

Verifiability, Privacy, and Coercion-Resistance: New Insights from a Case Study*

Ralf Küsters, Tomasz Truderung, and Andreas Vogt
University of Trier, Germany
{kuesters, truderung, vogt}@uni-trier.de

Abstract

In this paper, we present new insights into central properties of voting systems, namely verifiability, privacy, and coercion-resistance. We demonstrate that the combination of the two forms of verifiability considered in the literature—individual and universal verifiability—are, unlike commonly believed, insufficient to guarantee overall verifiability. We also demonstrate that the relationship between coercion-resistance and privacy is more subtle than suggested in the literature.

Our findings are partly based on a case study of prominent voting systems, ThreeBallot and VAV, for which, among others, we show that, unlike commonly believed, they do not provide any reasonable level of verifiability, even though they satisfy individual and universal verifiability. Also, we show that the original variants of ThreeBallot and VAV provide a better level of coercion-resistance than of privacy.

Keywords: voting; verifiability; coercion-resistance; privacy; protocol analysis

1 Introduction

Verifiability, privacy, and coercion-resistance are central security requirements for modern voting systems. *Privacy* is the most basic security requirement, which says that the way a particular voter voted is not revealed to anybody, see, e.g., [2, 4]. Intuitively, *verifiability* means that voters have a way of checking that their votes were actually counted and that the published result of the election is correct. In the literature, traditionally two forms of verifiability are considered: individual and universal verifiability [14, 1, 25, 12]. As stated, e.g., in [14], *individual verifiability* means that a voter can check that her own ballot appears on the bulletin board. *Universal verifiability* requires that anyone can check that the election outcome corresponds to the ballots published on the bulletin board. In addition, it is explicitly or implicitly required that each vote in the election outcome was cast by an eligible voter and that there is at most one vote per voter, a fact that in some voting protocols can be verified as well. To achieve (individual) verifiability, a voter typically obtains some kind of receipt which, together with additional data published in the election, she can use to check that her vote was counted. This, however, potentially opens up the possibility for vote buying and voter coercion. Besides verifiability, voting systems should therefore also provide so-called *coercion-resistance* [2, 13, 24].

*An abridged version of this work is published in [20].

In this paper, we present new insights into these central security properties. Our findings are partly based on a case study for prominent voting protocols, ThreeBallot and VAV [26]. More precisely, the contribution of this paper is as follows.

Contribution of this Paper. We demonstrate, using ThreeBallot and VAV as examples, that the combination of the two mentioned forms of verifiability considered in the literature—individual and universal verifiability—are, unlike commonly believed, insufficient to guarantee overall verifiability. More precisely, based on a definition of verifiability proposed in [17], we precisely measure the level of verifiability ThreeBallot and VAV provide. It turns out that, while ThreeBallot and VAV satisfy individual and universal verifiability, there is an attack on the verifiability of these protocols, which results in an insufficient level of verifiability. Our attack allows a dishonest bulletin board (or the scanner), collaborating with m dishonest voters, to turn m votes of honest voters for A into m votes for B . This goes unnoticed even if all honest voters check whether their receipts appear on the bulletin board and even if they check that the published result corresponds to the ballots shown on the bulletin board.

As for privacy and coercion-resistance, we first provide a game-based definition of privacy, along the lines of a game-based definition of coercion-resistance proposed in [16], since, as to the best of our knowledge, such a definition does not exist. There are, however, definitions of coercion-resistance in simulation-based settings (see, e.g., [22]) and in an abstract, Dolev-Yao style model [9]. We note that both our definition of privacy and the definition of coercion-resistance from [16] allow to measure the level of privacy/coercion-resistance a protocol provides: for privacy, it is measured how well external observers can distinguish whether an honest voter voted for candidate j or j' ; for coercion-resistance, the ability of coercers to distinguish whether coerced voters followed the coercer’s instructions or not is measured. This is important in order to make meaningful statements about protocols, as many voting protocols, in particular many paper-based protocols (e.g., ThreeBallot and VAV [26], Prêt à Voter [6], and Split-Ballot [23]), do not provide the ideal level of privacy/coercion-resistance. As discussed in [16], simulation-based security definitions, e.g., the one by Moran and Naor [22] provide a yes/no-answer, rather than measuring the level of privacy/coercion-resistance. Also, they are more demanding than game-based definitions and deem many reasonable protocols insecure.

One would expect that privacy and coercion-resistance are closely related: If the level of privacy is low, i.e., there is a good chance of correctly determining how a voter voted, then this should give the coercer leverage to coerce a voter. Some works in the literature indeed suggest a close connection. For example, the definition of coercion-resistance by Moran and Naor [22], being simulation-based, has privacy “built in”.

However, our case study, in which we precisely measured the level of privacy and coercion-resistance of different variants of ThreeBallot and VAV proposed in the literature,¹ demonstrates that the relationship between privacy and coercion-resistance is more subtle than what can be gathered from existing work.

Among others, it turns out that improving the level of privacy of a protocol in a natural way (e.g., by changing the way honest voters fill out ballots) can lead to a *lower* level of coercion-resistance. This is the case when going from the original variant of ThreeBallot to a “privacy enhanced” variant proposed by de Marneffe et al. [8]. Clearly, in general, one does not expect privacy to imply coercion-resistance. Still the effect is quite surprising.

¹We note that coercion-resistance of one of the variants of ThreeBallot considered in this paper has already been studied in [16].

A maybe even more important and unexpected finding that comes out of our case study is that the level of privacy of a protocol can be much *lower* than its level of coercion-resistance; this is the case for the original variant of ThreeBallot [26] and a natural variant of VAV. The reason behind this phenomenon is basically that it may happen that the counter-strategy a coerced voter may run to defend against coercion hides the behavior of the coerced voter, including her vote, better than the honest voting program.

To complete the picture, we also study the case in which the counter-strategy does not “outperform” the honest voting program in the above sense. For this case, we are able to prove a general theorem that states that if a voting system provides a certain level of coercion-resistance, it provides at least the same level of privacy. As discussed in Section 6.5, this theorem is applicable to many voting protocols.

Structure of this Paper. We first introduce some basic terminology and introduce the notion of a voting protocol. In Sections 3 and 4 we recall (the variants of) ThreeBallot and VAV, respectively. Our findings on verifiability are presented in Section 5 and those for privacy and coercion-resistance in Section 6. We conclude in Section 7. Some more details are provided in the appendix.

2 Preliminaries and Protocols

In this section, we introduce some basic terminology and the notion of a voting protocol.

Preliminaries. As usual, a function f from the natural numbers to the real numbers is *negligible*, if for every $c > 0$ there exists ℓ_0 such that $f(\ell) \leq \frac{1}{\ell^c}$ for all $\ell > \ell_0$. The function f is *overwhelming*, if the function $1 - f(\ell)$ is negligible. Let $\delta \in [0, 1]$. The function f is δ -*bounded* if f is bounded by δ plus a negligible function, i.e., for every $c > 0$ there exists ℓ_0 such that $f(\ell) \leq \delta + \frac{1}{\ell^c}$ for all $\ell > \ell_0$.

We use systems of probabilistic polynomial-time interactive Turing machines (ITMs) as our computational model (see, e.g., [15]). In a system of ITMs, also called a *process*, ITMs can communicate with other ITMs via input/output tapes, also called (*external*) *input/output* channels. If π and π' are processes (each with a set of external input/output tapes), then we write $\pi \parallel \pi'$ for the concurrent composition of π and π' . A process defines a family of probability distributions over runs, indexed by the security parameter.

Voting Protocols. A voting protocol P specifies the programs of the honest voters and authorities in a voting process. More precisely, let k be the number of candidates and q be the number of voters. Then, P specifies:

- A set $\{a_1, \dots, a_l\}$ of voting authorities and a program $\hat{\pi}_a$, for every voting authority a . The specification of $\hat{\pi}_a$ includes the specification of the interface of a to the rest of the voting process, i.e., the channels via which a is connected to other parties.
- A program (formally a process) $\hat{\pi}_v$, for every voter $v \in \{v_1, \dots, v_q\}$. The specification of $\hat{\pi}_v$ includes the specification of the interface of v to the rest of the voting process. The program $\hat{\pi}_v$ takes a choice $j \in \{0, \dots, k\}$, where $j = 0$ stands for abstention from voting, as parameter, indicating which candidate v votes for (if any).

In the following, we will consider a probability distribution $\vec{p} = p_0, \dots, p_k$ on the possible choices honest voters have, i.e., $p_0, \dots, p_k \in [0, 1]$ and $\sum_{i=0}^k p_i = 1$, where p_0 is the probability that a voter abstains from voting and p_i , $i \in \{1, \dots, k\}$, is the probability that a voter votes for candidate i .



Figure 1: Two ways of voting for the second candidate (candidate B) in the ThreeBallot protocol, where \times represents a marked position and o represents an unmarked position. All the other possibilities of voting for B can be obtained as permutations of these two.

We define $\hat{\pi}_v(\vec{p})$ to be the process which first chooses $j \in \{0, \dots, k\}$ according to \vec{p} and then runs $\hat{\pi}_v(j)$. We sometimes simply write $\hat{\pi}_v$ instead of $\hat{\pi}_v(\vec{p})$, if the distribution \vec{p} is clear from the context.

Because, as we will see, the level of privacy, coercion-resistance, and verifiability of a protocol P depends on several parameters, we consider protocol instantiations P^* of P , for which these parameters are fixed. The parameters are the following:

- (i) the set $A_H \subseteq \{a_1, \dots, a_l\}$ of honest voting authorities A_H ,
- (ii) the total number q of voters and the set $V_H \subseteq \{v_1, \dots, v_q\}$ of honest voters (static corruption),
- (iii) the number k of candidates, and
- (iv) the probability distribution \vec{p} , as described above.

Such a *protocol instantiation* will be denoted by $P^* = P(A_H, q, V_H, k, \vec{p})$. We note that in our theorems, only the number of honest (and dishonest) voters will matter, not the specific set V_H of honest voters. Therefore, we often simply write $P(A_H, q, n, k, \vec{p})$ with $n = |V_H|$.

3 The ThreeBallot Voting Scheme

In ThreeBallot [26], a voter is given a multi-ballot consisting of three simple ballots. On every simple ballot the candidates are printed in the same fixed order. In the secrecy of a voting booth, the voter is supposed to fill out all three simple ballots in the following way: She marks the candidate of her choice on exactly *two* simple ballots and every other candidate on exactly *one* simple ballot; Figure 1 shows two ways of voting for candidate B. After this, she feeds all three simple ballots to a voting machine (some kind of scanner) and indicates the simple ballot she wants to get as a receipt. The machine checks the well-formedness of the multi-ballot, prints secretly random numbers on each simple ballot, where length of these numbers is the length of the security parameter and where numbers on different simple ballots are chosen independently, and gives the voter a copy of the chosen simple ballot, with the random number printed on it. Note that the voter does not get to see the random numbers of the remaining two simple ballots. The scanner keeps all simple ballots (now separated) in a ballot box. We assume that clerks guarantee that only registered voters can vote and that every voter votes at most once.

In the tallying phase, the voting machine posts on the bulletin board (electronic copies of) all the cast simple ballots in a random order. From the ballots shown on the bulletin board the result can easily be computed: The number of votes for the i -th candidate is the number of simple ballots with the i -th position marked minus the total number of votes, which is the total number of simple ballots on the bulletin board divided by three.

Intuitively, the system is coercion-resistant (at least to some extent), as every simple ballot that a voter can take as a receipt can be part of a multi-ballot that forms a valid vote for any candidate. Also, ThreeBallot was meant to provide (some level of) verifiability. For this, a

crucial assumption, already made in the original paper [26], is that neither the scanner, the voting authority, nor the bulletin board knows which simple ballots have been taken as receipts by honest voters before all ballots were published. Now, as each voter may check whether the simple ballot she has taken as a receipt appears on the bulletin board, it should be risky for any party to remove or alter simple ballots in order to manipulate the result since the probability that the modification of k simple ballots goes undetected is merely $(\frac{2}{3})^k$. Unfortunately, as we will see in Section 5.2, this argument, found in the literature, is flawed.

As mentioned in the introduction, there are two variants of ThreeBallot which differ in the way an honest voter fills out the ballot: the original variant by Rivest [26] and a variant by de Marneffe et al. [8].

The original variant. In this variant of the protocol a voter first, for each candidate, randomly chooses a simple ballot on which she then marks the position corresponding to this candidate. Then, she randomly picks a simple ballot for which the position corresponding to the candidate of her choice is not yet marked, and she marks this position. Finally, she randomly chooses one ballot as a receipt.

The variant of de Marneffe et al. In this variant of the protocol a voter first, for each candidate, marks the position corresponding to this candidate on a randomly chosen simple ballot. Then, she randomly chooses one simple ballot to be taken as a receipt. Finally, she marks the position corresponding to the candidate of her choice on a randomly chosen simple ballot on which this position has not yet been marked and which is not the ballot chosen as a receipt; we remark that in some cases there will be only one such simple ballot.

The advantage of this procedure is that the receipt an honest voter gets is stochastically independent of the candidate the voter votes for, which in turn should give better privacy. We note that in [8], ThreeBallot was analyzed in a simulation-based setting, focusing on privacy. The analysis was based on the (only informally stated) assumption that the adversary, given a receipt, is not able to reconstruct the exact way the corresponding multi-ballot was filled out. However, this assumption is unjustified: Runs for which an adversary can reconstruct the multi-ballots occur with non-negligible probability, as illustrated by the following example:

It may happen (with *non-negligible* probability, depending only on the probability distribution \vec{p} and the number of voters) that *each* voter marks both positions on the first simple ballot, no position on the second one, exactly one position on the third ballot, and then take the last ballot as her receipt, as shown in Figure 1, (b) for the case that the voter votes for B . In this case, a receipt directly indicates the choice of the voter, which completely breaks privacy.

In what follows, we denote by $(\underline{x}, \underline{x}, \circ)$, $(\underline{\circ}, \underline{x}, \underline{x})$ etc. multi-ballots filled out by voters, where the underlined simple ballots (\underline{x} and $\underline{\circ}$, respectively) represent those simple ballots picked as receipts by the voters; we refer to these objects as *patterns*. A pattern does not fix the order of simple ballots, e.g., $(\underline{\circ}, \underline{x}, \underline{\circ})$ is considered to be the same pattern as $(\underline{x}, \underline{\circ}, \underline{\circ})$.

4 The VAV Voting Scheme

In this section, we describe the VAV voting scheme [26]. In VAV, a voter is given a multi-ballot consisting of three simple ballots. On every simple ballot the candidates are printed in the same fixed order. On the top of one of those simple ballots the letter A is printed; on the top of the remaining two simple ballots the letter V is printed. In the secrecy of a voting booth, the voter is supposed to fill out her multi-ballot in the following way: (S1) She marks the position next to

(a)	<table style="border-collapse: collapse; margin: 0 auto;"> <tr><td style="border-bottom: 1px solid black; padding: 2px 5px;">V</td></tr> <tr><td style="padding: 2px 5px;">1: x</td></tr> <tr><td style="padding: 2px 5px;">2: o</td></tr> <tr><td style="padding: 2px 5px;">3: o</td></tr> </table>	V	1: x	2: o	3: o	<table style="border-collapse: collapse; margin: 0 auto;"> <tr><td style="border-bottom: 1px solid black; padding: 2px 5px;">A</td></tr> <tr><td style="padding: 2px 5px;">1: x</td></tr> <tr><td style="padding: 2px 5px;">2: o</td></tr> <tr><td style="padding: 2px 5px;">3: o</td></tr> </table>	A	1: x	2: o	3: o	<table style="border-collapse: collapse; margin: 0 auto;"> <tr><td style="border-bottom: 1px solid black; padding: 2px 5px;">V</td></tr> <tr><td style="padding: 2px 5px;">1: x</td></tr> <tr><td style="padding: 2px 5px;">2: o</td></tr> <tr><td style="padding: 2px 5px;">3: o</td></tr> </table>	V	1: x	2: o	3: o
V															
1: x															
2: o															
3: o															
A															
1: x															
2: o															
3: o															
V															
1: x															
2: o															
3: o															

(b)	<table style="border-collapse: collapse; margin: 0 auto;"> <tr><td style="border-bottom: 1px solid black; padding: 2px 5px;">V</td></tr> <tr><td style="padding: 2px 5px;">1: x</td></tr> <tr><td style="padding: 2px 5px;">2: o</td></tr> <tr><td style="padding: 2px 5px;">3: o</td></tr> </table>	V	1: x	2: o	3: o	<table style="border-collapse: collapse; margin: 0 auto;"> <tr><td style="border-bottom: 1px solid black; padding: 2px 5px;">A</td></tr> <tr><td style="padding: 2px 5px;">1: o</td></tr> <tr><td style="padding: 2px 5px;">2: x</td></tr> <tr><td style="padding: 2px 5px;">3: o</td></tr> </table>	A	1: o	2: x	3: o	<table style="border-collapse: collapse; margin: 0 auto;"> <tr><td style="border-bottom: 1px solid black; padding: 2px 5px;">V</td></tr> <tr><td style="padding: 2px 5px;">1: o</td></tr> <tr><td style="padding: 2px 5px;">2: x</td></tr> <tr><td style="padding: 2px 5px;">3: o</td></tr> </table>	V	1: o	2: x	3: o
V															
1: x															
2: o															
3: o															
A															
1: o															
2: x															
3: o															
V															
1: o															
2: x															
3: o															

(c)	<table style="border-collapse: collapse; margin: 0 auto;"> <tr><td style="border-bottom: 1px solid black; padding: 2px 5px;">V</td></tr> <tr><td style="padding: 2px 5px;">1: x</td></tr> <tr><td style="padding: 2px 5px;">2: o</td></tr> <tr><td style="padding: 2px 5px;">3: o</td></tr> </table>	V	1: x	2: o	3: o	<table style="border-collapse: collapse; margin: 0 auto;"> <tr><td style="border-bottom: 1px solid black; padding: 2px 5px;">A</td></tr> <tr><td style="padding: 2px 5px;">1: o</td></tr> <tr><td style="padding: 2px 5px;">2: o</td></tr> <tr><td style="padding: 2px 5px;">3: x</td></tr> </table>	A	1: o	2: o	3: x	<table style="border-collapse: collapse; margin: 0 auto;"> <tr><td style="border-bottom: 1px solid black; padding: 2px 5px;">V</td></tr> <tr><td style="padding: 2px 5px;">1: o</td></tr> <tr><td style="padding: 2px 5px;">2: o</td></tr> <tr><td style="padding: 2px 5px;">3: x</td></tr> </table>	V	1: o	2: o	3: x
V															
1: x															
2: o															
3: o															
A															
1: o															
2: o															
3: x															
V															
1: o															
2: o															
3: x															

Figure 2: Three ways of voting for the candidate 1 in the VAV protocol, where x represents a marked position and o represents an unmarked position.

the candidate of her choice on one of the V-ballots and then (S2) she marks the position next to some randomly chosen candidate on the two remaining simple ballots (one V- and one A-ballot). Figure 2 shows all three ways of filling out the multi-ballot for candidate 1 in an election with three candidates. After this, she feeds all three simple ballots to a voting machine (some kind of scanner) and indicates the simple ballot she wants to get as a receipt. The machine checks the well-formedness of the multi-ballot, prints secretly random numbers on each simple ballot, where the length of these numbers is the length of the security parameter and numbers on different simple ballots are chosen independently, and gives the voter a copy of the chosen simple ballot, with the random number printed on it. Note that the voter does not get to see the random numbers of the remaining two simple ballots. The scanner keeps all simple ballots (now separated) in a ballot box.

In the tallying phase, the voting machine posts on the bulletin board (electronic copies of) all the cast simple ballots in a random order. From the ballots shown on the bulletin board the result can easily be computed: All A-ballots and the corresponding V-ballots (i.e. V-ballots marked at the same position) are removed. From the remaining V-ballots the number of votes for each candidate can directly be read off.

Intuitively, as in the case of ThreeBallot, the system is coercion-resistant (at least to some extent), as every simple ballot that a voter can take as a receipt can be part of a multi-ballot that forms a valid vote for any candidate. Similarly to ThreeBallot, VAV is supposed to provide (some level of) verifiability as it should be risky for any party to remove or alter simple ballots in order to manipulate the result.

In the description of VAV in [26] it is not specified how exactly a voter chooses the receipt. This, however, can heavily affect the properties we study in this paper. We therefore investigate two variants of this protocol.

Simple variant. In the simplest case, a voter could choose one of her three simple ballots with uniform probability as her receipt. We will call this the *simple variant* of VAV. For this variant, however, the receipt is not independent of the candidate of her choice, which, as we will see, has a significant negative effect on privacy and coercion-resistance. This is why we consider also

another variant, where privacy is enhanced.

Privacy enhanced variant. In this variant, a voter chooses one of the two simple ballots *not used* in step (S1), that is, one of those ballots where the random candidate has been marked. By this, the receipt is stochastically independent of the candidate the voter votes for.

While the privacy enhanced variant significantly improves the level of privacy and coercion-resistance, it decreases the level of verifiability: The voting machine is often able to determine one simple ballot in a multi-ballot that was certainly *not* taken as a receipt and consequently it can change this simple ballot without being detected. For instance, if an honest voter submits a multi-ballot as shown in Figure 2, (b) or (c), the machine knows that the left-most ballot cannot be chosen as a receipt. However, as we will see, the VAV voting scheme suffers from the same kind of attack the ThreeBallot protocol does, independently of how the receipts are chosen. Therefore, the level of verifiability is, in any case, very low.

In what follows, we will use, analogously to the case of ThreeBallot, the notion of a *pattern* which specifies how a multi-ballot is filled out and which simple ballot is taken as a receipt.

5 Verifiability

In this section, we first recall the definition of verifiability from [17], where, however, we use a slightly simplified definition which is sufficient for our setting. Next, we present our analysis of verifiability for ThreeBallot and VAV, including the mentioned attacks. We then conclude with remarks on the inadequacy of the notions of individual and universal verifiability demonstrated by our attacks.

5.1 Definition of Verifiability

The definition of verifiability in [17] assumes a *verifier*, also called a *judge*, who can be an honest regular protocol participant or an honest external observer. Now, informally speaking, verifiability says that if in a run of the voting protocol an important goal is not achieved — typically, the published result of the election is not correct, i.e., does not correspond to the votes actually cast by eligible voters —, then the verifier does not accept the run/the election. Conversely, if in a run certain parties which are supposed to make sure that the goal is achieved, such as (a subset of) the voting authorities, behave honestly, then the verifier accepts the run.

More formally, let $P^* = P(A_H, q, V_H, k, \vec{p})$ be a protocol instantiation. Given P^* , for each protocol participant a in P^* , we consider the set $\Pi(a)$ of all programs a may run. This set is defined as follows: If a is assumed to be honest (i.e. $a \in A_H \cup V_H$), then $\Pi(a) = \{\hat{\pi}_a\}$, i.e., $\Pi(a)$ consists only of the honest program of a as specified by the protocol. Otherwise, if a is not assumed to be honest, then $\Pi(a)$ consists of *all* processes limited only by a 's network interface, which is the network interface that $\hat{\pi}_a$ has. Note that in any case $\hat{\pi}_a \in \Pi(a)$.

Let $\Sigma = \{b_1, \dots, b_l\}$ be the set of all protocol participants of P^* . Then, a *process induced by P^** , also called an *instance*, is a process of the form $\pi = (\pi_{b_1} \parallel \dots \parallel \pi_{b_l})$, where $\pi_{b_i} \in \Pi(b_i)$. Such a process is called an *instance with honest $B \subseteq \Sigma$* if $\pi_{b_i} = \hat{\pi}_{b_i}$ for all $b_i \in B$. A *run of P^** is a run of some instance of P^* . Such a run is called a *run with honest B* if it is a run of an instance of P^* with honest B .

The definition of verifiability is parameterized by a goal γ , which, formally, is a set of runs of instances of P^* . In the context of voting, γ will typically contain all those runs in which the

published result of the election is correct, i.e., corresponds to the votes actually cast by eligible voters.

We say that a party a , playing the role of a verifier, *accepts a run*, if in this run a sends the message `accept` on some designated channel decision_a . Intuitively, a accepts a run if she believes that the goal γ has been achieved in this run.

For an instance π of P^* , by $\Pr[\pi(1^\ell) \mapsto (a : \text{accept})]$ we denote the probability that π , running with security parameter 1^ℓ , produces a run which is accepted by a . Similarly, by $\Pr[\pi(1^\ell) \mapsto \neg\gamma, (a : \text{accept})]$ we denote the probability that π , running with security parameter 1^ℓ , produces a run in which the goal has not been achieved, i.e., a run that does not belong to γ , but which nevertheless was accepted by a .

