

A Game-Based Definition of Coercion-Resistance and its Applications

Ralf Küsters Tomasz Truderung Andreas Vogt

University of Trier, Germany
{kuesters,truderun,vogt}@uni-trier.de

November 30, 2009

Abstract

Coercion-resistance is one of the most important and intricate security requirements for voting protocols. Several definitions of coercion-resistance have been proposed in the literature, both in cryptographic settings and more abstract, symbolic models. However, unlike symbolic approaches, only very few voting protocols have been rigorously analyzed within the cryptographic setting. A major obstacle is that existing cryptographic definitions of coercion-resistance tend to be complex and limited in scope: They are often tailored to specific classes of protocols or are too demanding.

In this paper, we therefore present a simple and intuitive cryptographic definition of coercion-resistance, in the style of game-based definitions. This definition allows to precisely measure the level of coercion-resistance a protocol provides. As the main technical contribution of this paper, we apply our definition to two voting systems, namely, the Bingo voting system and ThreeBallot. The results we obtain are out of the scope of existing approaches. We show that the Bingo voting system provides the same level of coercion-resistance as an ideal voting system. We also precisely measure the degradation of coercion-resistance of ThreeBallot in case the so-called short ballot assumption is not met and show that the level of coercion-resistance ThreeBallot provides is significantly lower than that of an ideal system, even in case of short ballots.

1 Introduction

Coercion-resistance is one of the most important and intricate security requirements for voting protocols [15, 20, 3]. Intuitively, a voting protocol is coercion-resistant if it prevents vote buying and voter coercion. Several definitions of coercion-resistance have been proposed in the literature (see, e.g., [15, 18, 7, 23, 11, 14, 10, 1, 17]), both based on cryptographic and symbolic models, where symbolic models take an idealized view on cryptography. However, in the cryptographic setting, only very few voting protocols have been analyzed rigorously w.r.t. coercion-resistance. A major obstacle

is that existing definitions tend to be complex and limited in scope: They are often tailored to a very specific class of protocols or are too demanding; some otherwise reasonable protocols are deemed insecure or can be shown to be secure only under stronger assumptions or using stronger cryptographic primitives (see Section 3 for more details). Even some relatively simple voting protocols are out of the scope of current cryptographic approaches. The situation is much better for symbolic approaches. Several quite complex voting protocols have been analyzed in this setting (see, e.g., [10, 1, 17]). For example, in [17] coercion-resistance of the Civitas voting system [8] was analyzed rigorously. However, being symbolic approaches, an idealized view on cryptography is taken and the level of coercion-resistance a protocol provides cannot be measured precisely.

Contribution of this paper. Inspired by a definition of coercion-resistance in the symbolic model [17], we present a definition of coercion-resistance in the cryptographic model, in the style of game-based definitions (rather than simulation-based definitions). The main idea is that a coercer should not be able to distinguish whether a coerced voter is following the instructions of the coerced voter (e.g., voting for a certain candidate) or whether the coerced voter is just trying to achieve her own goal (e.g., voting for her favorite candidate), by running a counter-strategy.

The resulting cryptographic definition has the following main features compared to other cryptographic definitions (see Section 3 for a detailed comparison with existing definitions and approaches): i) Our definition is simple and intuitive. ii) It allows to precisely measure the level of coercion-resistance a protocol provides. This quantitative approach is much preferable over approaches that provide only a yes/no answer: Not even an ideal voting protocol, which only reveals the result of the election, provides absolute coercion-resistance. Typically, the level of coercion-resistance depends on several parameters, including the number of voters and the number of choices voters have (e.g., the number of candidates in the election) as well as the probability distribution on these choices. iii) Our definition is applicable to a wide range of protocols, including protocols that are out of the scope of existing approaches, with less stringent security assumptions and weaker cryptographic primitives than some of the other approaches, as demonstrated by our case studies.

The case studies are the main technical contribution of our paper. Besides demonstrating the applicability of our definition, the results of our analysis are interesting in their own right as they constitute the first rigorous analyzes of the considered voting systems and introduce techniques that are applicable beyond our case studies.

We first provide a detailed analysis of an ideal voting system, which merely reveals the result of the election. The level of coercion-resistance such a system provides is a function of the number of honest voters in an election, the number of choices (e.g. candidates) voters have in the election, and a probability distribution on choices, which describes how honest voters vote. The analysis of the ideal voting system is a pure combinatorial argument. This analysis is motivated by the fact that the analysis of certain voting protocols, namely those that provide (almost) ideal coercion-

resistance, can often be divided into two parts: a combinatorial part corresponding to the ideal case and a cryptographic part. With the results presented in this paper, the combinatorial part does not have to be redone.

Based on the analysis of the ideal voting system, we show that the Bingo voting system [4], which has been used in practice [2], provides the same level of coercion-resistance as the ideal system (up to forced abstention attacks). This result is shown by a reduction to the ideal case, as explained above. It could not be obtained by previous approaches, as the Bingo voting system is either outside of the considered class of voting systems or, in case of simulation-based definitions, cannot be proven to be coercion-resistance, unless stronger security assumptions or more advanced cryptographic primitives are used (see Section 3 and 5 for more details).

We also provide a detailed analysis of the ThreeBallot voting system [21]. This system is known to leak partial information to a coercer. In particular, it is known that coercion-resistance cannot be obtained if the number of candidates in the election is high. In other words, coercion-resistance can at most be achieved under the so-called *short ballot assumption*. However, this assumption has so far not been defined or quantified within a formal framework. Using our definition, we rigorously measure the degradation of coercion-resistance as the number of candidates grows. Surprisingly, already with seven candidates and a few hundred voters the level of coercion-resistance ThreeBallot provides is insufficient. With ten candidates and two thousands of voters, ThreeBallot does virtually not provide any coercion-resistance. (Note that results of elections are often published per polling station and that a polling station often does not have more than a few hundred voters.) We also precisely analyze ThreeBallot in case the short ballot assumption is clearly met; more precisely, we consider the case of two candidates. Even in this case we find that the level of coercion-resistance ThreeBallot provides is significantly less than the ideal protocol. This analysis of ThreeBallot requires non-trivial combinatorial arguments as information can be leaked in subtle ways. As in case of the Bingo voting system, other approaches are unsuitable for the analysis of ThreeBallot (see Section 3 and 6 for more details).

Structure of this paper. In the following section, we present and discuss our definition of coercion-resistance. In Section 3 we provide a detailed comparison with other definitions and approaches. The analyzes of the three voting systems are carried out in Sections 4 to 6, with detailed proofs provided in the appendix.

2 Defining Coercion-Resistance

In this section, we present our definition of coercion-resistance. First, we introduce some notation and terminology.

2.1 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 $\ell > \ell_0$.

Our definition of coercion-resistance is based on a quite standard computational model, similar to models for simulation-based security (see, e.g., [5, 16]), in which *interactive Turing machines (ITMs)* communicate via tapes. The ITMs may perform probabilistic polynomial-time computations in the length of the security parameter and the input received so far. The details of the model are not essential for the rest of the paper. However, we fix some notation. A *system* \mathcal{S} of ITMs is a multi-set of ITMs, which we write as $\mathcal{S} = M_1 \parallel \dots \parallel M_l$, where M_1, \dots, M_l are ITMs. If \mathcal{S}_1 and \mathcal{S}_2 are systems of ITMs, then $\mathcal{S}_1 \parallel \mathcal{S}_2$ is a system of ITMs, assuming that the systems are connectible. As usual, the systems we consider are such that the length of a run is polynomially bounded in the length of the security parameter. Clearly, a run is uniquely determined by the random coins used by the ITMs in \mathcal{S} .

We assume that a system of ITMs has at most one ITM with a special output tape *decision*. For a system \mathcal{S} of ITMs and a security parameter ℓ , we write $\Pr[\mathcal{S}^{(\ell)} \mapsto 1]$ to denote the probability that \mathcal{S} outputs 1 (on tape decision) in a run with security parameter ℓ .

A *property* of a system \mathcal{S} is a subset of runs of \mathcal{S} . For a property γ of \mathcal{S} , we write $\Pr[\mathcal{S}^{(\ell)} \mapsto \gamma]$ to denote the probability that a run of \mathcal{S} , with security parameter ℓ , belongs to γ .

2.2 Voting Protocols

A *voting protocol* P specifies the programs (actions) carried out by honest voters and honest voting authorities, such as honest registration tellers, tallying tellers, bulletin boards, etc.

A voting protocol P , together with certain parameters, induces an *election system* $S = P(k, m, n, \vec{p})$. The parameters are as follows: k denotes the number of choices, an honest voter has in the election apart from abstaining from voting. In the simplest case, these choices can be the candidates the voter can vote for. Choices can also be preference lists of candidates, etc. In what follows, we often use the terms “candidate” and “choice” interchangeably. By m we denote the total number of voters and by n , with $n \leq m$, the number of honest voters. Honest voters follow the programs as specified in the protocol. The actions of dishonest voters and dishonest authorities are determined by the coercer, and hence, these participants can deviate from the protocol specification in arbitrary ways. We make the parameter n explicit since it is crucial for the level of coercion-resistance a system guarantees; intuitively the level of coercion-resistance increases with the number of honest voters. One can also think of n as the minimum number of voters the coercer may not corrupt. The vector $\vec{p} = p_0, \dots, p_k$ is a probability distribution on the possible choices, i.e., $p_0, \dots, p_k \in [0, 1]$ and $\sum_{i=0}^k p_i = 1$. Honest voters will abstain from voting with probability p_0 and vote for candidate i with probability p_i , $1 \leq i \leq k$. We make this distribution explicit, because it is realistic to

assume that the coercer knows this distribution (e.g., from opinion polls), and hence, uses it in his strategy, and because, as we will see later, the specific distribution is crucial for the level of coercion-resistance of a system.

An election system $S = P(k, m, n, \vec{p})$ specifies (sets of) ITMs for all participants, honest voters and authorities, the coercer, subsuming dishonest voters and dishonest authorities, and the coerced voter: (i) There are ITMs, say S_1, \dots, S_l , for all honest voting authorities. These ITMs run the programs as specified by the voting protocol. (ii) There is an ITM S_{v_i} , $i \in \{1, \dots, n\}$ for each of the honest voters. Every such ITM first makes a choice according to the probability distribution \vec{p} . Then, if the choice is not to abstain, it runs the program for honest voters according to the protocol specification with the candidate chosen before. (iii) The coercer is described by a set C_S of ITMs. This set contains all (probabilistic polynomial-time) ITMs, and hence, all possible coercion strategies the coercer can carry out. These ITMs are only constraint in their interface to the rest of the system. In particular, the ITMs can directly connect to the interface of dishonest voters and authorities. They can also communicate with the coerced voter. Moreover, they have access to all public information (e.g., bulletin boards) and possibly (certain parts of) the network. The precise interface of the ITMs in C_S depends on the specific protocol and the assumptions on the power of the coercer. iv) Similarly, the coerced voter is described by a set V_S of ITMs. Again, this set contains all (probabilistic polynomial-time) ITMs. This set represents all the possible programs the coercer can ask the coerced voter to run as well as all counter-strategies the coerced voter can run (see Section 2.3 for more explanation). The interface of these ITMs is typically the interface of an honest voter plus an interface for communication with the coercer. In particular, the set V_S 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. We note that a program in V_S can represent one coerced voter or a number of cooperating or independent coerced voters (see Section 2.3).

Given an election system $S = P(k, m, n, \vec{p})$, we denote by \mathbf{e}_S the system of ITMs containing all honest participants, i.e., $\mathbf{e}_S = (S_{v_1} \parallel \dots \parallel S_{v_n} \parallel S_1 \parallel \dots \parallel S_l)$, where, as explained above, $S_{v_1} \parallel \dots \parallel S_{v_n}$ are the ITMs modeling honest voters and $S_1 \parallel \dots \parallel S_l$ are the honest authorities. A system $(c \parallel v \parallel \mathbf{e}_S)$ of ITMs, with $c \in C_S$ and $v \in V_S$, is called an *instance of S* . We often implicitly assume a scheduler (modeled as an ITM) to be part of a system. Its role is to make sure that all components of the system are scheduled in a fair way, e.g., all voters get a chance to vote. For simplicity of notation, we do not state the scheduler explicitly. We define a *run of S* to be a run of some instance of S .

For an election system $S = P(k, m, n, \vec{p})$, 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 S , i.e., all other random bits used by honest voter 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 $\Pr_{\omega_2 \leftarrow \Omega_2}[\varphi]$.

A *property* of an election system $S = P(k, m, n, \vec{p})$ is defined to be a class γ of properties containing one property γ_T for each instance T of S . We will write $\Pr[T \mapsto \gamma]$ to denote the probability $\Pr[T \mapsto \gamma_T]$.

2.3 Coercion-Resistance

We can now present our definition of coercion-resistance, which, as already mentioned in the introduction is inspired by the symbolic definition of coercion-resistance in [17]. For now, we concentrate on the case that only a single voter is coerced. The case of multi-voter coercion-resistance is discussed later. In what follows let P be a voting protocol and $S = P(k, m, n, \vec{p})$ be an election system for P .

Our definition of coercion-resistance assumes that a coerced voter has a certain goal γ that she would try to achieve in absence of coercion. Formally, γ is a property of S . If, for example, 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 voted for this candidate and this vote is in fact counted. We note that in some cases such a goal cannot be achieved, e.g., in case ballots are sent over an unreliable channel or an election authority misbehaves in an observable way (e.g., fails to provide a valid proof of compliance) and as a result the election process is stopped. A more realistic goal γ would then be that the coerced voter successfully votes for a certain candidate, provided that the voters ballot is delivered in time and the election authority did not misbehave in an observable way.

