Who was that masked voter? The tally won't tell!

Peter Y. A. Ryan$^{1[0000-0002-1677-9034]}$, Peter B. Roenne$^{1[0000-0002-2785-8301]}$, Dimiter Ostrev$^{1[0000-0002-4098-0969]}$, Fatima-Ezzahra El Orche[1,3], Najmeh Soroush$^{1[0000-0002-9837-2376]}$, and Philip B. Stark$^{2[0000-0002-3771-9604]}$

[1] Interdisciplinary Centre for Security, Reliability, and Trust, SnT, University of Luxembourg
[2] Department of Statistics, University of California, Berkeley, CA, USA
[3] ENS, CNRS, PSL Research University, Paris, France
{firstname.lastname}@uni.lu, stark@stat.berkeley.edu, fatimaezzahra.elorche@ens.fr

Abstract. We consider elections that publish anonymised voted ballots or anonymised cast-vote records for transparency or verification purposes, investigating the implications for privacy, coercion, and vote selling and exploring how partially masking the ballots can alleviate these issues.
Risk Limiting Tallies (RLT), which reveal only a random sample of ballots, were previously proposed to mitigate some coercion threats. Masking some ballots provides coerced voters with plausible deniability, while risk-limiting techniques ensure that the required confidence level in the election result is achieved. Risk-Limiting Verification (RLV) extended this approach to masking a random subset of receipts or trackers.
Here we show how these ideas can be generalised and made more flexible and effective by masking at a finer level of granularity: at the level of the components of ballots. In particular, we consider elections involving complex ballots, where RLT may be vulnerable to pattern-based vote buying. We propose various measures of verifiability and coercion-resistance and investigate how several sampling/masking strategies perform against these measures. Using methods from coding theory, we analyse signature attacks, bounding the number of voters who can be coerced. We also define new quantitative measures for the level of coercion-resistance without plausible deniability and the level of vote-buying-resistance without "free lunch" vote sellers.
These results and the different strategies for masking ballots are of general interest for elections that publish ballots for auditing, verification, or transparency purposes.

1 Introduction

Some voting systems, including many end-to-end verifiable systems and some conventional elections, publish the (plaintext) ballots. If these ballots are suitably anonymised, by for example verifiable mixes published on a bulletin board, then this is typically quite safe. But in some contexts, revealing such information may be problematic: certain corner cases, such as unanimous votes or absence of any votes for a candidate and coercion threats, such as signature attacks.

In [4] the idea of Risk-Limiting Tallies (RLT) and Risk-Limiting Verification (RLV) was proposed to mitigate such threats. The idea is to shroud a proportion of the (anonymised) votes so voters can plausibly claim to have complied with the coercer, even though no votes appear for the candidate demanded by the coercer or no ballot with the pattern demanded by the coercer shows up in the tally. The proportion left shrouded can be adjusted using risk-limiting techniques to ensure that the confidence in the announced outcome achieves the required threshold, e.g, 99%. The idea extends to the verification aspects: shrouding some proportion of receipts or trackers. This proves particularly effective in for example the Selene scheme to counter the “sting in the tail”: the coercer claiming that the voter’s fake tracker is his own.

In this paper we note that, despite the pleasing features of the constructions of [4] there are still some drawbacks, in particular if the ballots are rather complex. While RLT may disincentivize *coercion*, there may still be an incentive for *vote buying*: the voter might still cast the required pattern vote in the hope that it will be revealed. Further, it has been suggested that RLT is arguably undemocratic in that some voters’ ballots do not contribute to the final tally. The second objection can be countered by arguing that every vote has an equal probability of being included in the count and that the outcome will be, with whatever confidence level required, a correct reflection of all votes cast. Nonetheless, it is an aspect that some people find troubling. A pleasing side effect of our construction is that all ballots are treated on an equal footing.

These observations suggest exploring different ways to apply RLT and RLV when ballots are complex: rather than shrouding entire ballots at random, we shroud, at random, some preferences on each ballot. In effect we are filtering the tally horizontally rather than vertically. This hits both of the issues above: the chance any given pattern remains identifiable after the filtering is reduced, and every ballot contributes to the outcome, albeit not necessarily to every contest. In the *full tally* construction below, every ballot contributes fully to the announced outcome, but we shroud the link between the tracker and some components of the ballots. For tracker-based schemes, the voters can verify some but not all of their selections. This paper seeks to quantify these effects and explore trade-offs among them.

Our techniques allow us to state and prove bounds on the number of voters an adversary is able to attack using pattern-based or “signature” attacks. Note that assigning the same, or similar, complex ballot pattern to many voters is counterproductive for the adversary: if even a few voters comply, the rest can point to the signature ballots that already appear and claim compliance. Thus, an adversary who wants to influence many voters with a signature attack must be able to produce many distinguishable ballot patterns. This observation motivates us to prove lower and upper bounds on the number of distinguishable patterns an adversary can construct. We prove these bounds using a connection to a well-studied problem in the theory of error-correcting codes.

This ballot-masking method and its privacy implications are interesting not only in for RLT and RLV but for all schemes where all or some ballots are published for auditing, verification, or transparency. As an example, Colorado is currently redacting cast-vote records (CVRs) by removing entire CVRs, e.g., for rare ballot styles; partial masking has been considered as an alternative. We note, however, that masking parts of the ballot might make it hard to detect ill-formed, e.g., over-votes etc.

We also note that this idea has similarities to the SOBA constructions for Risk-Limiting-Audits (RLAs), [1], which also publishes each audited ballot "disassembled" into different contests, whereas the auditors will see the intact ballot. The VAULT approach [2] also uses homomorphic encryption of the cast-vote records to achieve the SOBA goals more easily. (VAULT was used for the first time in a risk-limiting audit in Inyo County, California, in 2020.) The purpose and the underlying cryptographic constructions are quite different, but our analysis applies to these cases as well.

For some tally algorithms, we can separate ballots into their atomic parts and reveal these independently after anonymising them, which effectively counters signature attacks. However, that reduces public transparency and may reduce public confidence in the election result. For Selene, where voters verify their votes via trackers, this separation provides a method to verify without revealing individual ballots: we simply assign a distinct tracker to each element of the ballot. Voters can then verify some or all components of their ballot using those trackers. A coerced voter could use the Selene tracker-faking mechanism to assemble a ballot that matches the coercer's instructions. Technically this is straightforward but from a usability standpoint seems problematic. Moreover, even if the voter were prepared to go the effort of concocting such a fake ballot, the necessary ingredients might not be available, so coercion threats will remain, and the probability that one of atomic trackers is the same as the coercer's increases. Thus it makes sense to look for alternatives.