Definition 1 ([17], simplified²). *Let $P^* = P(A_H, q, V_H, k, \vec{p})$ be a protocol instantiation and let Σ be the set of protocol participants in P^* . Let $\delta \in [0, 1]$, $B \subseteq \Sigma$, $a \in A_H \cup V_H$ (playing the role of the verifier), and γ be a goal of P^* . Then, the goal γ is guaranteed in P^* by B and δ -verifiable by a if, for every instance π of P^* , the following conditions are satisfied:*

- (i) *If π is an instance with honest B , then $\Pr[\pi(1^\ell) \mapsto (a : \text{accept})]$ is overwhelming as a function of the security parameter.*
- (ii) *$\Pr[\pi(1^\ell) \mapsto \neg\gamma, (a : \text{accept})]$ is δ -bounded as a function of the security parameter.*

Condition (ii) guarantees that the probability that a accepts a run even though the goal has *not* been achieved (e.g., the published result of the election is incorrect) is “small”, i.e., bounded by δ . Condition (i) says that the protocol is sound w.r.t. a set B of agents in the following sense: If the agents in B are honest, then a accepts runs with overwhelming probability, which by Condition (ii) implies that in those runs the goal has indeed been achieved. Typically, the set B includes (a subset of) voting authorities/machines, i.e., those agents that suffice to guarantee that the goal is achieved. Note that without Condition (i) every protocol in which no runs are accepted by the verifier would be verifiable. Also note that requiring the probability in (ii) to be negligible, i.e., requiring $\delta = 0$, while highly desirable, would be too strong for many reasonable protocols. This is due to the fact that checks (by authorities and voters) are often imperfect and partial, as illustrated in subsequent sections. The value of δ determines the *level of verifiability* a protocol provides.

5.2 Verifiability of ThreeBallot and VAV

In this section, we study verifiability of ThreeBallot and VAV. We precisely measure the level of verifiability of these systems and show that, unlike commonly believed, these systems do not provide any reasonable level of verifiability. We start with the analysis of ThreeBallot.

ThreeBallot. We first describe the attack on the verifiability of ThreeBallot and then precisely state the level of verifiability of this system.

The Attack on the Verifiability of ThreeBallot. As mentioned in Section 3, in the literature the reasoning for the verifiability of ThreeBallot has so far been that, if voters check whether their receipt (a simple ballot) appears on the bulletin board, it should be risky for any party to remove or alter simple ballots since the probability that the modification of k simple ballots goes undetected is merely $(\frac{2}{3})^k$. However, the following attack shows that this reasoning is flawed.

²Definition 1 is a specific instance of the general definition presented in [17]. In [17], instead of the set B , we use a more general formalism to specify sets of protocol participants, namely positive boolean formulas.

Our attack assumes that there are dishonest voters and that one of the voting authorities, the voting machine or the bulletin board, is dishonest and collaborates with the dishonest voters. It is clearly realistic to assume dishonest voters and a dishonest voting authorities; defending against malicious authorities is the main point of verifiability. In what follows, we first consider the case of an election with two candidates and assume that the bulletin board is dishonest.

As already mentioned in the introduction, the effect of our attack is that m dishonest voters, collaborating with the dishonest bulletin board, can effectively vote for candidate B and, additionally, turn m votes of honest voters voting for A into votes for B . For instance, with 10 dishonest voters out of 101 voters, candidate B can win the election, even if 60 honest voters vote for A and only 31 honest voters vote for B . This goes unnoticed, provided no post-election audit based on paper ballots is performed, even if all honest voters check whether their receipts appear on the bulletin board and even if they check that the published result corresponds to the ballots shown on the bulletin board. Note that if no voter complains, then no post-election audit may be carried out. Moreover, for the post-election audit to be effective, additional trust assumptions are required.

The attack works as follows. Let us assume that there exists an honest voter who votes for candidate A and that the bulletin board, collaborating with some dishonest voter, wants to swap such a vote. To do so, the dishonest voter casts $\begin{pmatrix} x & o & o \\ o & x & x \end{pmatrix}$ and sends the serial number on her receipt to the bulletin board. Then, the bulletin board replaces the simple ballot with this serial number by o . The result of this manipulation is as if the dishonest voter had cast $\begin{pmatrix} o & o & o \\ x & x & x \end{pmatrix}$. The bulletin board remains consistent as these three simple ballots together with the multi-ballot submitted by the honest voter voting for A (which must be either $\begin{pmatrix} x & x & o \\ o & x & o \end{pmatrix}$ or $\begin{pmatrix} x & x & o \\ o & o & x \end{pmatrix}$) result in two valid votes for candidate B . Note that the multi-ballot of the honest voter remains unchanged, and hence, no voter or external observer will suspect any fraud.

By this attack, the bulletin board can safely change m votes of honest voters for one candidate to another candidate, where m is the number of dishonest voters.

A similar attack works for the case of multiple candidates. Here the simplest case is that the voting machine is dishonest. First observe that, given any multi-ballot of an honest voter voting for candidate i , it is easy to construct three simple ballots (which potentially do not form a valid multi-ballot) such that these simple ballots together with the multi-ballot of the honest voter form two valid multi-ballots for candidate j . Hence, for every dishonest voter, a voting machine can change the simple ballots of this voter in such a way that they, together with a multi-ballot of an honest voter, result in two valid votes for the candidate of the machine's choice. Note that several side channels are conceivable over which a voter could reveal himself as dishonest to the voting machine, e.g., voting at a specific time, pressing buttons in a specific unusual order, or in case of many candidates, using a pre-agreed pattern to fill out the ballot. Note that this attack works even if the voting machine does not know which simple ballots are taken as receipts.

The Precise Level of Verifiability of ThreeBallot. We now study the precise level of verifiability of both the original variant of ThreeBallot and the variant by de Marneffe et al., showing that only changing votes beyond the number of dishonest voters increases the risk of being detected.

In both cases, we assume that there is a protocol participant *ver* (a regular protocol participant or an external observer), the verifier, who *does not* accept a run iff some voter complains rightly (i.e. she has a receipt that does not appear correctly on the bulletin board) or the bulletin board is inconsistent (e.g., the number of simple ballots is not divisible by three, two serial numbers occur twice, a candidate got less marks than the number of voters, etc.). We assume that an honest voter checks that her receipt occurs on the bulletin board with probability p_{check} —it is realistic to assume that not all voters check their receipt. Clearly, this probability will affect the level of

verifiability. We also make the following assumptions:

1. Only eligible voters will be allowed to vote, and only once. Also, the number of voters who actually voted is properly counted. This is typically guaranteed by clerks. A polling station should at least have one honest clerk who oversees the actions of other clerks. This assumption prevents that the voting machine or the bulletin board can place extra ballots on the bulletin board.
2. Nobody involved in publishing the result, in particular, the voting machine and the bulletin board, should get to know which receipts honest voters chose before all ballots have been published. This assumption is clearly necessary in order to achieve any reasonable level of verifiability, as otherwise the voting machine and the bulletin board could safely change the ballots that were not taken as receipts, and hence, fabricate arbitrary outcomes.
3. The verifier ver behaves as described above.

We note that we neither assume the voting machine nor the bulletin board to be honest.

Let P_{TB}^o and P_{TB}^p denote the ThreeBallot protocol in the original variant and the variant by de Marneffe et al., respectively. Based on the assumptions made above, it is straightforward to formally define the protocol instantiations $S_{TB}^o = P_{TB}^o(\{ver\}, q, V_H, k, \vec{p})$ of P_{TB}^o and $S_{TB}^p = P_{TB}^p(\{ver\}, q, V_H, k, \vec{p})$ of P_{TB}^p , along with the sets $\Pi(a)$ for every protocol participant a as introduced in Section 5.1. Note that $A_H = \{ver\}$ does not include the voting machine or the bulletin board as they are not assumed to be honest. The verifier ver could also belong to V_H . Clerks are not modeled explicitly. The interface the voters have to the rest of the system guarantees assumption 1. above. We define $n = |V_H|$.

We consider the goal γ_ℓ , introduced in [17] (with full details presented in the technical report [18]), which, intuitively, states that at most ℓ votes of honest voters are changed, i.e., the published result is correct (1) up to votes of dishonest voters and (2) up to ℓ votes of honest voters. Note that for dishonest voters not much can be guaranteed as they might, for example, ignore the fact that their receipts are not shown or were modified on the bulletin board. More precisely, γ_ℓ is defined as follows: γ_ℓ contains all runs r for which there exist choices c_1, \dots, c_d , where d is the number of dishonest voters and a choice is either to abstain or to vote for one of the candidates, such that the result obtained from c_1, \dots, c_d together with the choices made by the honest voters³ in r differs only by at most ℓ votes from the published result in r (i.e., the result that can be computed from the simple ballots on the bulletin board). Note that c_1, \dots, c_d are simply quantified existentially, independently of the actions performed by dishonest voters in the considered run. We encourage the reader to consult [18] for more details and explanations about the definition of γ_ℓ .

The following theorem states the levels of verifiability of the two variants of ThreeBallot (see Appendix A for the proof), where q_ℓ denotes the probability of the event that in a run of the protocol there exists a candidate c such that the sum of all votes of honest voters for all candidates except c is at least ℓ . Note that if such an event has not occurred, then it is impossible to violate the goal $\gamma_{\ell-1}$, because, by assumption 1., dishonest parties cannot add new ballots to the bulletin board (as opposed to changing/replacing ballots). However, q_ℓ will typically be quite close to 1.

³Note that the choice made by an honest voter in a run is the parameter j in $\hat{\pi}_v(j)$ with which the voter ran the honest program $\hat{\pi}_v$ (see also Section 2).

Theorem 1. *Let B be the set containing the voting machine and the bulletin board and $x \in \{o, p\}$. Then, the goal γ_ℓ is guaranteed in S_{TB}^x by B and δ_{Ver}^x -verifiable by ver , where*

$$\delta_{Ver}^p = q_{\ell+1} \cdot \left(1 - \frac{1}{6} \cdot p_{check}\right)^{\ell+1-\min(\ell+1,m)}$$

and

$$\delta_{Ver}^o = q_{\ell+1} \cdot \left(1 - \frac{1}{3} \cdot p_{check}\right)^{\ell+1-\min(\ell+1,m)},$$

with m being the number of dishonest voters. Moreover, δ_{Ver}^x is optimal, i.e., there is no $\delta' < \delta_{Ver}^x$ for which the goal γ_ℓ is guaranteed in S_{TB}^x by B and δ' -verifiable by a .

Now, note that if $\ell + 1 \leq m$, i.e., the number of votes to be changed is at most the number of dishonest voters, then $\delta = \delta_{Ver}^p = \delta_{Ver}^o = q_{\ell+1}$. By the definition of δ -verifiability and the fact that δ is optimal, this means that there exists an attack—in fact, the one discussed—such that the probability that in a run more than ℓ votes of honest voters were changed and the verifier still accepted the run is δ . Note also that for $\ell > m$ the original variant of the protocol provides a better level of verifiability than the variant by de Marneffe et al., since in this case $\delta_{Ver}^o < \delta_{Ver}^p$.

VAV. An attack similar to the one for ThreeBallot also works for VAV: Let us assume that there exists an honest voter who votes for candidate j and that the bulletin board, collaborating with some dishonest voter, wants to switch such a vote to a vote for candidate i . To do so, the dishonest voter casts a multi-ballot where i is marked on all three ballots. Moreover, the dishonest voter takes the A -ballot as receipt and sends the serial number on her receipt to the bulletin board. Then, the bulletin board replaces the simple ballot with this serial number by an A -ballot where candidate j is marked. The bulletin board remains consistent, as the three simple ballots by the dishonest voter, with the A -ballot modified, together with the multi-ballot submitted by the honest voter voting for j constitute two valid votes for candidate i .

Hence, VAV too does not provide a reasonable level of verifiable. The exact (low) level of verifiability VAV achieves can be found in Appendix B.

5.3 Inadequacy of Individual and Universal Verifiability

In our analysis of verifiability of ThreeBallot and VAV above, we used Definition 1 as proposed by Küsters et al.[17]. This definition, applied to voting protocols, captures what is called *global verifiability* by Küsters et al. — in short, if the published result of the election is not correct, the verifier (a regular protocol participant or an external observer) should not accept the run, or only with small probability.

However, in the literature (see, e.g., [14, 1, 25, 12]), verifiability of voting protocols has traditionally been expressed by two forms of verifiability, as already mentioned in the introduction: *individual verifiability* (a voter can check that her own ballot appears on the bulletin board) and *universal verifiability* (anyone can check that the election outcome corresponds to the ballots published on the bulletin board). Note that, unlike global verifiability, these forms of verifiability assume some particular structure of the voting protocol. Also note that these forms of verifiability can be captured by Definition 1 using appropriate goals γ .

In the literature it was widely believed that individual and universal verifiability together achieve some form of global verifiability. However, our case study on ThreeBallot and VAV

shows that this is not the case. These protocols achieve both individual and universal verifiability, but as we proved, their level of (global) verifiability is completely insufficient: A voter can check whether her receipt appears on the bulletin board, which gives her relatively high assurance that all her simple ballots are unmodified and appear on the bulletin board; hence, we have individual verifiability. (More precisely, as explained in Sections 3 and 4, if fraud would be attempted, even only on a moderate scale, the probability that at least one voter would detect a problem with her receipt would be very high.) We also obviously have universal verifiability as the result of the election can be computed by everyone based on the information available on the bulletin board.

In general, what individual and universal verifiability ignore is that dishonest authorities/voters can break the integrity of ballots of honest voters by ill-formed ballots. Therefore, we advocate using global verifiability (see above) which directly captures the required property.

6 Privacy and Coercion-Resistance

In this section, we first introduce our definition of privacy. We also briefly recall the definition of coercion-resistance from [16]. We then present our analysis of privacy and coercion-resistance of the variants of ThreeBallot and VAV described in Sections 3 and 4. We conclude the section with a discussion of the relationship between privacy and coercion-resistance.

6.1 Definition of Privacy

For studying privacy of a protocol P , we assume that, besides the voting authorities and the voters, there is an additional party o called an *observer*. We denote by O the set of all programs an observer can run, i.e. all probabilistic polynomial-time ITMs with the following communication interface: An observer can directly connect to the interface of dishonest voters and authorities; in fact, the observer subsumes those parties. In addition, observers can observe publicly available information, such as messages posted by voting authorities. We also assume that, in a protocol instantiation $P^* = P(A_H, q, V_H, k, \vec{p})$, among the q voters, there is a voter who is under observation.

Now, a protocol instantiation $P^* = P(A_H, q, V_H, k, \vec{p})$, along with the set O of observer processes and a program π_v of the voter v under observation, induces a set of processes of the form $(\pi_o \parallel \pi_v \parallel e)$, where $\pi_o \in O$ and e denotes the concurrent composition of the processes $\hat{\pi}_{v'}$, $v' \in V_H$, of the honest voters and the processes $\hat{\pi}_a$, $a \in A_H$, of the honest authorities; recall that the dishonest voters and the dishonest authorities are subsumed by π_o . We denote by $\Pr[(\pi_o \parallel \pi_v \parallel e)^{(\ell)} \mapsto 1]$ the probability that π_o outputs 1 in a run of the process $(\pi_o \parallel \pi_v \parallel e)$ with security parameter 1^ℓ .

In the following definition, we formalize privacy to be the inability of the observer π_o to distinguish whether the voter v under observation voted for candidate j or candidate j' , where v runs her *honest* voting process $\hat{\pi}_v$ as specified by the voting protocol.

Definition 2. Let $P^* = P(A_H, q, V_H, k, \vec{p})$ be a protocol instantiation along with a set O of observer processes and a voter v under observation. Let $\delta \in [0, 1]$. We say that P^* achieves δ -privacy, if

$$\Pr[(\pi_o \parallel \hat{\pi}_v(j) \parallel e)^{(\ell)} \mapsto 1] - \Pr[(\pi_o \parallel \hat{\pi}_v(j') \parallel e)^{(\ell)} \mapsto 1] \quad (1)$$

is δ -bounded as a function of the security parameter 1^ℓ , for all $j, j' \in \{1, \dots, k\}$ and for all $\pi_o \in O$.

In the above definition we merely require (1) to be δ -bounded, instead of negligible, because there is always a non-negligible probability that an observer knows how a voter voted, e.g., in case

all (honest) voters and the voter under consideration voted for the same candidate. In general, and as we will see below and in Section 6.3, δ depends on the distribution \vec{p} , the number k of candidates, and the number $n = |V_H|$ of honest voters; the number of dishonest voters is typically not relevant. By δ , we can precisely measure the level of privacy a voting protocol offers.

We note that the above definition does not imply that an observer cannot distinguish whether or not a voter voted, i.e., abstention may be detected by an observer. If abstention should not be detectable, one can simply let j and j' range over $\{0, \dots, k\}$ instead of $\{1, \dots, k\}$ in the above definition. The above definition is motivated by the fact that for many voting protocols, including ThreeBallot and VAV, abstention can be detected by an observer, since, e.g., the observer is present at the polling station or the observer can see the receipts of all voters, and in particular, he can see whether a voter does not have a receipt.

The above definition could be generalized in the obvious way by letting the observer observe many voters at the same time and quantifying over two *tuples* of votes, instead over just j and j' . While, for simplicity, in our case study we consider the version with one observed voter, our findings and theorems extend to the case with multiple observed voters.

We note that the above (cryptographic) definition of privacy is close in spirit to a definition in an abstract, Dolev-Yao style model [9]. Simulation-based definitions (see, e.g., [22]) are stronger, but often too strong to be applicable (e.g., in case of ThreeBallot and VAV).

Privacy of the Ideal Protocol. As we have already mentioned, the level δ of privacy is bigger than zero for virtually every voting protocol, as some information is always leaked by the result of the election, e.g., if one candidate got all votes—an event with non-negligible probability—, it is clear that everybody voted for this candidate. In order to have a lower bound on δ for all voting protocols (where the results are of the form considered below), we now determine the optimal value of δ for the ideal voting protocol. An *ideal voting protocol* collects the votes of all voters and then correctly publishes the result, where we assume that a result reveals the number of votes for every candidate. The argument sketched below is similar to the one for determining the level of *coercion-resistance* of the ideal voting protocol in [16].

As we will see, the level of privacy of the ideal voting protocol, denoted by $\delta_{priv}(k, n, \vec{p})$, depends on the number k of candidates, on the number n of honest voters, and the probability distribution \vec{p} on the candidates.

To define this function, we need the following terminology. Since the observer knows the votes of the dishonest voters, he can simply subtract these votes from the final result and obtain what we call the *pure result* $\vec{r} = (r_0, \dots, r_k)$ of the election, where r_i , $i \in \{1, \dots, k\}$, is the number of votes for candidate i casted by honest voters, and r_0 is the number of honest voters who abstained from voting. Note that $r_0 + \dots + r_k = n + 1$ (n honest voters plus the observed voter). We denote by *Res* the set of all pure results. Let $A_{\vec{r}}^i$ denote the probability that the choices made by the honest voters yield the pure result \vec{r} , given that the voter under observation votes for the i -th candidate. (Clearly, $A_{\vec{r}}^i$ depends on \vec{p} . However, we omit this in the notation.) It is easy to see that

$$\begin{aligned} A_{\vec{r}}^i &= \binom{n}{r_0, \dots, r_{i-1}, r_i - 1, r_{i+1}, \dots, r_k} \cdot p_0^{r_0} \dots p_{i-1}^{r_{i-1}} p_i^{r_i - 1} p_{i+1}^{r_{i+1}} \dots p_k^{r_k} \\ &= \frac{n!}{r_0! \dots r_k!} \cdot p_0^{r_0} \dots p_k^{r_k} \cdot \frac{r_i}{p_i}, \end{aligned}$$

where $\binom{n}{m_0, \dots, m_k} = \frac{n!}{m_0! \dots m_k!}$. Moreover, let $M_{j, j'}^* = \{\vec{r} \in Res : A_{\vec{r}}^j \leq A_{\vec{r}}^{j'}\}$. Now, the intuition behind the definition of $\delta_{priv}(k, n, \vec{p})$ is as follows: If the observer, given a pure result \vec{r} , wants to

decide whether the observed voter voted for candidate j or j' , the best strategy of the observer is to opt for j' if $\vec{r} \in M_{j,j'}^*$, i.e., the pure result is more likely if the voter voted for candidate j' . This leads to the following definition:

$$\delta_{priv}(n, k, \vec{p}) = \max_{j, j' \in \{1, \dots, k\}} \sum_{\vec{r} \in M_{j,j'}^*} (A_{\vec{r}}^{j'} - A_{\vec{r}}^j).$$

The following theorem states that $\delta_{priv}(k, n, \vec{p})$ is indeed the optimal level of privacy, where VA denotes the trusted authority in the ideal voting protocol.

Theorem 2. *Let $S = P_{ideal}(\{VA\}, q, n, k, \vec{p})$ be an instantiation of the ideal protocol and $\delta = \delta_{priv}(n, k, \vec{p})$. Then S achieves δ -privacy. Moreover, S does not achieve δ' -privacy for any $\delta' < \delta$.*

The proof of this theorem can be found in Appendix C. Some values for $\delta_{priv}(n, k, \vec{p})$ are depicted in Figure 3 (see the values for the ideal protocol).

6.2 Definition of Coercion-Resistance

We now briefly recall the definition of coercion-resistance from [16]. Since the overall setting for coercion-resistance is similar to that of privacy, we highlight the differences to privacy.

For the definition of coercion-resistance, the voter under observation considered for privacy is now replaced by a *voter under coercion*, also called a *coerced voter*. Unlike a voter under observation, a coerced voter does not have to follow the honest voting procedure but can deviate from it. We denote by V the set of all programs the coerced voter v can run. This set includes all probabilistic polynomial-time ITMs where the communication interface is that of an honest voter plus an input and output channel for communication with the coercer (see below). In particular, the set V contains what we call a *dummy strategy* dum which simply forwards all the messages between the coercer and the interface the coerced voter has as an honest voter.

The observer in the case of privacy is now replaced by the *coercer*. We denote by C the set of all programs a coercer can run, i.e., all probabilistic polynomial-time ITMs with a communication interface similar to that of observers, where in addition the coercer can communicate with the coerced voter.

Before recalling the formal definition of coercion-resistance, we provide some intuition. We imagine that the coercer demands full control over the voting interface of the coerced voter, i.e., the coercer wants the coerced voter to run the dummy strategy $dum \in V$ instead of the program an honest voter would run. If the coerced voter in fact runs dum , the coercer can effectively vote on behalf of the coerced voter or decide to abstain from voting. Of course, the coercer is not bound to follow the specified voting procedure.

Now, informally speaking, a protocol is called coercion-resistant if the coerced voter, instead of running the dummy strategy, can run some *counter-strategy* $\tilde{v} \in V$ such that (i) by running this counter-strategy, the coerced voter achieves her own goal, e.g., votes for a specific candidate (see below), and (ii) the coercer is not able to distinguish whether the coerced voter followed his instructions (i.e., ran dum) or tried to achieve her own goal (by running \tilde{v}). If such a counter-strategy exists, then it indeed does not make sense for the coercer to try to influence a voter in any way, e.g., by offering money and/or threatening the voter: Even if the coerced voter tries to sell her vote by running dum , i.e., by following the coercer's instructions, the coercer is not able to tell whether the coerced voter is actually following the coercer's instructions or whether she

is just trying to achieve her own goal. For the same reason, the coerced voter can safely run the counter-strategy and achieve her own goal, even if she is coerced into running dum.

The *goal* of the coerced voter is formalized by a set γ of runs. For example, if γ is supposed to express that the coerced voter wants to vote for a certain candidate, then γ would contain all runs in which the coerced voter (successfully) voted for this candidate.

In the formal definition of coercion-resistance, we write, analogously to the case of privacy, $\Pr[(\pi_c \parallel \pi_v \parallel e)^{(\ell)} \mapsto 1]$ for the probability that π_c outputs 1 in a run of the process $(\pi_c \parallel \pi_v \parallel e)$ with security parameter 1^ℓ . We write $\Pr[(\pi_c \parallel \pi_v \parallel e)^{(\ell)} \mapsto \gamma]$ for the probability that a run of $(\pi_c \parallel \pi_v \parallel e)$, with security parameter 1^ℓ , belongs to γ .