In the definition of coercion-resistance 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` instead of the program an honest voter would run. As mentioned in Section 2.2, `dum` simply forwards all the messages between the coercer and the interface the coerced voter has as an honest voter. If the coerced voter 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; he can perform arbitrary coercion strategies: The coercer can send fake messages and depend his decisions on the information he has gathered so far. The intention of the coercer might even be to merely test whether the coerced voter follows his instructions, e.g., to find out whether this voter is “reliable”, and hence, is a good candidate for coercion in later elections. Also, the coercer is not necessarily bound to use the interface of the coerced voter in his coercion strategy. There may be other ways to vote on behalf of the coerced voter. However, for a protocol to be coercion-

resistance, there will always be at least one step in the protocol that the coercer cannot do all by himself, e.g., register, perform operations on a security token, or vote in a voting booth. For such actions, the coercer has to consult the coerced voter.

Now, for a protocol to be coercion-resistance our definition requires that there exists a *counter-strategy* \tilde{v} that the coerced voter can run instead of **dum** such that (i) the coerced voter achieves her own goal γ , with overwhelming probability, by running \tilde{v} and (ii) the coercer is not able to distinguish whether the coerced voter runs **dum** or \tilde{v} . More precisely, we will measure the ability of the coercer to distinguish between these two cases. Hence, \tilde{v} has to simulate **dum** while at the same time make sure that γ is achieved. 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 or threatening the voter, at least not from a technical point of view:¹ Even if the coerced voter tries to sell her vote, the coercer is not able to tell whether she is actually following the coercer’s instructions or just trying to achieve her own goal by running the counter-strategy. For the same reason, the coerced voter is safe, even if she wants to achieve her goal and therefore runs the counter-strategy.

Our formal definition of coercion-resistance is the following:

Definition 1. *Let P be a protocol and $S = P(k, m, n, \vec{p})$ be an election system. Let $\delta \in [0, 1]$, and γ be a property of S . The system S is δ -coercion-resistant w.r.t. γ , if there exists $\tilde{v} \in V_S$ such that for all $c \in C_S$ we have:*

- (i) $\Pr[(c \parallel \tilde{v} \parallel \mathbf{e}_S)^{(\ell)} \mapsto \gamma]$ is overwhelming, as a function of the security parameter.
- (ii) $\Pr[(c \parallel \mathbf{dum} \parallel \mathbf{e}_S)^{(\ell)} \mapsto 1] - \Pr[(c \parallel \tilde{v} \parallel \mathbf{e}_S)^{(\ell)} \mapsto 1]$ is δ -bounded, as a function of the security parameter.

Condition (i) says that by running the counter-strategy \tilde{v} the coerced voter achieves her goal with overwhelming probability, no matter which coercion-strategy the coercer performs.

Condition (ii) captures that the coercer is unable to distinguish whether the coerced voter run **dum** or \tilde{v} . More precisely, the coercer accepts a run (i.e., outputs 1 on tape **decision**) with almost the same probability no matter whether the coerced voter performs **dum** or \tilde{v} , where “almost the same” is formalized as δ -bounded, for some reasonably small δ (see below for more explanation). If the two probabilities are far apart, say for example, for some c , the probably of c accepting the run is 60% higher in case the coerced voter performs **dum**, this would give strong incentive to follow the instructions of the coercer, i.e., run **dum**: In case the coerced voter is threatened by the coercer, chances of being punished would be reduced significantly. In case the coerced voter wants to sell her vote, chances of being payed increase significantly.

In the rest of this section, we discuss further important aspects concerning the definition.

Negligible vs. δ -bounded. The reader might wonder why we require the difference in (ii) to be δ -bounded, rather than negligible. The reason is that negligibility is too

¹Of course, voters can be influenced psychologically.

strong. The difference, even for an ideal protocol, which merely reveals the result of the election, does not decrease with an increasing security parameter, but may depend on the number of choices, the distribution \vec{p} on these choices, and the number of honest voters: Imagine for example that a candidate did not get any vote in an election. Now, if the coercer asked the coerced voter to vote for this candidate, it is clear that the coerced voter did not follow the coercer’s instruction. The probability for this to happen is non-negligible and depends on \vec{p} and the number of voters; the larger the number of voters is, the more likely it is that a candidate gets a vote. In fact, in our examples (see Section 5 and 6), δ will depend on the number of candidates, \vec{p} , and the number of honest voters. Such a δ provides for a precise measure of the level of coercion-resistance, which is of practical relevance: It might, for example, indicate that a voting protocols does not have a sufficient level of coercion-resistance if the number of voters is below a certain threshold, the number of candidates is too big, or the probability distribution of the choices (e.g., according to opinion polls) is problematic in terms of coercion-resistance. These points will be illustrated in our examples.

Coercion strategies. In Definition 1, we assume that the coercer wants the coerced voter to run the dummy strategy **dum**. Alternatively, one could assume that the coercer wants the coerced voter to run some arbitrary coercion strategy $v \in V_S$. Then, one would demand that for every coercion strategy $v \in V_S$, there exists a counter-strategy v' such that (i) and (ii) are satisfied (with **dum** replaced by v and \tilde{v} replace by v'). However, it is easy to see that this formulation of coercion-resistance is not stronger than Definition 1: Intuitively, the reason is that the coercer can run v himself. More precisely, if there exists a counter-strategy \tilde{v} for **dum**, then it is easy to define a counter-strategy v' for a coercion strategy v , namely $v' = (v \parallel \tilde{v})$. Clearly, with this counter-strategy, (i) is satisfied, since for every $c \in C_S$, the system $(c \parallel (v \parallel \tilde{v}) \parallel e_S)$ behaves exactly as the system $((c \parallel v) \parallel \tilde{v} \parallel e_S)$ and $(c \parallel v)$ can be seen as a coercer $c' \in C_S$. By definition of \tilde{v} , we know that $\Pr[(c' \parallel \tilde{v} \parallel e_S)^{(\ell)} \mapsto \gamma]$ is overwhelming, as a function of the security parameter. Condition (ii) is satisfied as well, following a similar reasoning: The system $S_1 = (c \parallel v \parallel e_S)$ behave exactly the same as $S'_1 = ((c \parallel v) \parallel \mathbf{dum} \parallel e_S)$, since **dum** merely forwards messages. Moreover, the system $S_2 = (c \parallel (v \parallel \tilde{v}) \parallel e_S)$ is equivalent to $S'_2 = ((c \parallel v) \parallel \tilde{v} \parallel e_S)$. Now, as above, $(c \parallel v)$ can be considered to be a coercer $c' \in C_S$ and by definition of \tilde{v} , we know that $\Pr[S'_1 \mapsto 1] - \Pr[S'_2 \mapsto 1]$ is δ -bounded, as a function of the security parameter k , and hence, this is true for $\Pr[S_1 \mapsto 1] - \Pr[S_2 \mapsto 1]$.

We use Definition 1 since it simplifies proofs. Also, \tilde{v} is the strongest counter-strategy in that it can be used to construct counter-strategies for all coercion strategies (as shown above). Therefore, \tilde{v} should in fact be part of the protocol specification.

Specific voter goals. We have already pointed out that the flexibility of defining the voter goal γ is important to make reasonable statements about practical voting protocols. We illustrate this flexibility by another example: As already explained, in

elections where the probability for one candidate, say A , to get a vote is very low, the level of coercion-resistance can be quite low, i.e., δ can be quite big, because the coercer can tell the coerced voter to vote for A . Even in an ideal voting protocol the coerced voter has not much choice in such a situation than to vote for A . However, if there are two other candidates, B and C , say, with reasonably high probabilities, and the goal of the coerced voter is to vote for C , then γ could be defined as: If the coercer asks the coerced voter to vote for B (and the coerced voter can tell that this is the case), then the coerced voter votes for C . For such a (weakened) goal, δ would be smaller, saying that the level of coercion-resistance is high in case the coercer wants the coerced voter to vote for a candidate with high probability and the favorite candidate of the coerced voter is reasonably high as well.

Class of voter goals γ . Definition 1 is formulated w.r.t. a single goal γ the coerced voter tries to achieve. This can easily be generalized to a class of goals: A protocol is coercion-resistant for such a class if it is coercion-resistant for all goals in that class. The goals a coerced voter should be able to achieve should be all goals an honest voter, not under coercion, typically can achieve, e.g, vote for a certain candidate or abstain from voting.

Multi-voter coercion. We have so far focused on the case where only one voter is coerced. In reality a coercer can coerce many voters. From the point of view of a single coerced voter, say Alice, the behavior of other coerced voters may deviate in arbitrary ways from the prescribed protocol. Alice should be able to resist coercion, independently of the other coerced voters, whom Alice might not know anyway, and independently of their behavior. However, this is already captured by Definition 1 since other coerced voters can simply be considered to be dishonest voters, and hence, they are subsumed by the coercer. This makes the coercer only more powerful, since now he even fully dictates the behavior of other coerced voters in the coercion of Alice.

Conversely, coerced voters might want to team up, e.g., to have better chances to sell their votes. This can also be modeled using Definition 1 since \mathbf{dum} and \tilde{v} can represent a set of coerced voters. So, \mathbf{dum} would be a parallel composition of single dummy strategies, one for every coerced voter, and \tilde{v} would be either a joint counter-strategy or a parallel composition of independent counter-strategies.

3 Comparison with Other Definitions

One of the first rigorous definitions of coercion-resistance was presented by Juels et al. [15]. They defined coercion-resistance relative to an ideal system. However, their definition is tailored towards voting in a public-key setting, with protocols having a specific structure. In particular, neither the Bingo voting system nor ThreeBallot fall into the class of protocols considered by Juels et al. Conversely, the voting protocol proposed by Juels et al., and also the Civitas system [8] which generalizes the protocol by Juels et al., falls in our framework.

A rather general definition of coercion-resistance within the simulation-based approach was presented by Moran and Naor [18], based on a definition of coercion-resistance for multi-party computation by Canetti and Gennaro [6]. In this approach, a protocol is considered to be coercion-resistant, if it realizes an ideal voting functionality. The advantage of such definitions, compared to game-based definitions considered here, is that they provide composability by construction (see also [24]). However, this comes with a price: Some reasonable voting protocols cannot be proven secure due to the so-called commitment problem. This is, for example, the case for the Bingo voting system (see Section 5 for details). Other protocols are equipped with more advanced cryptographic primitives only in order for the security proofs in the simulation-based setting to go through (see, e.g., the split-ballot protocol [19]). Even if the commitment problem does not occur, the simulation-based definition is often too strong: It gives a yes/no answer—the difference between the ideal and real system is negligible or not—instead of measuring the level of coercion-resistance (as we do in our definition). Indeed for many protocols, such as paper-based protocols, the difference between a real and ideal system is non-negligible, but still reasonably small: For example, in some paper-based protocols there is a certain probability that a single fake ballot can be produced without being noticed (since, e.g., only partial auditing is done). If a coerced voter gets such a fake ballot, her vote might be revealed. However, the probability that a fixed coerced voter gets the fake ballot might be small (but non-negligible), e.g., approximately $\frac{1}{n}$, where n is the number of voters. Hence, the coercion level is increased by $\frac{1}{n}$, i.e., in our definition, δ is increased by $\frac{1}{n}$. This could be considered to be reasonably small, but is not captured by a yes/no answer as given in the simulation-based definition. In the simulation-based definition, one could replace negligibility by δ -boundedness. Unfortunately, in the definition of Moran and Naor δ might be quite big because the environment knows how honest voters vote, and hence, in situations like the above, it can tell with high probability whether it deals with the real or ideal system. So, replacing “negligible” by “ δ -boundedness” in the simulation-based definition often does not yield satisfactory results. In [19], Moran and Naor proposed and analyzed the paper-based voting protocol split-ballot which, in fact, is not perfect due to fake ballots. In this work, they indeed do not opt for δ -boundedness or the like, but change the ideal functionality. This approach can be problematic since it might not be clear whether the resulting functionality can be considered ideal. In particular, in their “ideal” functionality, they allow the adversary to *retroactively* change the votes of corrupted voters as a *function of the tally*, where the difference to the original tally is only bounded by the security parameter. For other “imperfect” protocols, such as ThreeBallot (see Section 6), finding a reasonable ideal functionality which is not too close to the protocol itself can be very challenging.

In [24], Unruh and Müller-Quade generalize the simulation-based framework of [18] and [6] for coercion-resistance. This paper was submitted to the Cryptology ePrint Archive only very recently. Independent of our work, this paper also presents a game-based definition of coercion-resistance which is similar to our definition. However, this

definition is not further applied, except as a means to illustrate the simulation-based framework: It is shown that their simulation-based definition implies their game-based definition. As such, our work is complementary to the work by Unruh and Müller-Quade. Our work directly uses our game-based definition, since, as explained above, the simulation-based approaches often cannot be used to analyze existing protocols. Also, while a main contribution of our work is the application of our game-based definition to existing voting protocols, Unruh and Müller-Quade do not apply their framework to published protocols. Finally, the game-based definition of Unruh and Müller-Quade misses two important parameters, which are crucial in the analysis of practical voting protocols: (i) While we have a parameter γ for the goal of a coerced voter, they fix a specific goal, requiring that the coerced voter has to vote for a specific candidate. As explained in Section 2.3, such a goal is too strong for most practical protocols (e.g., in presence of network delays or observable misbehavior). (ii) While we have a parameter δ for specifying the level of coercion-resistance, they fix δ to be the level of coercion-resistance an ideal protocol guarantees plus a negligible function. As argued before, many reasonable protocols, such as some paper-based protocols, do not achieve this level of coercion-resistance.