Below, we present the main ideas and analyse differences in privacy, coercion-resistance, and receipt-freeness for the different methods. Section 2 introduces the idea of partially masking ballots. Section 3 describes how it can be used in masked RLT and RLV. Section 4 defines a distinguishing distance between randomly masked ballots, establishes a connection to the Hamming distance, characterizes the class of masking strategies for which this connection holds, and proves bounds on the number of voters that can be approached with a pattern-based attack. It provides another application of the distinguishing distance: to quantify the effect of masking on individual verifiability. Section 5 considers quantitative game-based notions of privacy, coercion-resistance, and receipt-freeness. Section 6 concludes.

2 Masking Complex Ballots

Many elections use simple plurality voting: the voter selects at most one candidate from a set, in the simplest case, a referendum, a choice between "yes" and "no." The next level of complexity is single-winner plurality, aka "first past the post."

More complex social choice functions and correspondingly more complex ballots are common. Perhaps the next level in complexity are *approval voting* in which the voter can cast votes for several candidates for a single office, and multi-winner plurality, in which a voter can vote for up to k candidates for k offices. In some cases voters may have a quota of votes and is allowed to cast more than one vote for a given candidate, up to some limit. Some methods allow voters to give a preference ranking to the candidates.

Common to all of these social choice functions, if the ballots are published, is that they are vulnerable to signature attacks (also known as "Italian" attacks), i.e. a coercer chooses a particular, unlikely, pattern, instructs the victim to mark a ballot with that pattern and checks whether a ballot with that pattern appears in the tally.

Let us assume that the ballots are of the form $(v_1, v_2, \ldots, v_k)$ with k the number of candidates and v_i taking values from a specified set $\mathcal{V}$. $\mathcal{V}$ might for example just be $\{0, 1\}$ or a set of integers plus a blank: $\{1, ..., s\} \bigcup \{\text{blank}\}$ etc.

In many types of elections, these ballot-level selections, or subsets thereof, will reappear as part of the tally procedure (e.g. in electronic mixnet tallies), as part of an audit trail or for transparency (electronic scans of paper ballots), in Risk-Limiting Audits using samples of votes, or verification procedures (e.g. in tracker-based schemes such as Selene). In order to preserve privacy, the mapping between the published votes and the voter is normally anonymised.

As mentioned above, revealing these ballots may endanger the receipt-freeness of the election. With *Masked Tallies*, introduced here, only parts of each ballot are revealed:

$$(\, \text{mask}_{i1}(v_1^{(i)}), \text{mask}_{i2}(v_2^{(i)}), \, \ldots \, , \text{mask}_{ik}(v_k^{(i)}) \,) \quad \text{for } i = 1, \ldots, n.$$

The functions mask_{ij} are either the identity, displaying the component of the vote, or a constant, .e.g. $* \; (\notin \mathcal{V})$, masking the component. n is the number of ballots cast.

Risk-Limiting Tallies [4], involved unmasking only as many randomly selected ballots as are needed to determine the election result with a chosen risk limit. The remaining ballots were kept completely masked. Here we suggest a generalization, allowing partial masking of the ballots, and we will discuss the impact on risk limits, privacy, coercion-resistance, and resistance to vote-buying.

3 Partially Masked RLTs and RLVs

We reprise risk-limiting tallies and verification, RLT and RLV [4], before extending these to general masks. First we recapitulate the idea of tracker-based verification in terms of Selene.

Outline of Selene Selene [8] enables verification by posting the votes in the clear on the BB along with private tracking numbers. Voters are only notified of their tracker some time after the vote/tracker pairs have been publicly posted, giving a coerced voter the opportunity to choose an alternative tracker to placate the coercer. The voter is able to fake the tracker and related cryptographic data using a secret trapdoor key. The notification of the trackers is carefully designed to provide assurance to the voter that it is their correctly assigned tracker, i.e. unique to them, while being deniable to any third party.

Assuming that votes are encrypted component-wise, at the end of the mixing we will have encrypted votes and trackers on the bulletin board:

$$(\{tr_i\}_{PK}, (\{v_1^{(i)}\}_{PK}, \{v_2^{(i)}\}_{PK}, \ldots\ldots \{v_k^{(i)}\}_{PK}))$$

where $\{\cdot\}_{PK}$ denotes encryption under the public key PK. These ballots can now be verifiably decrypted to reveal the vote/tracker pairs that can be checked by the voters, and anyone can compute the tally directly on the plaintext votes.

Risk-Limiting Tallies and Verification with Partially Masked Ballots In the original approach to RLT (where ballots are without trackers) and RLV (with trackers for individual verification), see [4], the idea was to only decrypt a random subset of the ballots. The number decrypted being controlled by a risk-limit that bounds the probability that the announced election result will be wrong.

In the new masked RLV and RLT approach, we instead reveal randomly selected components of the ballots (and the trackers for RLV). If there is more than one contest on the ballot, the contests can be treated independently. How much we reveal will again be governed by a specified risk limit, as in [4]. A natural choice is to first decrypt m of the k entries in each ballot at random, and to increase m if necessary to meet the risk limit. This is simplest and will be used in the analysis below. In practice, it may make sense to dynamically change the rate of openings per candidate, e.g. if a candidate is popular we might be able to decrease the rate of unmasking of votes for that candidate, maintaining the risk limit while improving coercion-resistance.

Using this masked approach for RLV with tracker verification, the masking means that only parts of the ballot can be verified, but unlike to the original RLV every voter can verify *something*. We will quantify how much.

Full Tally with Partial Verification (FTPV) A social choice function is *separable* if, for the purposes of tallying, the components of each vote can be considered separately. Plurality, approval, and Borda count are separable; instant-runoff voting and single transferrable vote are not. For separable social choice functions, it is possible to compute the full tally, i.e. achieve 100% confidence in the outcome while partially masking selections. For each ballot, we randomly select some components. All selected components for all ballots are gathered in another part of the BB and subjected to a full, componentwise shuffling before decryption. Their positions in the original ballots are replaced by $*$. Thus, the way that these selected components appeared in the original ballots is lost.

The FTPV approach above might still hit corner cases, for instance if no vote was cast for a particular candidate. This suggests using a hybrid approach in which we use the approach above but reveal a random subset of the components separated out from the ballots. Thus we reveal enough of each ballot linked to the tracker to make verification meaningful while mitigating coercion threats, while a larger portion of the ballots is revealed without a link to the trackers to attain the required risk limit for the tally.