Definition 3 ([16]). Let $P^* = P(A_H, q, V_H, k, \vec{p})$ be a protocol instantiation and let V and C be sets of processes as above. Let $\delta \in [0, 1]$ and let γ be a goal. Then, P^* is δ -coercion-resistant w.r.t. γ , if there exists $\tilde{v} \in V$ such that the following conditions are satisfied:

- (i) $\Pr[(\pi_c \parallel \tilde{v} \parallel e)^{(\ell)} \mapsto \gamma]$ is overwhelming, as a function of ℓ , for every $\pi_c \in C$.
- (ii) $\Pr[(\pi_c \parallel \text{dum} \parallel e)^{(\ell)} \mapsto 1] - \Pr[(\pi_c \parallel \tilde{v} \parallel e)^{(\ell)} \mapsto 1]$ is δ -bounded, as a function of ℓ , for every $\pi_c \in C$.

Similar to the case of privacy, in Condition (ii) the difference is required to be δ -bounded instead of negligible, since there is always a non-negligible chance for the coercer to know for sure whether the coerced voter followed his instructions or not. For example, if one candidate got all votes, but the coercer told the coerced voter to vote for a different candidate, then the coercer knows that the coerced voter did not follow his instructions. In general, as in the case of privacy, δ is a function of \vec{p} , k , and the number $n = |V_H|$ of honest voters (see below and Section 6.4). Clearly, a small δ is preferable. Let us illustrate the meaning of δ by the following example. Assume that if π_c outputs 1, i.e., the coercer thinks that the coerced voter is following his instructions, then the coercer pays \$100 to the coerced voter, and otherwise, if π_c outputs 0, i.e., the coercer thinks that the coerced voter did not follow his instructions, he might harm the coerced voter. Now, if $\delta = 0.8$, then this means that if the coerced voter follows the instructions of the coercer, the coerced voter increases her chances of getting payed (not being harmed) by up to 80%. Conversely, by following the counter-strategy, the coerced voter drastically decreases her chances of getting payed and increases her chances of being harmed. This might be a strong incentive for the coerced voter to follow the instructions of the coercer.

While here we concentrated on the case for one coerced voter, the above definition in fact also applies to the setting of multiple coerced voters (see [16]).

Coercion-resistance of the Ideal Protocol. Since we will refer to the level of coercion-resistance of the ideal protocol in this paper, we recall the optimal level of coercion-resistance established in [16]. Similarly to the case of privacy, let $A_{\vec{r}}^i$ denote the probability that the choices made by the honest voters and the coerced voter yield the pure result $\vec{r} = (r_0, \dots, r_k)$, given that the coerced voter votes for the i -th candidate. Also, let $M_{i,j}^* = \{\vec{r} \in \text{Res} : A_{\vec{r}}^i \leq A_{\vec{r}}^j\}$ and

$$\delta_{\min}^i(n, k, \vec{p}) = \max_{j \in \{1, \dots, k\}} \sum_{\vec{r} \in M_{i,j}^*} (A_{\vec{r}}^j - A_{\vec{r}}^i).$$

Let γ_i be the goal of the coerced voter which is achieved if the coerced voter votes for candidate i , in case she is instructed by the coercer to vote (for some candidate). Note that coercion-resistance w.r.t. this goal does not imply that forced abstention attacks are prevented: If the coercer

wants the coerced voter to abstain from voting, the coerced voter, when running her counter-strategy, does not need to vote in order to fulfil the goal. While for the ideal protocol a stronger goal which says that the coerced voter in any case votes for i could be considered, for ThreeBallot and VAV such a goal, which requires that forced abstention attacks are not possible, is too strong (see Section 6.4). The following theorem states that $\delta_{min}^i(n, k, \vec{p})$ is optimal for γ_i :

Theorem 3 ([16]). *Let $S = P_{ideal}(\{VA\}, q, n, k, \vec{p})$. Then, S is δ -coercion-resistant w.r.t. γ_i , where $\delta = \delta_{min}^i(n, k, \vec{p})$. Moreover, S is not δ' -coercion-resistant for any $\delta' < \delta$.*

We note that the level of privacy of the ideal protocol coincides with the level of coercion-resistance of the ideal protocol, if the goal of the coerced voter is to vote for the candidate with the smallest probability according to \vec{p} .

6.3 Privacy of ThreeBallot and VAV

In this section, we analyze the level of privacy provided by all variants of ThreeBallot and VAV described in Sections 3 and 4. In all cases, the presentation of the results follows the same structure: First, we define what we call an *essential view of the observer*, where we abstract away from some parts of the observer’s full view in a given protocol run. Based on the notion of an essential view, we define the optimal level of privacy, δ , and state the result. Due to the similar structure, we introduce the necessary terminology and present the results “in parallel” for all protocol variants. We start with our modeling and security assumptions, which are largely the same for all variants.

Modeling and Security Assumptions. In our analysis, we assume that the observer can see whether a voter enters the voting booth. We also assume that an honest voter may reveal her (paper) receipt to the observer, after the voting phase is finished. However, to measure how much information an observer gains from the receipts of honest voters, we will also consider the case that the observer does not get to see the receipts of honest voters.

Moreover, we assume that the voting machine (the scanner) is honest; the bulletin board may be dishonest. Note that this assumption is indeed necessary for privacy: without this assumption, the observer gets to know how voters vote, as the voters disclose their votes to the machine. Even though the machine, at the moment a voter votes, might not know who the voter is, this information could be reconstructed from the order in which voters voted.

In our analysis of ThreeBallot, we focus on the case with two candidates, i.e., a case where the so-called short ballot assumption is fulfilled. It is well-known that without this assumption, ThreeBallot does not have a sufficient level of privacy and coercion-resistance (see, e.g., [27, 11]). The degradation of the level of coercion-resistance of the variant of ThreeBallot by Marneffe et al. in the multi-candidate case was formally studied in [16]. However, in our analysis of VAV we do not restrict the number of candidates.

By P_{VAV}^{s+} and P_{VAV}^{s-} we denote the simple variant of the VAV protocol (modeled as a protocol in the sense of Section 2), where ‘+’ and ‘-’ indicate whether or not the honest voters reveal their receipts. Similarly, P_{VAV}^{p+} and P_{VAV}^{p-} denote the privacy enhanced variant of VAV, with/without receipts being revealed. As for ThreeBallot, we use P_{TB}^{o+} and P_{TB}^{o-} for the original variant and P_{TB}^{p+} and P_{TB}^{p-} for the variant by Marneffe et al.

Following our modeling and security assumptions, we consider instantiations of P_{VAV}^{s+} , P_{VAV}^{s-} , P_{VAV}^{p+} , P_{VAV}^{p-} , P_{TB}^{o+} , P_{TB}^{o-} , P_{TB}^{p+} and P_{TB}^{p-} , where the parameters are chosen as follows: (i) the set $A_H = \{M\}$ of honest authorities contains the voting machine M only, (ii) the number q of (honest

and dishonest) voters, (iii) some number n of honest voters, (iv) some number k of candidates, and (v) some probability distribution \vec{p} on the candidates. In case of ThreeBallot, we assume that $k = 2$. For brevity of notation, we, for instance, simply write P_{VAV}^{s+} , instead of $P_{\text{VAV}}^{s+}(A_H, q, n, k, \vec{p})$.

We denote the set of all the instantiations described above by \mathcal{S} .

Views and essential views. The *view* of the observer in a protocol run contains (1) the random coins generated by the observer, (2) optionally, depending on the case under consideration, the receipts of the honest voters, after the voting phase is finished, and (3) all messages received from the interface of the dishonest parties (which the observer controls). The latter includes all dishonest voters and the bulletin board, containing the shuffled simple ballots with serial numbers. Note that the observer cannot directly see the information the honest voters obtain or the actions they perform in the voting booth.

In an *essential view* of the observer we abstract away from those parts of his view which are not relevant for distinguishing how voters vote, e.g., the serial numbers on the simple ballots, the order of the simple ballots on the bulletin board or the simple ballots of the dishonest voters (which are determined by the observer). The crucial part of the proof of Theorem 4, stated below, is to show that, indeed, the observer can without loss of generality base his decision solely on such essential views.

More precisely, if the observer cannot see the receipts of the voters, the essential view is defined to be the pure result of the election, as defined in Section 6.1. If the observer can see the receipts, the definition of an essential view depends on the system under consideration:

- *ThreeBallot*: The essential view of the observer consists of (i) three integers n_x^s , n_o^s , and n_o^o indicating the number of the respective simple ballots on the bulletin board cast by honest voters, including the observed voter, and (ii) the receipt r of the voter under observation and (iii) integers r_x^s , r_o^s , and r_o^o , indicating the number of the respective receipts taken by the honest voters.

Note that from these numbers the number of o -ballots on the bulletin board and the number of o -receipts can easily be computed.

- *Privacy enhanced variant of VAV*: The essential view is just the pure result.

- *Simple variant of VAV*: The essential view of the observer consists of (i) integers n_j^A , n_j^V , for each candidate j , indicating the number of A- and V-ballots, respectively, on the bulletin board on which candidate j is marked, (ii) the receipt r of the voter under observation and (iii) integers r_j^A and r_j^V , indicating the number of A- and V-ballots, respectively, taken by the honest voters as receipts on which candidate j is marked.

By EV^S we denote the set of all essential views of the observer for the instantiation $S \in \mathcal{S}$.

Level of privacy. Let $S \in \mathcal{S}$. For an essential view $\rho \in EV^S$, and a candidate i , let $A_{\rho,i}^S$ denote the probability that, in a run of S , the essential view of the observer is ρ , given that the observed voter votes for i . For $i, j \in \{1, \dots, k\}$, let $M_{j,i}^S = \{\rho \in EV^S : A_{\rho,j}^S \leq A_{\rho,i}^S\}$. Similar to the case of privacy for the ideal protocol, the intuition behind the definition of δ , given below, is the following: If the observer, given an essential view ρ , wants to decide whether the observed voter voted for candidate i or j , the best strategy for the observer is to opt for i , if $\rho \in M_{j,i}^S$, i.e. his (essential) view is more likely if the voter voted for candidate i .




Figure 3: Level of privacy (the smaller δ the higher privacy) for the considered variants of Three-Ballot with two candidates, $p_0 = 0.3$, $p_1 = 0.35$, $p_2 = 0.35$.

Now, we are ready to express the function representing the level of privacy in the instantiation $S \in \mathcal{S}$:

$$\delta_{priv}^S(n, k, \vec{p}) = \max_{i,j=1,\dots,k} \sum_{\rho \in M_{j,i}^S} (A_{\rho,i}^S - A_{\rho,j}^S).$$

The following theorem shows that $\delta_{priv}^S = \delta_{priv}^S(n, k, \vec{p})$ indeed is the optimal level of privacy achieved by the considered instantiations S of ThreeBallot and VAV.

Theorem 4. *Let $S \in \mathcal{S}$. The instantiation S achieves δ_{priv}^S -privacy. Moreover, S does not achieve δ' -privacy for any $\delta' < \delta_{priv}^S$.*

The proof of this theorem can be found in Appendix D. While the representation of δ_{priv}^S is the same for every instantiation (except that the definitions of essential views differ), the proofs of Theorem 4 for the various instantiations differ significantly. In each case, we first show that the view of the observer can indeed be reduced to the corresponding essential view, and then, by combinatorial arguments, we show that δ_{priv}^S is optimal. Due to space limitations, we do not present this proof here.

For every $S \in \mathcal{S}$, we have developed concrete formulas for $A_{\rho,i}^S$, which, in some cases, involved non-trivial combinatorial reasoning (see Appendix E). These formulas allowed us to compute concrete values for δ_{priv}^S , as depicted in Figure 3 for the case of two candidates. Note that election results are often published per polling station, with just a few hundred voters each.

As can be seen from Figure 3 and as we formally prove in Appendix G, the variants of the protocols where the observer does not get to see receipts of voters provide the ideal level of privacy. The privacy enhanced variant of VAV with receipts being revealed is ideal too: Intuitively, the reason for this is that the receipts taken by the honest voters are picked independently of the chosen candidates. Furthermore, the bulletin board does not leak any information about how a given voter voted, except for the bare result—the ballots of type A and V which cancel out are chosen independently of the voters' choices. The variant of ThreeBallot by Marneffe et al. is close to ideal. However, the level of privacy of the original variant of ThreeBallot and the simple variant

of VAV is unacceptable. This is due to the receipts which, for these variants of the protocols, leak a lot of information about a vote. In case of VAV, for example, it is easy to see that with probability $\frac{1}{3}$, a voter takes the simple ballot as a receipt which exactly shows her choice, hence, δ can never drop below $\frac{1}{3}$.

6.4 Coercion-Resistance of ThreeBallot and VAV

In this section, we analyze the level of coercion-resistance provided by all variants of ThreeBallot and VAV described in Sections 3 and 4. We note that the level of coercion-resistance of the variant by Marneffe et al. [8] has already been established in [16]. However, the results for VAV and the one for the original variant of ThreeBallot are new.

As in the case of privacy, the presentation of the results follows the same structure for all protocol variants, which is why we again introduce the necessary terminology and present the results “in parallel” for all these variants. For coercion-resistance, we also use the notion of an essential view (although defined differently). In addition, we have to define the goal of the coerced voter and the counter-strategy.

Modeling and security assumptions. We make the same modeling and security assumptions as in the case of privacy and consider the same set \mathcal{S} of concrete instantiations.

The goal of the coerced voter. Our analysis is w.r.t. the goal γ_i , for $i \in \{1, \dots, k\}$, which is met if the coerced voter votes for candidate i , in case she is instructed by the coercer to vote for some candidate. Note that if the coerced voter is not instructed to vote, she cannot vote, as this fact would be observed by the coercer, who sees if the voter enters the voting booth (forced-abstention attack). Recall that for ThreeBallot, we assume $k = 2$.

Counter-strategy. We define the counter-strategy of the coerced voter for an instance $S \in \mathcal{S}$ as follows: The counter-strategy coincides with the dummy strategy *dum* with one exception: If the coerced voter is requested to fill out her ballot and cast it according to a certain pattern Z , then the coerced voter will, instead, fill out the ballot according to $C^S(Z, i)$, as defined next. (Recall that the goal of the coerced voter is to vote for i .) We define $C^S(Z, i)$ to be Z , if the pattern Z forms a valid vote for i . Otherwise, the definition of $C^S(Z, i)$ depends on the protocol under consideration:

- *ThreeBallot*: We define $C^S(Z, i)$ in such a way that it yields the same receipt as Z does. Moreover, it adjusts the two remaining ballots in such a way that the resulting multi-ballot is a valid vote for candidate i . By this requirement, $C^S(Z, i)$ is uniquely determined, except for two cases: $C^S((\begin{smallmatrix} x & o \\ o & x \end{smallmatrix}, x), 1)$ and $C^S((\begin{smallmatrix} o & x \\ x & o \end{smallmatrix}, 2)$. In the former case, for instance, one can take $(\begin{smallmatrix} x & x \\ o & o \end{smallmatrix}, x)$ or $(\begin{smallmatrix} x & o \\ o & x \end{smallmatrix}, o)$, or randomly pick one of the two, possibly based on some further information. For these cases, we define $C^S(Z, i)$ to choose one of the possible patterns uniformly at random.
- *VAV*: Similarly as in the case for ThreeBallot, $C^S(Z, i)$ is defined in such a way that it yields the same receipt as Z does. It adjusts the two remaining ballots in such a way that the resulting multi-ballot is a valid vote for candidate i . By this requirement, $C^S(Z, i)$ is uniquely determined, except for the case where the coercer demands a V -receipt for candidate i . In this case, one can mark an arbitrary candidate on the remaining two ballots. We demand that $C^S(Z, i)$ then marks candidate i also on the remaining ballots.

We use these strategies in the proof of Theorem 5. From the proof of this theorem it follows that these counter-strategies achieve the maximal level of coercion-resistance and, in this sense, are optimal for the coerced voter.

Essential views. The *essential view of the coercer* is defined as follows:

- *ThreeBallot*: If the coercer can see the receipts of honest voters, the essential view is defined just like the essential view of the observer in case of privacy (see Section 6.3), except that it does not contain the receipt of the coerced voter (as the coerced voter returns always the receipt demanded by the coercer). If the coercer cannot see the receipts, the essential view consists of the integers n_x^s , n_o^s , and n_z^s , representing the numbers of the respective simple ballots on the bulletin board.
- *Privacy enhanced variant of VAV*: Regardless of whether the coercer can see the receipts or not, an essential view of the coercer consists of two integers n_j^A and n_j^V for each candidate j , indicating the number of A- and V-ballots on the bulletin board, respectively, with candidate j marked.
- *Simple variant of VAV*: The essential view of the coercer consists of integers n_j^A and n_j^V , for each candidate j , as defined above and, if the coercer can see the receipts, additionally, two integers r_j^A and r_j^V , for each candidate j , indicating the number of A- and V-ballots, respectively, taken by the honest voters as receipts with candidate j marked.

By EV^S we denote the set of all essential views of the coercer for the instantiation $S \in \mathcal{S}$.

Level of coercion-resistance. Let $S \in \mathcal{S}$ and $\rho \in EV^S$. We define $A_{\rho,Z}^S$ to be the probability that, in a run of S , the choices made by the honest voters and the coerced voter result in the essential view ρ , given that the coerced voter chooses the pattern Z . The definition of the level of coercion-resistance, δ , now follows the same idea as in the case of privacy. We define $M_{Z,i}^S = \{\rho \in EV^S : A_{\rho,C(Z,i)}^S \leq A_{\rho,Z}^S\}$ to be the set of those essential views for which the coercer should accept the run and we define

$$\delta_{CR}^S(n, k, \vec{p}) = \max_Z \sum_{\rho \in M_{Z,i}^S} (A_{\rho,Z}^S - A_{\rho,C(Z,i)}^S). \quad (2)$$

The following theorem shows that $\delta_{CR}^S = \delta_{CR}^S(n, k, \vec{p})$ indeed is the optimal level of coercion-resistance for the instantiation S , where the case in which S is the variant of ThreeBallot by Marneffe et al. was shown in [16].

Theorem 5. *Let $S \in \mathcal{S}$. Then S is δ_{CR}^S -coercion-resistant. Moreover, S is not δ' -coercion-resistant for any $\delta' < \delta_{CR}^S$.*

Similar to the case of privacy, the details of the proofs for the different variants S of the protocols differ significantly. The proof of this theorem can be found in Appendix F.

We developed concrete formulas for $A_{\rho,i}^S$ and $A_{\rho,C(Z,i)}^S$, which involved non-trivial combinatorial arguments, but allowed us to compute concrete values for δ_{CR}^S , as depicted in Figure 4 for the case of two candidates (see Appendix E). To put these values in context, we present also the corresponding values for the variant of ThreeBallot and the ideal protocol studied in [16]. As we can see in Figure 4, for all the considered protocols, the level of coercion-resistance is lower if the coercer can see the receipts.⁴ Altogether, under the same assumptions, VAV provides a better level of coercion-resistance than ThreeBallot, but both are still worse than the ideal protocol.

Some selected values for the multi-candidate case, namely 10 candidates, are depicted in Figure 5. To put these values in context, we present also the corresponding values for the ideal

⁴Note that in [20], incorrect values of the level of coercion-resistance for P_{VAV}^{p+} are given.

protocol and the variant of ThreeBallot by Marneffe et al. as studied in [16]. We can see that the (privacy enhanced) variant of VAV handles the case of multiple candidates much better than ThreeBallot, which for 10 candidate basically does not provide any coercion-resistance. (Recall that δ close to 1 means that the coercer can tell almost for sure whether the coerced voter followed his instructions or not.)




Figure 4: Level of coercion (δ) for different protocols with two candidates, $p_0 = 0.3$, $p_1 = p_2 = 0.35$. The goal of the coerced voter is to vote for candidate 1.

6.5 Relationship between Privacy and Coercion-resistance

As already mentioned in the introduction, one would expect that privacy and coercion-resistance are closely connected: If a protocol provides a poor level of privacy, i.e., an observer has a good chance of distinguishing whether an honest voter voted for candidate j or j' , then this should give the coercer leverage to distinguish whether the coerced voter followed the coercer's instructions or not. Indeed some works suggest a close connection between coercion-resistance and privacy, most notable the already mentioned work by Moran and Naor [22] and work in an abstract, Dolev-Yao style approach [10], which, however, puts strong restrictions on the coercer and counter-strategies. The definitions used in these works yield yes/no-answers, rather than measuring the level of coercion-resistance and privacy.

Our case study of ThreeBallot and VAV demonstrates that the connection between privacy and coercion-resistance, in particular when it comes to measuring the level of privacy and coercion-resistance, is more subtle than what can be gathered from existing work.

One observation that comes out of our case study is that improving the level of privacy of a protocol in a natural way (e.g., by changing the way honest voters fill out ballots) can lead to a *worse* level of coercion-resistance. This is the case when going from the original variant of ThreeBallot to the “privacy enhanced” variant by de Marneffe et al., as can be seen from the results in Sections 6.3 and 6.4 (compare the levels of privacy of the systems S_{TB}^{o+} and S_{TB}^{p+} , given in Figure 3, with those for coercion-resistance of the same systems in Figure 4). Clearly, in general, one does not expect privacy to imply coercion-resistance. One might, however, expect that improving privacy also improves coercion-resistance. In this sense, the illustrated effect is




Figure 5: The lower-bound of coercion-resistance (δ) for ThreeBallot in the variant by Marneffe et al. and the precise values for VAV in the privacy enhanced variant and the ideal voting protocol with 10 candidates, where an honest voter abstains from voting with probability $p_0 = 0.3$ and she chooses a candidate with probability $((1 - p_0)/10)$.

surprising. At the end of this section, we propose another (though artificial) variant of ThreeBallot which better explains and amplifies the described effect.

Another finding that comes out of our case study, which maybe more unexpected, is that the level of privacy of a protocol can be much lower than its level of coercion-resistance. This is so for the original variant of ThreeBallot and the simple variant of VAV, as can be seen from the results in Sections 6.3 and 6.4 (compare the level of privacy of S_{TB}^{o+} , given in Figure 3, with the level of coercion resistance of this system, given in Figure 4; similarly for the system S_{VAV}^{+}). The reason behind this phenomenon is basically that the counter-strategy hides the behavior of the coerced voter, including her vote, better than the honest voting program hides the vote. Conversely, one could say that the honest voting program is “suboptimal” in hiding the way the voter voted. In the original variant of ThreeBallot and the simple variant of VAV, a receipt an honest voter obtains indeed discloses more information than necessary. The following simple, but unlike ThreeBallot and VAV, artificial example, carries this effect to extremes: Consider the ideal voting protocol which collects all votes and publishes the correct result. Now imagine a voting protocol in which voters use the ideal voting protocol to cast their vote, but where half of the voters publish how they voted (e.g., based on a coin flip). Clearly, the level of privacy this protocol provides is very low, namely $\delta \geq \frac{1}{2}$. However, a coerced voter can be more clever and simply lie about how she voted. This protocol indeed provides a high level of coercion-resistance, namely $\delta \approx \delta_{min}^i(n/2, k, \vec{p})$ (see Section 6.2). Below we also provide a slightly more subtle example based on ThreeBallot.

In case the counter-strategy does not “outperform” the honest voting program (or conversely, the honest voting program does not leak more information than the counter-strategy), one would expect that if a voting system provides a certain level of coercion-resistance, then it provides at least the same level of privacy. We now show that this is indeed true.

We first have to define what it means for the counter-strategy to outperform the honest voting program.

For this purpose, let $P^* = P(A_H, q, V_H, k, \vec{p})$ be a protocol instantiation, with sets V and C of

processes as in Definition 3. As usual, with $\hat{\pi}_v(j)$ we denote the honest voting program voting for candidate j and with γ_i we denote the goal which contains all runs in which the coerced voter voted for candidate i (or the weaker goal, where this is required, only if the voter is instructed by the coercer to vote for some candidate). Let \tilde{v}_i be a counter-strategy that tries to achieve this goal. Let π_c^j be a process of the coercer which only connects to dum, but does not use any other part of the interface the coercer can connect to, and which simply simulates the program of an honest voter voting for candidate j . Clearly, the systems $(\pi_c^j \parallel \text{dum})$ and $\hat{\pi}_v(j)$ are identical from the point of view of an external observer.

Now, informally speaking, for \tilde{v}_i to *not* outperform $\hat{\pi}_v(i)$ we require that from the point of view of an external observer, \tilde{v}_i , if instructed to vote for some candidate j by following the honest program, behaves like $\hat{\pi}_v(i)$. Recall that, the program π_o of an *external observer* may output 0 or 1 on some designated channel and may use the same communication interface as the coercer, except for connecting to dum, i.e., there is no direct communication between an external observer and the coerced voter. Now, formally we say that \tilde{v}_i *does not outperform* $\hat{\pi}_v(i)$, if for all programs π_o of the external observer we have that

$$\Pr[(\pi_o \parallel \pi_c^j \parallel \tilde{v}_i \parallel e)^{(\ell)} \mapsto 1] - \Pr[(\pi_o \parallel \hat{\pi}_v(i) \parallel e)^{(\ell)} \mapsto 1]$$

is negligible as a function in the security parameter 1^ℓ .