Teague et al. [23] proposed a definition of coercion-resistance which takes a quantitative approach. However, this definition has the following limitations: (i) It is intended to be used for ideal protocols, combined, as the authors suggest, with a simulation-based definition. (ii) The coercer may only use a specific strategy to decide whether to punish the coerced voter or not. Also, the class of counter-strategies available to the coerced voter is limited. (iii) Only the probability that a cheating voter gets punished is considered, ignoring the possibility that a voter might try to sell her vote by following the instructions of the coercer.

A recent definition of coercion-resistance by Gardner et al. [11] is specifically tailored to the protocol considered by the authors. It also considers only a very restricted part of an election process, denying, for example, the coercer access to information in the tallying phase. In particular, the Bingo voting system and ThreeBallot are not in the scope of this definition.

As already mentioned in the introduction, several definitions of coercion-resistance were proposed in symbolic models (see, e.g., [10, 1, 17]), where, as mentioned, our game-based definition is inspired by the definition in [17].

4 Analyzing the Ideal Protocol

In this section, we analyze an ideal voting protocol and precisely establish the level of coercion-resistance this protocol provides. More precisely, we determine the optimal (i.e. minimal) constant δ_{min} for which the ideal protocol is coercion-resistant. In particular, no real protocol can be δ -coercion-resistant for any $\delta < \delta_{min}$. As already explained in the introduction, the results of this section are motivated by the fact that the analysis of real voting protocols can often be reduced to the ideal case (see Section 5 for an example).

We consider here the most common tallying function. It returns the number of votes each candidate gets. This kind of tallying function is used in the protocols that we analyze in the following sections. We note that the level of coercion-resistance depends on the tallying functions used in an election.

The ideal protocol. In the ideal protocol, denoted by P_{ideal} , a voter sends her choice directly to the fully trusted election process. The election process properly counts the votes and outputs the result, without revealing any additional information. Here we consider a result to be a $(k + 1)$ -tuple indicating the number of abstaining voters and the number of votes each of the k candidates got.

More precisely, let $S = P_{\text{ideal}}(k, m, n, \vec{p})$ denote the election system defined as follows. The system S contains exactly one voting authority. The program of an honest voter randomly picks a choice according to the distribution \vec{p} and either abstains from voting or, in the other case, sends the chosen candidate on some untappable channel to the voting authority. In particular, only the voter and the voting authority know whether the voter abstained and, if the voter did not abstain, the chosen candidate, unless the voter is dishonest. The program of the voting authority simply collects the votes received on the untappable channels from the voters (one vote for each voter) and then outputs the result of the election.

The coercer completely controls the dishonest voters and can also send messages to the coerced voter. In fact, by definition of the ideal protocol, the only reasonable message the coercer can send to the coerced voter and on the untappable channels of the dishonest voters are the desired candidates; everything else would be ignored by the voting authority. Since the protocol does not output messages to voters, the coercer does not expect to receive messages either. Hence, the view of the coercer merely consists of his own random coins and the result of the election.

A coerced voter, running the dummy strategy or emulating it by running the counter-strategy, can receive a message from the coercer and send her choice on the untappable channel to the voting authority.

Goals of the coerced voter. We will consider goals γ_i of the coerced voter, for $i \in \{1, \dots, k\}$, defined as follows: γ_i is satisfied in a run, if whenever the coerced voter has sent her candidate to the voting authority, she has successfully voted for the i -th candidate. This implies that if the coerced voter is not instructed by the coercer to vote, i.e., the coercer does not send his candidate to the coerced voter, and hence, effectively wants the coerced voter to abstain from voting, the coerced voter does not have to vote in order to fulfill γ_i . In other words, by γ_i abstention attacks are not prevented.

Alternatively, we could consider a stronger and simpler goal γ'_i which requires the coerced voter to vote for i , even if the coercer wants the coerced voter to abstain. In fact, for this goal we obtain very similar results. However, γ'_i is too strong for most practical protocols, including the ones we consider in this paper. For reasons of uniformity, we therefore restrict ourselves to the goal γ_i .

We also note that, for the ideal protocol, we could consider abstention to be a goal of the coerced voter. But again, this goal cannot be achieved in most of practical protocols in which a voter is given a receipt, as such receipts can be used by the coercer to verify that the voter has actually voted.

The optimal constant δ_{min} . Now, we establish the optimal constant δ_{min} mentioned above. As this constant depends on the number of candidates, on the number of honest voters, and the probability distribution \vec{p} , we will denote it by $\delta_{min}(k, n, \vec{p})$.

This constant will be achieved if the counter-strategy \tilde{v} of the coerced voter is as follows: If the coerced voter receives a candidate from the coercer, then the coerced voter sends the i -th candidate to the voting authority; otherwise, she abstains from voting. Clearly, this counter-strategy guarantees that γ_i is met.

To determine δ_{min} , we need some terminology and notation.

Since the coercer knows the votes of dishonest voters, he can simply subtract these votes from the final result and obtain what we will call the *pure result* of the election. The pure results only depends on the votes by the n honest voters and the coerced voter. Hence, a pure result is a tuple $\vec{r} = (r_0, \dots, r_k)$ of non-negative integers such that $r_0 + \dots + r_k = n + 1$, where r_i , for $i \in \{1, \dots, k\}$, is the number of votes for the i -th candidate and r_0 denotes the number of voters who abstained from voting. As already mentioned above, the coercer has to base his decision—accept or reject—solely on such a pure result \vec{r} . We will denote the set of pure results by Res .

In the definition of $\delta_{min}(k, n, \vec{p})$, we will use the probability $A_{\vec{r}}^i$ 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. 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 $\binom{n}{m_0, \dots, m_k} = \frac{n!}{m_0! \dots m_k!}$.

The intuition behind the definition of $\delta_{min}(k, n, \vec{p})$ is the following: If the coercer wants the coerced voter to vote for j and the coerced voter wants to vote for i , for some $i, j \in \{1, \dots, k\}$, then, as we will show, the best strategy of the coercer to distinguish whether the coerced voter has voted for j or i is to accept a run if the pure result \vec{r} of the election in this run is such that $A_{\vec{r}}^i \leq A_{\vec{r}}^j$. Let $M_{i,j}^* = \{\vec{r} \in Res : A_{\vec{r}}^i \leq A_{\vec{r}}^j\}$ be the set of those results, for which – according to his best strategy – the coercer should accept the run.

The following lemma yields a convenient and intuitive characterization of the set $M_{i,j}^*$. It says that a result should be accepted by the coercer iff the actual ratio $\frac{r_j}{r_i}$ of the number of votes for j to the number of votes for i is bigger than the expected ratio $\frac{p_j}{p_i}$.

Lemma 1. $A_{\vec{r}}^i \leq A_{\vec{r}}^j$ iff $\frac{r_j}{r_i} \geq \frac{p_j}{p_i}$, and therefore $M_{i,j}^* = \{\vec{r} \in Res : \frac{r_j}{r_i} \geq \frac{p_j}{p_i}\}$.




Figure 1: Level of coercion-resistance (δ) for the ideal protocol. The goal of the coerced voter is, in each case, to vote for candidate 1.

Proof. We have the following equation:

$$A_r^j - A_r^i = \frac{n!}{r_0! \cdots r_k!} \cdot p_0^{r_0} \cdots p_k^{r_k} \cdot \left(\frac{r_j}{p_j} - \frac{r_i}{p_i} \right).$$

This term is ≥ 0 if and only if $\frac{r_j}{r_i} \geq \frac{p_j}{p_i}$. \square

Now, we are ready to define the constant δ_{min}^i , which we will show to be optimal:

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

In the definition of this constant, we take into account all possible candidates $1, \dots, k$ that the coercer can wish the coerced voter to vote for, excluding abstention, as in this case the counter-strategy coincides with the dummy strategy. We take the worst possible case, i.e., the index j for which the sum in the expression above is maximal.

The following theorem shows that δ_{min} is indeed optimal (see Appendix A).

Theorem 1. *Let $S = P_{ideal}(k, m, n, \vec{p})$. Then S is δ -coercion-resistant with respect to γ_i , where $\delta = \delta_{min}^i(n, k, \vec{p})$. Moreover, S is not δ' -coercion-resistant for any $\delta' < \delta$.*

In Figure 1, we depict values of $\delta = \delta_{min}$ for some selected cases. These values illustrate that the level of coercion-resistance heavily depends on the number of honest voters, the number of candidates, and the probability distribution on the choices. Note that even for national elections, it is realistic to assume that the number of voters is small, since results are often published per polling station and the number of voters who voted in one polling station is often not more than a few hundred.

The following example illustrates differences in the level of coercion-resistance depending on the parameters.

Example 1. Consider two elections that use the ideal protocol. In both cases, we assume that the goal of the coerced voter is γ_1 (to vote for 1) and that the coercer is willing to pay \$50 to a coerced voter if (using some decision procedure) he decides that the voter followed his instructions.

In the first election, we assume 2000 honest voters, two candidates, and probabilities $p_0 = 0.3$, $p_1 = 0.35$, $p_2 = 0.35$ that an honest voter abstains from voting, chooses candidate 1, or chooses candidate 2, respectively. By Theorem 1 we know that this system is (0.021)-coercion-resistant w.r.t. γ_1 . This means that if the coerced voter runs her counter-strategy to vote for her own candidate, then she will be paid—in the worst possible case—with probability only 2.1% less, and thus will earn, on average, \$1.05 less, compared to the case when she follows the instructions of the coercer. Hence, in this case, the coerced voter has only little incentive to follow the instructions of the voter. Conversely, by running the counter-strategy, chances of being accused of not following the instructions of the coercer are not much bigger than in case the coerced voter would actually follow the coercer’s instructions.

In the second election, we take 100 honest voters, five candidates, and probabilities $p_0 = 0.3$, $p_1 = 0.2$, $p_2 = 0.05$, and $p_3 = p_4 = p_5 = 0.15$. In this case the system is only (0.198)-coercion resistant w.r.t. γ_1 , which means that the coerced voter earns, on average, \$9.9 less when she runs her counter-strategy, which might give sufficient incentives to follow the instructions of the coercer. Also, chances of being accused of not following the coercer’s instructions are now much higher when running the counter-strategy compared to following the coercer’s instructions.

5 Analyzing Bingo Voting

In this section, we analyze the Bingo voting system [4]. We prove that this system enjoys the same level of coercion-resistance as the ideal protocol, except for forced-abstention attacks.

5.1 Description of the System

We describe the Bingo Voting system, which we will denote by P_{Bingo} .

In addition to the voters, the participants in this system are the following: (i) A *voting machine*, which is the main component in the voting process. (ii) A *trusted random number generator (RNG)*, which is an independent source of randomness, with its own display, and which is connected to the voting machine. (iii) A *bulletin board*. iv) Some number of *auditors* who will contribute randomness in a distributed way used for randomized partial checking (RPC) in zero-knowledge proofs provided by the voting machine. While in our analysis we concentrate on the case of one voting machine, the analysis easily carries over to the case of several voting machines, as they are independent.

The election consists of three phases described below: initialization, voting, and tallying.

Initialization phase. In this phase, the voting machine, for every candidate j , generates n random numbers x_1^j, \dots, x_n^j (where n is the number of voters), along with an unconditionally hiding commitment $\text{comm}(j, x_l^j)$ for each pair (j, x_l^j) ; more precisely, Pedersen commitments are used. All commitments are then shuffled and published on the bulletin board. Moreover, zero-knowledge proofs are published to guarantee that the same number n of commitments is created for every candidate (see below for details).

Voting phase. In this phase, a voter enters the voting booth to indicate the candidate of her choice, say j , to the voting machine, by pressing a button corresponding to j . Note that a voter can of course also abstain from voting. Then, the RNG creates a fresh random number which is displayed to the voter and transferred to the voting machine. The machine then prints a receipt consisting of the candidate names along with the following numbers next to them: The number next to the chosen candidate is the fresh random number, where the voter is expected to check that this number is the same as the one displayed by the RNG. Next to the other candidate names the machine prints a so far unused number x_l^j , for some l .

Tallying phase. In this phase, the voting machine publishes the result of the election as well as all the receipts given to voters (in a lexicographical order). The machine also opens the commitments to all pairs (j, x_l^j) where the number x_l^j is unused, i.e., x_l^j has not been printed on any receipt.

Moreover, the machine provides zero-knowledge proofs to show that (i) for each unopened commitment on a pair of the form (j, x_l^j) , the number x_l^j occurs on exactly one receipt, and (ii) every receipt contains $(k - 1)$ numbers x_l^j for distinct candidates j . (The k -th number is the one provided by the RNG.) These zero-knowledge proofs are described below.

The zero-knowledge proofs are checked as described below. If they are valid, every observer can verify the correctness of the result: the number of votes for candidate j should be the number of opened commitments of the form $\text{comm}(j, x_l^j)$, for some x_l^j , minus the number of abstaining voters.

Zero-knowledge proofs. Now, we describe the zero-knowledge proofs used both in the tallying phase and the initialization phase.