4 Distinguishing distance and applications to signature attacks and individual verifiability

In this section, we define a metric on the set of complex ballots that characterizes how well pairs of strings can be distinguished under random masking. We then observe that in some cases this metric is a monotone transformation of the Hamming distance used in coding theory. We also precisely characterize the cases when this occurs. Next, we use the connection to coding theory to answer the following question: how many simultaneous signature attacks can a coercer and/or vote-buyer launch? Finally, we give another application of the distinguishing distance: we use it to quantify the effect of a masking strategy on individual verifiability.

Throughout this section, we consider complex ballots with k components taken from the set $\mathcal{V}$; thus, the set of possible ballots is $\mathcal{V}^k$. We ignore here any constraints on what constitute valid ballots. For $x \in \mathcal{V}^k$ and $S \subset \{1, \ldots, k\}$, we denote by x_S the substring of x on the positions in S.

4.1 Definition and basic properties of distinguishing distance

How distinguishable are pairs of elements of $\mathcal{V}^k$ under masking? For every probability distribution p_S over subsets of $\{1, \ldots, k\}$, for every $x \in \mathcal{V}^k$ there is an induced probability distribution q_{S,x_S} of the pair (S, x_S), given by $q_{S,x_S}(s, \alpha) = p_S(s)\delta_{x_s,\alpha}$. If we keep p_S fixed and consider a pair $x, y \in \mathcal{V}^k$, we can define the distance between x and y as the statistical distance of q_{S,x_S}, q_{S,y_S}; thus, we take

$$d_{p_S}(x,y) = \frac{1}{2}\|q_{S,x_S} - q_{S,y_S}\|_1 = \sup_D |\Pr(D(S,x_S)=1) - \Pr(D(S,y_S)=1)|, \quad (1)$$

where the supremum is over distinguishers D. We can obtain the following formula for d_{p_S}:

Proposition 1. *For all distributions p_S, for all $x, y \in \mathcal{V}^k$,*

$$d_{p_S}(x,y) = \sum_{s:x_s \neq y_s} p_S(s) = \sum_s p_S(s)\mathbb{I}(s \cap t \neq \emptyset)$$

where t is the set of positions on which x, y differ and the operator $\mathbb{I}$ transforms the true/false value of a statement to $1, 0$ respectively.

Proof.

$$\begin{aligned} d_{p_S}(x,y) =& \frac{1}{2}\|q_{S,x_S} - q_{S,y_S}\|_1 = \sum_{(s,\alpha):q_{S,x_S}(s,\alpha)>q_{S,y_S}(s,\alpha)} (q_{S,x_S}(s,\alpha) - q_{S,y_S}(s,\alpha)) \\ &= \sum_{(s,\alpha):q_{S,x_S}(s,\alpha)>q_{S,y_S}(s,\alpha)} (p_S(s)\delta_{x_s,\alpha} - p_S(s)\delta_{y_s,\alpha}) \\ &= \sum_{s:x_s\neq y_s} p_S(s) = \sum_s p_S(s)\mathbb{I}(s\cap t\neq\emptyset) \end{aligned}$$

□

Under the mild assumption that each position is revealed with strictly positive probability, d_{p_S} is a metric on $\mathcal{V}^k$.

Proposition 2. *For all p_S, d_{P_S} is symmetric, satisfies the triangle inequality and satisfies $\forall x, d_{P_S}(x,x) = 0$. If in addition $\forall i, \Pr(i\in S) > 0$, then $d_{p_S}(x,y) = 0 \implies x = y$.*

Proof. The first three claims follow directly from (1). For the last claim, take any i, any $v \in \mathcal{V}$, any x, y with $d_{p_S}(x,y) = 0$. Consider the distinguisher D given by "On input s, α, if i is among the revealed positions and the corresponding entry is v output 1, else output zero." Then,

$$\Pr(i\in S)\delta_{x_i,v} = \Pr(D(S,x_S)=1) = \Pr(D(S,y_S)=1) = \Pr(i\in S)\delta_{y_i,v}.$$

Therefore, $\forall i \forall v, x_i = v \iff y_i = v$, so $x = y$. □

Now, we look at another question: how to find an optimal distinguisher between a pair of strings. For each $x \in \mathcal{V}^k$, define distinguisher D_x by "On input (s,α), if $x_s = \alpha$, output 1, else output 0." This is optimal regardless of the particular p_S, and regardless of the particular second element y.

Proposition 3. *For all distributions p_S, for all $x, y \in \mathcal{V}^k$,*

$$d_{p_S}(x,y) = \Pr(D_x(S,x_S)=1) - \Pr(D_x(S,y_S)=1)$$

Proof.

$$\begin{aligned} &\Pr(D_x(S,x_S)=1) - \Pr(D_x(S,y_S)=1) \\ &\qquad = \sum_s p_S(s)(\Pr(D_x(s,x_s)=1) - \Pr(D_x(s,y_s)=1)) \\ &\qquad\qquad = \sum_s p_S(s)(1-\delta_{x_s,y_s}) = \sum_{s:x_s\neq y_s} p_S(s) = d_{p_S}(x,y) \end{aligned}$$

□

4.2 Distinguishing distance and Hamming distance

From Proposition 1, we see that for any p_S, $d_{p_S}(x,y)$ does not depend on all details of the strings x, y, but only on the set of positions where x, y differ. It turns out that there is a class of distributions p_S such that d_{p_S} does not even depend on all details of the set of positions where x, y differ, but only on the Hamming distance between x and y, $d_H(x, y) = |\{i : x_i \neq y_i\}|$. This class of probability distributions is precisely those that assign equal weight to subsets of equal size.

Theorem 1. *For all p_S, the following are equivalent:*

1. *There exists a probability vector $(r(0), \ldots r(k))$ such that $\forall s, p_S(s) = \frac{r(|s|)}{\binom{k}{|s|}}$*
2. *There exists a function f_{p_S} such that for all $x, y \in \mathcal{V}^k$, $d_{p_S}(x, y) = f_{p_S}(d_H(x, y))$.*

We prove the forward direction of Theorem 1 by computing an explicit formula for the function f_{p_S}.