Now, we can state that, under the assumption that the counter-strategy does not outperform the honest voting program, we have that if a protocol is δ -coercion-resistant, then it also achieves δ -privacy. In the following theorem, we say that P^* is *δ -coercion-resistant w.r.t. γ_i and \tilde{v}_i* , if \tilde{v}_i can be used to show δ -coercion-resistant of P^* w.r.t. γ_i . For simplicity, we focus here on the case of single-voter coercion-resistance/privacy, but this result can easily be lifted to the multi-voter case.

Theorem 6. *Let P^* be δ -coercion-resistant w.r.t. γ_i and \tilde{v}_i , for every candidates i . Then, if \tilde{v}_i does not outperform $\hat{\pi}_v(i)$, for every candidate i , then P^* achieves δ -privacy.*

Proof sketch. Let $j, j' \in \{1, \dots, k\}$ and let π_o be a program of the observer, in the sense of Section 6.1. We define $\pi_c = (\pi_o \parallel \pi_c^j)$ to be a coercer program. Recall from Section 6.2 that the communication interface of a coercer is that of an observer, except that a coercer can also communicate with the coerced voter. Therefore, π_c indeed is a coercer program. Since P^* is δ -coercion-resistant w.r.t. $\gamma_{j'}$ and $\tilde{v}_{j'}$, we know that $\Pr[T \mapsto 1] - \Pr[\tilde{T} \mapsto 1]$ is δ -bounded for $T = (\pi_c \parallel \text{dum} \parallel e)^{(\ell)}$ and $\tilde{T} = (\pi_c \parallel \tilde{v}_{j'} \parallel e)^{(\ell)}$.

Now, by the definition of π_c^j , we know that $\Pr[T \mapsto 1] = \Pr[T' \mapsto 1]$, for $T' = (\pi_o \parallel \hat{\pi}_v(j) \parallel e)^{(\ell)}$. Moreover, because $\tilde{v}_{j'}$ does not outperform $\hat{\pi}_v(j')$, we have that $\Pr[\tilde{T} \mapsto 1] - \Pr[\tilde{T}' \mapsto 1]$ is negligible for $\tilde{T}' = (\pi_o \parallel \hat{\pi}_v(j') \parallel e)^{(\ell)}$. It follows that $\Pr[T' \mapsto 1] - \Pr[\tilde{T}' \mapsto 1]$ is δ -bounded, which proves that P^* achieves δ -privacy. \square

It turns out that for many voting protocol which have been analyzed with respect to coercion-resistance, the (optimal) counter-strategies indeed do not outperform the honest voting program of the respective protocol. In particular, it is not hard to check that at least the following protocols satisfy the condition in Theorem 6: the Bingo voting system [3] (see [16]), Scantegrity II [5] (see [19]), the JCJ voting protocol [13] and the Civitas voting system [7] (see also [21]). Also the “privacy enhanced” variants of ThreeBallot and VAV satisfy the condition (but clearly not the

other variants of ThreeBallot and VAV we considered, since, as mentioned, for these variants the level of coercion-resistance is higher— δ is smaller—than the level of privacy.)

For such protocols, once δ -coercion-resistance is proven, by Theorem 6 we obtain δ -privacy for free. In case the level of coercion-resistance corresponds to the ideal one—as, e.g., proven for the Bingo voting system and Scantegrity II in [16] and [19]—, by Theorems 2 and 3, the level of privacy is ideal as well. However, in general, the actual level of privacy might be better than what can be concluded from the theorem, with the “privacy enhanced” variants of ThreeBallot and VAV being examples.

We conclude this section with the postponed examples mentioned above.

Example (Improving Privacy Significantly Lowers the Level of Coercion-Resistance). We consider the following variant of ThreeBallot. An honest voter is supposed to submit either $(\underline{x}, \underline{x}, \circ)$ or $(\underline{x}, \circ, \circ)$ (according to her favorite candidate). This scheme is ideal in terms of privacy, because the bulletin board and the receipts do not leak any information apart from the pure result of the election. However, this scheme provides no coercion-resistance whatsoever: When the coerced voter is instructed to submit $(\circ, \underline{x}, \circ)$ (which is allowed but never done by honest voters), but wants to vote for candidate A, she would have to submit $(\circ, \underline{x}, \underline{x})$. But then, as all the honest voters submit $(\underline{x}, \underline{x}, \circ)$ or $(\underline{x}, \circ, \circ)$, the coercer could easily detect that he was cheated, by counting the number of ballots of type \circ on the bulletin board.

Example (Improving Coercion Resistance Significantly Lowers the Level of Privacy). We consider the following variant of ThreeBallot. In order to vote for candidate A, an honest voter is supposed to submit $(\underline{x}, \underline{x}, \circ)$ or $(\underline{x}, \circ, \circ)$, each with probability $\frac{1}{2}$. Analogously for a voter that wants to vote for candidate 2. By this, in 50% of the cases, an honest voter reveals her vote, namely in the case where she does not have \underline{x} as receipt. That means that this variant provides a very low level of privacy ($\delta \geq 0.5$), i.e. the observer can with quite high probability tell which candidate a given voter voted for. However, this variant is not that bad in terms of coercion-resistance, as here, the coerced voter is not bound to follow the honest strategy and can choose patterns in a more clever way. In fact, we can use here the same counter-strategy we used previously—take the receipt required by the coercer and adjust the remaining ballots to form a valid vote for the favorite candidate. Note that, following this strategy, the coerced voter may submit patterns which are valid but never chosen by the program of an honest voter. For this counter-strategy, although the coercer might learn approximatively half of the votes of the honest voters, the actual vote of the coerced voter is still hidden behind the votes of the honest voters that submitted $(\underline{x}, \underline{x}, \circ)$ or $(\underline{x}, \circ, \circ)$, i.e. that did not reveal their votes to the coercer. This results in a reasonably small δ .

7 Conclusion

In this paper, we presented new insights into central security properties, namely verifiability, privacy, and coercion-resistance. Our findings, in part, come from a case study, in which we precisely measure the level of verifiability, privacy, and coercion-resistance of different variants of ThreeBallot and VAV proposed in the literature.

For verifiability we have demonstrated that the combination of individual and universal verifiability is, unlike commonly believed, insufficient to provide overall/global verifiability. Our case study shows that the main problem with individual and universal verifiability is that these notions ignore that dishonest authorities/voters can break the integrity of ballots of honest voters

by ill-formed ballots. We therefore advocate the concept of *global verifiability*, as captured by the definition of verifiability in [17] and used in the present paper.

We also demonstrated that the relationship between privacy and coercion-resistance is more subtle than what can be gathered from the literature. Our case study highlighted interesting phenomena for existing protocols: i) improving privacy may degrade the level of coercion-resistance, ii) the level of coercion-resistance may be higher than the level of privacy. The latter is due to the fact that the counter-strategy a coerced voter uses maybe “smarter” in hiding information than the honest voting program. For the case that this is not so, we were able to prove that δ -coercion-resistance implies δ -privacy. As discussed in Section 6.5, for many protocols, the counter-strategy does indeed not outperform the honest voting program. We conjecture that if it does, then it should be possible to improve the honest voting program.

Besides these general findings on verifiability, privacy, and coercion-resistance, our case study also provides the first comprehensive picture on the security of prominent voting systems, Three-Ballot and VAV.

Acknowledgment. This work was carried out within the Priority Programme 1496 *Reliably Secure Software Systems – RS³* supported by *Deutsche Forschungsgemeinschaft* (DFG) under Grant KU 1434/6-1.

References

- [1] B. Adida and C.A. Neff. Ballot Casting Assurance. In *USENIX/ACCURATE Electronic Voting Technology (EVT 2006)*, 2006.
- [2] J. C. Benaloh and D. Tuinstra. Receipt-free secret-ballot elections (extended abstract). In *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing (STOC 1994)*, pages 544–553. ACM Press, 1994.
- [3] J.-M. Bohli, J. Müller-Quade, and S. Röhrich. Bingo Voting: Secure and Coercion-Free Voting Using a Trusted Random Number Generator. In A. Alkassar and M. Volkamer, editors, *E-Voting and Identity (VOTE-ID 2007)*, volume 4896 of *Lecture Notes in Computer Science*, pages 111–124. Springer, 2007.
- [4] D. Chaum. Elections with unconditionally-secret ballots and disruption equivalent to breaking rsa. In *Advances in Cryptology – Eurocrypt ’88*, volume 330 of *Lecture Notes in Computer Science*, pages 177–182. Springer, 1988.
- [5] D. Chaum, R. Carback, J. Clark, A. Essex, S. Popoveniuc, R. L. Rivest, P. Y. A. Ryan, E. Shen, and A. T. Sherman. Scantegrity II: End-to-End Verifiability for Optical Scan Election Systems using Invisible Ink Confirmation Codes. In *USENIX/ACCURATE Electronic Voting Technology (EVT 2008)*. USENIX Association, 2008. See also <http://www.scantegrity.org/elections.php>.
- [6] D. Chaum, P.Y.A. Ryan, and S. Schneider. A Practical, Voter-verifiable Election Scheme. In *Proceedings of the 10th European Symposium on Research in Computer Security (ESORICS 2005)*, volume 3679 of *Lecture Notes in Computer Science*, pages 118–139. Springer, 2005.

- [7] M. R. Clarkson, S. Chong, and A. C. Myers. Civitas: Toward a Secure Voting System. In *2008 IEEE Symposium on Security and Privacy (S&P 2008)*, pages 354–368. IEEE Computer Society, 2008.
- [8] O. de Marneffe, O. Pereira, and J.-J. Quisquater. Simulation-Based Analysis of E2E Voting Systems. In A. Alkassar and M. Volkamer, editors, *E-Voting and Identity (VOTE-ID 2007)*, volume 4896 of *Lecture Notes in Computer Science*, pages 137–149. Springer, 2007.
- [9] S. Delaune, S. Kremer, and M. D. Ryan. Verifying Privacy-type Properties of Electronic Voting Protocols. *Journal of Computer Security*, 17(4):435–487, 2009.
- [10] S. Delaune, S. Kremer, and M.D. Ryan. Coercion-Resistance and Receipt-Freeness in Electronic Voting. In *Proceedings of the 19th IEEE Computer Security Foundations Workshop (CSFW'06)*, pages 28–39. IEEE Computer Society Press, 2006.
- [11] Kevin Henry, Douglas R. Stinson, and Jiayuan Sui. The Effectiveness of Receipt-based Attacks on ThreeBallot. *IEEE Transactions on Information Forensics and Security*, 4(4):699–707, 2009.
- [12] M. Hirt and K. Sako. Efficient receipt-free voting based on homomorphic encryption. In B. Preneel, editor, *Advances in Cryptology – EUROCRYPT 2000*, volume 1807 of *Lecture Notes in Computer Science*, pages 539 – 556. Springer, 2000.
- [13] A. Juels, D. Catalano, and M. Jakobsson. Coercion-resistant Electronic Elections. In *Proceedings of Workshop on Privacy in the Electronic Society (WPES 2005)*, pages 61–70. ACM Press, 2005.
- [14] Steve Kremer, Mark Ryan, and Ben Smyth. Election Verifiability in Electronic Voting Protocols. In Dimitris Gritzalis, Bart Preneel, and Marianthi Theoharidou, editors, *15th European Symposium on Research in Computer Security (ESORICS2010)*, volume 6345 of *Lecture Notes in Computer Science*, pages 389–404. Springer, 2010.
- [15] R. Küsters. Simulation-Based Security with Inexhaustible Interactive Turing Machines. In *Proceedings of the 19th IEEE Computer Security Foundations Workshop (CSFW-19 2006)*, pages 309–320. IEEE Computer Society, 2006. See <http://eprint.iacr.org/2013/025/> for a full and revised version.
- [16] Ralf Küsters, Tomasz Truderung, and Andreas Vogt. A Game-based Definition of Coercion-Resistance and its Applications. In *23th IEEE Computer Security Foundations Symposium, CSF 2010*, pages 122–136. IEEE Computer Society, 2010.
- [17] Ralf Küsters, Tomasz Truderung, and Andreas Vogt. Accountability: Definition and Relationship to Verifiability. In *Proceedings of the 17th ACM Conference on Computer and Communications Security (CCS 2010)*, pages 526–535. ACM, 2010.
- [18] Ralf Küsters, Tomasz Truderung, and Andreas Vogt. Accountabiliy: Definition and Relationship to Verifiability. Technical Report 2010/236, Cryptology ePrint Archive, 2010. <http://eprint.iacr.org/>.

- [19] Ralf Küsters, Tomasz Truderung, and Andreas Vogt. Proving Coercion-Resistance of Scantegrity II. In Miguel Soriano, Sihang Qing, and Javier López, editors, *Proceedings of the 12th International Conference on Information and Communications Security (ICICS 2010)*, volume 6476 of *Lecture Notes in Computer Science*, pages 281–295. Springer, 2010.
- [20] Ralf Küsters, Tomasz Truderung, and Andreas Vogt. Verifiability, Privacy, and Coercion-Resistance: New Insights from a Case Study. In *32nd IEEE Symposium on Security and Privacy (S&P 2011)*, pages 538–553. IEEE Computer Society, 2011.
- [21] Ralf Küsters and Tomasz Truderung. An Epistemic Approach to Coercion-Resistance for Electronic Voting Protocols. In *2009 IEEE Symposium on Security and Privacy (S&P 2009)*, pages 251–266. IEEE Computer Society, 2009.
- [22] T. Moran and M. Naor. Receipt-Free Universally-Verifiable Voting With Everlasting Privacy. In C. Dwork, editor, *Advances in Cryptology - CRYPTO 2006, 26th Annual International Cryptology Conference, Proceedings*, volume 4117 of *Lecture Notes in Computer Science*, pages 373–392. Springer, 2006.
- [23] T. Moran and M. Naor. Split-ballot voting: everlasting privacy with distributed trust. In P. Ning, S. De Capitani di Vimercati, and P. F. Syverson, editors, *Proceedings of the 2007 ACM Conference on Computer and Communications Security, CCS 2007*, pages 246–255. ACM, 2007.
- [24] T. Okamoto. Receipt-Free Electronic Voting Schemes for Large Scale Elections. In B. Christianson, B. Crispo, T. M. A. Lomas, and M. Roe, editors, *Proceedings of the 5th International Workshop on Security Protocols*, volume 1361 of *Lecture Notes in Computer Science*, pages 25–35. Springer, 1997.
- [25] B. Riva and A. Ta-Shma. Bare-Handed Electronic Voting with Pre-processing. In *USENIX/ACCURATE Electronic Voting Technology (EVT 2007)*, 2007.
- [26] R. L. Rivest and W. D. Smith. Three Voting Protocols: ThreeBallot, VAV and Twin. In *USENIX/ACCURATE Electronic Voting Technology (EVT 2007)*, 2007.
- [27] Charlie E. M. Strauss. A critical review of the triple ballot voting system, part 2: Cracking the triple ballot encryption. <http://www.cs.princeton.edu/~appel/voting/Strauss-ThreeBallotCritique2v1.5.pdf>, October 8, 2006. Draft V1.5.

A Proof Sketch of Theorem 1

In the proof we will use the following fact:

Lemma 1. *Let $[b_1, b_2, b_3]$ be the simple ballots submitted by an honest voter in a protocol run (where the order of these simple ballots matters). Then, for $i = 1, 2, 3$, the probability that b_i is taken as the receipt is $\frac{1}{3}$ in the system S_{TB}^o and bigger than or equal to $\frac{1}{6}$ in the system S_{TB}^p .*

Moreover, if $[b_1, b_2, b_3]$ is a vote for candidate j (for $j = 1, \dots, k$) and b_l (for $l = 1, 2, 3$) is a ballot where the j -th position is marked, then the probability that b_l is taken as a receipt in the system S_{TB}^p is $[b_1, b_2, b_3]$ is $\frac{1}{6}$.

Proof. As the receipts are chosen uniformly in S_{TB}^o , we immediately get the result for this system. Let us now consider the system S_{TB}^p . Without loss of generality, we can assume that $[b_1, b_2, b_3]$ is a vote for candidate 1 and that the first position is marked on b_1 and b_2 (but not on b_3). What we need to compute is, in fact, the conditional probability that an honest voter takes b_i as her receipt, given she submits $[b_1, b_2, b_3]$.

We start with computing the probability that an honest voter submits $[b_1, b_2, b_3]$. By the definition of the voting procedure of the system S_{TB}^p , an honest voter can obtain (and submit) $[b_1, b_2, b_3]$ in two ways: right before the voter has chosen her receipt, she must have either $[b'_1, b_2, b_3]$ or $[b_1, b'_2, b_3]$, where b'_1 and b'_2 are as b_1 and b_2 with the difference that the first position is not marked. From $[b'_1, b_2, b_3]$ the voter obtains $[b_1, b_2, b_3]$, if she chooses b_3 as the receipt (which happens with probability $\frac{1}{3}$) and then marks candidate 1 on the first ballot (which happens with probability p_1) or if she chooses b_2 as the receipt (which happens with probability $\frac{1}{3}$) and then marks the first position of b_1 (which happens with probability $\frac{1}{2} \cdot p_1$). Now, the probability for having $[b'_1, b_2, b_3]$ before the receipt is picked is $p = \frac{1}{3k}$. In the same way we can handle the case when $[b_1, b_2, b_3]$ is obtained from $[b_1, b'_2, b_3]$. From this, we can compute the probability that the honest voter submits $[b_1, b_2, b_3]$ as $2 \cdot p \cdot p_1 \cdot (\frac{1}{3} + \frac{1}{6}) = p \cdot p_1$.

Now, the probability that an honest voter submits $[b_1, b_2, b_3]$ and takes b_1 as the receipt is $p_1 \cdot p \cdot \frac{1}{6}$ (this is only possible if she starts with $[b_1, b'_2, b_3]$, picks b_1 as the receipt and then marks candidate 1 on the second ballot). From this, we conclude that the conditional probability that an honest voter takes b_1 as the receipt, given she submits $[b_1, b_2, b_3]$, is $\frac{1}{6}$. In a very similar way, we can show that the probability that an honest voter takes b_2 as the receipt, given she submits $[b_1, b_2, b_3]$, is $\frac{1}{6}$, too. Finally, from the above we easily obtain that the probability that an honest voter takes b_3 as the receipt, given she submits $[b_1, b_2, b_3]$, is $\frac{2}{3}$. \square

Now we prove Theorem 1. For the first condition of Definition 1, we have to show that whenever, in a run of the system, the machine and the bulletin board are honest, *ver* accepts this run with overwhelming probability. This is easy to see: If the machine and the bulletin board are honest in a run of the system, only well-formed ballots are sent to the bulletin board, (with overwhelming probability) no serial numbers occur twice, and the bulletin board correctly displays the ballots received from the machine. Now, by the definition of the verifier, it follows that *ver* accepts such a run.

For the second condition of Definition 1, we have to show that the probability that the system produces a run which is accepted by *ver*, even though the goal is violated, is bounded by δ_{ver}^p or δ_{ver}^o , respectively. In such a run, since it is accepted by *ver*, the bulletin board must be consistent. Furthermore, because the goal is violated, there must exist a candidate, say candidate i , such that the sum of all votes of honest voters for all candidates except i is at least $\ell + 1$.

As we have already shown, the machine can change m votes without any risk of this being detected. Therefore, in order to violate the goal γ_ℓ , it remains to change $k' = \ell + 1 - \min(\ell + 1, m)$ votes of honest voters that did not vote for candidate i . Let us consider the following procedure a dishonest party (the voting machine or the bulletin board) can use to do this: If a voter votes for candidate $j \neq i$, candidate j is marked on two simple ballots, whereas candidate i is marked only on one ballot. The dishonest party chooses one simple ballot of such a voter with marked j -th position, but not i -th position, and replaces it by a similar ballot, but with the markings on the i -th and j -th position swapped. By Lemma 1 we know that every time this is done, the probability that this is detected by an honest voter is $\frac{1}{6} \cdot p_{\text{check}}$ in the system S_{TB}^p and $\frac{1}{3} \cdot p_{\text{check}}$ in S_{TB}^o . As it must be done k' times and, as mentioned, there must exist $\ell + 1$ voters who voted not for candidate i ,

we conclude that, when this strategy is used, the probability that the goal γ_ℓ is violated and the observer accepts the run is δ_{Ver}^p and δ_{Ver}^o , respectively. This shows that these bounds are optimal (i.e. the considered goal is not δ' -verifiable for smaller δ').

To prove that the goal γ_ℓ is δ_{Ver}^x -verifiable, we need to show that the strategy described above is optimal, i.e. for any other strategy of the dishonest party, the probability that the run is accepted even though the goal is not satisfied is smaller. If a strategy does not change m votes using m dishonest voters, as described earlier, it clearly is not optimal (recall that doing this does not involve any risk of being detected). If it does it, there are still k' votes, as defined above, to be changed. So, let us consider any way k' votes are changed. First, let us observe that to change k' votes the dishonest party must change at least k' simple ballots of honest voters. If exactly one simple ballot belonging to some honest voter is changed, the probability that this goes undetected is, by Lemma 1, at most $(1 - \frac{1}{6} \cdot p_{check})$. Moreover, it is always worse to change $t > 1$ simple ballots of one honest voter, than to change t ballots of t separate voters: If all three simple ballots of an honest voter are changed, the probability that this is not detected is 0. In case two ballots of an honest voter are changed, this probability is at least $1 - \frac{2}{6} \cdot p_{check}$, which is, in fact, less than $(1 - \frac{1}{6} \cdot p_{check})^2$ (which expresses the probability that the cheating is detected if two simple ballots of two different honest voters are changed). Altogether, we can conclude that probability that k votes are changed and the run is accepted is at most δ_{Ver}^p and δ_{Ver}^o , respectively.

B Verifiability of VAV

Let P_{VAV}^s and P_{VAV}^p denote the VAV protocol in the simple variant and the privacy enhanced variant, respectively. Based on analogous assumptions as those for ThreeBallot (see Section 5.2), it is straightforward to formally define the protocol instantiations $S_{VAV}^s = P_{VAV}^s(\{ver\}, q, V_H, k, \vec{p})$ of P_{VAV}^s and $S_{VAV}^p = P_{TB}^p(\{ver\}, q, V_H, k, \vec{p})$ of P_{VAV}^p .

To state the following theorem, we need to introduce the following notation. For a given run of the protocol, we denote by A the set of those candidates j for which the sum of all votes of honest voters for all candidates except j is at least $\ell + 1$. Intuitively, if a candidate does not belong to A , then it is impossible for a dishonest party to violate the considered goal by increasing the number of votes for this candidate, because there is not enough votes for other candidates that can be changed.

Next, by X_j we denote the number of multi-ballots in a run submitted by honest voters for which the following holds:

- (i) the multi-ballot forms a vote for a candidate different from j , and
- (ii) on the multi-ballot, not the same candidate is marked on all three simple ballots.

These multi-ballots can safely be changed to votes for j , as explained in Section 4.

Finally, we define $p_r = \Pr[A \neq \emptyset \text{ and } \max_{j \in A} X_j = r]$, where the probability is over runs of the protocol. (Note that p_r only depends on choices made by honest voters.) As one can see, p_r is the probability that it is possible in principle to violate the goal ($A \neq \emptyset$) and that r is the candidate such that the number of votes that we can safely change (as explained in Section 4) to votes for r is maximal.

Theorem 7. *Let B be the set containing the voting machine and the bulletin board and $x \in \{s, p\}$. The goal γ_ℓ is guaranteed in S_{VAV}^x by B and δ_{VerVAV}^x -verifiable by a , where $\delta_{VerVAV}^s = \delta_{Ver}^o$ with δ_{Ver}^o*

as in Theorem 1 and

$$\delta_{VerVAV}^p = \sum_{r=0}^n p_r \left(1 - \frac{1}{4} \cdot p_{check}\right)^{\max(\ell+1-r-m, 0)},$$

where m is the number of dishonest voters and p_r is defined as above. Moreover, δ_{VerVAV}^x is optimal, i.e. there is not $\delta' < \delta_{VerVAV}^x$ for which the goal is guaranteed in S_{VAV}^x by B and δ_{VerVAV}^x -verifiable by a .

Proof sketch. The statement for S_{VAV}^s follows as in the proof of Theorem 1.