ZK-proofs in the tallying phase. To prove conditions (i) and (ii) in the tallying phase, the following steps are performed for every receipt: First, the voting machine generates a new commitment on the pair (j, r) , where j is the chosen candidate and r is the number generated by the RNG and printed next to j . Then, all the commitments for the receipt are published: one of them is the commitment just described, the other $(k - 1)$ commitments are unopened commitments published on the bulletin board in the initialization phase, where for different receipts, different commitments are taken from the bulletin board. An observer can verify that this is the case. Next, these

commitments are re-randomized and shuffled twice; both the intermediate and the final set of commitments are published. The final commitments are opened. Now an observer can check that there is exactly one commitment for each candidate. Finally, the auditors choose a random bit in some distributed way. Depending on the value of this bit, the voting machine publishes the random factors for the first or for the second re-randomization step.

If the voting machine would try to cheat, this would be detected with a probability of 50%; this probability can be increased by repeating the procedure.

ZK-proofs in the initialization phase. This proof was not precisely defined in [4], but it can be implemented by randomized partial checking similarly to the zero-knowledge proof in the tallying phase. To this end, we assume that a commitment $\text{comm}(j, x_l^j)$ on a pair (j, x_l^j) is implemented as the pair $(C_{jl}, D_l^j) = (\text{comm}(j), \text{comm}(x_l^j))$, where the commitments on the single components are Pederson commitments. Now, to show that among the published commitments there are exactly n of the form $\text{comm}(j, x_l^j)$ for every candidate j , the zero-knowledge proof proceeds similarly as in the tallying phase, except that it only uses the first component C_{jl} of a commitment.

5.2 Modeling and Security Assumptions

The modeling of the Bingo voting system as an election system $S = \text{P}_{\text{Bingo}}(k, m, n, \vec{p})$ is straightforward. However, we highlight some modeling issues, and most importantly, our security assumptions.

Voting authorities. We assume that the voting machine and the random number generator are honest; the bulletin board may be dishonest. This assumption is necessary for the Bingo voting system to be coercion-resistant. (However, for voter verifiability the voting machine does not need to be honest.) We also assume that at least one of the auditors is honest; all others may be dishonest.

Honest voters. As usual, an honest voter first makes a choice according to the probability distribution \vec{p} . If the choice is to abstain from voting, she abstains, otherwise follows the procedure described for the voting phase. After the voting phase is finished, an honest voter reveals her (paper) receipt, e.g., mails it to an organization to ask it to verify the correctness of the voting process w.r.t. her receipt or to publish it on some bulletin board. In particular, the coercer will get to see all receipts of honest voters, and hence, knows whether a voter voted or not. The assumption that the paper receipts are revealed after the voting phase is reasonable. Also, the (presumably small) fraction of honest voters for which the coercer manages to get hold of the receipt earlier, could be considered to be dishonest. In any case, the assumption helps in the proof and we believe that our results also hold without that assumption.

The coerced voter. A coerced voter, running the dummy strategy or emulating it by running a counter-strategy, can communicate with the coercer and send her candidate on an untappable channel to the voting authority.

The coercer. The coercer can freely communicate with dishonest participants (voters and authorities) as well as with the coerced voter; in fact, dishonest participants are considered to be part of the coercer program. In a run of the system the coercer can see the following: (v1) his random coins, (v2) all messages published by the voting machine, both in the initialization phase and the tallying phase, (v3) receipts of all honest voters, as already explained above and (v4) the messages received from the coerced voter (and dishonest parties), including the receipt of the coerced voter. However, the coercer cannot directly see the information the coerced voter obtains in the voting booth. In particular, the coerced voter can lie about what she sees and does in the voting booth, such as the random number shown by the RNG or the candidate she picked. So, while talking with the coercer on the phone would be allowed, taking pictures or videos should be prohibited (unless they can be manipulated on-the-fly, which, however, is unrealistic).

5.3 Coercion-Resistance of the System

We now show that the Bingo voting system enjoys the same level of coercion-resistance as the ideal protocol. However, since we assume that the coercer can see all receipts of voters who voted, the coerced voter can be forced to abstain from voting. Hence, the coerced voter can only achieve goal γ_i , but not γ'_i (see Section 4).

More precisely, the goal γ_i , $i \in \{1, \dots, k\}$, is satisfied in a run, if whenever the coerced voter has indicated her candidate to the voting machine, she has successfully voted for the i -th candidate.

We prove the following theorem:

Theorem 2. *Let $S = \mathsf{P}_{\text{Bingo}}(k, m, n, \vec{p})$. Then S is δ -coercion-resistant with respect to γ_i , where $\delta = \delta_{\min}^i(n, k, \vec{p})$.*

As already mentioned in Section 3, other approaches are unsuitable for the analysis of the Bingo voting system. We note that the simulation-based definitions [18, 24] cannot be applied due to the commitment problem. However, they might be applicable if we weakened the security assumptions, assuming that all the auditors are honest. In this case a simulator can simulate these auditors, which allows it to fake the zero-knowledge proofs in the tallying phase, as it “knows” the challenges. Another alternative could be to consider more advanced commitments, as, e.g., in [19]. The game-based definition in [24] could be adapted to deal with the Bingo voting system (see also Section 3). However, the simulation-based approach taken in [24] to prove coercion-resistance would, as explained, not work.

The remainder of this section is devoted to the proof of Theorem 2. First, we define the counter-strategy \tilde{v} of the coerced voter: \tilde{v} coincides with the dummy strategy dum , with the following exceptions:

1. \tilde{v} votes for candidate i , i.e., the coerced voter presses the button for candidate i , if the coercer instructs the coerced voter to vote for some candidate j .
2. If dum would forward the number that is shown on the display of the random number generator to the coercer, \tilde{v} forwards the number next to the candidate j , as shown on her receipt.

It is easy to see that if the coerced voter runs the counter-strategy \tilde{v} , then condition (i) of Definition 1 is satisfied for every $c \in C_S$. Note that if the coercer does not instruct the coerced voter to vote for some candidate j (abstention attack), then following the counter-strategy the coerced voter abstains from voting, which is in accordance with γ_i .

It remains to prove condition (ii) of Definition 1. For this purpose, let us fix a program c of the coercer. We need to prove that $\Pr[T \mapsto 1] - \Pr[\tilde{T} \mapsto 1] \leq \delta$, where $T = (\text{dum} \parallel c \parallel \mathbf{e}_S)$ and $\tilde{T} = (\tilde{v} \parallel c \parallel \mathbf{e}_S)$. The rest of the proof consists of the two parts already mentioned in the introduction, a cryptographic and a combinatorial part. The cryptographic part is Lemma 2. Using Lemma 2, the combinatorial part is merely a reduction to the ideal case, as studied in the previous section; it does not have to be redone.

As introduced in Section 2.2, by $\omega_1 \in \Omega_1$ we denote a vector of choices made by the honest voters and by $\omega_2 \in \Omega_2$ we denote all the remaining random coins of a system. We denote by ρ a view of the coercer, as described in Section 5.2, (v1)–(v4). We use the notion of a *pure result* $\vec{r} = (r_0, \dots, r_k)$ as introduced in Section 4. In particular, it holds that $r_0 + \dots + r_k = n + 1$ and the coercer can compute this result from his view, by subtracting the votes of dishonest voters from the result of the election. We will denote the pure result determined by a view ρ of the coercer by $\text{res}(\rho)$. A pure result determined by ω_1 and the choice j of the coerced voter will be denoted by $\text{res}(\omega_1, j)$.

As mentioned before, the coercer can derive from his view which voters abstained from voting. Given a view ρ of the coercer, we denote by $\text{abst}(\rho)$ the set of voters who abstained from voting, among the honest voters and the coerced voter; the number of such voters is referred to by $r_0(\rho) = |\text{abst}(\rho)|$. As this set/number depends only on ω_1 , we will sometimes write $\text{abst}(\omega_1)/r_0(\omega_1)$.

For a coercer view ρ in a run of the system, we denote by $f(\rho)$ the candidate the coercer wants the coerced voter to vote for; if the coercer does not instruct the coerced voter to vote, then $f(\rho)$ is undefined. Note that the coercer has to provide the coerced voter with $f(\rho)$ before the end of the election. Consequently, all messages the coercer has seen up to this point only depend on ω_2 and are independent of the choices made by honest voters, which are determined by ω_1 . Therefore, we sometimes write $f(\omega_2)$ for the candidate the coercer wants the coerced voter to vote for in runs that use the random coins ω_2 .

For a coercer view ρ , let φ_ρ be a predicate over Ω_1 such that $\varphi_\rho(\omega_1)$ holds iff $\text{res}(\omega_1, f(\rho)) = \text{res}(\rho)$ and $\text{abst}(\omega_1) = \text{abst}(\rho)$, i.e., the choices ω_1 of the honest voter are consistent with the view of the coercer, as far as the result of the election and the set

of abstaining voters is concerned, in case the coerced voter runs the dummy strategy. Analogously, for the counter-strategy, we define that $\tilde{\varphi}_\rho(\omega_1)$ holds iff $\text{res}(\omega_1, i) = \text{res}(\rho)$ and $\text{abst}(\omega_1) = \text{abst}(\rho)$.

For a coercer view ρ , by $T(\omega_1, \omega_2) \mapsto \rho$, or simply $T \mapsto \rho$, we denote the fact that the system T , when run with ω_1, ω_2 , produces the view ρ (similarly for \tilde{T}). For a set M of views, we write $T(\omega_1, \omega_2) \mapsto M$ if $T(\omega_1, \omega_2) \mapsto \rho$ for some $\rho \in M$.

The following lemma is the key fact used in the proof of Theorem 2 (see Appendix B for the proof). It constitutes the cryptographic part of the proof of Theorem 2.

Lemma 2. *Let ρ be a coercer view such that $f(\rho)$ is defined. Let ω_1^ρ and $\tilde{\omega}_1^\rho$ be some fixed elements of Ω_1 such that $\varphi_\rho(\omega_1^\rho)$ and $\tilde{\varphi}_\rho(\tilde{\omega}_1^\rho)$, respectively. Then, the following equations hold true:*

$$\Pr[T \mapsto \rho] = \Pr_{\omega_1}[\varphi_\rho(\omega_1)] \cdot \Pr_{\omega_2}[T(\omega_1^\rho, \omega_2) \mapsto \rho] \quad (1)$$

$$\Pr[\tilde{T} \mapsto \rho] = \Pr_{\omega_1}[\tilde{\varphi}_\rho(\omega_1)] \cdot \Pr_{\omega_2}[\tilde{T}(\tilde{\omega}_1^\rho, \omega_2) \mapsto \rho] \quad (2)$$

$$\Pr_{\omega_2}[T(\omega_1^\rho, \omega_2) \mapsto \rho] = \Pr_{\omega_2}[\tilde{T}(\tilde{\omega}_1^\rho, \omega_2) \mapsto \rho] . \quad (3)$$

Intuitively, the lemma says 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 result of the election given in this view.

Now, using this lemma, we can link the level of coercion-resistance the Bingo voting system provides with the optimal bound δ_{min} established in Section 4. Clearly, if $f(\rho)$ is defined, we have:

$$\Pr_{\omega_1}[\varphi_\rho(\omega_1)] = A_{\text{res}(\rho)}^{f(\rho)} \cdot \Pr_{\omega_1}[\text{abst}(\omega_1) = \text{abst}(\rho) \mid \text{res}(\omega_1, f(\rho)) = \text{res}(\rho)]$$

and

$$\Pr_{\omega_1}[\tilde{\varphi}_\rho(\omega_1)] = A_{\text{res}(\rho)}^i \cdot \Pr_{\omega_1}[\text{abst}(\omega_1) = \text{abst}(\rho) \mid \text{res}(\omega_1, i) = \text{res}(\rho)].$$

Furthermore, we have

$$\begin{aligned} \Pr_{\omega_1}[\text{abst}(\omega_1) = \text{abst}(\rho) \mid \text{res}(\omega_1, f(\rho)) = \text{res}(\rho)] &= \\ &= \Pr_{\omega_1}[\text{abst}(\omega_1) = \text{abst}(\rho) \mid r_0(\omega_1) = r_0(\rho)] \\ &= \Pr_{\omega_1}[\text{abst}(\omega_1) = \text{abst}(\rho) \mid \text{res}(\omega_1, i) = \text{res}(\rho)], \end{aligned}$$

as the set of abstaining voters depends only on the number of abstaining voters.