Theorem 2. *Suppose $\exists(r(0), \ldots r(k)) \forall s, p_S(s) = \frac{r(|s|)}{\binom{k}{|s|}}$ Then,*

$$d_{p_S}(x, y) = \sum_{i=1}^{d_H(x,y)} \sum_{j=0}^{k-d_H(x,y)} \frac{\binom{d_H(x,y)}{i}\binom{k-d_H(x,y)}{j} r(i+j)}{\binom{k}{i+j}}$$

Proof (Theorem 2). Take any x, y and let t be the subset of positions where x, y differ. Then,

$$d_{p_S}(x, y) = \sum_{s : x_s \neq y_s} p_S(s) = \sum_{s : s \cap t \neq \emptyset} p_S(s) = \sum_{i=1}^{|t|} \sum_{j=0}^{k-|t|} \frac{\binom{|t|}{i}\binom{k-|t|}{j} r(i+j)}{\binom{k}{i+j}} \quad \square$$

To prove the reverse direction of Theorem 1, we think of the $2^k - 1$ dimensional vector space over $\mathbb{C}$ with entries indexed by non-empty subsets of $\{1, \ldots k\}$, we think of the subspace

$$W = \{w \in \mathbb{C}^{2^k - 1} : |s| = |t| \implies w(s) = w(t)\}$$

and we also think of the $(2^k - 1) \times (2^k - 1)$ matrix M with entries $M(s, t) = \mathbb{I}(s \cap t \neq \emptyset)$ indexed by non-empty subsets of $\{1, \ldots, k\}$.

From Theorem 2, we see that $w \in W \implies Mw \in W$, that is, M leaves the subspace W invariant. Next, we observe that M is self-adjoint, and that M is also invertible:

Theorem 3. *For all $k \in \mathbb{N}$, the matrix M_k with entries $M_k(s, t) = \mathbb{I}(s \cap t \neq \emptyset)$ indexed by non-empty subsets of $\{1, \ldots k\}$ is invertible.*

a fact that we will prove at the end of this subsection. From this, we see that M^{-1} also leaves subspace W invariant.

Now, assume $\exists f_{p_S}, \forall x, y : d_{p_S}(x, y) = f_{p_S}(d_H(x, y))$. Form the vector $w \in W$ with entries $w(t) = f_{p_S}(|t|)$. The relation $d_{p_S}(x, y) = f_{p_S}(d_H(x, y))$ and Proposition 1 imply $\forall t, w(t) = \sum_{s \neq \emptyset} M(t, s) p_S(s)$. Then, $(p_S(s))_{s \neq \emptyset} = M^{-1} w \in W$, so p_S assigns equal weight to subsets of equal size. This completes the proof of Theorem 1, assuming Theorem 3 holds.

It remains to prove Theorem 3. The proof is by induction on k. When $k = 1$, $M_1 = (1)$ is invertible. Assume now M_k is invertible and consider M_{k+1}. We order subsets according to the following: a subset corresponds to a string of 0s and 1s, and this encodes an integer between 1 and $2^{k+1}-1$. With this ordering of the subsets, the matrix M_{k+1} has the following block form:

$$\begin{pmatrix} M_k^{(2^k-1)\times(2^k-1)} & 0^{(2^k-1)\times 1} & M_k^{(2^k-1)\times(2^k-1)} \\ 0^{1\times(2^k-1)} & 1^{1\times 1} & 1^{1\times(2^k-1)} \\ M_k^{(2^k-1)\times(2^k-1)} & 1^{(2^k-1)\times 1} & 1^{(2^k-1)\times(2^k-1)} \end{pmatrix}$$

where the sizes of the blocks are indicated in the superscript, and a 0 or 1 indicates that all entries of that block are 0 or 1.

Now we consider the following elementary row operations: subtract the middle row from all the bottom rows, then subtract the top block of rows from the bottom block of rows. We arrive at the matrix

$$\begin{pmatrix} M_k^{(2^k-1)\times(2^k-1)} & 0^{(2^k-1)\times 1} & M_k^{(2^k-1)\times(2^k-1)} \\ 0^{1\times(2^k-1)} & 1^{1\times 1} & 1^{1\times(2^k-1)} \\ 0^{(2^k-1)\times(2^k-1)} & 0^{(2^k-1)\times 1} & (-M_k)^{(2^k-1)\times(2^k-1)} \end{pmatrix}$$

and this is invertible by the inductive hypothesis. Hence, M_{k+1} is also invertible.

4.3 Bounds on the number of simultaneous signature attacks

We consider a coercer and/or vote buyer who wants to launch signature attacks on multiple voters simultaneously. Thus, the adversary chooses r signatures $x_1, \ldots, x_r \in \mathcal{V}^k$ and approaches many voters requiring each to submit one of the signature ballots.

What is the largest number r_{max} of different signatures that a coercer can use subject to the natural constraint that the strings $x_1, \ldots, x_r$ are pairwise distinguishable under random masking? We use the connection to coding theory from subsection 4.2 to answer this question.

First, we prove some properties of the function f_{p_S} from Theorem 1.

Lemma 1. *For every p_S that satisfies $\exists(r(0), \ldots, r(k))\forall s, p_S(s) = \frac{r(|s|)}{\binom{k}{|s|}}$, the function f_{p_S} is non-decreasing, $f_{p_S}(0) = 0$, and $f_{p_S}(k) = 1 - p_S(\emptyset)$.*

Proof. Take any $i < j \in \{0, \ldots, k\}$. Take $x, y \in \mathcal{V}^k$ that differ in the first i positions and $x', y' \in \mathcal{V}^k$ that differ in the first j positions. Using Proposition 1 we get $f_{p_S}(i) = f_{p_S}(d_H(x, y)) = d_{p_S}(x, y) = \sum_s p_S(s)\mathbb{I}(s \cap \{1, \ldots, i\} \neq \emptyset)$ $\leq \sum_s p_S(s)\mathbb{I}(s \cap \{1, \ldots, j\} \neq \emptyset) = d_{p_S}(x', y') = f_{p_S}(d_H(x', y')) = f_{p_S}(j)$.

For the other two claims, take $z, w \in \mathcal{V}^k$ that differ in all positions. Then,

$$f_{p_S}(0) = f_{p_S}(d_H(z, z)) = d_{p_S}(z, z) = 0$$

$$f_{p_S}(k) = f_{p_S}(d_H(z, w)) = d_{p_S}(z, w) = \sum_s p_S(s)\mathbb{I}(s \cap \{1, \ldots, k\} \neq \emptyset) = 1 - p_S(\emptyset) \quad \square$$

The properties of f_{p_S} established in Lemma 1 allow us to define a partial inverse of f_{p_S}. Take $g_{p_S} : [0, 1 - p_S(\emptyset)] \to \{0, 1, \ldots k\}$ given by

$$g_{p_S}(q) = \min\{d \in \{0, 1, \ldots, k\} : f_{p_S}(d) \geq q\}$$

so that we have

$$f_{p_S}(d) \geq q \iff d \geq g_{p_S}(q) \tag{2}$$

Now, we are ready to state and prove our bounds on the number of simultaneous signature attacks under a pairwise distinguishability constraint.