The intuition behind the statement for S_{VAV}^p is the following. The best strategy of the adversary for violating the goal in a given run is as follows: First, he determines those candidates j such that the number of submitted multi-ballots not for j is bigger than ℓ . (Those candidates form the set A introduced above). Second, among those candidates, he determines a candidate j for which the number X_j is maximal. The probability of this number being r is p_r . Now, if $X_j = r$, then the adversary can safely change r votes. If there are some further votes to be changed (i.e. if $r < \ell + 1$), then the adversary can use dishonest voters to change additional m votes, as described in Section 5.2. Only if there are still some votes to be changed (i.e. $r + m < \ell + 1$), the adversary changes further $q = \ell + 1 - r - m$ ballots of q different voters, which is detected with probability $\frac{1}{4} \cdot p_{check}$ for each ballot (see Lemma 2 below). Hence, the probability that the adversary goes undetected when changing $\ell + 1 - r - m$ ballots is $(1 - \frac{1}{4} \cdot p_{check})^{\ell+1-r-m}$. Now, considering all possible values of r and the corresponding probabilities p_r , we obtain the overall probability that ℓ changes goes undetected, as in the theorem.

The theorem follows from the observation that the strategy described above is the best strategy of the adversary. It can be shown in the similar way as in the case of Theorem 1. Here, we use the following lemma. \square

Lemma 2. *Let $[b_1, b_2, b_3]$ be the simple ballots submitted by an honest voter in a protocol run of the system S_{VAV}^p (where the order of these simple ballots matters) such that some candidate j is marked on all b_1, b_2 and b_3 (recall that only those ballots cannot be changed safely in S_{VAV}^p as described earlier). Then, for $i = 1, 2, 3$, the probability that b_i is taken as the receipt is $\frac{1}{4}$ if b_i is a V-ballot and $\frac{1}{2}$ if b_i is an A-ballot.*

Proof. Let $[b_1, b_2, b_3]$ be as in the lemma. Without loss of generality, we can assume that b_1 is an A-ballot and b_2 and b_3 are V-ballots. It follows directly from the definition of the voting procedure of S_{VAV}^p that b_1 is picked as the receipt with probability $\frac{1}{2}$. Moreover, b_1 and b_2 are picked as the receipt with the same probability which must be then $\frac{1}{4}$. \square

C Proof of Theorem 2

Our goal is to prove that $S = P_{ideal}(\{VA\}, q, n, k, \vec{p})$ achieves δ -privacy where $\delta = \delta_{Priv}(n, k, \vec{p})$. We begin with some auxiliary definitions and facts.

For a protocol instantiation P^* , we denote by $\Omega_1 = \{0, \dots, k\}^n$ the set of all possible combinations of choices made by the honest voters, with the corresponding probability distribution μ_1 derived from $\vec{p} = p_0, p_1, \dots, p_k$. All other random bits used by ITMs in an instance of P^* , i.e., all other random bits used by honest voters as well as all random bits used by honest authorities, the coercer, and the coerced voter, are uniformly distributed. We take μ_2 to be this

distribution over the space Ω_2 of random bits. Formally, this distribution depends on the security parameter. We can, however, safely ignore it in the notation without causing confusion. We define $\Omega = \Omega_1 \times \Omega_2$ and $\mu = \mu_1 \times \mu_2$, i.e., μ is the product distribution obtained from μ_1 and μ_2 . For an event φ , we will write $\Pr_{\omega_1, \omega_2 \leftarrow \Omega}[\varphi]$, $\Pr_{\omega_1, \omega_2}[\varphi]$, or simply $\Pr[\varphi]$ to denote the probability $\mu(\{(\omega_1, \omega_2) \in \Omega : \varphi(\omega_1, \omega_2)\})$. Similarly, $\Pr_{\omega_1 \leftarrow \Omega_1}[\varphi]$ or simply $\Pr_{\omega_1}[\varphi]$ will stand for $\mu_1(\{\omega_1 \in \Omega_1 : \varphi(\omega_1)\})$; analogously for $\Pr_{\omega_2 \leftarrow \Omega_2}[\varphi]$.

Let

$$\Delta_{ij} = \sum_{\vec{r} \in M_{i,j}^*} (A_{\vec{r}}^j - A_{\vec{r}}^i). \quad (3)$$

So, we have $\delta_{priv}(n, k, \vec{p}) = \max_{i,j \in \{1, \dots, k\}} \Delta_{ij}$. By $\text{res}(\omega_1, i)$, where $\omega_1 \in \Omega_1$ and $i \in \{1, \dots, k\}$, we denote the pure result of the election (i.e. an element of Res) obtained when the honest voters vote according to ω_1 and the voter under observation votes for i . Therefore, we have

$$A_{\vec{r}}^i = \Pr_{\omega_1}[\text{res}(\omega_1, i) = \vec{r}].$$

By definition of $M_{i,j}^*$, it is easy to see that for every $i, j \in \{1, \dots, k\}$ and every set $M \subseteq \text{Res}$ of pure results, the following inequality holds:

$$\sum_{\vec{r} \in M} (A_{\vec{r}}^j - A_{\vec{r}}^i) \leq \sum_{\vec{r} \in M_{i,j}^*} (A_{\vec{r}}^j - A_{\vec{r}}^i) = \Delta_{ij}. \quad (4)$$

Let $\pi_o \in \mathcal{O}$ be an arbitrary observer process and M be the set of views accepted by π_o . Note that the view of the observer in a run of the system consists only of his random coins $\omega_2 \in \Omega_2$ and the result of the election. Therefore, each element of M can be represented as (ω_2, \vec{r}) , where $\omega_2 \in \Omega_2$ and \vec{r} is a pure result (note that, given the program c , the complete result is determined by \vec{r} and ω_2). For $\omega_2 \in \Omega_2$, we define M^{ω_2} to be $\{\vec{r} \in \text{Res} : (\omega_2, r) \in M\}$. With this, we obtain:

$$\begin{aligned} & \Pr[(\pi_o \parallel \hat{\pi}_v(j) \parallel e)^{(\ell)} \mapsto 1] - \Pr[(\pi_o \parallel \hat{\pi}_v(i) \parallel e)^{(\ell)} \mapsto 1] = \\ &= \Pr[(\omega_2, \text{res}(\omega_1, j)) \in M] - \Pr[(\omega_2, \text{res}(\omega_1, i)) \in M] \\ &= \sum_{\omega_2 \in \Omega_2} \left(\mu_2(\omega_2) \cdot \Pr_{\omega_1}[\text{res}(\omega_1, j) \in M^{\omega_2}] - \mu_2(\omega_2) \cdot \Pr_{\omega_1}[\text{res}(\omega_1, i) \in M^{\omega_2}] \right) \\ &= \sum_{\omega_2 \in \Omega_2} \mu_2(\omega_2) \sum_{\vec{r} \in M^{\omega_2}} \left(\Pr_{\omega_1}[\text{res}(\omega_1, j) = \vec{r}] - \Pr_{\omega_1}[\text{res}(\omega_1, i) = \vec{r}] \right) \\ &= \sum_{\omega_2 \in \Omega_2} \mu_2(\omega_2) \cdot \sum_{\vec{r} \in M^{\omega_2}} (A_{\vec{r}}^j - A_{\vec{r}}^i) \\ &\stackrel{\text{(by (4))}}{\leq} \sum_{\omega_2 \in \Omega_2} \mu_2(\omega_2) \cdot \sum_{\vec{r} \in M_{i,j}^*} (A_{\vec{r}}^j - A_{\vec{r}}^i) \\ &\stackrel{\text{(by (3))}}{=} \sum_{\omega_2 \in \Omega_2} \mu_2(\omega_2) \cdot \Delta_{i,j} \\ &\leq \sum_{\omega_2 \in \Omega_2} \mu_2(\omega_2) \cdot \max_{i,j \in \{1, \dots, k\}} \Delta_{i,j} \\ &= \max_{j \in \{1, \dots, k\}} \Delta_{i,j} = \delta_{priv}^i(n, k, \vec{p}). \end{aligned}$$

This implies that $\Pr[(c \parallel \hat{\pi}_v(j) \parallel e)^{(\ell)} \mapsto 1] - \Pr[(c \parallel \hat{\pi}_v(i) \parallel e)^{(\ell)} \mapsto 1]$ is δ -bounded, for $\delta = \delta_{\text{priv}}(n, k, \vec{p})$.

It remains to show that δ is optimal. We observe that in the above inequalities we obtain equality for an observer program $\pi_o = \pi_o^*$ which, for i, j , accepts a run only if the pure result \vec{r} in his view belongs to $M_{i,j}^*$ and for i, j such that $\Delta_{i,j} = \max_{l,l' \in \{1, \dots, k\}} \Delta_{l,l'}$. Hence, for $\pi_o = \pi_o^*$

$$\Pr[(\pi_o \parallel \hat{\pi}_v(j) \parallel e)^{(\ell)} \mapsto 1] - \Pr[(\pi_o \parallel \hat{\pi}_v(i) \parallel e)^{(\ell)} \mapsto 1] = \delta$$

and, therefore S does not achieve δ' -privacy, for any $\delta' < \delta$.

D Proof of Theorem 4

In this section, we prove Theorem 4. We only deal with the case where the observer can see the receipt of the voter under observation. The other case is analogous (and simpler). Let $S \in \{P_{\text{TB}}^{\rho+}, P_{\text{TB}}^{\rho+}, P_{\text{VAV}}^{\rho+}, P_{\text{VAV}}^{\rho+}\}$

We first introduce some notation. We will assume that an element $\omega_1 \in \Omega_1 = \Omega_1^S$ now determines the choices of the honest voters plus the random coins the honest voters *and* the voter under observation use to determine, together with their chosen candidates, the exact pattern they submit (see Section 3 and Section 4 for the definition of a pattern). We define the following random variables on Ω_1 .

Let $\text{rec}(\omega_1) = \text{rec}^S(\omega_1)$ be the vector of numbers of receipts of honest voters add: and the voter under observation of the different types, i.e., \times , \times , \circ , \times , or \circ in P_{TB} and in P_{VAV} , for every candidate i , an A- and a V-ballot where candidate i is marked. Let $\text{Rec}(\omega_1) = \text{Rec}^S(\omega_1)$ be the vector (r_0, r_1, \dots, r_n) , where r_0 is the receipt of the voter under observation and r_i ($1 \leq i \leq n$) is the receipt of the i -th honest voter (without a serial number).

We can represent an element ω_2 of the space of random bits Ω_2 used in a run of a system, in addition to the random choices ω_1 , as a tuple $\omega_2 = (\alpha, \vec{r}, \pi)$, where α is a sequence of random coins of the observer, $\vec{r} = (r_{ij})_{i \in \{0, \dots, m\}, j \in \{1, 2, 3\}}$, where r_{ij} is the serial number printed by the voting machine on the j -th ballot cast by the i -th voter (where the 0-th voter is the observed voter), and π is a permutation applied to the set of ballots before publishing. (Note that serial number r_{ij} , for $j \in \{1, 2, 3\}$, are not printed, if the i -th voter does not vote.)

We will use the letter η to range over views of the observer. Further, ρ is used to denote an essential view of the observer, as defined in Section 6.3.

Let i, j be arbitrary candidates. By $\rho(\eta)$ we will denote the essential view determined by η . By $\rho(i, \omega_1)$, for a candidate i and $\omega_1 \in \Omega_1$, we denote the essential view obtained when the voter under observation casts simple ballots for candidate i in a way determined by the randomness in ω_1 and the honest voters casts ballots determined by ω_1 . By $\text{Rec}(\eta)$ we denote the receipts *without serial numbers* that the observed voter and the honest voters give to the observer in run η .

For an observer view η , let φ_η be a predicate over Ω_1 such that $\varphi_\eta(\omega_1)$ iff $\rho(j, \omega_1) = \rho(\eta)$ and $\text{Rec}(\omega_1) = \text{Rec}(\eta)$. Analogously, we define a predicate $\bar{\varphi}_\eta(\omega_1)$ which holds iff $\rho(i, \omega_1) = \rho(\eta)$ and $\text{Rec}(\omega_1) = \text{Rec}(\eta)$.

Let c be a program of the observer and let $T = (c \parallel \hat{\pi}_v(j) \parallel e_S)$ and $\tilde{T} = (c \parallel \hat{\pi}_v(i) \parallel e_S)$. Note that the behavior of the systems T and \tilde{T} is completely determined by $(\omega_1, \omega_2) \in \Omega_1 \times \Omega_2$, as the exact pattern of the voter under observation submits in $\hat{\pi}_v(j)$ and $\hat{\pi}_v(i)$ is determined by ω_1 .

We now show that the view of the observer is information-theoretically independent of the choices of the honest voters as long as these choices are consistent with the essential view and the order of the receipts. This is formulated in Lemma 4, with the core stated in the following lemma.

Lemma 3. *Let η be a view of the observer and $\omega_1, \omega'_1, \omega''_1, \omega'''_1$ be arbitrary elements of Ω_1 with $\varphi_\eta(\omega_1), \varphi_\eta(\omega'_1), \tilde{\varphi}_\eta(\omega''_1)$ and $\tilde{\varphi}_\eta(\omega'''_1)$. Then the sets*

$$\begin{aligned} A &= \{\omega_2 : T(\omega_1, \omega_2) \mapsto \eta\}, & B &= \{\omega_2 : T(\omega'_1, \omega_2) \mapsto \eta\}, \\ C &= \{\omega_2 : \tilde{T}(\omega''_1, \omega_2) \mapsto \eta\}, & D &= \{\omega_2 : \tilde{T}(\omega'''_1, \omega_2) \mapsto \eta\} \end{aligned}$$

have the same cardinality, and hence, have the same probability.

Proof. We will show how to construct a bijection $h : A \rightarrow B$. The proof for the remaining cases are very similar.

Let $I = \{0, \dots, n\} \times \{1, 2, 3\}$. For $(m, l) \in I$, by $b_{m,l}$ and $b'_{m,l}$ we denote the marking on the l -th ballot cast by the m -th voter according to ω_1 and ω'_1 , respectively. Because $\varphi_\eta(\omega_1)$ and $\varphi_\eta(\omega'_1)$, we know that there exists a permutation $\sigma : I \rightarrow I$ such that $b'_{(m,l)} = b_{\sigma(m,l)}$. Moreover, we can assume that σ preserves receipts of voters, i.e. if the m -th voter picks the l -th ballot as her receipt according to ω_1 and she picks the l' -th ballot as a receipt according to ω'_1 , then $\sigma(m, l') = (m, l)$. Note that, in this case, $b_{(m,l)} = b'_{(m,l')}$.

Let $(\alpha, \vec{r}, \pi) \in \omega_2$. We define $h(\alpha, \vec{r}, \pi) = (\alpha, \vec{r}', \pi')$, where $r'_{(m,p)} = r_{\sigma(m,p)}$ and $\pi'(m, l) = \pi(\sigma(m, l))$ (recall that π determines the position $\pi(m, p)$ of the ballot $b_{(m,p)}$ on the bulletin board). It is easy to check that h is a bijection from A to B . \square

Using the lemma above, we can obtain the following fact.

Lemma 4. *Let η be an observer view. Let ω_1^η and $\tilde{\omega}_1^\eta$ be some fixed elements of Ω_1 such that $\varphi_\eta(\omega_1^\eta)$ and $\tilde{\varphi}_\eta(\tilde{\omega}_1^\eta)$, respectively. Then, the following equations hold true:*

$$\Pr[T \mapsto \eta] = \Pr_{\omega_1}[\varphi_\eta(\omega_1)] \cdot \Pr_{\omega_2}[T(\omega_1^\eta, \omega_2) \mapsto \eta] \quad (5)$$

$$\Pr[\tilde{T} \mapsto \eta] = \Pr_{\omega_1}[\tilde{\varphi}_\eta(\omega_1)] \cdot \Pr_{\omega_2}[\tilde{T}(\tilde{\omega}_1^\eta, \omega_2) \mapsto \eta] \quad (6)$$

$$\Pr_{\omega_2}[T(\omega_1^\eta, \omega_2) \mapsto \eta] = \Pr_{\omega_2}[\tilde{T}(\tilde{\omega}_1^\eta, \omega_2) \mapsto \eta] . \quad (7)$$

Proof. Using Lemma 3 we obtain:

$$\begin{aligned} \Pr[T \mapsto \eta] &= \Pr[\varphi_\eta(\omega_1), T(\omega_1, \omega_2) \mapsto \eta] \\ &= \sum_{\omega'_1: \varphi_\eta(\omega'_1)} \Pr[\omega_1 = \omega'_1, T(\omega'_1, \omega_2) \mapsto \eta] \\ &= \sum_{\omega'_1: \varphi_\eta(\omega_1)} \Pr_{\omega_1}[\omega_1 = \omega'_1] \cdot \Pr_{\omega_2}[T(\omega'_1, \omega_2) \mapsto \eta] \\ &= \sum_{\omega'_1: \varphi_\eta(\omega_1)} \Pr_{\omega_1}[\omega_1 = \omega'_1] \cdot \Pr_{\omega_2}[T(\omega_1^\eta, \omega_2) \mapsto \eta] \\ &= \Pr_{\omega_1}[\varphi_\eta(\omega_1)] \cdot \Pr_{\omega_2}[T(\omega_1^\eta, \omega_2) \mapsto \eta] . \end{aligned}$$

This proves (5). One can prove (6) in an analogous way. The equation (7) follows directly from Lemma 3. \square

Now, using Lemma 4, we can link the level of privacy ThreeBallot/VAV provides with the optimal bound δ_{priv}^S stated in Section 6.3. Clearly we have:

$$\Pr_{\omega_1}[\varphi_\eta(\omega_1)] = A_{\rho(\eta),j}^S \cdot \Pr_{\omega_1}[Rec(\omega_1) = Rec(\eta) \mid \rho(j, \omega_1) = \rho(\eta)]$$

and

$$\Pr_{\omega_1}[\tilde{\varphi}_\eta(\omega_1)] = A_{\rho(\eta),i}^S \cdot \Pr_{\omega_1}[Rec(\omega_1) = Rec(\eta) \mid \rho(i, \omega_1) = \rho(\eta)].$$

Furthermore, it is easy to show that given two essential views with the same number of receipts of every type (and otherwise possibly different information on the bulletin board), the probability of obtaining a specific vector of receipts (which links receipts and voters) stays the same. From this it follows:

$$\begin{aligned} \Pr_{\omega_1}[Rec(\omega_1) = Rec(\eta) \mid \rho(j, \omega_1) = \rho(\eta)] &= \\ &= \Pr_{\omega_1}[Rec(\omega_1) = Rec(\eta) \mid \rho(i, \omega_1) = \rho(\eta)] \\ &= \Pr_{\omega_1}[Rec(\omega_1) = Rec(\eta) \mid rec(\omega_1) = rec(\eta)]. \end{aligned}$$

Together with Lemma 4, we immediately obtain for all ω_1^η with $\varphi_\eta(\omega_1^\eta)$

$$\begin{aligned} \Pr[T \mapsto \eta] - \Pr[\tilde{T} \mapsto \eta] &= (A_{\rho(\eta),j}^S - A_{\rho(\eta),i}^S) \cdot \\ &\cdot \Pr_{\omega_2}[T(\omega_1^\eta, \omega_2) \mapsto \eta] \cdot \\ &\cdot \Pr_{\omega_1}[Rec(\omega_1) = Rec(\eta) \mid rec(\omega_1) = rec(\eta)]. \end{aligned}$$

Note that if there does not exist $\tilde{\omega}_1^\eta$ such that $\tilde{\varphi}_\eta(\tilde{\omega}_1^\eta)$, then $A_{\rho(\eta),i}^S = 0$ and $\Pr[\tilde{T} \mapsto \eta] = 0$.

Now, we are ready to prove that the system S , as defined in Theorem 4 in the case the observer can see the receipts of honest voters, achieves δ -privacy for $\delta = \delta_{priv}^S$.

Let M be the set of views for which the observer program c outputs 1 (i.e., informally, where the observer thinks that the voter under observation voted for j). In what follows, let ρ range over all essential views, η over all views, and Rec over all possible vectors of receipts. Finally, let $M^{\rho, Rec} = \{\eta \in M : Rec(\eta) = Rec, \text{ and } \rho(j, \eta) = \rho\}$. For ρ, Rec with $M^{\rho, Rec} \neq \emptyset$ let $\omega_1^{\rho, Rec}$ be an arbitrary element of Ω_1 with $\rho(j, \omega_1^{\rho, Rec}) = \rho$ and $Rec(\omega_1^{\rho, Rec}) = Rec$. Then we have $\varphi_\eta(\omega_1^{\rho, Rec})$ for all $\eta \in M^{\rho, Rec}$. We have:

$$\begin{aligned} \Phi &= \Pr[T \mapsto 1] - \Pr[\tilde{T} \mapsto 1] \\ &= \Pr[T \mapsto M] - \Pr[\tilde{T} \mapsto M] \\ &= \sum_{\rho} \sum_{Rec} \sum_{\eta \in M^{\rho, Rec}} (\Pr[T \mapsto \eta] - \Pr[\tilde{T} \mapsto \eta]) \\ &= \sum_{\rho} (A_{\rho,j}^S - A_{\rho,i}^S) \sum_{Rec} \sum_{\eta \in M^{\rho, Rec}} \Pr_{\omega_2}[T(\omega_1^{\rho, Rec}, \omega_2) \mapsto \eta] \\ &\quad \cdot \Pr_{\omega_1}[Rec(\omega_1) = Rec \mid rec(\omega_1) = rec(\rho)]. \end{aligned}$$

We get:

$$\begin{aligned}
\Phi &\leq \sum_{\rho \in M_{i,j}^S} (A_{\rho,j}^S - A_{\rho,i}^S) \sum_{Rec} \sum_{\eta \in M^{\rho, Rec}} \Pr_{\omega_2} [T(\omega_1^{\rho, Rec}, \omega_2) \mapsto \eta] \\
&\quad \cdot \Pr_{\omega_1} [Rec(\omega_1) = Rec \mid rec(\omega_1) = rec(\rho)] \\
&\leq \sum_{\rho \in M_{i,j}^S} (A_{\rho,j}^S - A_{\rho,i}^S) \sum_{Rec} \Pr_{\omega_1} [Rec(\omega_1) = Rec \mid rec(\omega_1) = rec(\rho)] \\
&\quad \cdot \sum_{\eta \in M^{\rho, Rec}} \Pr_{\omega_2} [T(\omega_1^{\rho, Rec}, \omega_2) \mapsto \eta] \\
&\leq \sum_{\rho \in M_{i,j}^S} (A_{\rho,j}^S - A_{\rho,i}^S) \sum_{Rec} \Pr_{\omega_1} [Rec(\omega_1) = Rec \mid rec(\omega_1) = rec(\rho)] \\
&\leq \sum_{\rho \in M_{i,j}^S} (A_{\rho,j}^S - A_{\rho,i}^S) \leq \delta_{Priv}^S.
\end{aligned}$$

This shows that S achieves δ_{Priv}^S -privacy. It remains to show that δ_{Priv}^S is optimal.

Let us consider the program c of the observer which accepts a view η only if $\rho(\eta)$ is in $M_{i,j}^S$. With this program c we have, for each essential view ρ :

$$\Pr[T \mapsto \rho] = A_{\rho,j}^S \quad \text{and} \quad \Pr[\tilde{T} \mapsto \rho] = A_{\rho,i}^S.$$

For i, j such that $\sum_{\rho \in M_{i,j}^S} (A_{\rho,j}^S - A_{\rho,i}^S) = \max_{l,l'=1,\dots,k} \sum_{\rho \in M_{l,l'}^S} (A_{\rho,l'}^S - A_{\rho,l}^S)$ we immediately obtain:

$$\Phi = \sum_{\rho \in M_{i,j}^S} (\Pr[T \mapsto \rho] - \Pr[\tilde{T} \mapsto \rho]) = \sum_{\rho \in M_{i,j}^S} (A_{\rho,j}^S - A_{\rho,i}^S) = \delta. \quad (8)$$

This concludes the proof of Theorem 4.

E Formulas

In this section we will sometimes use the same symbol to denote a random variable and a constant, when the context will allow to distinguish between these cases. For instance, we may use ρ to represent a random function over Ω_1 (such that $\rho(\omega_1)$ is the essential view determined by ω_1) and fix some essential view ρ . Then, in contexts like $\Pr_{\omega_1} [\rho(\omega_1) = \rho]$, the first instance of ρ denotes the random variable and the second instance denotes the fixed essential view.