Together with Lemma 2, we immediately obtain for all ω_1^ρ with $\varphi_\rho(\omega_1^\rho)$:

$$\begin{aligned} \Pr[T \mapsto \rho] - \Pr[\tilde{T} \mapsto \rho] &= \\ (A_{\text{res}(\rho)}^{f(\rho)} - A_{\text{res}(\rho)}^i) \cdot \Pr_{\omega_2}[T(\omega_1^\rho, \omega_2) \mapsto \rho] \cdot \Pr_{\omega_1}[\text{abst}(\omega_1) = \text{abst}(\rho) \mid r_0(\omega_1) = r_0(\rho)]. \end{aligned}$$

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 j range over the set of candidate names $\{1, \dots, k\}$, $\vec{r} = (r_0, \dots, r_k)$ over all the pure results and S over all subsets of $\{1, \dots, n\}$. Let $M_j^{\vec{r}, S} = \{\rho \in M : f(\rho) = j, \text{ abst}(\rho) = S \text{ and } \text{res}(\rho) = \vec{r}\}$. Further, let $\omega_1^{j, \vec{r}, S}$ be an arbitrary element, such that $\text{res}(\omega_1^{j, \vec{r}, S}, j) = \vec{r}$ and $\text{abst}(\omega_1^{j, \vec{r}, S}) = S$. Then we have $\varphi_\rho(\omega_1^{j, \vec{r}, S})$ for all $\rho \in M_j^{\vec{r}, S}$. We have

$$\begin{aligned}
\Phi &= \Pr[T \mapsto 1] - \Pr[\tilde{T} \mapsto 1] \\
&= \Pr[T \mapsto M] - \Pr[\tilde{T} \mapsto M] \\
&= \sum_j \sum_{\vec{r}} \sum_S \sum_{\rho \in M_j^{\vec{r}, S}} (\Pr[T \mapsto \rho] - \Pr[\tilde{T} \mapsto \rho]) \\
&= \sum_j \sum_{\vec{r}} \sum_S \sum_{\rho \in M_j^{\vec{r}, S}} (A_{\vec{r}}^j - A_{\vec{r}}^i) \cdot \Pr_{\omega_2}[T(\omega_1^{j, \vec{r}, S}, \omega_2) \mapsto \rho] \\
&\quad \cdot \Pr_{\omega_1}[\text{abst}(\omega_1) = S | r_0(\omega_1) = r_0] \\
&= \sum_j \sum_{\vec{r}} (A_{\vec{r}}^j - A_{\vec{r}}^i) \sum_S \sum_{\rho \in M_j^{\vec{r}, S}} \Pr_{\omega_2}[T(\omega_1^{j, \vec{r}, S}, \omega_2) \mapsto \rho] \cdot \Pr_{\omega_1}[\text{abst}(\omega_1) = S | r_0(\omega_1) = r_0].
\end{aligned}$$

Let $M_{i,j}^* = \{\vec{r} : A_{\vec{r}}^j \geq A_{\vec{r}}^i\}$. Then, we obtain

$$\begin{aligned}
\Phi &\leq \sum_j \sum_{\vec{r} \in M_{i,j}^*} (A_{\vec{r}}^j - A_{\vec{r}}^i) \sum_S \sum_{\rho \in M_j^{\vec{r}, S}} \Pr_{\omega_2}[T(\omega_1^{j, \vec{r}, S}, \omega_2) \mapsto \rho] \\
&\quad \cdot \Pr_{\omega_1}[\text{abst}(\omega_1) = S | r_0(\omega_1) = r_0].
\end{aligned}$$

Next, we use that, by the definition of $M_j^{\vec{r}, S}$, for $\rho \in M_j^{\vec{r}, S}$ we have $f(\rho) = j$ and, because $f(\rho)$ depends only on ω_2 , $T(\omega_1^{j, \vec{r}, S}, \omega_2) \mapsto \rho$ implies $f(\omega_2) = j$. With this, we obtain:

$$\Pr_{\omega_2}[T(\omega_1^{\rho}, \omega_2) \mapsto \rho] = \Pr_{\omega_2}[f(\omega_2) = j] \cdot \Pr_{\omega_2}[T(\omega_1^{j, \vec{r}, S}, \omega_2) \mapsto \rho \mid f(\omega_2) = j]$$

for $\rho \in M_j^{\vec{r}, S}$. Now, we can conclude

$$\begin{aligned}
\Phi &\leq \sum_j \sum_{\vec{r} \in M_{i,j}^*} (A_{\vec{r}}^j - A_{\vec{r}}^i) \sum_S \Pr_{\omega_1}[\text{abst}(\omega_1) = S | r_0(\omega_1) = r_0] \\
&\quad \sum_{\rho \in M_j^{\vec{r}, S}} \Pr_{\omega_2}[f(\omega_2) = j] \\
&\quad \cdot \Pr_{\omega_2}[T(\omega_1^{j, \vec{r}, S}, \omega_2) \mapsto \rho \mid f(\omega_2) = j]
\end{aligned}$$



Figure 2: Two ways of voting for the second candidate (candidate B) in the ThreeBallot protocol, where x 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.

$$\begin{aligned}
&\leq \sum_j \Pr_{\omega_2}[f(\omega_2) = j] \sum_{r \in M_{i,j}^*} (A_r^j - A_r^i) \\
&\qquad \sum_S \Pr_{\omega_1}[\text{abst}(\omega_1) = S | r_0(\omega_1) = r_0] \\
&\leq \sum_j \Pr_{\omega_2}[f(\omega_2) = j] \sum_{\vec{r} \in M_{i,j}^*} (A_{\vec{r}}^j - A_{\vec{r}}^i) \\
&\leq \sum_j \Pr_{\omega_2}[f(\omega_2) = j] \cdot \delta_{\min}^i(n, k, \vec{p}) \leq \delta_{\min}^i(n, k, \vec{p}) = \delta.
\end{aligned}$$

This concludes the proof of Theorem 2.

6 Analyzing ThreeBallot

In this section, we study the ThreeBallot voting system [22]. As already mentioned in the introduction, based on our definition of coercion-resistance, we measure the degradation of coercion-resistance of ThreeBallot as the number of candidates grows, i.e., in case the so-called short ballot assumption is not met. We also show that the level of coercion-resistance ThreeBallot provides is significantly lower than that of an ideal system, even in case of short ballots. We first recall the ThreeBallot voting system and state our security assumptions.

6.1 Description of the System

In ThreeBallot, a voter is given a multi-ballot consisting of three simple ballots, where the candidates are written in a 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 2 shows two ways of voting for candidate B in an election with two candidates. After this, she feeds all three simple ballots to a 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 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 ballots.

In the tallying phase, all the cast simple ballots are shuffled by an voting authority and published on a bulletin board. From this publicly available information, the result

can be easily 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 for any candidate.

In our analyzes we consider the variant of ThreeBallot as proposed in [9]. In this variant a specific way of filling out the ballots is proposed: 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 some simple ballot, excluding the one chosen as a receipt (if there is more than one possibility, one of the two possible simple ballots is chosen randomly). The advantage of this procedure is that the receipt an honest voter gets is stochastically independent from the candidate the voter votes for, which gives better privacy. We note that in [9], ThreeBallot was analyzed in simulation-based setting, focussing on privacy. The analysis was based on the (only informally stated) assumption that the adversary is not able to reconstruct the multi-ballot corresponding to a receipt. However, this assumption is unjustified: Runs for which an adversary can reconstruct the multi-ballots occur with non-negligible probability (see Section 6.3).

6.2 Modeling and Security Assumptions

The modeling of ThreeBallot as an election system $S = \mathbf{P}_{\text{ThreeBallot}}(k, m, n, \vec{p})$ is straightforward. Here, we only highlight some modeling issues and our security assumptions.

Voting authorities. We assume that the scanner and the authorities in charge of shuffling the ballots are honest; the bulletin board may be dishonest. Without this assumption, coercion would easily be possible.

Honest voters. As usual, an honest voter first makes a choice according to the probability distribution \vec{p} . If the choice is to abstain from voting, she abstains, otherwise follows the procedure described for the voting phase. After the voting phase is finished, an honest voter may reveal her (paper) receipt. However, to measure how much information a coercer gains from the receipts of honest voters, we will also consider the case that the coercer does not get to see the receipts of honest voters.

The coerced voter. A coerced voter, running the dummy strategy or emulating it by running a counter-strategy, can communicate with the coercer. Just as an honest voter, she can also fill out a multi-ballot, feed it to the scanner and pick a receipt. If the coerced voter follows the dummy strategy, she will carry out these steps following the instructions of the coercer. Of course, if she follows the counter-strategy she can deviate from these instructions.

The coercer. As usual, the coercer subsumes dishonest voters and can freely communicate with the coerced voter. In a run of the system, the coercer can see the following: (v1) his random coins, (v2) the bulletin board consisting of the shuffled simple ballots with serial numbers of all voters (v3) optionally, depending on the case under consideration, the receipts of the honest voters, after the voting-phase is finished, (v4) the messages received from the coerced voter, including the receipt of the coerced voter. As in case of Bingo voting, the coercer cannot directly see the information the coerced voter obtains or the actions she performs in the voting booth.

6.3 ThreeBallot with Two Candidates

Based on our definition, we now precisely measure the level of coercion-resistance ThreeBallot provides and show that it is significantly lower than that of an ideal system, even in case of short ballots, and hence, under the so-called short ballot assumption (see, e.g., [21]). More precisely, we analyze the case of two candidates. The case for multiple candidates will be studied in Section 6.4.

As a warming up, we note that the bulletin board and the receipts potentially reveal more information to the coercer than just the result of the election: It may, for instance, happen, that the multi-ballots of all voters are of the form $(\underline{\circ}, \underline{x}, \circ)$ or $(\underline{\circ}, \underline{x}, \circ)$, where the underlined ballots ($\underline{\circ}$ and \underline{x} , respectively) are picked as receipts. In this case, a receipt directly indicates the choice of the voter, which allows for successful coercion.

In what follows, we often use the above notation for multi-ballots and the receipt picked, and refer to this object as a *pattern*. A pattern does not fix the order of simple ballots, e.g., $(\underline{\circ}, \underline{x}, \circ)$ is considered to be the same pattern as, say, $(\underline{x}, \underline{\circ}, \circ)$.

As before, our analysis is w.r.t. the goal γ_i , for $i \in \{1, 2\}$, which is met if the voter votes for candidate i , in case she is instructed by the coercer to vote for some candidate.

We proceed as follows: First, we define a counter-strategy, which is optimal for the coerced voter. Second, we define the constant δ , which describes the optimal level of coercion-resistance ThreeBallot achieves. For this, we introduce what we call an essential view of the coercer which abstracts away from some details of the actual view of the coercer. Finally, we state the main result of this section, namely δ -coercion-resistance of ThreeBallot and the optimality of δ . This is done both for the case where the coercer gets to see all receipts of voters and for the case where receipts of honest voters are hidden from the coercer, resulting in two constants δ_{TB+} and δ_{TB-} .

Counter-strategy. We define the *counter-strategy* of the coerced voter to coincide with the dummy strategy 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 $C(Z, i)$, as defined next. (Recall that the goal of the coerced voter is to vote for i).

We define $C(Z, i)$ in such a way that it yields the same receipt as Z does and adjusts the two remaining ballots in such a way that the resulting multi-ballot is a vote for

candidate i . By this requirement, $C(Z, i)$ is uniquely determined, except for two cases: $C(\binom{x}{\circ}, \binom{\circ}{x}, \binom{\circ}{x}, 1)$ and $C(\binom{\circ}{x}, \binom{x}{\circ}, \binom{x}{\circ}, 2)$. In the former case, for instance, one can take $(\binom{x}{\circ}, \binom{x}{\circ}, \binom{\circ}{x})$, $(\binom{x}{\circ}, \binom{x}{\circ}, \binom{\circ}{\circ})$, or randomly pick one of the two, possibly based on further information. For these cases, we define $C(\binom{x}{\circ}, \binom{\circ}{x}, \binom{\circ}{x}, 1) = (\binom{x}{\circ}, \binom{x}{\circ}, \binom{\circ}{\circ})$ and $C(\binom{\circ}{x}, \binom{x}{\circ}, \binom{x}{\circ}, 2) = (\binom{\circ}{x}, \binom{\circ}{\circ}, \binom{x}{\circ})$. We use this strategy in the proof of Theorem 3. From the proof of this theorem it follows that this counter-strategy achieves the maximal level of coercion-resistance and, in this sense, is optimal for the coerced voter.

Essential views. In the essential view of the coercer, we abstract away from the following information: the serial numbers on the simple ballots, the order of the simple ballots on the bulletin board, the order of the receipts of the honest voters (if considered), the random coins of the coercer (i.e., randomness does not help the coercer), the receipt of the coerced voter (as both in the dummy strategy and the counter-strategy as defined above, she returns what the coercer expects her to return) and the simple ballots of the dishonest voters (which are as expected by the coercer).

More precisely, an *essential view* of the coercer consists only of (i) three integers $n_x, n_{\circ}, n_{\times}$, indicating the number of the respective simple ballots on the bulletin board and (ii) in case the coercer can see the receipts of honest voters, three integers $r_x, r_{\circ}, r_{\times}$, indicating the number of the respective receipts taken by the those voters. Note that from these numbers the number of (\circ) -ballots on the bulletin board and the number of (\circ) -receipts can be derived by the coercer.

By V^+ and V^- we denote the set of all essential views of the coercer, when he can or cannot see the receipts of the honest voters, respectively.

The constants δ_{TB-}^i and δ_{TB+}^i . To define these constants we use the probability A_{ρ}^Z that the choices made by the honest voters and the coerced voter result in an essential view ρ , given that the coerced voter chooses the pattern Z .

The intuition behind the result given below is similar to the one for the ideal protocol (Section 4): If the coercer wants the coerced voter to choose the pattern Z and the coerced voter wants to vote for candidate i , then the best strategy of the coercer to distinguish whether the coerced voter has chosen Z or $C(Z, i)$ is to accept a run if the essential view ρ in this run is such that $A_{\rho}^{C(Z, i)} \leq A_{\rho}^Z$. Let $M_{Z, i}^- = \{\rho \in V^- : A_{\rho}^{C(Z, i)} \leq A_{\rho}^Z\}$ and $M_{Z, i}^+ = \{\rho \in V^+ : A_{\rho}^{C(Z, i)} \leq A_{\rho}^Z\}$ be the sets of those essential views for which – according to his best strategy – the coerced should accept the run.