Theorem 4. *For every finite set $\mathcal{V}$, for every $k \in \mathbb{N}$, for every probability distribution p_S on subsets of $\{1, \ldots, k\}$ satisfying $\exists (r(0), \ldots, r(k)) \forall s, p_S(s) = \frac{r(|s|)}{\binom{k}{|s|}}$, for every $q \in [0, 1 - p_S(\emptyset)]$, let $r_{max}(\mathcal{V}, k, p_S, q)$ denote the size of the largest collection $\{x_1, \ldots x_r\}$ with the property $\forall i \neq j, d_{p_S}(x_i, x_j) \geq q$. Then*

$$\frac{|\mathcal{V}|^k}{\sum_{j=0}^{g_{p_S}(q)-1} \binom{k}{j} (|\mathcal{V}| - 1)^j} \leq r_{max}(\mathcal{V}, k, p_S, q) \leq \frac{|\mathcal{V}|^k}{\sum_{j=0}^{\lfloor (g_{p_S}(q)-1)/2 \rfloor} \binom{k}{j} (|\mathcal{V}| - 1)^j}$$

Proof. We use the same argument that is used in coding theory to establish the Gilbert-Varshamov lower bound and the Hamming upper bound on the maximum number of codewords subject to a pairwise Hamming distance constraint.

First, we observe that a collection $\{x_1, \ldots x_r\}$ satisfies $\forall i \neq j, d_{p_S}(x_i, x_j) \geq q$ if and only if it satisfies $\forall i \neq j, d_H(x_i, x_j) \geq g_{p_S}(q)$. This follows from the relation $d_{p_S}(x_i, x_j) = f_{p_S}(d_H(x_i, x_j))$ and the property (2) of the partial inverse g_{p_S}.

Now, take a collection $\{x_1, \ldots x_{r_{max}(\mathcal{V},k,p_S,q)}\}$ with the maximum number of elements subject to the constraint $\forall i \neq j, d_H(x_i, x_j) \geq g_{p_S}(q)$. To prove the upper bound, note that the Hamming balls of radius $\lfloor (g_{p_S}(q) - 1)/2 \rfloor$ around $x_1, \ldots, x_{r_{max}}$ must be disjoint, that each such ball contains $\sum_{j=0}^{\lfloor (g_{p_S}(q)-1)/2 \rfloor} \binom{k}{j} (|\mathcal{V}| - 1)^j$ elements, and that the total number of elements in all these balls must not exceed the size of the whole set $\mathcal{V}^k$.

To prove the lower bound, note that the Hamming balls of radius $g_{p_S}(q) - 1$ around $x_1, \ldots, x_{r_{max}}$ must completely cover $\mathcal{V}^k$, or else another element could be found that has Hamming distance $\geq g_{p_S}(q)$ to all of $x_1, \ldots, x_{r_{max}}$ and this would contradict the choice of $\{x_1, \ldots x_{r_{max}(\mathcal{V},k,p_S,q)}\}$ as having the maximum number of elements. Now, we have r_{max} Hamming balls with $\sum_{j=0}^{g_{p_S}(q)-1} \binom{k}{j} (|\mathcal{V}| - 1)^j$ elements each and their total number of elements must exceed $|\mathcal{V}|^k$, giving the lower bound on r_{max}. □

These upper and lower bounds are exemplified in Figure 1 for an election with $k = 5$ candidates and $|\mathcal{V}| = 2$ (like the student election example in next section). We have $g_{p_S}(q) = k - m + 1$ when g is applied to a uniform distribution over m-element subsets (m openings) evaluated at $q = 1$ (perfect distinguishability).

4.4 Quantifying the effect of masking on individual verifiability

We would like to quantify the effect of a particular masking strategy, specified by the probability distribution p_S, on individual verifiability. We propose the following quantity:

$$IV(p_S) = \inf_{x \neq y \in \mathcal{V}^k} d_{p_S}(x, y)$$

This quantity takes values between 0 and 1, where $IV(p_S) = 1$ means that the masking strategy p_S leaves the individual verifiability of the underlying voting protocol invariant, while $IV(p_S) = 0$ means that the masking strategy p_S destroys any individual verifiability that was present in the underlying voting protocol.

The motivation for choosing the quantity $IV(p_S)$ is the following: a voter who has voted x obtains a pair (s, α) where $s \subset \{1, \ldots, k\}$ and $\alpha \in \mathcal{V}^{|s|}$ and must decide whether this revealed vote was obtained from his submitted vote x or from some $y \neq x$. Taking the infimum over $x \neq y$ corresponds to considering the worst case over voter choices x and modifications of the voter choice y.

One attractive feature of this setup is that an individual voter does not need to know the distribution p_S or the modification y in order to apply the optimal verification strategy; indeed the optimal strategy for a voter who has chosen x is to apply the distinguisher D_x considered in Proposition 3.

For distributions p_S that satisfy $\exists(r(0), \ldots, r(k)) \forall s, p_S(s) = \frac{r(|s|)}{\binom{k}{|s|}}$, Theorem 2 gives a simple formula for $IV(p_S)$:

$$IV(p_S) = \sum_{j=0}^{k-1} \frac{\binom{k-1}{j} r(j+1)}{\binom{k}{j+1}} = \sum_{l=1}^{k} \frac{l}{k} r(l)$$

where we have used the fact that the transformation from Hamming to distinguishing distance is non-decreasing (Lemma 1), and so the smallest distinguishing distance is between x, y such that $d_H(x, y) = 1$.

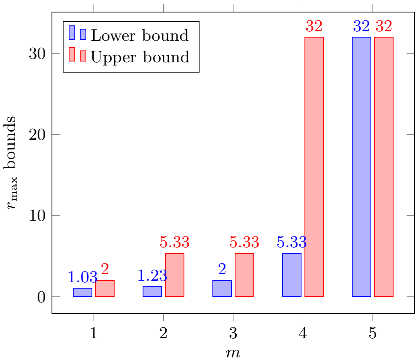

Fig. 1. Example for $|\mathcal{V}| = 2$ and $k = 5$

$m \setminus p$	p_{col}	$(1 - p_{col})^n$
1	0.16	$1.46 \cdot 10^{-79}$
2	0.018	$8.3 \cdot 10^{-9}$
3	0.0005	0.60
4	$9.7 \cdot 10^{-6}$	0.99
5	$1.6 \cdot 10^{-7}$	0.9998

Fig. 2. The probability, p_{col} that a single (resp. no) honest voter casts a ballot which after masking equals the mask of $v_0^O = (0, 1, 1, 1, 1)$ for the student election.