E.1 Formulas for P_{TB}

In this section, we begin by stating some auxiliary facts. Then we develop the constants expressing the level of privacy and coercion resistance of ThreeBallot.

E.1.1 Auxiliary Facts

Lemma 5. *In the original version of ThreeBallot, the following formulas are true:*

1. Consider the case when the honest voters keep their receipts secret. Let $\rho = (n_x^{\times}, n_x^{\circ}, n_o^{\circ})$ be an essential view produced by n honest voters, as defined in Section 6.4. Then the probability A_{ρ}^{o-} that n honest voters produce the essential view ρ is as follows:

$$A_{\rho}^{o-} = \binom{n}{N} \binom{N}{R} \cdot p_0^{n-N} \cdot p_1^R \cdot p_2^{N-R} \cdot \binom{N}{n_x^{\times}} \left(\frac{2}{3}\right)^{n_x^{\times}} \left(\frac{1}{3}\right)^{N-n_x^{\times}},$$

where $N = (2n_x^{\times} + n_x^{\circ} + n_o^{\circ})/3$ denotes the total number of non-abstaining voters and $R = (n_x^{\times} + n_x^{\circ}) - N$ denotes the number of votes for candidate 1.

2. Consider the case when the honest voters publish their receipts secret. Let $\rho = (n_x^{\times}, n_x^{\circ}, n_o^{\circ}, r_x^{\times}, r_x^{\circ}, r_o^{\circ})$ be an essential view of the coercer, as defined in Section 6.4. Then the probability A_{ρ}^{o+} that n honest voters produce the essential view ρ is:

$$\begin{aligned} A_{\rho}^{o+} = & \binom{n}{N} \binom{N}{R} p_0^{n-N} p_1^R p_2^{N-R} \sum_{\tau_1, \tau_2} \binom{R}{\tau_1} \binom{R-\tau_1}{\tau_2} \left(\frac{4}{9}\right)^{\tau_1} \left(\frac{1}{9}\right)^{\tau_2} \\ & \binom{N-R}{r_x^{\circ} - \tau_1} \binom{N-R-r_x^{\circ} + \tau_1}{r_o^{\circ} - \tau_2} \left(\frac{1}{9}\right)^{r_x^{\circ} - \tau_1} \left(\frac{4}{9}\right)^{r_o^{\circ} - \tau_2} \\ & \binom{N-r_x^{\circ} - r_o^{\circ}}{r_x^{\times}} \left(\frac{2}{9}\right)^{N-r_x^{\circ} - r_o^{\circ}} \cdot \binom{\tau_1 + r_x^{\circ} - \tau_2}{n_x^{\times} - r_x^{\times} - (N - r_x^{\times} - r_x^{\circ} - r_o^{\circ})} \left(\frac{1}{2}\right)^{\tau_1 + r_x^{\circ} - \tau_2} \end{aligned}$$

with N and R as above and τ_1 and τ_2 ranging over the set $\{0, \dots, n\}$. We use the convention that $\binom{m}{l} = 0$ for $m < 0$.

In the proof of Lemma 5 (see below), we will use the following easy to prove facts.

Lemma 6. Consider honest, non-abstaining voter in the original variant of ThreeBallot. The following facts hold true.

1. The probability that a voter takes receipt \times_x^{\times} is $\frac{2}{9}$.
2. The probability that a voter takes receipt \circ_o° is $\frac{2}{9}$.
3. The probability that a voter that votes for candidate 1 takes receipt \times_o^{\times} is $\frac{4}{9}$ and $\frac{1}{9}$ for receipt \times_x° .
4. The probability that a voter that votes for candidate 2 takes receipt \circ_x° is $\frac{4}{9}$ and $\frac{1}{9}$ for receipt \times_o^{\times} .
5. The probability that a voter produces a \times_x^{\times} -ballot is $\frac{1}{2}$ in either of the following cases: (a) if we assume that she votes for candidate 1 and takes \times_o^{\times} as a receipt, and (b) if she votes for candidate 2 and takes \circ_x° as a receipt.

Proof of Lemma 5. We only prove 2., 1. is analogous (and simpler).

We will use the following random variables on Ω_1 : $F(\omega_1)$ denotes the number of (\times_x^{\times}) -ballots, $R(\omega_1)$ is the number of votes of honest voters for candidate 1, $N(\omega_1)$ is the number of non-abstaining honest voters, $\tau_1(\omega_1)$ denotes the number of voters that vote for 1 and take \times_o^{\times} as receipt, and $\tau_2(\omega_1)$ is the number of voters that vote for 1 and take \circ_x° as receipt.

Let $\rho = (n_x^{\times}, n_x^{\circ}, n_o^{\circ}, r_x^{\times}, r_x^{\circ}, r_o^{\circ})$. Let $N = (2n_x^{\times} + n_x^{\circ} + n_o^{\circ})/3$ denote the total number of non-abstaining voters and $R = (n_x^{\times} + n_x^{\circ}) - N$ denote the votes for candidate 1. Then we have the following equality, where τ_1 and τ_2 range over $\{0, \dots, n\}$.

$$A_{\rho}^{o+} = \Pr_{\omega_1}[\rho(\omega_1) = \rho] = \sum_{\tau_1, \tau_2} \Pr_{\omega_1}[\rho(\omega_1) = \rho, \tau_1(\omega_1) = \tau_1, \tau_2(\omega_1) = \tau_2].$$

Moreover, since n_x^\times , N , and R determine n_o^\times and n_o° and vice versa, we have for all τ_1, τ_2

$$\begin{aligned} & \Pr_{\omega_1}[\rho(\omega_1) = \rho, \tau_1(\omega_1) = \tau_1, \tau_2(\omega_1) = \tau_2] \\ &= \Pr_{\omega_1}[F(\omega_1) = n_x^\times, R(\omega_1) = R, \tau_1(\omega_1) = \tau_1, \tau_2(\omega_1) = \tau_2, \text{rec}(\omega_1) = (r_x^\times, r_o^\times, r_o^\circ), N(\omega_1) = N] \\ &= \Pr_{\omega_1}[F(\omega_1) = n_x^\times \mid R(\omega_1) = R, \tau_1(\omega_1) = \tau_1, \tau_2(\omega_2) = \tau_2, \\ & \quad \text{rec}(\omega_1) = (r_x^\times, r_o^\times, r_o^\circ), N(\omega_1) = N]. \end{aligned} \quad (9)$$

$$\begin{aligned} & \cdot \Pr_{\omega_1}[r_x^\times(\omega_1) = r_x^\times \mid r_o^\times(\omega_1) = r_o^\times, r_o^\circ(\omega_1) = r_o^\circ, \tau_1(\omega_1) = \tau_1, \tau_2(\omega_2) = \tau_2, \\ & \quad R(\omega_1) = R, N(\omega_1) = N]. \end{aligned} \quad (10)$$

$$\cdot \Pr_{\omega_1}[r_o^\times(\omega_1) = r_o^\times, r_o^\circ(\omega_1) = r_o^\circ, \mid \tau_1(\omega_1) = \tau_1, \tau_2(\omega_2) = \tau_2, R(\omega_1) = R, N(\omega_1) = N]. \quad (11)$$

$$\cdot \Pr_{\omega_1}[\tau_1(\omega_1) = \tau_1, \tau_2(\omega_1) = \tau_2 \mid R(\omega_1) = R, N(\omega_1) = N]. \quad (12)$$

$$\cdot \Pr_{\omega_1}[R(\omega_1) = R, N(\omega_1) = N]. \quad (13)$$

For (13), we have to distribute $n - N$ abstaining voters and R voters voting for candidate 1 among the n voters, which gives us

$$\Pr_{\omega_1}[R(\omega_1) = R, N(\omega_1) = N] = \binom{n}{N} \binom{N}{R} p_0^{n-N} p_1^R p_2^{N-R}.$$

For (12), we have to distribute τ_1 , respectively τ_2 , voters that take receipt $^\times$, respectively $^\circ$, among the R voters that vote for candidate 1. Using Lemma 6, we get

$$\begin{aligned} & \Pr_{\omega_1}[\tau_1(\omega_1) = \tau_1, \tau_2(\omega_1) = \tau_2 \mid R(\omega_1) = R, N(\omega_1) = N] = \\ &= \binom{R}{\tau_1} \binom{R - \tau_1}{\tau_2} \left(\frac{4}{9}\right)^{\tau_1} \left(\frac{1}{9}\right)^{\tau_2} \left(\frac{4}{9}\right)^{R - \tau_1 - \tau_2}. \end{aligned}$$

For (11), we have to distribute the remaining $r_o^\times - \tau_1$ receipts of type $^\times$ among the $N - R$ voters that do not vote for 1 and the remaining $r_o^\circ - \tau_2$ receipts of type $^\circ$ among the remaining $N - R - (r_o^\times - \tau_1)$ voters that do not vote for 1. With Lemma 6 we obtain:

$$\begin{aligned} & \Pr_{\omega_1}[r_o^\times(\omega_1) = r_o^\times, r_o^\circ(\omega_1) = r_o^\circ, \mid \tau_1(\omega_1) = \tau_1, \tau_2(\omega_2) = \tau_2, R(\omega_1) = R, N(\omega_1) = N] = \\ &= \binom{N - R}{r_o^\times - \tau_1} \binom{N - R - r_o^\times + \tau_1}{r_o^\circ - \tau_2} \left(\frac{1}{9}\right)^{r_o^\times - \tau_1} \left(\frac{4}{9}\right)^{r_o^\circ - \tau_2} \left(\frac{4}{9}\right)^{N - R - r_o^\times + \tau_1 - r_o^\circ + \tau_2} \end{aligned}$$

For (10), we have to distribute r_x^\times receipts of type $^\times$ among the $N - r_o^\times - r_o^\circ$ voters who take $^\times$ or $^\circ$ as receipt. With Lemma 6 we obtain:

$$\begin{aligned} & \Pr_{\omega_1}[r_x^\times(\omega_1) = r_x^\times \mid r_o^\times(\omega_1) = r_o^\times, r_o^\circ(\omega_1) = r_o^\circ, \tau_1(\omega_1) = \tau_1, \tau_2(\omega_2) = \tau_2, R(\omega_1) = R, N(\omega_1) = N] = \\ &= \binom{N - r_o^\times - r_o^\circ}{r_x^\times} \left(\frac{1}{2}\right)^{N - r_o^\times - r_o^\circ} \end{aligned}$$

Once $r_x^\times, r_o^\times, r_o^\circ, \tau_1$ and τ_2 are fixed, we know that the number of $^\times$ -ballots submitted by the honest voters who have taken $^\times$ or $^\circ$ as the receipt is $r_x^\times + (N - r_x^\times - r_o^\times - r_o^\circ)$. So, for (9), we have to distribute the remaining $n_x^\times - r_x^\times - (N - r_x^\times - r_o^\times - r_o^\circ)$ voters that submit a $^\times$ ballot among the

remaining $\tau_1 + r_{\times}^{\circ} - \tau_2$ voters who do not take \times nor \circ as the receipt, but still may submit \times (note that a voter who votes for 1 and takes \circ as the receipt cannot submit \times -ballot; similarly for a voter who votes for 2 and takes \times as the receipt). With Lemma 6 we obtain:

$$\Pr_{\omega_1} [F(\omega_1) = n_{\times}^{\times} \mid r_{\times}^{\times}(\omega_1) = r_{\times}^{\times}, r_{\circ}^{\times}(\omega_1) = r_{\circ}^{\times}, r_{\times}^{\circ}(\omega_1) = r_{\times}^{\circ}, \tau_1(\omega_1) = \tau_1, \tau_2(\omega_2) = \tau_2, \\ R(\omega_1) = R, N(\omega_1) = N] = \binom{\tau_1 + r_{\times}^{\circ} - \tau_2}{n_{\times}^{\times} - r_{\times}^{\times} - (N - r_{\times}^{\times} - r_{\circ}^{\times} - r_{\times}^{\circ})} \left(\frac{1}{2}\right)^{\tau_1 + r_{\times}^{\circ} - \tau_2}$$

By combining and rewriting the above formulas, we obtain the formula in Lemma 5. \square

The following lemma is proved in [16].

Lemma 7. *In the version of [8], the following formulas are true:*

1. *Consider the case where the honest voters keep their receipts secret. Let ρ be an essential view of the form $(n_{\times}^{\times}, n_{\circ}^{\times}, n_{\times}^{\circ})$ produced by n honest voters. Then, the probability A_{ρ}^{p-} that n honest voters produce the essential view ρ is as follows:*

$$A_{\rho}^{p-} = \binom{n}{N} \binom{N}{R} \cdot p_0^{n-N} \cdot p_1^R \cdot p_2^{N-R} \cdot \binom{N}{n_{\times}^{\times}} \left(\frac{2}{3}\right)^{n_{\times}^{\times}} \left(\frac{1}{3}\right)^{N-n_{\times}^{\times}},$$

with $N = (2n_{\times}^{\times} + n_{\circ}^{\times} + n_{\times}^{\circ})/3$ denoting the total number of non-abstaining voters and $R = (n_{\times}^{\times} + n_{\circ}^{\times}) - N$ denoting the number of votes for candidate 1.

2. *Consider the case where the honest voters publish their receipts. Let ρ be an essential view of the form $(n_{\times}^{\times}, n_{\circ}^{\times}, n_{\times}^{\circ}, r_{\times}^{\times}, r_{\circ}^{\times}, r_{\times}^{\circ})$ produced by n honest voters. Then, the probability A_{ρ}^{p+} that n honest voters produce the essential view ρ is as follows:*

$$A_{\rho}^{p+} = \binom{n}{N} p_0^{n-N} p_1^R p_2^{N-R} \cdot \binom{N}{r_{\times}^{\times}, r_{\circ}^{\times}, r_{\times}^{\circ}} \cdot \left(\frac{1}{9}\right)^{r_{\times}^{\times}} \left(\frac{2}{9}\right)^{r_{\circ}^{\times} + r_3} \left(\frac{4}{9}\right)^{N - r_{\times}^{\times} - r_{\circ}^{\times} - r_{\times}^{\circ}} \\ \cdot \sum_{\tau_1, \tau_2} \binom{r_{\circ}^{\times} + r_{\times}^{\circ} - \tau_1 - \tau_2}{n_{\times}^{\times} - N + r_{\circ}^{\times} + r_{\times}^{\circ}} \left(\frac{1}{2}\right)^{r_{\circ}^{\times} + r_{\times}^{\circ} - \tau_1 - \tau_2} \\ \cdot \binom{N - r_{\circ}^{\times} - r_{\times}^{\circ}}{R - (r_{\times}^{\times} - \tau_1) - \tau_2} \binom{r_{\times}^{\times}}{\tau_1} \binom{r_{\circ}^{\times}}{\tau_2}.$$

with N and R as above and τ_1 and τ_2 ranging over the set $\{0, \dots, n\}$. We use the convention that $\binom{m}{l} = 0$ for $m < 0$.

E.1.2 Formulas for Privacy

For an essential view ρ as defined in Section 6.3 (that is in the setting of privacy), i.e., with indication of the receipt r of the voter under observation, let $\tilde{\rho}$ denote the view without the receipt of the voter under observation. Furthermore, for a pattern Z and an essential view ρ , let $\rho - Z$ denote the essential view obtained by removing the ballots of Z from ρ .

In the following we write $A_{\rho, i}^{o+}$ for $A_{\rho, i}^S$ with $S = P_{\text{TB}}^{o+}(A_H, q, n, k, \vec{p})$ and analogously for $A_{\rho, i}^{p+}$

Lemma 8. *1. For the original variant of ThreeBallot we have*

$$1. A_{\rho,1}^{o+} = \begin{cases} \frac{4}{9} \cdot A_{\tilde{\rho}_{-x\circ\circ}}^{o+}, & r = \begin{smallmatrix} x \\ x \end{smallmatrix} \text{ or } \begin{smallmatrix} \circ \\ \circ \end{smallmatrix} \\ \frac{2}{9} \cdot A_{\tilde{\rho}_{-x\circ\circ}}^{o+} + \frac{2}{9} \cdot A_{\tilde{\rho}_{-x\circ x}}^{o+}, & r = \begin{smallmatrix} x \\ \circ \end{smallmatrix} \\ \frac{1}{9} \cdot A_{\tilde{\rho}_{-x\circ x}}^{o+}, & r = \begin{smallmatrix} \circ \\ x \end{smallmatrix} \end{cases}$$

$$2. A_{\rho,2}^{o+} = \begin{cases} \frac{4}{9} \cdot A_{\tilde{\rho}_{-x\circ\circ}}^{o+}, & r = \begin{smallmatrix} x \\ x \end{smallmatrix} \text{ or } \begin{smallmatrix} \circ \\ \circ \end{smallmatrix} \\ \frac{1}{9} \cdot A_{\tilde{\rho}_{-x\circ\circ}}^{o+}, & r = \begin{smallmatrix} x \\ \circ \end{smallmatrix} \\ \frac{2}{9} \cdot A_{\tilde{\rho}_{-x\circ\circ}}^{o+} + \frac{2}{9} \cdot A_{\tilde{\rho}_{-x\circ x}}^{o+}, & r = \begin{smallmatrix} \circ \\ x \end{smallmatrix} \end{cases}$$

2. For the variant of [8] we have

$$1. A_{\rho,1}^{p+} = \begin{cases} \frac{5}{9} \cdot A_{\tilde{\rho}_{-x\circ\circ}}^{p+}, & r = \begin{smallmatrix} x \\ x \end{smallmatrix} \text{ or } \begin{smallmatrix} \circ \\ \circ \end{smallmatrix} \\ \frac{1}{9} \cdot A_{\tilde{\rho}_{-x\circ\circ}}^{p+} + \frac{1}{9} \cdot A_{\tilde{\rho}_{-x\circ x}}^{p+}, & r = \begin{smallmatrix} x \\ \circ \end{smallmatrix} \\ \frac{2}{9} \cdot A_{\tilde{\rho}_{-x\circ x}}^{p+}, & r = \begin{smallmatrix} \circ \\ x \end{smallmatrix} \end{cases}$$

$$2. A_{\rho,2}^{p+} = \begin{cases} \frac{5}{9} \cdot A_{\tilde{\rho}_{-x\circ\circ}}^{p+}, & r = \begin{smallmatrix} x \\ x \end{smallmatrix} \text{ or } \begin{smallmatrix} \circ \\ \circ \end{smallmatrix} \\ \frac{2}{9} \cdot A_{\tilde{\rho}_{-x\circ\circ}}^{p+}, & r = \begin{smallmatrix} x \\ \circ \end{smallmatrix} \\ \frac{1}{9} \cdot A_{\tilde{\rho}_{-x\circ\circ}}^{p+} + \frac{1}{9} \cdot A_{\tilde{\rho}_{-x\circ x}}^{p+}, & r = \begin{smallmatrix} \circ \\ x \end{smallmatrix} \end{cases}$$

Proof. We now focus on $A_{\rho,1}^{o+}$. The proofs for the remaining cases are similar.

Let us consider the case when $r = \begin{smallmatrix} x \\ \circ \end{smallmatrix}$. A voter that votes for candidate 1 submits either the pattern $\begin{smallmatrix} x\circ\circ \\ x\circ\circ \end{smallmatrix}$ or $\begin{smallmatrix} x\circ\circ \\ \circ\circ x \end{smallmatrix}$. An elementary calculation yields that the probability that a voter who votes for candidate 1 submits $\begin{smallmatrix} x\circ\circ \\ x\circ x \end{smallmatrix}$ and takes $\begin{smallmatrix} x \\ \circ \end{smallmatrix}$ as receipt is $\frac{2}{9}$, which is also the probability that a voter that votes for candidate 1 submits $\begin{smallmatrix} x\circ\circ \\ \circ\circ x \end{smallmatrix}$ and takes $\begin{smallmatrix} x \\ \circ \end{smallmatrix}$ as receipt.

In the following, we will use the following random variables on Ω_1 : $\rho(\omega_1)$ is the essential view determined by ω_1 , $\tilde{\rho}(\omega_1)$ is the essential view without the receipt of the voter under observation, $Z_0(\omega_1)$ is the pattern of the voter under observation (as determined by ω_1), and $v_0(\omega_1)$ is the candidate chosen by this voter. Then we obtain

$$\begin{aligned} A_{\rho,1}^{o+} &= \Pr_{\omega_1}[\rho(\omega_1) = \rho \mid v_0(\omega_1) = 1] + \\ &= \Pr_{\omega_1}[\rho(\omega_1) = \rho, Z_0(\omega_1) = \begin{smallmatrix} x\circ\circ \\ x\circ\circ \end{smallmatrix} \mid v_0(\omega_1) = 1] + \\ &\quad + \Pr_{\omega_1}[\rho(\omega_1) = \rho, Z_0(\omega_1) = \begin{smallmatrix} x\circ\circ \\ \circ\circ x \end{smallmatrix} \mid v_0(\omega_1) = 1] = \\ &= \Pr_{\omega_1}[\tilde{\rho}(\omega_1) = \rho, Z_0(\omega_1) = \begin{smallmatrix} x\circ\circ \\ x\circ\circ \end{smallmatrix} \mid v_0(\omega_1) = 1] + \\ &\quad + \Pr_{\omega_1}[\tilde{\rho}(\omega_1) = \rho, Z_0(\omega_1) = \begin{smallmatrix} x\circ\circ \\ \circ\circ x \end{smallmatrix} \mid v_0(\omega_1) = 1] = \\ &= \Pr_{\omega_1}[\tilde{\rho}(\omega_1) = \tilde{\rho} \mid Z_0(\omega_1) = \begin{smallmatrix} x\circ\circ \\ x\circ\circ \end{smallmatrix}] \cdot \Pr_{\omega_1}[Z_0(\omega_1) = \begin{smallmatrix} x\circ\circ \\ x\circ\circ \end{smallmatrix} \mid v_0(\omega_1) = 1] + \\ &\quad + \Pr_{\omega_1}[\tilde{\rho}(\omega_1) = \tilde{\rho} \mid Z_0(\omega_1) = \begin{smallmatrix} x\circ\circ \\ \circ\circ x \end{smallmatrix}] \cdot \Pr_{\omega_1}[Z_0(\omega_1) = \begin{smallmatrix} x\circ\circ \\ \circ\circ x \end{smallmatrix} \mid v_0(\omega_1) = 1] = \\ &= \frac{2}{9} \cdot A_{\tilde{\rho}_{-x\circ\circ}}^{o+} + \frac{2}{9} \cdot A_{\tilde{\rho}_{-x\circ x}}^{o+}. \end{aligned}$$

The cases $r = \begin{smallmatrix} x \\ x \end{smallmatrix}$, $r = \begin{smallmatrix} \circ \\ \circ \end{smallmatrix}$, or $r = \begin{smallmatrix} \circ \\ x \end{smallmatrix}$ are analogously to that (and simpler), as then the pattern used by the voter under observation is uniquely determined. \square

Now let us consider the case when the voters keep their receipts secret. In this case an essential view is of the form $\rho = (n_x, n_\circ, n_x)$.