Now, we are ready to define the constants expressing the (optimal) level of coercion-resistance of ThreeBallot, for the case that the coercer cannot see the receipts of the honest voters:

$$\delta_{TB-}^i(n, \vec{p}) = \max_Z \sum_{\rho \in M_{Z, i}^-} (A_{\rho}^Z - A_{\rho}^{C(Z, i)}) , \quad (4)$$

and for the case that the coercer can see these receipts:

$$\delta_{TB^+}^i(n, \vec{p}) = \max_Z \sum_{\rho \in M_{Z,i}^+} (A_\rho^Z - A_\rho^{C(Z,i)}). \quad (5)$$

The following theorem shows that the two constants (more precisely, functions) just defined in fact capture the optimal level of coercion-resistance provided by ThreeBallot in case of an election with two candidates.

Theorem 3. *Let $S = P_{\text{ThreeBallot}}(2, m, n, \vec{p})$. Then:*

1. *If the coercer cannot see the receipts of the honest voters, then $P_{\text{ThreeBallot}}$ is δ -coercion resistant with respect to γ_i for $\delta = \delta_{TB^-}^i(n, \vec{p})$.*
2. *Similarly, if the coercer can see the receipts of the honest voters, but only after the voting phase, then $P_{\text{ThreeBallot}}$ is δ -coercion resistant with respect to γ_i for $\delta = \delta_{TB^+}^i(n, \vec{p})$.*

Moreover, in both cases the system is not δ' -coercion-resistant for any $\delta' < \delta$.

The proof of this theorem is given in Appendix C. The main part of this proof is to show that the additional information given in a full view of the coercer, and omitted in an essential view, can safely be discarded. This is similar to the proof of Theorem 2, where we reduced the analysis of the Bingo voting system to the ideal protocol, although the technical details differ and are simpler for ThreeBallot.

For ThreeBallot the bigger challenge is to come up with explicit formulas for the probabilities A_ρ^Z , which allow to compute the level of coercion-resistance for concrete parameters. In particular, this is so for the case where the coercer can see the receipts of honest voters. The formulas are stated in the following two lemmas.

Lemma 3. *Consider the case when the coercer cannot see the receipts of the honest voters. Let $\rho = (n_x, n_o, n_z)$ be an essential view. Then we have $A_\rho^Z = A_{\rho-Z}$, where $\rho - Z$ denotes the view we get by removing the ballots of Z from ρ and*

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

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

While the above formula can be obtained relatively easily, the following formula is harder to obtain (see Appendix C).

Lemma 4. *Consider the case when the coercer can see the receipts of the honest voters. Let $\rho = (n_x, n_o, n_z, r_x, r_o, r_z)$ be an essential view of the coercer. Then we have*

$A_\rho^Z = A_{\rho-Z}$ where

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

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

Using these formulas, we have computed the level of coercion-resistance for concrete values (see Figure 3). In order to be able to compare this level with the one for the ideal protocol, we depict the corresponding values also for this protocol.

As can be seen from the diagram (a) in Figure 3, the level of coercion-resistance of the ideal protocol is about double the level provided by ThreeBallot, in case receipts of honest voters can be seen by the coercer, i.e., the value for δ for the ideal protocol is half of the value for ThreeBallot (if the number of honest voters is at least five). This difference is quite significant. It means that in case of ThreeBallot the expected gain when trying to sell ones vote (by following the instructions of the coercer instead of running the counter-strategy) is twice as high as in the ideal protocol. Conversely, by running the counter-strategy (instead of following the instructions of the coercer) the expected growth in the risk of being caught is twice as big as in the ideal protocol.

The difference between ThreeBallot and the ideal protocol decreases in case the coercer cannot see the receipts of honest voters. We found that it also decreases if the probability distribution for the candidates is less uniform as the deficiency of ThreeBallot then becomes less significant (see the diagram (b) in Figure 3).

As already pointed out in Section 3, existing cryptographic approaches are unsuitable for analyzing coercion-resistance of ThreeBallot.

6.4 ThreeBallot with multiple candidates

We now analyze the degradation of coercion-resistance of ThreeBallot as the number of candidates grows, i.e., the case where the so-called small ballot assumption is not met. The degradation itself is not surprising since certain patterns become very unlikely to occur. This has been noted for variants of ThreeBallot, e.g., in [12, 13]. However, our definition allows us to measure the degradation rigorously in the context of coercion-resistance, showing that the level of coercion-resistance ThreeBallot provides is completely insufficient already with five to seven candidates and a few hundred voters.

More precisely, in this section we state negative results for ThreeBallot with multiple candidates by providing *lower bounds* for the level of coercion of ThreeBallot, i.e.

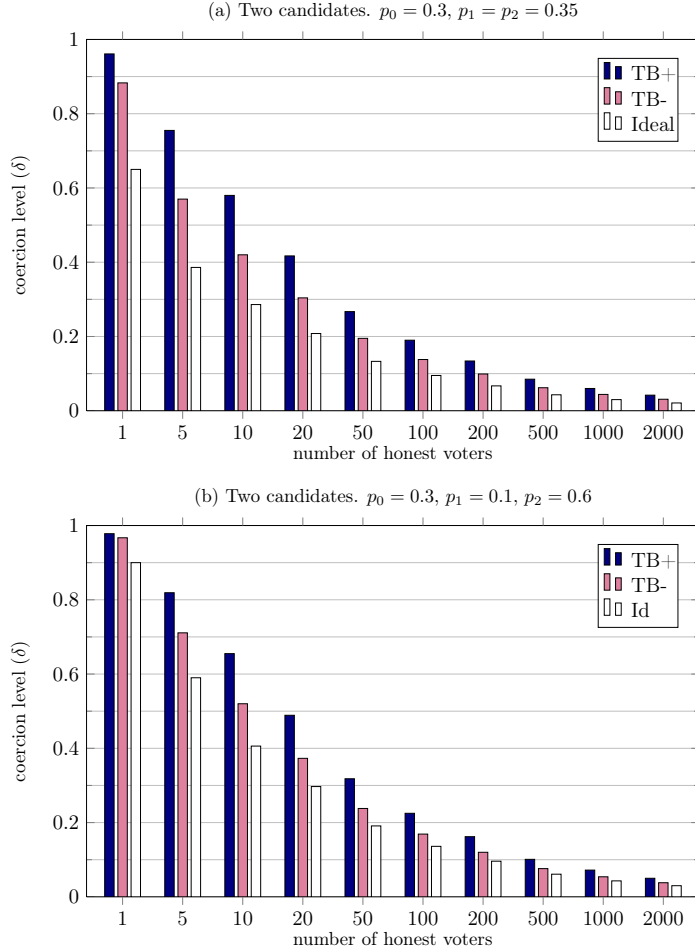


Figure 3: Level of coercion resistance (δ) for the ideal protocol (Id), ThreeBallot without revealing receipts of the honest voters (TB-) and with revealed receipts of the honest voters (TB+).

we show that the studied systems are *not* δ -coercion-resistant for any δ smaller than the lower bound. (Unlike the two candidate case, we do not show that the systems *are* coercion-resistant for the given lower bound.) These results apply both to the case with and without receipts.

Our method is the following. We consider a restricted class $\mathcal{C} \subseteq C_S$ of programs of the coercer. Then we define a counter-strategy of the coerced voter (which is optimal for \mathcal{C}), and apply the definition of coercion-resistance with the set of coercer programs restricted to the class \mathcal{C} .

The class \mathcal{C} is defined as follows. In every program $c \in \mathcal{C}$, the coercer instructs the coerced voter to vote for some candidate j by marking all the position on one single ballot (we will call such a single ballot *fully marked*), the j -th position on the second ballot, and no position on the third ballots. The coercer then asks the coerced

voter to bring the second ballot (the one with one position marked) as a receipt. The program c , which decides whether to accept a run, only uses the following parts of its view: (v1) the receipt given by the coerced voter, (v2) the pure result $\vec{r} = (r_1, \dots, r_k)$, as introduced in Section 4, and (v3) the number u of all fully marked ballots on the bulletin board cast by honest voters and the coerced voter. A tuple $\rho = (\vec{r}, u)$ will be called a *restricted view* of the coercer. For the same reason as in the two candidate case, the receipt of the coerced voter is not part of the view; the counter-strategy will always return the expected receipt.

Now, let γ_i denote the goal as specified in the two candidate case. We define the counter-strategy v^* as follows: The coerced voter, when instructed to vote as determined by the coercer, fills out the multi-ballot in such a way that (a) she votes for i and (b) one of the single ballots is the required receipt. This can be done in possibly many ways; v^* just fixes one of them.

This counter-strategy is optimal for \mathcal{C} because any two strategies satisfying these conditions produce exactly the same restricted views (since they do not use fully marked ballots), and it is clear that any successful counter-strategy has to satisfy (a) and (b).

Now, the technique for obtaining the lower bound is very similar to the one used for the case with two candidates without receipts.

Let k, \vec{p} and γ_i be as usual. Let $\rho = (\vec{r}, u)$ be a restricted view. We will denote by $A_\rho^{i,o}$ ($A_\rho^{i,c}$) the probability that the choices of the honest voters and the coerced voter result in the restricted view ρ , given that the coerced voter votes for candidate i with (without) one fully marked ballot. Note that if the coerced voter obeys the instructions of the coercer her multi-ballot contains a fully marked ballot; otherwise, it does not. These probabilities are given by

$$A_\rho^{i,o} = A_{\vec{r}}^i \cdot \binom{n-r_0}{u-1} q^{u-1} (1-q)^{n-r_0-u+1}$$

and

$$A_\rho^{i,c} = A_{\vec{r}}^i \cdot \binom{n-r_0}{u} q^u (1-q)^{n-r_0-u},$$

respectively, where $A_{\vec{r}}^i$ is defined as in Section 4 and $q = \frac{2}{3^{k-1}}$ is the probability that an honest, non-abstaining voter produces a fully marked ballot.

Let $M_{i,j}$ be the set of those restricted views ρ for which $A_\rho^{j,o} > A_\rho^{i,c}$ and let

$$\delta_i(n, k, \vec{p}) = \max_{j \in \{1, \dots, k\}} \sum_{\rho \in M_{i,j}} (A_\rho^{j,o} - A_\rho^{i,c}). \quad (6)$$

Then, we obtain the following result (see Appendix D for the proof):

Theorem 4. *Let $S = \text{P}_{\text{ThreeBallot}}(k, m, n, \vec{p})$. Then S is not δ -coercion resistant w.r.t. γ_i for any $\delta < \delta_i(n, k, \vec{p})$.*

This result allows us to compute lower bounds for the level of coercion of ThreeBallot for different numbers of candidates. Figure 4 depicts some of these lower bounds.

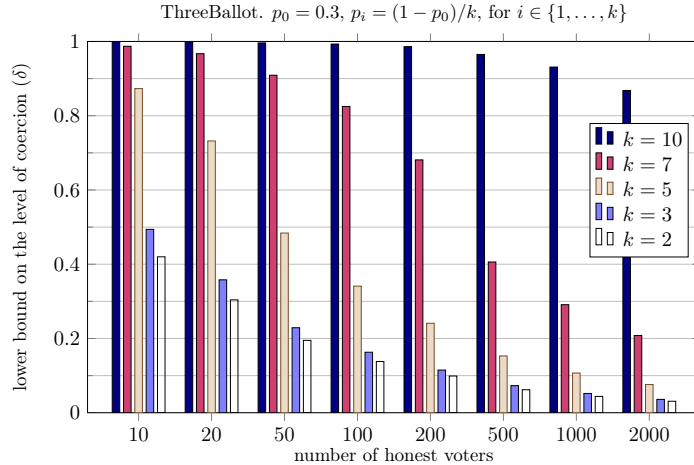


Figure 4: The lower-bound of coercion resistance (δ) for the ThreeBallot system with k candidates with the probability that an honest voter abstains $p_0 = 0.3$ and equal probability of choosing a candidate $((1 - p_0)/k)$.

As already mentioned, the figure shows that the level of coercion-resistance ThreeBallot provides is completely insufficient already with five to seven candidates and a few hundred voters. Note that the actual levels of coercion can even be higher than depicted in this figure.

References

- [1] M. Backes, C. Hritcu, and M. Maffei. Automated Verification of Remote Electronic Voting Protocols in the Applied Pi-Calculus. In *Proceedings of the 21st IEEE Computer Security Foundations Symposium (CSF 2008)*, pages 195–209. IEEE Computer Society, 2008.
- [2] M. Bär, C. Henricj, J. Müller-Quade, S. Röhrich, and C. Stüber. Real World Experiences with Bingo Voting and a Comparison of Usability. In *IAVoSS Workshop On Trustworthy Elections (WOTE 2008)*, 2008.
- [3] 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.
- [4] 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.
- [5] R. Canetti. Universally Composable Security: A New Paradigm for Cryptographic Protocols. Technical report, Cryptology ePrint Archive, December 2005. Online available at <http://eprint.iacr.org/2000/067.ps>.