5 Quantitative Privacy-Type Properties

We now want to measure and compare privacy-properties for different masked tally methods. When computing concrete values we will consider approval voting with k candidates only 0 or 1 is allowed for each candidate, without any overall constraint, $(v_1, \ldots, v_k) \in \{0, 1\}^k$. For the n honest voters we assume for simplicity that the probability to vote $v_i = 1$ is p_i and these probabilities are independent. As a special concrete case we consider a student election with $n = 1001$ voters (one voter is under observation), $k = 5$ candidates with probabilities $(0.6, 0.4, 0.01, 0.01, 0.01)$, i.e. two popular candidates and three unpopular.

5.1 Privacy

In order to compare the different approaches we first consider the quantitative δ-privacy definition from [5]. The main other quantitative privacy definition is [3], but it is less suited considering signature attacks. The parties are an observer O, who can use public data, n_h honest voters and an additional voter under observation V_{obs}, whose vote the observer tries to guess.

Definition 1 (δ-privacy). *Let P be a voting protocol and V_{obs} be the voter under observation. We say that P achieves δ-privacy if*

$$\Pr[(\pi_O||\pi_{V_{obs}}(v_0^O)||\pi_v)^{(l)} \rightarrow 1] - \Pr[(\pi_O||\pi_{V_{obs}}(v_1^O)||\pi_v)^{(l)} \rightarrow 1]$$

is δ-bounded as a function of the security parameter ℓ for all vote choices v_0^O and v_1^O of the observed voter. Here π_O, $\pi_{V_{obs}}$ and π_v are respectively the programs run by the observer O, the voter under observation V_{obs} and all the honest voters.

The value δ will depend on the chosen vote distribution, and we see that it is especially relevant to penalize signature attacks: if we assume that there is a vote choice $v^* = (v_1^*, \dots, v_k^*)$ which rarely gets selected and has a probability close to zero, then an unmasked tally which reveals all cast plaintext ballots, even in anonymised form, will have $\delta = 1$ – the adversary simply checks if v^* appears.

Full ballot disclosure When we reveal all ballots, we can consider the case where the observer tries to distinguish a voter casting the most unpopular vote vs the most popular vote, as in a signature attack. That is, in the definition we let $v_0^O = (v_1, \dots, v_k)$ with $v_i = 1$ if $p_i \leq 1/2$ and $v_i = 0$ if $p_i > 1/2$, and we have $v_1^O = (1 - v_1, \dots, 1 - v_k)$. Denote the corresponding probability p_{min}. Now a good strategy is simply to check if at least one $(v_1, \dots, v_k)$ appears in the disclosed ballots, and the algorithm then outputs "1". This means $\Pr[(\pi_O||\pi_{V_{obs}}(v_0^O)||\pi_v)^{(l)} \rightarrow 1] = 1$ but $(\pi_O||\pi_{V_{obs}}(v_1^O)||\pi_v)$ will also output "1" if another voter chooses v_0^O. This happens with probability $1 - (1 - p_{min})^{n_h}$. We conclude that $\delta \geq (1 - p_{min})^{n_h}$. For the case of the student election we have that $v_0^O = (0, 1, 1, 1, 1)$ with $p_{min} = 0.4^2 \cdot 0.01^3 = 1.6 \cdot 10^{-7}$. Thus for $n_h = 1000$ we have $\delta \geq (1 - p_{min})^{n_h} \approx 0.99984$, i.e. close to 1.

Result Only We now consider the case where we only reveal the overall result $r = (r_1, \dots, r_k)$. In this case we can follow an analysis close to [5,7] for calculating δ. For every possible result r we calculate the probability that the result happened if the observed voter cast v_0^O or v_1^O. The algorithm will then output one if the former probability is larger. We get $\delta = \sum_{r \in M^*_{v_0^O, v_1^O}} (A_r^{v_0^O} - A_r^{v_1^O})$ where $M^*_{v_0^O, v_1^O} = \{r \in \mathbb{R} : A_r^{v_1^O} \leq A_r^{v_0^O}\}$, $\mathbb{R}$ is the set of all possible results of the election and A_r^v denotes the probability that the choices of the honest voters yield the result r given that V_{obs}'s choice is v. These probabilities can explicitly be calculated since each candidate count from the honest voters, X_i, is binomially distributed, $X_i \sim BD(n_h, p_i)$. We thus have $A_r^v = \mathbb{P}(X_1 = r_1 - v_1) \cdots \mathbb{P}(X_k = r_k - v_k) = \prod_{i=1}^k \binom{n-1}{r_i - v_i} p_i^{r_i - v_i} (1 - p_i)^{n - r_i + v_i - 1}$.

RLT In the original RLT method we keep a certain fraction, f_{blind}, of the ballots hidden, that is $(1 - f_{blind})n$ ballots are published. If we consider the optimal algorithm from the full ballot disclosure and the corresponding δ_{full} we see that $\delta = (1 - f_{blind})\delta_{full}$ since the probability that observed voter's ballot is hidden is $(1 - f_{blind})$.

Masked RLT We now consider the case of masked RLTs where the we release all ballots but with only m out of k components unmasked. A good strategy to lower bound δ is to count the number N_b of colliding ballots v which satisfy $\text{mask}_v\, v = \text{mask}_v\, v_b^O$ for $b = 0, 1$. We choose v_0^O as the most unlikely ballot, as above and take v_1^O as the opposite ballot to discriminate optimally between the two counts. The main distinguishing power comes from N_0, and we let the distinguishing algorithm output "1" if the probability of the honest voters casting $N_0 - 1$ colliding votes is higher than getting N_0 collisions. The probability for each honest voter to have a collision is $p_{col} = 1/\binom{k}{m} \cdot \sum_{1 \leq i_1 < i_2 < \ldots < i_m \leq k} p_{i_1} \cdots p_{i_m}$ and $N_0 \sim BD(n_h, p)$, where p_i is the probability of a match in the ith candidate. In Fig. 2 we have displayed the probabilities for the student election example. The algorithm above will then simply give the probability at the mode of the binomial distribution with p_{col}. For $m = 3$ we find $\delta \geq 0.6$ for the student election.

5.2 Coercion-Resistance

In [6] the authors present a definition of quantitative coercion-resistance following similar ideas as in Definition 1. We will here use their strategy version and not go into all details. We let S denote the election system with specified number candidates, honest (n_h) and dishonest voters (mostly neglected here) and a ballot distribution, and attacker, C_S, and voter, V_S, interactive Turing machine models. We let γ denote a property defining the goal of the coerced voter, e.g. to vote for a specified candidate.