Lemma 9. For both the original variant of ThreeBallot and the variant from [8], i.e., for $S \in \{P_{\text{TB}}^{o-}, P_{\text{TB}}^{p-}\}$, we have:

$$1. A_{\rho,1}^S = \frac{2}{3} A_{\rho_{-x\circ\circ}}^S + \frac{1}{3} A_{\rho_{-x\circ x}}^S$$

$$2. A_{\rho,2}^S = \frac{2}{3}A_{\rho-\begin{smallmatrix} \times \times \circ \\ \times \times \circ \end{smallmatrix}}^S + \frac{1}{3}A_{\rho-\begin{smallmatrix} \circ \circ \times \\ \circ \circ \times \end{smallmatrix}}^S.$$

Proof. The proof is analogous to the proof of Lemma 8, where one uses the fact that the probability that an honest voter that votes for candidate 1 (respectively 2) submits the pattern $\begin{smallmatrix} \times \times \circ \\ \times \times \circ \end{smallmatrix}$ (respectively $\begin{smallmatrix} \circ \circ \times \\ \circ \circ \times \end{smallmatrix}$) is $\frac{2}{3}$. \square

E.1.3 Formulas for Coercion-Resistance

Lemma 10. *Let $S \in \{p+, p-, o+, o-\}$. Let ρ denote an essential view as defined in Section 6.4. Then we have $A_{\rho,Z}^S = A_{\rho-Z}^S$.*

Proof. This lemma follows directly from the definition of $A_{\rho,Z}^S$ and $A_{\rho-Z}^S$. \square

E.2 Formulas for VAV

E.2.1 Auxiliary Facts

Lemma 11. *Let $S \in \{P_{\text{VAV}}^{p-}, P_{\text{VAV}}^{s-}\}$ and $\rho = (n_1^A, \dots, n_k^A, n_1^V, \dots, n_k^V)$ be an essential view in S , as defined in Section 6.4, produced by n honest voters. Let $c_i = n_i^V - n_i^A$ denote the number of votes for candidate i and $N = \sum_{i=1}^k c_i$ the number of non-abstaining voters. Then, the probability A_{ρ}^S that n honest voters produce the view ρ is as follows:*

$$A_{\rho}^S = \binom{n}{n-N, c_1, \dots, c_k} p_0^{n-N} p_1^{c_1} \cdots p_k^{c_k} \binom{N}{n_1^A, \dots, n_k^A} \left(\frac{1}{k}\right)^N. \quad (14)$$

Proof. For (14), we have to distribute the abstaining voters $n - N$ and the voters that vote for candidate c_i ($1 \leq i \leq k$) among the n honest voters, which gives us the first part of (14). For the second part, we have to distribute the A-ballots among the voters. Note that the probability that a voter submits a particular A-ballot is $\frac{1}{k}$, independent of the chosen candidate. As the chosen candidate of a voter and the A-ballot she submitted uniquely determines the pattern she submitted, the essential view is uniquely determined by the chosen candidates and the A-ballots. From this (14) follows. \square

Lemma 12. *Let $\rho = (n_1^A, \dots, n_k^A, n_1^V, \dots, n_k^V, r_1^A, \dots, r_k^A, r_1^V, \dots, r_k^V)$ be an essential view in $S = P_{\text{VAV}}^{p+}$, as defined in Section 6.4, produced by n honest voters. Let $c_i = n_i^V - n_i^A$ denote the number of votes for candidate i and $N = \sum_{i=1}^k c_i$ the number of non-abstaining voters. Then, the probability A_{ρ}^S that n honest voters produce the view ρ is as follows:*

$$A_{\rho}^S = \binom{n}{n-N, c_1, \dots, c_k} p_0^{n-N} p_1^{c_1} \cdots p_k^{c_k} \binom{N}{r_1^A, \dots, r_k^A, r_1^V, \dots, r_k^V} \left(\frac{1}{2k}\right)^N. \quad (15)$$

Proof. The proof is analogously to the proof of Lemma 11, with the difference that, instead of distributing the A-ballots among the voters, we now distribute the receipts among the voters. As every receipt is taken with probability $\frac{1}{2k}$ (independent of the chosen candidate), and the receipt together with the chosen candidate uniquely determines the pattern, (15) follows (analogously to the proof of Lemma 11). \square

For the simple variant of VAV, in case the voters publish their receipts, we restrict ourselves to the case of two candidates.

Lemma 13. Let $\rho = (n_1^A, n_2^A, n_1^V, n_2^V, r_1^A, r_2^A, r_1^V, r_2^V)$ be an essential view in $S = P_{\text{VAV}}^{s+}$, as defined in Section 6.4, produced by n honest voters. Let $c_i = n_i^V - n_i^A$ denote the number of votes for candidate i and $N = c_1 + c_2$ the number of non-abstaining voters. Then, the probability A_ρ^S that n honest voters produce the view ρ is:

$$A_\rho^S = \sum_{\tau_{11}^A, \tau_{22}^A, \tau_2, \tau_{22}^V} \binom{n}{n-N, c_1, c_2} p_0^{c_1} p_1^{c_2} p_2^{c_2} \binom{c_1}{\tau_{11}^V, \tau_{11}^A, \tau_{12}^V, \tau_{12}^A, \tau_1} \binom{c_2}{\tau_{21}^V, \tau_{21}^A, \tau_{22}^V, \tau_{22}^A, \tau_2} 2^{\tau_{11}^V + \tau_{22}^V} \left(\frac{1}{6}\right)^N$$

Proof. We define the following random variables on Ω_1 : Let $\tau_{11}^V(\omega_1)$ denote the number of voters that submit a pattern, where candidate 1 is marked on every ballot and a V-ballot is taken as receipt. Similarly, let $\tau_{11}^A(\omega_1)$ denote the number of voters that submit the same pattern, but the A ballot is taken as receipt. Let $\tau_{12}^V(\omega_1)$ denote the number of voters that vote for candidate 1 but submit an A-ballot for candidate 2 and take the V-ballot for candidate 1 as receipt. Let $\tau_{12}^A(\omega_1)$ denote the number of voters that submit the same pattern but take the A-ballot as receipt. Further, let $\tau_1(\omega_1)$ denote the number of voters that submit the same pattern, but take the V-ballot for candidate 2 as receipt.

In the similar way we handle the symmetric case: Let $\tau_{22}^V(\omega_1)$ denote the number of voters that submit a pattern, where candidate 2 is marked on every ballot and a V-ballot is taken as receipt. Let $\tau_{22}^A(\omega_1)$ be the number of voters that submit the same pattern, but the A ballot is taken as receipt. Let $\tau_{21}^V(\omega_1)$ denote the number of voters that vote for candidate 2 but submit an A-ballot for candidate 1 and take the V-ballot for candidate 2 as receipt. Let $\tau_{21}^A(\omega_1)$ denote the number of voters that submit the same pattern but take the A-ballot as receipt. Finally, let $\tau_2(\omega_1)$ denote the number of voters that submit the same pattern, but take the V-ballot for candidate 1 as receipt.

Now, we have

$$\begin{aligned} A_\rho^S &= \Pr[\rho(\omega_1) = \rho] \\ &= \sum_{\tau_{11}^A, \tau_{22}^A, \tau_2, \tau_{22}^V} \Pr[\rho(\omega_1) = \rho, \tau_{11}^A(\omega_1) = \tau_{11}^A, \tau_{22}^A(\omega_1) = \tau_{22}^A, \tau_2(\omega_1) = \tau_2, \tau_{22}^V(\omega_1) = \tau_{22}^V] \\ &= \sum_{\tau_{11}^A, \tau_{22}^A, \tau_2, \tau_{22}^V} \Pr[M^\rho, \tau_{11}^A, \tau_{22}^A, \tau_2, \tau_{22}^V] \end{aligned}$$

where

$$\begin{aligned} M^\rho, \tau_{11}^A, \tau_{22}^A, \tau_2, \tau_{22}^V &= \\ &= \{\omega_1 : \rho(\omega_1) = \rho, \tau_{11}^A(\omega_1) = \tau_{11}^A, \tau_{22}^A(\omega_1) = \tau_{22}^A, \tau_2(\omega_1) = \tau_2, \tau_{22}^V(\omega_1) = \tau_{22}^V\}. \end{aligned}$$

By taking a look at the definition of the random variables $\tau_{12}^A(\omega_1), \dots$ above, one can easily verify

that the following equations hold true:

$$\begin{aligned}
n_1^V(\omega_1) &= 2\tau_{11}^V(\omega_1) + \tau_{12}^V(\omega_1) + \tau_1(\omega_1) + \tau_{21}^V(\omega_1) + \tau_{12}^A(\omega_1) + \tau_{21}^A(\omega_1) + 2\tau_{11}^A(\omega_1) + && + \tau_2(\omega_1) \\
n_1^A(\omega_1) &= \tau_{11}^V(\omega_1) + && + \tau_{21}^V(\omega_1) + && + \tau_{21}^A(\omega_1) + \tau_{11}^A(\omega_1) + && + \tau_2(\omega_1) \\
n_2^V(\omega_1) &= && \tau_{12}^V(\omega_1) + \tau_1(\omega_1) + \tau_{21}^V(\omega_1) + \tau_{12}^A(\omega_1) + \tau_{21}^A(\omega_1) + && + 2\tau_{22}^V(\omega_1) + 2\tau_{22}^A(\omega_1) + \tau_2(\omega_1) \\
n_2^A(\omega_1) &= && \tau_{12}^V(\omega_1) + \tau_1(\omega_1) + && + \tau_{12}^A(\omega_1) && + 2\tau_{22}^V(\omega_1) + 2\tau_{22}^A(\omega_1) \\
r_1^V(\omega_1) &= \tau_{11}^V(\omega_1) + \tau_{12}^V(\omega_1) + && && && + \tau_2(\omega_1) \\
r_1^A(\omega_1) &= && && + \tau_{21}^A(\omega_1) + \tau_{11}^A(\omega_1) \\
r_2^V(\omega_1) &= && \tau_1(\omega_1) + && + \tau_{21}^V(\omega_1) + && + \tau_{22}^V(\omega_1) \\
r_2^A(\omega_1) &= && && \tau_{12}^A(\omega_1) + && + \tau_{22}^A(\omega_1)
\end{aligned}$$

Using the Gaussian elimination algorithm (and the fact that $n_1^V(\omega_1) + n_2^V(\omega_1) = 2n_1^A(\omega_1) + 2n_2^A(\omega_1) = 2r_1^V(\omega_1) + 2r_2^V(\omega_1) + 2r_1^A(\omega_1) + r_2^A(\omega_1)$, which is true, as every expression obviously equals two times the number of voters), we get that the above system of linear equations is equivalent to the following system:

$$\begin{aligned}
\tau_{12}^A(\omega_1) &= r_2^A(\omega_1) - \tau_{22}^A(\omega_1) \\
\tau_{21}^A(\omega_1) &= r_1^A(\omega_1) - \tau_{11}^A(\omega_1) \\
\tau_{21}^V(\omega_1) &= -n_1^V(\omega_1) + 2n_1^A(\omega_1) + n_2^A(\omega_1) - \tau_{21}^A(\omega_1) - \tau_{22}^A(\omega_1) - \tau_2(\omega_1) - \tau_{22}^V(\omega_1) \\
\tau_1(\omega_1) &= n_1^V(\omega_1) - n_1^A(\omega_1) - \tau_{11}^A(\omega_1) - \tau_{12}^A(\omega_1) - r_1^V(\omega_1) + \tau_2(\omega_1) \\
\tau_{12}^V(\omega_1) &= n_1^V(\omega_1) - 2n_1^A(\omega_1) - \tau_{12}^A(\omega_1) + \tau_{21}^A(\omega_1) + \tau_2(\omega_1) + \tau_{21}^V(\omega_1) - \tau_1(\omega_1) \\
\tau_{11}^V(\omega_1) &= n_1^A(\omega_1) - \tau_2(\omega_1) - \tau_{21}^A(\omega_1) - \tau_{21}^V(\omega_1) - \tau_{11}^A(\omega_1).
\end{aligned}$$

whose equations, therefore, also hold true.

From this, we immediately get that every element of $M^{\rho, \tau_{11}^A, \tau_{22}^A, \tau_2, \tau_{22}^V}$ is included in the set

$$\begin{aligned}
\tilde{M}^{\rho, \tau_{11}^A, \tau_{22}^A, \tau_2, \tau_{22}^V} &= \{ \omega_1 : \tau_{11}^A(\omega_1) = \tau_{11}^A, \tau_{22}^A(\omega_1) = \tau_{22}^A, \tau_2(\omega_1) = \tau_2, \tau_{22}^V(\omega_1) = \tau_{22}^V \\
&\tau_{12}^A(\omega_1) = \tau_{12}^A, \tau_{21}^A(\omega_1) = \tau_{21}^A, \tau_{21}^V(\omega_1) = \tau_{21}^V, \tau_1(\omega_1) = \tau_1, \\
&\tau_{12}^V(\omega_1) = \tau_{12}^V, \tau_{11}^V(\omega_1) = \tau_{11}^V \}
\end{aligned}$$

where

$$\begin{aligned}
\tau_{12}^A &= r_2^A - \tau_{22}^A \\
\tau_{21}^A &= r_1^A - \tau_{11}^A \\
\tau_{21}^V &= -n_1^V + 2n_1^A + n_2^A - \tau_{21}^A - \tau_{22}^A - \tau_2 - \tau_{22}^V \\
\tau_1 &= n_1^V - n_1^A - \tau_{11}^A - \tau_{12}^A - r_1^V + \tau_2 \\
\tau_{12}^V &= n_1^V - 2n_1^A - \tau_{12}^A + \tau_{21}^A + \tau_2 + \tau_{21}^V - \tau_1 \\
\tau_{11}^V &= n_1^A - \tau_2 - \tau_{21}^A - \tau_{21}^V - \tau_{11}^A.
\end{aligned}$$

Furthermore, one can use the above equations and the fact that $n_1^V + n_2^V = 2n_1^A + 2n_2^A = 2r_1^V + 2r_2^V + 2r_1^A + r_2^A$, to show that every element ω_1 of $\tilde{M}^{\rho, \tau_{11}^A, \tau_{22}^A, \tau_2, \tau_{22}^V}$ is also an element of $M^{\rho, \tau_{11}^A, \tau_{22}^A, \tau_2, \tau_{22}^V}$. For example, to show that $r_2^A(\omega_1) = r_2^A$ we use the equation $r_2^A(\omega_1) = \tau_{12}^A(\omega_1) + \tau_{22}^A(\omega_1)$ and the facts $\tau_{12}^A(\omega_1) = \tau_{12}^A = r_2^A - \tau_{22}^A$ and $\tau_{22}^A(\omega_1) = \tau_{22}^A$.

Therefore

$$A_{\rho}^S = \sum_{\tau_{11}^A, \tau_{22}^A, \tau_2, \tau_{22}^V} \Pr[\tilde{M}^{\rho, \tau_{11}^A, \tau_{22}^A, \tau_2, \tau_{22}^V}]. \quad (16)$$

We now compute $\Pr[\tilde{M}^{\rho, \tau_{11}^A, \tau_{22}^A, \tau_2, \tau_{22}^V}]$. First, a simple calculation yields that the probability that a voter marks the same candidate on all three simple ballots and takes V-ballot as receipt is $\frac{1}{3}$; the probability for all other patterns is $\frac{1}{6}$. We obtain:

$$\Pr[\tilde{M}^{\rho, \tau_{11}^A, \tau_{22}^A, \tau_2, \tau_{22}^V}] = \binom{n}{n-N, c_1, c_2} p_0^{c_1} p_1^{c_2} p_2^{c_2} \quad (17)$$

$$\cdot \binom{c_1}{\tau_{11}^V, \tau_{11}^A, \tau_{12}^V, \tau_{12}^A, \tau_1} \left(\frac{1}{3}\right)^{\tau_{11}^V} \left(\frac{1}{6}\right)^{\tau_{11}^A} \left(\frac{1}{6}\right)^{\tau_{12}^V} \left(\frac{1}{6}\right)^{\tau_{12}^A} \left(\frac{1}{6}\right)^{\tau_1} \quad (18)$$

$$\cdot \binom{c_2}{\tau_{21}^V, \tau_{21}^A, \tau_{22}^V, \tau_{22}^A, \tau_2} \left(\frac{1}{3}\right)^{\tau_{22}^V} \left(\frac{1}{6}\right)^{\tau_{21}^V} \left(\frac{1}{6}\right)^{\tau_{21}^A} \left(\frac{1}{6}\right)^{\tau_{22}^A} \left(\frac{1}{6}\right)^{\tau_2} \quad (19)$$

Here, for (17) we distributed the abstaining voters, the voters for candidate 1 and the voters for candidate 2 among the n voters. For (18), we distributed the different possible patterns for candidate 1 among c_1 voters for candidate 1 and analogously for (19) for candidate 2.

From this and equation (16), we obtain the formula in Lemma 13. \square

E.2.2 Formulas for privacy

From the above lemmas, we get the following formulas for privacy.

In the following, we denote by Z_j^i the pattern (a filled out multi ballot) which is a vote for candidate i , where candidate j is marked on A-ballot. Let $\rho - Z_j^i$ be the essential view obtained from ρ by removing Z_j^i , that is by removing a V-ballot for candidate i , a V-ballot for candidate j , and an A-ballot for j .

Lemma 14. *Let $S \in \{P_{\text{VAV}}^-, P_{\text{VAV}}^s\}$ and $\rho = (n_1^A, \dots, n_k^A, n_1^V, \dots, n_k^V)$ be an essential view in S produced by n honest voters and the voter under observation. Then we have*

$$A_{\rho, i}^S = \frac{1}{k} \sum_{j=1}^k A_{\rho - Z_j^i}^S.$$

Proof. Clearly, $A_{\rho - Z_j^i}^S$ is the probability that the essential view is ρ , given that the voter under observation votes for candidate i and submits an A-ballot for j . The probability that the voter under observation submits an A-ballot for candidate j is $\frac{1}{k}$. Hence, $\sum_{j=1}^k \frac{1}{k} A_{\rho - Z_j^i}^S$ gives us the probability for the essential view ρ , given that the voter under observation votes for candidate i (and submits an arbitrary A-ballot). \square

Lemma 15. *Let $\rho = (n_1^A, \dots, n_k^A, n_1^V, \dots, n_k^V, r_1^A, \dots, r_k^A, r_1^V, \dots, r_k^V, r)$ be an essential view in $S = P_{\text{VAV}}^+$ produced by n honest voters and the voter under observation. Let j be the candidate marked on the ballot r . Then we have*

$$A_{\rho, i}^S = \frac{1}{2k} A_{\rho - Z_j^i}^S.$$

Proof. Clearly, $A_{\rho-Z_j^i}^S$ is the probability that the essential view is ρ , given that the voter under observation votes for candidate i and has a receipt where candidate j is marked (note that, in the system P_{VAV}^{s+} , the candidate marked on the A-ballot is j). The probability that the voter under observation has the receipt r is $\frac{1}{2k}$. From this, the formula follows. \square

In the following lemma, we denote by A_x° an A-ballot with marked first candidate. In the analogous way, we define A_x° , V_x° , and V_x° . Also, $\tilde{\rho}$ denotes the essential view without the receipt of the voter under observation.

Lemma 16. *Let $\rho = (n_1^A, n_2^A, n_1^V, n_2^V, r_1^A, r_2^A, r_1^V, r_2^V, r)$ be an essential view in $S = P_{\text{VAV}}^{s+}$ produced by n honest voters and the voter under observation. Let r be a ballot where candidate j is marked. Then we have*

$$A_{\rho,1}^S = \begin{cases} \frac{1}{6}A_{\tilde{\rho}-Z_1^1}^S, & r = A_x^\circ \\ \frac{1}{6}A_{\tilde{\rho}-Z_2^1}^S, & r = A_x^\circ \\ \frac{1}{6}A_{\tilde{\rho}-Z_2^1}^S, & r = V_x^\circ \\ \frac{1}{3}A_{\tilde{\rho}-Z_1^1}^S + \frac{1}{6}A_{\tilde{\rho}-Z_2^1}^S, & r = V_x^\circ \end{cases} \quad A_{\rho,2}^S = \begin{cases} \frac{1}{6}A_{\tilde{\rho}-Z_2^2}^S, & r = A_x^\circ \\ \frac{1}{6}A_{\tilde{\rho}-Z_1^2}^S, & r = A_x^\circ \\ \frac{1}{6}A_{\tilde{\rho}-Z_1^2}^S, & r = V_x^\circ \\ \frac{1}{3}A_{\tilde{\rho}-Z_2^2}^S + \frac{1}{6}A_{\tilde{\rho}-Z_1^2}^S, & r = V_x^\circ \end{cases}$$

where $\tilde{\rho} - Z_j^i$ is obtained from ρ by removing the receipt of the voter under observation and a V-ballot for candidate i , a V-ballot for candidate j , and an A-ballot for candidate j .

Proof. We now focus on $A_{\rho,1}^S$, the case of $A_{\rho,2}^S$ is similar.

Note that, if $r = V_x^\circ$, then the pattern is not uniquely determined. Here, either on all ballots candidate 1 is marked or on one V-ballot and on the A-ballot, candidate 2 is marked. The probability that a voter that votes for candidate 1 takes the receipt V_x° and marks candidate 1 on all three ballots is $\frac{1}{3}$, whereas the probability that a voter that votes for candidate 1 takes V_x° as receipt and marks candidate 2 on the remaining ballots is $\frac{1}{6}$.

In the following, we will use the following random variables on Ω_1 : $\rho(\omega_1)$ is the essential view determined by ω_1 , $\tilde{\rho}(\omega_1)$ is the essential view without the receipt of the voter under observation, $r(\omega_1)$ is the receipt of the voter under observation determined by ω_1 , $Z_0(\omega_1)$ is the pattern of the voter under observation (as determined by ω_1), and $v_0(\omega_1)$ is the candidate chosen by this voter. We have

$$\begin{aligned} A_{\rho,1}^{s+} &= \Pr_{\omega_1}[\rho(\omega_1) = \rho \mid v_0(\omega_1) = 1] \\ &= \Pr_{\omega_1}[\rho(\omega_1) = \rho, Z_0(\omega_1) = Z_1^1 \mid v_0(\omega_1) = 1] + \\ &\quad + \Pr_{\omega_1}[\rho(\omega_1) = \rho, Z_0(\omega_1) = Z_2^1 \mid v_0(\omega_1) = 1] = \\ &= \Pr_{\omega_1}[\tilde{\rho}(\omega_1) = \tilde{\rho}, Z_0(\omega_1) = Z_1^1, r(\omega_1) = V_x^\circ \mid v_0(\omega_1) = 1] + \\ &\quad + \Pr_{\omega_1}[\tilde{\rho}(\omega_1) = \tilde{\rho}, Z_0(\omega_1) = Z_2^1, r(\omega_1) = V_x^\circ \mid v_0(\omega_1) = 1] = \\ &= \Pr_{\omega_1}[\tilde{\rho}(\omega_1) = \tilde{\rho} \mid Z_0(\omega_1) = Z_1^1] \cdot \Pr_{\omega_1}[Z_0(\omega_1) = Z_1^1, r(\omega_1) = V_x^\circ \mid v_0(\omega_1) = 1] + \\ &\quad + \Pr_{\omega_1}[\tilde{\rho}(\omega_1) = \tilde{\rho} \mid Z_0(\omega_1) = Z_2^1] \cdot \Pr_{\omega_1}[Z_0(\omega_1) = Z_2^1, r(\omega_1) = V_x^\circ \mid v_0(\omega_1) = 1] = \\ &= \frac{1}{3} \cdot A_{\tilde{\rho}-Z_1^1}^{s+} + \frac{1}{6} \cdot A_{\tilde{\rho}-Z_2^1}^{s+}. \end{aligned}$$

The cases $r = A_x^\circ$, $r = A_x^\circ$, and $r = V_x^\circ$ are analogous (and simpler, as in these cases the pattern is uniquely determined). \square

E.2.3 Formulas for Coercion-Resistance

From Lemma 11, we also get the following formula for coercion-resistance:

Lemma 17. *Let $S \in \{P_{\text{VAV}}^{p-}, P_{\text{VAV}}^{s-}, P_{\text{VAV}}^{p+}, P_{\text{VAV}}^{s+}\}$ and $\rho = (n_1^A, \dots, n_k^A, n_1^V, \dots, n_k^V)$ be an essential view in S produced by n honest voters and the coerced voter. Let Z be an arbitrary pattern. Then we have*

$$A_{\rho, Z}^S = A_{\rho-Z}^S,$$

where $\rho - Z$ results from ρ by removing the ballots of Z .

Proof. This formula follows directly from the definition. □

F Proofs of Theorem 5

In this section, we prove Theorem 5. The proof for the ThreeBallot protocol (that is for P_{TB}^{o+} , P_{TB}^{o-} , P_{TB}^{p+} , and P_{TB}^{p-}) is given in [16]. Therefore, we only need to focus on VAV (P_{VAV}^{s+} , P_{VAV}^{s-} , P_{VAV}^{p+} , and P_{VAV}^{p-}). In the following, we will only deal with the case where the coercer can see the receipt of the voters (i.e. we will consider the systems P_{VAV}^{s+} and P_{VAV}^{p+}). Proofs for P_{VAV}^{s-} and P_{VAV}^{p-} are similar (and simpler).

For a coercer's view η , let $f(\eta)$ denotes the pattern that, in a run with η , the coercer requires the coerced voter to use, if any (otherwise, $f(\eta)$ is undefined). Since in the systems we consider, this patterns depends only on ω_2 , we will also denote it by $f(\omega_2)$.

We use the same notations as in Appendix D with the following exceptions: $\omega_1 \in \Omega_1$ now only determines the choices and the submitted patterns of the honest voters (recall that in Appendix D, ω_1 also determined the way the voter under observation votes for candidate j in $\hat{\pi}_v(j)$). The receipt of the coerced voter is not included in Rec and rec . Further, let φ_η be a predicate over Ω_1 such that $\varphi_\eta(\omega_1)$ iff $\rho(f(\eta), \omega_1) = \rho(\eta)$ and $Rec(\omega_1) = Rec(\eta)$, and, similarly, $\tilde{\varphi}_\eta(\omega_1)$ be a predicate which holds iff $\rho(C(f(\eta), i), \omega_1) = \rho(\eta)$ and $Rec(\omega_1) = Rec(\eta)$, Recall that, according to the considered goal γ_i , i is the candidate the coerced voter wants to vote for.