- [6] R. Canetti and R. Gennaro. Incoercible Multiparty Computation (extended abstract). In *37th Annual Symposium on Foundations of Computer Science (FOCS '96)*, pages 504–513. IEEE Computer Society, 1996.
- [7] B. Chevallier-Mames, P.-A. Fouque, D. Pointcheval, J. Stern, and J. Traoré. On Some Incompatible Properties of Voting Schemes. In *IAVoSS Workshop On Trustworthy Elections (WOTE 2006)*, 2006.
- [8] 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.
- [9] 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.
- [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] R. Gardner, S. Garera, and A. D. Rubin. Coercion Resistant End-to-end Voting. In Roger Dingledine and Philippe Golle, editors, *13th International Conference on Financial Cryptography (FC 2009)*, volume 5628 of *Lecture Notes in Computer Science*, pages 344–361. Springer, 2009.
- [12] Viliam Geffert, Juhani Karhumäki, Alberto Bertoni, Bart Preneel, Pavol Návrat, and Mária Bieliková, editors. *Short Ballot Assumption and Threeballot Voting Protocol*, volume 4910 of *Lecture Notes in Computer Science*. Springer, 2008.
- [13] Kevin Henry, Douglas R. Stinson, and Jiayuan Sui. The effectiveness of receipt-based attacks on threeballot. <http://www.cacr.math.uwaterloo.ca/~dstinson/papers/ThreeBallot-Jan.30.pdf>, January 30, 2008.
- [14] H.L. Jonker and W. Pieters. Receipt-Freeness as a special case of Anonymity in Epistemic Logic. In *IAVoSS Workshop On Trustworthy Elections (WOTE 2006)*, 2006.
- [15] A. Juels, D. Catalano, and M. Jakobsson. Coercion-resistant electronic elections. In *Proceedings of Workshop on Privacy in the Eletronic Society (WPES 2005)*. ACM Press, 2005.
- [16] 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.
- [17] R. Küsters and T. 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.
- [18] T. Moran and M. Naor. Receipt-Free Universally-Verifiable Voting With Everlasting Privacy. In *Advances in Cryptology—CRYPTO 2006, 26th Annual International Cryptology Conference*, 2006. To appear.
- [19] T. Moran and M. Naor. Split-ballot voting: everlasting privacy with distributed trust. In *ACM Conference on Computer and Communications Security (CSS 2007)*, pages 246–255, 2007.

- [20] 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.
- [21] R. L. Rivest and W. D. Smith. Three Voting Protocols: ThreeBallot, VAV and Twin. In *USENIX/ACCURATE Electronic Voting Technology (EVT 2007)*, 2007.
- [22] Ron Rivest. The threeballot voting system. people.csail.mit.edu/rivest/Rivest-TheThreeBallotVotingSystem.pdf, October 1, 2006.
- [23] V. Teague, K. Ramchen, and L. Naish. Coercion-Resistant tallying for STV voting. In *IAVoSS Workshop On Trustworthy Elections (WOTE 2008)*, 2008.
- [24] D. Unruh and J. Müller-Quade. Universally Composable Incoercibility. Cryptology ePrint Archive, Report 2009/520, 2009. <http://eprint.iacr.org/>. Submitted October 27, 2009.

A Proof of Theorem 1

Our goal is to prove that $S = P_{\text{ideal}}(k, m, n, \vec{p})$ is δ -coercion-resistant w.r.t. γ_i , where $i \in \{1, \dots, k\}$ and $\delta = \delta_{\min}^i(n, k, \vec{p})$. To show δ -coercion-resistance, we take the counter-strategy \tilde{v} which, when the coerced voter is instructed to vote (for some candidate), votes for the i -th candidate. This strategy clearly meets condition (i) of Definition 1, for every $c \in C_S$.

We need to show that condition (ii) of this definition is satisfied. We begin with some auxiliary definitions and facts. Let

$$\Delta_{ij} = \sum_{\vec{r} \in M_{i,j}^*} (A_{\vec{r}}^j - A_{\vec{r}}^i). \quad (7)$$

So, we have $\delta_{\min}^i(n, k, \vec{p}) = \max_{j \in \{1, \dots, k\}} \Delta_{ij}$.

By $\text{res}(\omega_1, i)$, where $\omega_1 \in \Omega_1$ (recall that Ω_1 is the set of candidate choices made by honest voters) 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 coerced voter v_0 votes for i . Therefore, we have

$$A_{\vec{r}}^i = \Pr_{\omega_1}[\text{res}(\omega_1, i) = \vec{r}]. \quad (8)$$

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:

$$\sum_{\vec{r} \in M} (A_{\vec{r}}^j - A_{\vec{r}}^i) \leq \sum_{\vec{r} \in M_{i,j}^*} (A_{\vec{r}}^j - A_{\vec{r}}^i) = \Delta_{ij}. \quad (9)$$

Now, to prove condition (ii) of Definition 1, let $c \in C_S$. Recall that the view of the coercer in a run of the system consists only of his random coins $\omega_2 \in \Omega_2$ and the result of the election.

The only action of the coerced voter, besides receiving a candidate name from the coercer, is to indicate the candidate of choice to the voting machine. Therefore, the dummy strategy of the coerced voter also only needs to forward one message, namely the candidate name; no other message is going to be accepted by the voting machine. Moreover, this message has to be sent before the result of the election is published, in order for the voting machine to accept the message. Therefore, if the coercer demands that the coerced voter votes for candidate j , he has to do this before the result is published. In particular, the coercer has to determine j – the candidate he wants the coerced voter to vote for – based solely on his random coins ω_2 , independently of the result of the election. Hence j is a function of ω_2 , which we denote by $f(\omega_2)$; this function can be undefined if the coercer does not want the coerced voter to vote. Note that $f(\omega_2) \in \{1, \dots, k\}$. Hence, the view of the coercer if the coerced voter runs the dummy strategy \mathbf{dum} is the random variable $view(\mathbf{dum}, c)$, where

$$view(\mathbf{dum}, c)(\omega_1, \omega_2) = (\omega_2, \text{res}(\omega_1, f(\omega_2)))$$

for every $\omega = (\omega_1, \omega_2) \in \Omega$. It is

$$view(\tilde{v}, c)(\omega_1, \omega_2) = (\omega_2, \text{res}(\omega_1, i)),$$

if the coerced voter runs the counter-strategy \tilde{v} to vote for i .

Now, let M_c be the set of views accepted by the machine c . Each element of M_c is of the form (ω_2, \vec{r}) , where $\omega_2 \in \Omega_2$ and \vec{r} is a pure result. For $\omega_2 \in \Omega_2$, we define $M_c^{\omega_2}$ to be $\{\vec{r} \in \text{Res} : (\omega_2, \vec{r}) \in M_c\}$. Moreover, we define $\Omega'_2 = \{\omega_2 \in \Omega_2 \mid f(\omega_2) \text{ is defined}\}$. Note that for $\omega_2 \notin \Omega'_2$, the counter-strategy behaves exactly like the dummy strategy, namely, abstains from voting. With this, we obtain:

$$\begin{aligned} & \Pr[(c \parallel \mathbf{dum} \parallel e_S) \mapsto 1] - \Pr[(c \parallel \tilde{v} \parallel e_S) \mapsto 1] \\ &= \Pr[view(\mathbf{dum}, c) \in M_c] - \Pr[view(\tilde{v}, c) \in M_c] \\ &= \sum_{\omega_2 \in \Omega'_2} \sum_{\vec{r} \in M_c^{\omega_2}} \left(\Pr[view(\mathbf{dum}, c) = (\omega_2, \vec{r})] - \Pr[view(\tilde{v}, c) = (\omega_2, \vec{r})] \right) \\ &= \sum_{\omega_2 \in \Omega'_2} \sum_{\vec{r} \in M_c^{\omega_2}} \left(\mu_2(\omega_2) \cdot \Pr_{\omega_1}[\text{res}(\omega_1, f(\omega_2)) = \vec{r}] - \mu_2(\omega_2) \cdot \Pr_{\omega_1}[\text{res}(\omega_1, i) = \vec{r}] \right) \\ &= \sum_{\omega_2 \in \Omega'_2} \mu_2(\omega_2) \cdot \sum_{\vec{r} \in M_c^{\omega_2}} \left(A_{\vec{r}}^{f(\omega_2)} - A_{\vec{r}}^i \right) \\ &\leq \sum_{\omega_2 \in \Omega'_2} \mu_2(\omega_2) \cdot \sum_{\vec{r} \in M_{i, f(\omega_2)}^*} \left(A_{\vec{r}}^{f(\omega_2)} - A_{\vec{r}}^i \right) \quad (\text{by (9)}) \end{aligned}$$

$$\begin{aligned}
&= \sum_{\omega_2 \in \Omega'_2} \mu_2(\omega_2) \cdot \Delta_{i,f(\omega_2)} && \text{(by (7))} \\
&\leq \sum_{\omega_2 \in \Omega'_2} \mu_2(\omega_2) \cdot \max_{j \in \{1, \dots, k\}} \Delta_{i,j} \\
&= \max_{j \in \{1, \dots, k\}} \Delta_{i,j} = \delta_{min}^i(n, k, \vec{p}).
\end{aligned}$$

This implies that $\Pr[(c \parallel \mathbf{dum} \parallel \mathbf{e}_S) \mapsto 1] - \Pr[(c \parallel \tilde{v} \parallel \mathbf{e}_S) \mapsto 1]$ is δ -bounded, for $\delta = \delta_{min}^i(n, k, \vec{p})$. So, Condition (ii) of Definition 1 follows.

It remains to show that δ is optimal. First, we observe that in the above inequalities we obtain equality for a coercer program c i) which always instructs the coerced voter to vote for candidate j_0 , where, for the fixed i , $\Delta_{i,j}$ takes its maximum for $j = j_0$, and hence, $f(\omega_2) = j_0$ for all $\omega_2 \in \Omega_2$, and ii) accepts a run only if the pure result \vec{r} in his view belongs to M_{i,j_0}^* .

Second, we observe that if the coercer wants the coerced voter to abstain from voting, then the best counter-strategy is to abstain as well, because in this case the dummy strategy and the counter-strategy behave the same, and the counter-strategy satisfies γ_i ; if a counter-strategy would not abstain in this case, then this could possibly be detected by the coercer.

Conversely, if the coercer wants the coerced voter to vote, the counter-strategy has to vote for i with overwhelming probability, as it has to achieve γ_i with overwhelming probability. Hence, using a different counter-strategy than \tilde{v} the term $\Pr[(c \parallel \mathbf{dum} \parallel \mathbf{e}_S) \mapsto 1] - \Pr[(c \parallel \tilde{v} \parallel \mathbf{e}_S) \mapsto 1]$ would at most be negligibly smaller than δ . In particular, there is no constant δ' such that $\delta' < \delta$ and S is δ' -coercion-resistant.

B Proof of Lemma 2

The core of Lemma 2 is stated in the following lemma.

Lemma 5. *Let ρ be an arbitrary view such that $f(\rho)$ is defined. Let $\omega_1, \omega'_1, \omega''_1, \omega'''_1$ be arbitrary, fixed elements of Ω_1 with $\varphi_\rho(\omega_1), \varphi_\rho(\omega'_1), \tilde{\varphi}_\rho(\omega''_1)$, and $\tilde{\varphi}_\rho(\omega'''_1)$. Then the sets*

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

have the same cardinality, and hence, $\mu_2(A) = \mu_2(B) = \mu_2(C) = \mu_2(D)$.

To prove this lemma, we use Lemma 6. To state Lemma 6, we use the following notation. By \tilde{T}_j we denote the system $(\tilde{v}_j \parallel c \parallel \mathbf{e}_S)$, where \tilde{v}_j is defined like \tilde{v} but votes for j instead of i . So we have $\tilde{v} = \tilde{v}_i$ and $\tilde{T} = \tilde{T}_i$. Moreover, for each view ρ of the coercer, for which $f(\rho)$ is defined, we clearly have: $T(\omega_1, \omega_2) \mapsto \rho$ iff $\tilde{T}_{f(\rho)}(\omega_1, \omega_2) \mapsto \rho$. A permutation σ on a tuple $(v_0, \dots, v_n) \in \{0, 1, \dots, k\}^{n+1}$ is a permutation on the set of indices $\{0, \dots, n\}$. We write $\sigma(v_0, \dots, v_n)$ for the tuple $(v_{\sigma(0)}, \dots, v_{\sigma(n)})$. For

simplicity of notation, we sometimes write $\sigma(v_i)$ instead of $v_{\sigma(i)}$. We say that σ does not change the abstaining votes of (v_0, \dots, v_n) if $\sigma(j) = j$ for every $j \in \{0, \dots, k\}$ with $v_j = 0$. For $j \in \{1, \dots, k\}$ and $\omega_1 \in \Omega_1 (= \{0, 1, \dots, k\}^n)$, we consider (j, ω_1) to be an $(n+1)$ -tuple over $\{0, 1, \dots, k\}$. If σ is a permutation on (j, ω_1) , we may apply σ to ω_1 , written $\sigma(\omega_1)$, with the obvious meaning. With this and the above conventions, we have that $\sigma(j, \omega_1) = (\sigma(j), \sigma(\omega_1))$.

Lemma 6. *For every $j \in \{1, \dots, k\}$, every $\omega_1 \in \Omega_1$ and every permutation σ^0 on (j, ω_1) that does not change the abstaining votes, there is a bijective function $h = h^{j, \omega_1, \sigma^0}$ from Ω_2 to Ω_2 such that for all ω_2 we have that $\tilde{T}_j(\omega_1, \omega_2)$ yields the same view as $\tilde{T}_{\sigma^0(j)}(\sigma^0(\omega_1), h(\omega_2))$.*

We postpone the proof of this lemma to the end of this section. Now, Lemma 5 follows directly from Lemma 6: Given the assumptions of Lemma 5, there are permutations σ_1^0 , σ_2^0 , and σ_3^0 such that $(f(\rho), \omega_1) = \sigma_1^0(f(\rho), \omega_1') = \sigma_2^0(i, \omega_1'') = \sigma_3^0(i, \omega_1''')$. Moreover, $T(\omega_1, \omega_2) \mapsto \rho$ iff $\tilde{T}_{f(\rho)}(\omega_1, \omega_2) \mapsto \rho$ and $\tilde{T}(\omega_1, \omega_2) \mapsto \rho$ iff $\tilde{T}_i(\omega_1, \omega_2) \mapsto \rho$. From this and Lemma 6 we obtain that the functions $h^{f(\rho), \omega_1, (\sigma_1^0)^{-1}}$, $h^{f(\rho), \omega_1, (\sigma_2^0)^{-1}}$, and $h^{f(\rho), \omega_1, (\sigma_3^0)^{-1}}$, are bijections between A and B , A and C , and A and D , respectively.