Definition 2. *S achieves δ^{cr}-coercion-resistance if for all dictated coerced strategies $\pi_{V_{co}} \in V_S$ there exists a counter-strategy $\tilde{\pi}_{V_{co}} \in V_S$ s.t. for all coercer programs $\pi_c \in C_S$:*

- $\Pr[(\pi_c || \tilde{\pi}_{V_{co}} || \pi_v)^{(l)} \mapsto \gamma]$ *is overwhelming.*
- $\Pr[(\pi_c || \pi_{V_{co}} || \pi_v)^{(l)} \mapsto 1] - \Pr[(\pi_c || \tilde{\pi}_{V_{co}} || \pi_v)^{(l)} \mapsto 1]$ *is δ^{cr}-bounded.*

With bounded and overwhelming defined in the security parameter. The first part says that the voter is able to achieve her goal (e.g. vote for a specific candidate) and the second part says that the coercer's distinguishing power is bounded by δ^{cr}. This level of coercion-resistance depends on several parameters especially the probability distribution on the candidates.

Whereas this definition gives a level of coercion-resistance, it does not tell the full story. To see this let us consider two different election systems. System A outputs voter names and corresponding votes with probability 1/2, completely breaking privacy, and otherwise it only outputs the election result. Neglecting the information from the election result we get $\delta^A = 1/2$. In system B the voter secretly gets a signed receipt of her vote with probability 1/2 and otherwise the protocol works ideally. In this case a coerced voter can always cast her own choice and claim that no receipt was received. A voter following the coercer's instruction will with probability 1/2 give the corresponding receipt, i.e. we again have $\delta^B = 1/2$. However, the two systems are very different from the point of view of the voter: in system A the coerced voter gets caught cheating with probability 1/2, whereas in system B, the voter always has plausible deniability.

Since plausible deniability is an essential factor for the usability of coercion-resistance mechanisms, we need a new definition to be able to measure this aspect.

5.3 No Deniability

The level of plausibility of a voter claiming to have followed the coercer, while actually following the counter strategy, relates to the probability of false positives when the coercer tries to determine if the voter disregarded the instructions. In the following we assume without loss of generality that the coercer outputs 1 when blaming the voter. We now want to define the maximal probability of getting caught without any deniability, i.e. we consider the case where $\Pr[(\pi_c||\pi_{V_{co}}||\pi_v)^{(l)} \mapsto 1] = 0$ or negligible, i.e. the coercer only uses strategies where he never blames an honest voter.

Definition 3. *S achieves $\delta^{cr,no-d}$-coercion-resistance if for all dictated coerced strategies $\pi_{V_{co}} \in V_S$ there exists a counter-strategy $\tilde{\pi}_{V_{co}} \in V_S$ s.t. for all coercer programs $\pi_c \in C_S$:*

- $\Pr[(\pi_c||\tilde{\pi}_{V_{co}}||\pi_v)^{(l)} \mapsto \gamma]$ *is overwhelming.*
- $\Pr[(\pi_c||\tilde{\pi}_{V_{co}}||\pi_v)^{(l)} \mapsto 1]$ *is $\delta^{cr,no-d}$-bounded and* $\Pr[(\pi_c||\pi_{V_{co}}||\pi_v)^{(l)} \mapsto 1]$ *is negligible.*

Note that the coercer's optimal strategy to obtain this $\delta^{cr,no-d}$ and the voter's strategy might be different from the ones in Definition 2 but $\delta^{cr,no-d} \leq \delta^{cr}$.

The no deniability probability clearly separates the RLT approaches. The original RLT always has plausible deniability if we choose to keep some ratio of ballots shrouded and the voter can claim her ballot was not revealed. This is e.g. important for RLV giving deniability against an attack where the coercer provides a ciphertext to cast and asks for its decrypted vote.

In the case of masked ballots, there can be a chance of getting caught undeniably. This will depend strongly on the number of revealed ballot components m, the vote distribution and the voter's goal. For the student election analysed above, the worst case when the goal of the voter is to cast $(1, 0, 0, 0, 0)$. The coercer's optimal strategy is then to demand a vote for $(0, 1, 1, 1, 1)$. The coercer will blame the voter if there is no matching masked ballot, i.e. if no honest voters produce a collision which happens with probability $(1 - p_{col})^{n_h+1}$ computed Fig. 2. The probability of no deniability is then $p = 8 \cdot 10^{-9}$ for $m = 2$ but jumps abruptly to $p = 0.6$ for $m = 3$.

An interesting case is when the voter has a relaxed goal allowing to cast a signature part or not, and when the vote distribution has some ballots strictly zero probability. Let us consider a three candidate 0/1 election with 1-vote probabilities $(1/2, 1/2, 0)$. The voter's goal is to cast a 1 for the first candidate. The coercer's optimal strategy is to demand a signature ballot $(0, 0, 1)$. The voter has two counter-strategies: 1) cast a vote $(1, 0, 0)$ without the 0 probability signature part or 2) casting a vote $(1, 0, 1)$ with the signature part. For 1) the there is no deniability if no other voter casts a matching ballot and the coerced voter's ballot does not match either. For $m = 1$ this happens with $p = (2/3)^{n_h+1}$ and for $m = 2$ with $p = (11/12)^{n_h}$, both are small if we have many voters. For 2) there will always be a matching vote if the first part of the coerced voter's ballot is masked. However, if the last part is revealed the coercer can deduce this ballot comes from the coerced voter since this candidate had probability 0, and if the 1 vote in the first part is revealed as well then the voter is caught with no deniability. Thus is no deniability with probability $1/3(2/3)^{n_h}$ for $m = 1$ and $1/3 + 1/3(11/12)^{n_h}$ for $m = 2$. Thus for $m = 1$ strategy 2) is always better, but for $m = 2$ strategy 1) is better when we have more than 13 voters. In some cases the voter strategy thus depends on m, which might not be know beforehand.

Finally, it is also natural to define the level of plausability we can provide. The average plausability that a voter has e.g. in Definition 2 is a useful quantity for the voter, but it would be more useful to guarantee that the voter always has a certain level for coercion-resistance. We leave a precise definition for future work.

5.4 Receipt-Freeness

Following [6], definition 2 also covers receipt-freeness. However, we again argue that modelling some variants is useful. The following definition is based on a swap of $\pi_{V_{co}}$ and $\pi_{\tilde{V_{co}}}$ in Definition 3, and models vote buyers who do not want to pay a "free lunch" to vote sellers who follow their own goal. The voter goal γ can here be to cast a specified vote or set of votes.

