathcal{E}}}[\upsilon, W]$, we start with the equation

$$V_{\neg\alpha_{\mathcal{E}}}[\upsilon \triangleright W] = 1 - \frac{1}{m^n} \sum_{w\in\mathcal{C}} \min_{c\in\mathcal{C}} \sum_{b\in\mathcal{B}} \alpha_{\mathcal{E}}((c,0),b) \times W_{bw}$$

which could be derived in the same way as in Appendix B. Now, using the fact that $W = T{\cdot}R$ and a bit of rewriting and counting, we have

$$\begin{array}{rll} & V_{\neg\alpha_{\mathcal{E}}}[\upsilon\triangleright W] & \\ = & 1 - \frac{1}{m^n} \sum_{w\in\mathcal{C}} \min_{c\in\mathcal{C}} \sum_{y\in\mathbb{N}^{\mathcal{C}}} \sum_{b\in\mathcal{B}} (1 \text{ if } b_0 = c \text{ else } 0) \times \sum_{y\in\mathbb{N}^{\mathcal{C}}} T_{by} \times R_{yw} & \text{“Def. } \alpha_{\mathcal{E}}, W = T{\cdot}R\text{”} \\ = & 1 - \frac{1}{m^n} \sum_{w\in\mathcal{C}} \min_{c\in\mathcal{C}} \sum_{y\in\mathbb{N}^{\mathcal{C}}} |\{b | b_0 = c \text{ and } T_{by} = 1 \text{ and } R_{yw} = 1\}| & \text{“Rewriting”} \\ = & 1 - \frac{1}{m^n} \sum_{w\in\mathcal{C}} \min_{c\in\mathcal{C}} \sum_{y\in\mathbb{N}^{\mathcal{C}}}^{R_{yw}=1} \frac{y_c}{n} \binom{n}{y} & \text{“Counting”} \end{array}$$

This last formula is faster to compute since it does not involve any sum over $\mathcal{B}$ anymore.

References

[1] Alvim MS, Chatzikokolakis K, McIver A, Morgan C, Palamidessi C, Smith G. Additive and multiplicative notions of leakage, and their capacities. In: IEEE 27th computer security foundations symposium, CSF 2014, Vienna, Austria, 19–22 July, 2014. IEEE; 2014. p. 308–22.

[2] Alvim MS, Chatzikokolakis K, Palamidessi C, Smith G. Measuring information leakage using generalized gain functions. Proceedings of the 25th IEEE computer security foundations symposium (CSF 2012); 2012. p. 265–79.

[3] Benaloh J, Moran T, Naish L, Ramchen K, Teague V. Shuffle-sum: coercion-resistant verifiable tallying for stv voting. IEEE Trans Inf Forensics Secur 2009;4(4):685–98.

[4] Bernhard D, Cortier V, Pereira O, Warinschi B. Measuring vote privacy, revisited. In: Yu T, Danezis G, Gligor VD, editors. The ACM conference on computer and communications security, CCS'12, Raleigh, NC, USA, October 16–18, 2012. ACM; 2012. p. 941–52.

[5] Czumaj A, Stemann V. Randomized allocation processes. In: Proceedings of the 38th annual symposium on foundations of computer science. FOCS '97, pages 194–, Washington, DC, USA; 1997. IEEE Computer Society

[6] Cosmo R.D. On privacy and anonymity in electronic and non electronic voting: the ballot-as-signature attack2007.

[7] Dwork C. Differential privacy. Automata, languages and programming, 33rd international colloquium, ICALP 2006, Venice, Italy, July 10–14, 2006, Proceedings, Part II; 2006. p. 1–12.

[8] Gonnet GH. Expected length of the longest probe sequence in hash code searching. J ACM 1981;28:289–304.

[9] Johnson NL, Kotz S. Urn models and their application : an approach to modern discrete probability theory. Wiley series in probability and mathematical statistics. Wiley, New York; 1977.

[10] Lindeman M, Stark PB. A gentle introduction to risk-limiting audits. IEEE Secur Privacy 2012;10(5):42–9.

[11] McIver A, Meinicke L, Morgan C. Compositional closure for Bayes Risk in probabilistic noninterference. Automata, languages and programming, 37th international colloquium, ICALP 2010, Bordeaux, France, July 6–10, 2010, Proceedings, Part II; 2010. p. 223–35.

[12] Raab M, Steger A. "Balls into bins" – a simple and tight analysis. In: Proceedings of the second international workshop on randomization and approximation techniques in computer science, RANDOM '98. London, UK, UK: Springer-Verlag; 1998. p. 159–70.

[13] Smith G. On the foundations of quantitative information flow. Proc. 12th international conference on foundations of software science and computational structures (FoSSaCS '09), volume 5504 of lecture notes in computer science; 2009. p. 288–302.

[14] Wen R. Online elections in terra australis. School of Computer Science and Engineering, The University of New South Wales; 2010. Phd thesis.

[15] Wen R, McIver A, Morgan C. Towards a Formal Analysis of Information Leakage for Signature Attacks in Preferential Elections. FM 2014: Formal Methods - 19th International Symposium, Singapore, May 12–16, 2014; 2014. p. 595–610.