Further, let c be now a program of the coercer. Let $T = (c \parallel \text{dum} \parallel e_S)$ and $\tilde{T} = (c \parallel \tilde{v} \parallel e_S)$.

We now show that the view of the coercer is information-theoretically independent of the choices of honest voters and the coerced voter as long as these choices are consistent with the essential view and the order of the receipts. This is formally stated in Lemma 19 below. The following lemma, which is the counterpart of Lemma 3 from Appendix D, is the core fact used to prove Lemma 19.

Lemma 18. *Let η be a view of the coercer such that $f(\eta)$ is defined. Let $\omega_1, \omega'_1, \omega''_1, \omega'''_1$ be arbitrary elements of Ω_1 with $\varphi_\eta(\omega_1)$, $\varphi_\eta(\omega'_1)$, $\tilde{\varphi}_\eta(\omega''_1)$ and $\tilde{\varphi}_\eta(\omega'''_1)$. Then the sets*

$$\begin{aligned} A &= \{\omega_2 : T(\omega_1, \omega_2) \mapsto \eta\}, & B &= \{\omega_2 : T(\omega'_1, \omega_2) \mapsto \eta\}, \\ C &= \{\omega_2 : \tilde{T}(\omega''_1, \omega_2) \mapsto \eta\}, & D &= \{\omega_2 : \tilde{T}(\omega'''_1, \omega_2) \mapsto \eta\} \end{aligned}$$

have the same cardinality, and hence, have the same probability.

The proof of this lemma is very similar to the proof of Lemma 3. Using Lemma 18, one can prove the following lemma (the proof is very similar to the proof of Lemma 4).

Lemma 19. *Let η be a coercer view such that $f(\eta)$ is defined. Let ω_1^η and $\tilde{\omega}_1^\eta$ be some fixed elements of Ω_1 such that $\varphi_\eta(\omega_1^\eta)$ and $\tilde{\varphi}_\eta(\tilde{\omega}_1^\eta)$, respectively. Then, the following equations hold true:*

$$\Pr[T \mapsto \eta] = \Pr_{\omega_1}[\varphi_\eta(\omega_1)] \cdot \Pr_{\omega_2}[T(\omega_1^\eta, \omega_2) \mapsto \eta] \quad (20)$$

$$\Pr[\tilde{T} \mapsto \eta] = \Pr_{\omega_1}[\tilde{\varphi}_\eta(\omega_1)] \cdot \Pr_{\omega_2}[\tilde{T}(\tilde{\omega}_1^\eta, \omega_2) \mapsto \eta] \quad (21)$$

$$\Pr_{\omega_2}[T(\omega_1^\eta, \omega_2) \mapsto \eta] = \Pr_{\omega_2}[\tilde{T}(\tilde{\omega}_1^\eta, \omega_2) \mapsto \eta] \quad (22)$$

Now, using Lemma 19, we can link the level of coercion-resistance VAV provides with the optimal bound δ_{CR}^S stated in Section 6.4. Clearly we have:

$$\begin{aligned} \Pr_{\omega_1}[\varphi_\eta(\omega_1)] &= \Pr_{\omega_1}[\rho(f(\eta), \omega_1) = \rho(\eta), \text{Rec}(\omega_1) = \text{Rec}(\eta)] \\ &= \Pr_{\omega_1}[\rho(f(\eta), \omega_1) = \rho(\eta)] \cdot \Pr_{\omega_1}[\text{Rec}(\omega_1) = \text{Rec}(\eta) \mid \rho(f(\eta), \omega_1) = \rho(\eta)] \\ &= A_{\rho(\eta)}^{f(\eta)} \cdot \Pr_{\omega_1}[\text{Rec}(\omega_1) = \text{Rec}(\eta) \mid \rho(f(\eta), \omega_1) = \rho(\eta)] \end{aligned}$$

and, similarly,

$$\Pr_{\omega_1}[\tilde{\varphi}_\eta(\omega_1)] = A_{\rho(\eta)}^{C(f(\eta), i)} \cdot \Pr_{\omega_1}[\text{Rec}(\omega_1) = \text{Rec}(\eta) \mid \rho(C(f(\eta), i), \omega_1) = \rho(\eta)].$$

Furthermore, it is easy to show that given two essential views with the same number of receipts of every type (and otherwise possibly different information on the bulletin board), the probability of obtaining a specific vector of receipts (which links receipts and voters) stays the same. From this it follows:

$$\begin{aligned} \Pr_{\omega_1}[\text{Rec}(\omega_1) = \text{Rec}(\eta) \mid \rho(f(\eta), \omega_1) = \rho(\eta)] &= \\ &= \Pr_{\omega_1}[\text{Rec}(\omega_1) = \text{Rec}(\eta) \mid \rho(C(f(\eta), i), \omega_1) = \rho(\eta)] \\ &= \Pr_{\omega_1}[\text{Rec}(\omega_1) = \text{Rec}(\eta) \mid \text{rec}(\omega_1) = \text{rec}(\eta)]. \end{aligned}$$

Together with Lemma 19, we immediately obtain for all η such that $f(\eta)$ is defined and for all ω_1^η with $\varphi_\eta(\omega_1^\eta)$:

$$\begin{aligned} \Pr[T \mapsto \eta] - \Pr[\tilde{T} \mapsto \eta] &= (A_{\rho(\eta)}^{f(\eta)} - A_{\rho(\eta)}^{C(f(\eta), i)}) \cdot \\ &\cdot \Pr_{\omega_2}[T(\omega_1^\eta, \omega_2) \mapsto \eta] \cdot \Pr_{\omega_1}[\text{Rec}(\omega_1) = \text{Rec}(\eta) \mid \text{rec}(\omega_1) = \text{rec}(\eta)]. \end{aligned}$$

Note that if there does not exist $\tilde{\omega}_1^\eta$ such that $\tilde{\varphi}_\eta(\tilde{\omega}_1^\eta)$, then $A_{\rho(\eta)}^{C(f(\eta), i)} = 0$ and $\Pr[\tilde{T} \mapsto \eta] = 0$.

Now, we are ready to prove that $S \in \{P_{VAV}^{s+}, P_{VAV}^{p+}\}$ is δ -coercion resistant w.r.t. \tilde{V}_i for $\delta = \delta_{CR}^S(n, k, \vec{p})$.

Let M be the set of views that are accepted by the program c of the coercer, i.e., for which the coercer outputs 1. In what follows, let Z range over the set of all possible patterns, ρ over all essential views, η over all views of the coercer, and Rec over all possible vectors of receipts. We abbreviate $C(Z, i)$ by $C(Z)$. Finally, let $M_Z^{\rho, \text{Rec}} = \{\eta \in M : f(\eta) = Z, \text{Rec}(\eta) = \text{Rec}, \text{ and } \rho(\eta) = \rho\}$. For Z, ρ, Rec with $M_Z^{\rho, \text{Rec}} \neq \emptyset$ let $\omega_1^{Z, \rho, \text{Rec}}$ be an arbitrary element of Ω_1 with $\rho(Z, \omega_1^{Z, \rho, \text{Rec}}) =$

ρ and $Rec(\omega_1^{Z,\rho,Rec}) = Rec$. Then we have $\varphi_\eta(\omega_1^{Z,\rho,Rec})$ for all $\eta \in M_Z^{\rho,Rec}$. We have:

$$\begin{aligned}
\Phi &= \Pr[T \mapsto 1] - \Pr[\tilde{T} \mapsto 1] \\
&= \Pr[T \mapsto M] - \Pr[\tilde{T} \mapsto M] \\
&= \sum_Z \sum_\rho \sum_{Rec} \sum_{\eta \in M_Z^{\rho,Rec}} (\Pr[T \mapsto \eta] - \Pr[\tilde{T} \mapsto \eta]) \\
&= \sum_Z \sum_\rho (A_\rho^Z - A_\rho^{C(Z)}) \sum_{Rec} \sum_{\eta \in M_Z^{\rho,Rec}} \Pr_{\omega_2}[T(\omega_1^{Z,\rho,Rec}, \omega_2) \mapsto \eta] \\
&\quad \cdot \Pr_{\omega_1}[Rec(\omega_1) = Rec \mid rec(\omega_1) = rec(\rho)].
\end{aligned}$$

For the third equation we use the fact that $\Pr[T \mapsto \eta] - \Pr[\tilde{T} \mapsto \eta] = 0$, if $f(\eta)$ is not defined, as in this case T and \tilde{T} coincide. With $M_Z^S = M_{Z,i}^S$ as defined in Section 6.4, we get:

$$\begin{aligned}
\Phi &\leq \sum_Z \sum_{\rho \in M_Z^S} (A_{\rho,Z}^S - A_{\rho,C(Z)}^S) \sum_{Rec} \sum_{\eta \in M_Z^{\rho,Rec}} \Pr_{\omega_2}[T(\omega_1^{Z,\rho,Rec}, \omega_2) \mapsto \eta] \\
&\quad \cdot \Pr_{\omega_1}[Rec(\omega_1) = Rec \mid rec(\omega_1) = rec(\rho)]
\end{aligned}$$

Now, by the definition of $M_Z^{\rho,Rec}$, for $\eta \in M_Z^{\rho,Rec}$, we have $f(\eta) = Z$ and, because $f(\eta)$ depends only on ω_2 , we know that $T(\omega_1^{Z,\rho,Rec}, \omega_2) \mapsto \eta$ implies $f(\omega_2) = Z$. Therefore, we have $\sum_{\eta \in M_Z^{\rho,Rec}} \Pr_{\omega_2}[T(\omega_1^{Z,\rho,Rec}, \omega_2) \mapsto \eta] \leq \Pr_{\omega_2}[f(\omega_2) = Z]$. With this, we obtain:

$$\Phi \leq \sum_Z \Pr_{\omega_2}[f(\omega_2) = Z] \sum_{\rho \in M_Z^S} (A_\rho^Z - A_\rho^{C(Z)}) \tag{23}$$

$$\sum_{Rec} \Pr_{\omega_1}[Rec(\omega_1) = Rec \mid rec(\omega_1) = rec(\rho)] \leq \tag{24}$$

$$\leq \sum_Z \Pr_{\omega_2}[f(\omega_2) = Z] \sum_{\rho \in M_Z^S} (A_{\rho,Z}^S - A_{\rho,C(Z)}^S) \tag{25}$$

$$\leq \sum_Z \Pr_{\omega_2}[f(\omega_2) = Z] \cdot \delta \tag{26}$$

$$\leq \delta. \tag{27}$$

This shows that S is δ -coercion resistant w.r.t. \tilde{V}_i . It remains to show that δ is optimal.

Let us consider the program c of the coercer which requests the coerced voter to vote using Z^* and accepts a view η only if $\rho(\eta)$ is in $M_{Z^*,i}^S$, where $M_{Z^*,i}^S$ is as defined in Section 6.4, and Z^* is a pattern with

$$\max_Z \sum_{\rho \in M_{Z^*,i}^S} (A_{\rho,Z}^S - A_{\rho,C(Z)}^S) = \sum_{\rho \in M_{Z^*,i}^S} (A_{\rho,Z^*}^S - A_{\rho,C(Z^*)}^S).$$

With this program c of the coercer we have, for each essential view ρ :

$$\Pr[T \mapsto \rho] = A_{\rho,Z^*}^S \quad \text{and} \quad \Pr[\tilde{T} \mapsto \rho] = A_{\rho,C(Z^*)}^S.$$

We immediately obtain:

$$\Phi = \sum_{\rho \in M_{Z^*,i}^S} (\Pr[T \mapsto \rho] - \Pr[\tilde{T} \mapsto \rho]) = \sum_{\rho \in M_{Z^*,i}^S} (A_{\rho,Z^*}^S - A_{\rho,C(Z^*)}^S) = \delta, \tag{28}$$

which shows that S is not δ' -coercion resistant for any $\delta' < \delta$, in case the counter-strategy \tilde{v} is used. To complete the proof of Theorem 5, we need to show that every other counter-strategy \tilde{v}' does not yield a smaller δ .

First, note that every reasonable counter-strategy \tilde{v}' should, up to a negligible set of runs, (a) cast ballots only when instructed by the coercer, (b) in case instructed by the coercer to cast a ballot, cast a ballot for candidate i , and (c) take the receipt requested by the coercer. Failing to meet (b) would mean that $\tilde{v}' \notin \tilde{V}_i$. Conversely, to guarantee that $\tilde{v}' \in \tilde{V}_i$, the coerced voter only needs to vote if instructed by the coercer. Therefore, it is clear that in order to be as indistinguishable from the dummy strategy as possible, a counter-strategy should only cast a ballot if instructed to do so by the coercer, which explains (a). As for (c), it is clear that if a counter-strategy takes a receipt different from the one requested by the coercer, the coercer can easily distinguish this strategy from the dummy strategy. Therefore, \tilde{v}' must be like \tilde{v} , up to the response in case it is instructed to vote for some $j \neq i$ according to a pattern where the coercer demands a V-receipt for candidate i . Let \mathcal{Z} be the set of such patterns. For $Z_0 \in \mathcal{Z}$, a counter-strategy may mark an arbitrary candidate on the remaining two ballots (one of these responses can be chosen randomly, according to some strategy). Recall from Section 6.4 that the response of the counter-strategy \tilde{v} , $C(Z_0)$, marks in this case candidate i also on the remaining ballots. Let $C'(Z) = C'(Z, i)$ denotes the response of the counter-strategy \tilde{v}' to a pattern Z . We also know that $C'(Z) = C(Z)$ for every $Z \notin \mathcal{Z}$. As just explained, for $Z_0 \in \mathcal{Z}$, $C'(Z_0)$ has k choices which could be chosen randomly. For simplicity of the argument, we assume that $C'(Z_0)$ always marks some candidate j on the remaining ballots; the case of a randomized choice can be treated similarly.

Let c be the program of the coercer which requests the coerced voter to vote using Z^* and accepts a view η only if $\rho(\eta)$ is in $\tilde{M}_{Z^*,i}^S$, where $\tilde{M}_{Z^*,i}^S = \{\rho : A_{\rho,Z}^S \geq A_{\rho,C'(Z^*)}^S\}$, and Z^* is a pattern such that

$$\max_Z \sum_{\rho \in \tilde{M}_{Z,i}^S} (A_{\rho,Z}^S - A_{\rho,C'(Z)}^S) = \sum_{\rho \in \tilde{M}_{Z^*,i}^S} (A_{\rho,Z^*}^S - A_{\rho,C'(Z^*)}^S) .$$

With this, analogously to (28), we have:

$$\tilde{\Phi} = \Pr[(c \parallel \text{dum} \parallel e_S) \mapsto 1] - \Pr[(c \parallel \tilde{v}' \parallel e_S) \mapsto 1] = \sum_{\rho \in \tilde{M}_{Z^*,i}^S} (A_{\rho,Z}^S - A_{\rho,C'(Z^*)}^S).$$

Hence it remains to show that

$$\max_Z \sum_{\rho \in \tilde{M}_{Z,i}^S} (A_{\rho,Z}^S - A_{\rho,C'(Z)}^S) \geq \max_Z \sum_{\rho \in \tilde{M}_{Z,i}^S} (A_{\rho,Z}^S - A_{\rho,C(Z)}^S) = \delta.$$

For $Z_0 \in \mathcal{Z}$, let Z'_0 be as Z_0 with the only difference that now the A-ballot for candidate i has to be taken as receipt. Then $C'(Z'_0)$ is uniquely determined and equal to $C(Z'_0)$. As the receipt of the coerced voter is *not* part of the essential view, we have, for all ω_1 :

$$\rho(Z'_0, \omega_1) = \rho(Z_0, \omega_1) \text{ and } \rho(C(Z'_0), \omega_1) = \rho(C(Z_0), \omega_1).$$

It follows

$$\sum_{\rho \in \tilde{M}_{Z'_0,i}^S} (A_{\rho,Z'_0}^S - A_{\rho,C'(Z'_0)}^S) = \sum_{\rho \in \tilde{M}_{Z_0,i}^S} (A_{\rho,Z_0}^S - A_{\rho,C(Z_0)}^S). \quad (29)$$

Now, we obtain:

$$\begin{aligned}
\max_Z \sum_{\rho \in \mathcal{M}_{Z,i}^S} (A_{\rho,Z}^S - A_{\rho,C'(Z)}^S) &\geq \max_{Z \notin \mathcal{Z}} \sum_{\rho \in \mathcal{M}_{Z,i}^S} (A_{\rho,Z}^S - A_{\rho,C'(Z)}^S) \\
&= \max_{Z \notin \mathcal{Z}} \sum_{\rho \in \mathcal{M}_{Z,i}^S} (A_{\rho,Z}^S - A_{\rho,C(Z)}^S) \\
&= \max_Z \sum_{\rho \in \mathcal{M}_{Z,i}^S} (A_{\rho,Z}^S - A_{\rho,C(Z)}^S),
\end{aligned}$$

where the last equation follows from (29). This concludes the proof of Theorem 5.

G Relationship between some constants

The following theorem states that the level of privacy of ThreeBallot and VAV in the case

$S \in \{P_{\text{TB}}^{o-}, P_{\text{TB}}^{p-}, P_{\text{VAV}}^{s-}, P_{\text{VAV}}^{p-}, P_{\text{VAV}}^{p+}\}$ is as good as the level of privacy of the ideal protocol (which only reveals the result of the election).

Theorem 8. *Let $S \in \{P_{\text{TB}}^{o-}, P_{\text{TB}}^{p-}, P_{\text{VAV}}^{s-}, P_{\text{VAV}}^{p-}, P_{\text{VAV}}^{p+}\}$. Then we have $\delta_{\text{Priv}}^S = \delta_{\text{Priv}}$, where δ_{Priv} is the level of privacy the ideal protocol achieves.*

Proof. Let $S = P_{\text{TB}}^{o-}$ and let $\rho = (n_x^-, n_x^+, n_x^0)$ be an essential view of S . By $f(\rho)$ we denote the pure result determined by ρ . Let $(N, R, r) = f(\rho)$. Note that N is the number of abstaining voters and R is the number of votes for candidate 1.

1. We show that $A_{\rho,2}^S \geq A_{\rho,1}^S \Leftrightarrow A_{f(\rho)}^2 \geq A_{f(\rho)}^1$ where $A_{f(\rho)}^i$ is as in Section 6.1. By Lemma 9 and Lemma 5, we have

$$\begin{aligned}
A_{\rho,1}^S &= \frac{2}{3}A_{\rho-\times\times\times}^S + \frac{1}{3}A_{\rho-\times\times\times}^S = \\
&= \binom{n}{N-1} \binom{N-1}{R-1} p_0^{n-N+1} p_1^{R-1} p_2^{N-R} \frac{2}{3} \binom{N-1}{n_x^- - 1} \left(\frac{2}{3}\right)^{n_x^- - 1} \left(\frac{1}{3}\right)^{N - n_x^-} + \\
&\quad + \binom{n}{N-1} \binom{N-1}{R-1} p_0^{n-N+1} p_1^{R-1} p_2^{N-R} \frac{1}{3} \binom{N-1}{n_x^+} \left(\frac{2}{3}\right)^{n_x^+} \left(\frac{1}{3}\right)^{N - n_x^+ - 1} \\
&= \binom{n}{N-1} \binom{N-1}{R-1} p_0^{n-N+1} p_1^{R-1} p_2^{N-R} \left(\frac{2}{3}\right)^{n_x^-} \left(\frac{1}{3}\right)^{N - n_x^-} \left(\binom{N-1}{n_x^- - 1} + \binom{N-1}{n_x^+} \right) = \\
&= \binom{n}{N-1} \binom{N-1}{R-1} p_0^{n-N+1} p_1^{R-1} p_2^{N-R} \binom{N}{n_x^-} \left(\frac{2}{3}\right)^{n_x^-} \left(\frac{1}{3}\right)^{N - n_x^-}
\end{aligned}$$

and

$$A_{\rho,2}^S = \binom{n}{N-1} \binom{N-1}{R} p_0^{n-N+1} p_1^R p_2^{N-R-1} \binom{N}{n_x^+} \left(\frac{2}{3}\right)^{n_x^+} \left(\frac{1}{3}\right)^{N - n_x^+}.$$

From this, we get

$$A_{\rho,2}^S - A_{\rho,1}^S = \binom{N}{n_x} \left(\frac{2}{3}\right)^{n_x} \left(\frac{1}{3}\right)^{N-n_x} \cdot \left(\binom{n}{N-1} \binom{N-1}{R-1} p_0^{n-N+1} p_1^{R-1} p_2^{N-R} - \binom{n}{N-1} \binom{N-1}{R} p_0^{n-N+1} p_1^R p_2^{N-R-1} \right) \geq 0$$

if and only if

$$\binom{n}{N-1} \binom{N-1}{R-1} p_0^{n-N+1} p_1^{R-1} p_2^{N-R} - \binom{n}{N-1} \binom{N-1}{R} p_0^{n-N+1} p_1^R p_2^{N-R-1} \geq 0,$$

i.e., $A_{f(\rho)}^2 \geq A_{f(\rho)}^1$.

2. Now, we show that for all essential views $\vec{r} = (N^*, R^*, r^*)$, we have $\sum_{\rho: f(\rho)=\vec{r}} A_{\rho,i}^S = A_{\vec{r}}^i$. Indeed, we have

$$\begin{aligned} \sum_{\rho: f(\rho)=\vec{r}} A_{\rho,1}^S &= \sum_{\rho: f(\rho)=\vec{r}} \binom{n}{N^*-1} \binom{N^*-1}{R^*-1} p_0^{n-N^*+1} p_1^{R^*-1} p_2^{N^*-R^*} \cdot \binom{N^*}{n_x} \left(\frac{2}{3}\right)^{n_x} \left(\frac{1}{3}\right)^{N^*-n_x} = \\ &= \binom{n}{N^*-1} \binom{N^*-1}{R^*-1} p_0^{n-N^*+1} p_1^{R^*-1} p_2^{N^*-R^*} \cdot \sum_{\rho: f(\rho)=\vec{r}} \binom{N^*}{n_x} \left(\frac{2}{3}\right)^{n_x} \left(\frac{1}{3}\right)^{N^*-n_x} = \\ &= \binom{n}{N^*-1} \binom{N^*-1}{R^*-1} p_0^{n-N^*+1} p_1^{R^*-1} p_2^{N^*-R^*} = A_{\vec{r}}^1. \end{aligned}$$

Here, we used that, for every $n_x \in \{0, \dots, N^*\}$, the equation $f(n_x, n_x, n_x) = \vec{r}$ has exactly one solution $n_x(n_x, N^*, R^*)$, $n_x^o(n_x, N^*, R^*)$, and it has no solutions with $n_x \notin \{0, \dots, N^*\}$ and, therefore

$$\sum_{\rho: f(\rho)=\vec{r}} \binom{N^*}{n_x} \left(\frac{2}{3}\right)^{n_x} \left(\frac{1}{3}\right)^{N^*-n_x} = \sum_{n_x=0}^{N^*} \binom{N^*}{n_x} \left(\frac{2}{3}\right)^{n_x} \left(\frac{1}{3}\right)^{N^*-n_x} = 1.$$

3. We conclude:

$$\begin{aligned} \sum_{\rho \in M_{j,i}^S} (A_{\rho,i}^S - A_{\rho,j}^S) &= \sum_{\rho: f(\rho) \in M_{j,i}^*} (A_{\rho,i}^S - A_{\rho,j}^S) = \\ &= \sum_{\vec{r} \in M_{j,i}^*} \sum_{\rho: f(\rho)=\vec{r}} (A_{\rho,i}^S - A_{\rho,j}^S) = \sum_{\vec{r} \in M_{j,i}^*} (A_{\vec{r}}^i - A_{\vec{r}}^j) \end{aligned}$$

By taking the maximum on both sides, the proof is done for $S = P_{\text{TB}}^{\circ-}$.

For $S \in \{P_{TB}^{p-}, P_{VAV}^{s-}, P_{VAV}^{p-}, P_{VAV}^{p+}\}$ the proofs are analogous, where, in step 1, we have to use the multinomial equivalent to

$$\binom{N-1}{\ell-1} + \binom{N-1}{\ell} = \binom{N}{\ell},$$

namely

$$\sum_{j=1}^k \binom{N-1}{\ell_1, \dots, \ell_j-1, \dots, \ell_k} = \binom{N}{\ell_1, \dots, \ell_k},$$

where $\ell_1 + \dots + \ell_k = N$.

□