Definition 4 (Weak Vote Buying Resistance). *For a given small p_{fl}, S achieves δ^{wvb}-coercion-resistance if for all dictated coerced strategies $\pi_{V_{co}} \in V_S$ there exists a counter-strategy $\tilde{\pi}_{V_{co}} \in V_S$ s.t. for all coercer programs $\pi_c \in C_S$:*

- $\Pr[(\pi_c||\tilde{\pi}_{V_{co}}||\pi_v)^{(l)} \mapsto \gamma]$ *is overwhelming.*
- $\Pr[(\pi_c||\pi_{V_{co}}||\pi_v)^{(l)} \mapsto 1] - \Pr[(\pi_c||\tilde{\pi}_{V_{co}}||\pi_v)^{(l)} \mapsto 1]$ *is δ^{wvb}-bounded and* $\Pr[(\pi_c||\tilde{\pi}_{V_{co}}||\pi_v)^{(l)} \mapsto 1]$ *is p_{fl}-bounded.*

We here interpret outputting "1" as paying the vote seller and this definition bounds how often an instruction-following vote seller gets paid by a vote-buyer (by $\delta^{wvb} + p_{\mathrm{fl}}$), but under the condition that a voter who casts another vote is only paid with a (very) small probability p_{fl}. This is a weakened vote-buyer model but interesting since a vote buyer should avoid vote sellers going for a "free lunch". If the probability of an honest vote seller getting paid is low, it would help curb vote selling (even though the vote buyer could increase the price and create a "vote selling lottery"). In this definition, it also makes sense to drop the quantification over the coercer's strategies to see the resistance to vote buying for different vote choices.

RLT In the original RLT a signature ballot will get revealed with probability $1 - f_{blind}$. If the vote buyer sees this he can pay the vote seller and will only pay the voter seller wrongly with a small probability p_{fl} equal to the probability that one of the honest voters cast the signature ballot, i.e. $\delta^{vb} \simeq 1 - f_{blind}$ which can be rather high and protects badly against vote buying.

Masked RLT For the masked ballots we can however choose m such that several ballots will have the same masking as the signature ballot and makes it hard for the vote buyer to assess if the signature ballot was cast. For the student election we see from Fig. 2 that the number of matches with the optimal signature ballot $(0, 1, 1, 1, 1)$ is binomially distributed with an expectation value of 18.4 colliding ballots and a standard deviation of around 4.

For a more precise example, we can consider the three candidate election with probabilites $(1/2, 1/2, 0)$ as above and assume that the goal of the voter is to cast 0 for candidate 1 and $p_{\text{fl}} = 0$. For $m = 1$ we will have $\delta^{vb} = 0$, but for $m = 2$ the vote-buyer can demand a vote for candidate 1 and 3 and pay out if he sees $(1, *, 1)$. Any counter-strategy with 0 for candidate 1 gives $\delta^{vb} = 1/3$.

We note that the new quantitative definitions for no deniability coercion-resistance (Def. 3), the weak vote buying resistance (Def. 4) and the original δ^{cr}-coercion-resistance (Def. 2) are considering different aspects of coercion-resistance and stating the three different δ-values gives a more nuanced description of the security of a given voting protocol. Also note that the δ values are calculated using potentially different strategies for the coercer and voter, and finding unified strategies optimising the parameters is an interesting line of future work. Finally, there are natural, more fine-grained, definitions extending these which should be also considered in the future.

6 Conclusion

We have shown that the idea of risk-limiting tallies and risk-limiting verification can be applied effectively to complex ballots. By partially masking each ballot rather than simply masking a subset of the ballots as in the original RLT and RLV we gain far greater flexibility in terms of masking strategies. This will be explored further in order to optimise the trade-offs between the various measures defined here in future work.

The approach is more robust against any claims of being undemocratic: all ballots are counted, and indeed in the full tally/partial verification option, all are counted fully. The only compromise then is some reduction in the level of verifiability, but this can be adjusted and is probably acceptable. If we compare this with ThreeBallot, there the chance of detecting a manipulated ballot is 1/3, assuming that the attacker does not learn which ballot was retained by the voter. In our case we can achieve a good level of coercion mitigation with say a shrouding of $\pm 1/2$ of each ballot. Finally, we did a preliminary analysis of the quantitative privacy for the different tally methods, and the coercion-resistance, in particular, the probability a coerced voter gets undeniably caught. The new masked tallies however, are more appropriate for receipt-freeness, in particular with upper bounds on the number of vote sellers, whereas the old RLT provides good plausible deniability to coerced voters. This suggests combining both methods when possible, but future work is needed to define the precise level of vote-buying resistance.

Acknowledgements

This research was funded in part by the Luxembourg National Research Fund (FNR) grant references STV C18/IS/12685695, Q-CoDe CORE17/IS/11689058 and PRIDE15/10621687/ SPsquared.

References

1. Benaloh, J., Jones, D.W., Lazarus, E., Lindeman, M., Stark, P.B.: SOBA: Secrecy-preserving observable ballot-level audit. EVT/WOTE **11** (2011)
2. Benaloh, J., Stark, P.B., Teague, V.: VAULT: Verifiable audits using limited transparency. E-Vote-ID 2019 p. 69 (2019)
3. Bernhard, D., Cortier, V., Pereira, O., Warinschi, B.: Measuring vote privacy, revisited. In: Proceedings of the 2012 ACM conference on Computer and communications security. pp. 941–952 (2012)
4. Jamroga, W., Roenne, P.B., Ryan, P.Y., Stark, P.B.: Risk-limiting tallies. In: International Joint Conference on Electronic Voting. pp. 183–199. Springer (2019)
5. Küsters, R., Truderung, T., Vogt, A.: Verifiability, privacy, and coercion-resistance: New insights from a case study. In: 2011 IEEE Symposium on Security and Privacy. pp. 538–553. IEEE (2011)
6. Küsters, R., Truderung, T., Vogt, A.: A game-based definition of coercion resistance and its applications. J. Comput. Secur. **20**(6), 709–764 (2012). https://doi.org/10.3233/JCS-2012-0444
7. Liedtke, J., Küsters, R., Müller, J., Rausch, D., Vogt, A.: Ordinos: A verifiable tally-hiding electronic voting protocol. In: IEEE 5th European Symposium on Security and Privacy (EuroS&P 2020) (2020)
8. Ryan, P.Y.A., Rønne, P.B., Iovino, V.: Selene: Voting with transparent verifiability and coercion-mitigation. In: International Conference on Financial Cryptography and Data Security. pp. 176–192. Springer (2016)