Now with Lemma 5 we can easily complete the proof of Lemma 2:

$$\begin{aligned}
\Pr[T \mapsto \rho] &= \Pr_{\omega_1, \omega_2} [\varphi_\rho(\omega_1), T(\omega_1, \omega_2) \mapsto \rho] \\
&= \sum_{\omega_1': \varphi_\rho(\omega_1')} \Pr_{\omega_1, \omega_2} [\omega_1 = \omega_1', T(\omega_1', \omega_2) \mapsto \rho] \\
&= \sum_{\omega_1': \varphi_\rho(\omega_1')} \mu_1(\omega_1') \cdot \Pr_{\omega_1, \omega_2} [T(\omega_1', \omega_2) \mapsto \rho \mid \omega_1 = \omega_1'] \\
&= \sum_{\omega_1': \varphi_\rho(\omega_1)} \mu_1(\omega_1') \cdot \Pr_{\omega_2} [T(\omega_1', \omega_2) \mapsto \rho] \\
&= \sum_{\omega_1': \varphi_\rho(\omega_1)} \mu_1(\omega_1') \cdot \Pr_{\omega_2} [T(\omega_1^\rho, \omega_2) \mapsto \rho] \\
&= \Pr_{\omega_1} [\varphi_\rho(\omega_1)] \cdot \Pr_{\omega_2} [T(\omega_1^\rho, \omega_2) \mapsto \rho].
\end{aligned}$$

This proves (1). The proof for (2) is analogous. Statement (3) follows immediately from Lemma 5.

Proof of Lemma 6. To prove Lemma 6, we first introduce notation for the components (cryptographic operations, random numbers, etc.) of the Bingo Voting protocol.

The cryptographic components. For the sake of simplicity, we omit the description of the zero-knowledge proofs in the initialization phase. However, the proof of Lemma 6 can easily be extended to deal with these proofs as these proofs can be dealt with very similarly to the zero-knowledge proofs in the tallying phase.

We first describe in detail the structure of the sequence $\omega_2 \in \Omega_2$ of random coins. In the following, by $\text{comm}(a)^r$ we denote the commitment on a with randomness r .

- (a) α — the random coins of the coercer.
- (b) x_i^j and r_i^j , for $i \in \{0, \dots, m\}$ and $j \in \{1, \dots, k\}$ — the random numbers and the randomness used in the commitments $c_i^j = \text{comm}(j, x_i^j)^{r_i^j}$.
- (c) π — the permutation used by the machine to shuffle the commitments c_i^j .
- (d) x_i , for $i \in \{0, \dots, m\}$ — the random number generated by RNG for the i -th voter.
- (e) π_j , for every candidate $j \in \{1, \dots, k\}$ — a permutation of $\{0, \dots, m\}$, such that $x_{\pi_j(i)}^j$ is the number (taken from the pool of random numbers generated for the j -th candidate) assigned by the machine in the voting booth to the i -th voter (if necessary, i.e. if the i -th voter does not abstain and does not vote for j).
- (f) r_i , for every candidate $i \in \{0, \dots, m\}$ who does not abstain — a random number used by the voting machine to create a commitment $c_i = \text{comm}(v_i, x_i)^{r_i}$.
- (g) σ_i^0 , for every candidate $i \in \{0, \dots, m\}$ who does not abstain — a permutation used by the machine to shuffle the commitments associated with the receipt R_i of the i -th voter (see (B3) below).
- (h) $\tau_{i,j}^1$ and σ_i^1 , for every $j \in \{1, \dots, k\}$ and every candidate $i \in \{0, \dots, m\}$ who does not abstain — random numbers and permutations used for masking and shuffling commitments in C_{left}^i (see (B5) below).
- (i) $\tau_{i,j}^2$ and σ_i^2 , for every $j \in \{1, \dots, k\}$ and every candidate $i \in \{0, \dots, m\}$ who does not abstain — random numbers and permutations used for masking and shuffling commitments in C_{middle}^i (see (B6) below).
- (j) Random values contributed by the auditors to compute a challenge $s \in \{1, 2\}$.

A view ρ of the coercer, depending on ω_2 and the choices v_0, \dots, v_n taken by the voters, consists of the following parts:

- (B1) α — random coins of the coercer.
- (B2) The commitments c_i^j shuffled with π .
- (B3) R_i — the receipt of the i -th candidate, for every non-abstaining candidate i . Such a receipt is of the form s_1, \dots, s_k , where $s_j = (j, x_{\pi_j(i)}^j)$, for $j \neq v_i$, and $s_{v_i} = (v_i, x_i)$.
- (B4) The values $x_{\pi_j(i)}^j$ and $r_{\pi_j(i)}^j$ for opening the unused commitments $c_{\pi_j(i)}^j$, for all $j \in \{0, \dots, k\}$ and $i \in \{0, \dots, m\}$ such that $v_i = 0$ or $v_i = j$.

In the following items, i ranges over all the non-abstaining voters $i \in \{0, \dots, m\}$:

- (B5) The list of commitments $C_{left}^i = d_i^1, \dots, d_i^k$ shuffled with σ_i^0 , where $d_i^j = c_i$, if $j = v_i$, and $d_i^j = c_{\pi_j(i)}^j$, otherwise,
- (B6) The list of commitments $C_{middle}^i = \bar{d}_i^1, \dots, \bar{d}_i^k$ shuffled with $(\sigma_i^1 \circ \sigma_i^0)$, where $\bar{d}_i^j = \text{comm}(v_i, x_i)^{r_i + \tau_{i,\sigma_i^0(v_i)}^1}$, if $j = v_i$, and $\bar{d}_i^j = \text{comm}(j, x_{\pi_j(i)}^j)^{r_{\pi_j(i)}^j + \tau_{i,\sigma_i^0(j)}^1}$, otherwise.
- (B7) The list of commitments $C_{right}^i = \hat{d}_i^1, \dots, \hat{d}_i^k$ shuffled with $(\sigma_i^2 \circ \sigma_i^1 \circ \sigma_i^0)$, where $\hat{d}_i^j = \text{comm}(v_i, x_i)^{r_i + \tau_{i,\sigma_i^0(v_i)}^1 + \tau_{i,\sigma_i^1(\sigma_i^0(v_i))}^2}$, if $j = v_i$, and

$$\hat{d}_i^j = \text{comm}(j, x_{\pi_j(i)}^j)^{r_{\pi_j(i)}^j + \tau_{i, \sigma_i^1(j)}^1 + \tau_{i, \sigma_i^1(\sigma_i^0(j))}^2}, \text{ otherwise.}$$

(B8) The values $r_i + \tau_{i, \sigma_i^0(v_i)}^1 + \tau_{i, \sigma_i^1(\sigma_i^0(v_i))}^2$ and $r_{\pi_j(i)}^j + \tau_{i, \sigma_i^0(j)}^1 + \tau_{i, \sigma_i^1(\sigma_i^0(j))}^2$, for $j \in \{1, \dots, k\}$, $j \neq i$, for opening the commitment in C_{right}^i .

(B9) The challenge s along with masking factors σ_i^s and permutations $\tau_{i,j}^s$.

Proof. Because every permutation is the finite composition of permutations that switch only two successive positions, it suffices to consider the case where σ flips the positions l and $l+1$; the rest follows from composing permutations and bijections. Let $\tilde{v}_0, \dots, \tilde{v}_n$ be such that

$$\sigma(v_0, \dots, v_n) = (\tilde{v}_0, \dots, \tilde{v}_n) = (v_0, \dots, v_{l+1}, v_l, \dots, v_n)$$

Further, we assume that $v_l = y \neq z = v_{l+1}$, as the case that $\sigma^0(v_0, \dots, v_n) = (v_0, \dots, v_n)$ is trivial. Recall that, by assumption, we have that $y, z \neq 0$.

Let ω_2 be any element of Ω_2 and let $\alpha, x_i^j, r_i^j, \pi, x_i, \pi_j, r_i, \sigma_i^0, \tau_{i,j}^1, \sigma_i^1, \tau_{i,j}^2, \sigma_i^2$ and s be the parts of ω_2 defined as above. Here, i ranges over $0, \dots, m$ and j over $1, \dots, k$. We will denote the corresponding parts of $h(\omega_2)$ by $\tilde{\alpha}, \tilde{x}_i^j$, and so on. We define $h(\omega_2)$ as follows:

- $\tilde{\alpha} = \alpha$. As one can see, (B1) remains unchanged.
- $\tilde{\pi}_j = \pi_j$, for all j .
- \tilde{x}_i^j are defined like x_i^j , except for:

$$\tilde{x}_l = x_{\pi_z(l)}^z, \quad \tilde{x}_{\pi_y(l)}^y = x_l, \quad (10)$$

$$\tilde{x}_{l+1} = x_{\pi_y(l+1)}^y, \quad \tilde{x}_{\pi_z(l+1)}^z = x_{l+1}, \quad (11)$$

$$\tilde{x}_{\pi_z(l)}^z = x_{\pi_z(l+1)}^z, \quad \tilde{x}_{\pi_y(l+1)}^y = x_{\pi_y(l)}^y. \quad (12)$$

One can check that, by (10) and (11), the receipts (B3) remain unchanged.

- \tilde{r}_i^j are defined like r_i^j , except for: $\tilde{r}_{\pi_z(l)}^z = r_{\pi_z(l+1)}^z$ and $\tilde{r}_{\pi_y(l+1)}^y = r_{\pi_y(l)}^y$ (which, together with (12) implies that (B4) remains unchanged) and, furthermore, $\tilde{r}_{\pi_y(l)}^y$ and $\tilde{r}_{\pi_z(l+1)}^z$ are (uniquely) defined in such a way that

$$\text{comm}(y, \tilde{x}_{\pi_y(l)}^y)^{\tilde{r}_{\pi_y(l)}^y} = \text{comm}(z, x_{\pi_z(l)}^z)^{r_{\pi_z(l)}^z} \quad (13)$$

$$\text{comm}(z, \tilde{x}_{\pi_z(l+1)}^z)^{\tilde{r}_{\pi_z(l+1)}^z} = \text{comm}(y, x_{\pi_y(l+1)}^y)^{r_{\pi_y(l+1)}^y}. \quad (14)$$

(Note that Pedersen commitments used in this protocol guarantee that for each a and b there exists exactly one r such that $\text{comm}(a)^r = b$.)

- $\tilde{\pi}$ is as π with the straightforward adjustment such that the list of published commitments (B2) in both cases (i.e. for ω_2 and $h(\omega_2)$) is exactly the same (it can be easily done, because, as one can check, the produced commitments in both cases are, up to the ordering, the same).

- \tilde{r}_i are like r_i with the two following exceptions: \tilde{r}_l and \tilde{r}_{l+1} are (uniquely) defined in such a way that

$$\begin{aligned}\text{comm}(z, \tilde{x}_l)^{\tilde{r}_l} &= \text{comm}(y, x_l)^{r_l} \\ \text{comm}(y, \tilde{x}_{l+1})^{\tilde{r}_{l+1}} &= \text{comm}(z, x_{l+1})^{r_{l+1}}.\end{aligned}$$

- $\tilde{\sigma}_i^0$ are like σ_i^0 with the following exceptions:

$$\begin{aligned}\tilde{\sigma}_l^0(y) &= \sigma_l^0(z), & \tilde{\sigma}_{l+1}^0(y) &= \sigma_{l+1}^0(z), \\ \tilde{\sigma}_l^0(z) &= \sigma_l^0(y), & \tilde{\sigma}_{l+1}^0(z) &= \sigma_{l+1}^0(y).\end{aligned}$$

One can verify that (B5) remains unchanged.

In the following, we assume that $s = 1$; the case for $s = 2$ is very similar.

- Let $\tilde{s} = s = 1$, $\tilde{\sigma}_i^1 = \sigma_i^1$, and $\tilde{\tau}_{i,j}^1 = \tau_{i,j}^1$ for all i, j . Therefore (B9) remains the same. One can also check that (B6) remains the same (this is because (B5) remains the same and (B6) is obtained from it using the same permutations and masking factors).
- Let $\tilde{\sigma}_i^2$ be like σ_i^2 with the following exceptions:

$$\tilde{\sigma}_l^2(\tilde{\sigma}_l^1(\tilde{\sigma}_l^0(y))) = \sigma_l^2(\sigma_l^1(\sigma_l^0(y))) \quad \text{and} \quad \tilde{\sigma}_l^2(\tilde{\sigma}_l^1(\tilde{\sigma}_l^0(z))) = \sigma_l^2(\sigma_l^1(\sigma_l^0(z)))$$

and analogously for $(l + 1)$.

- Let $\tilde{\tau}_{i,j}^2$ be like $\tau_{i,j}^2$, except for $\tilde{\tau}_{l,\tilde{\sigma}_l^1(\tilde{\sigma}_l^0(z))}^2$ and $\tilde{\tau}_{l,\tilde{\sigma}_l^1(\tilde{\sigma}_l^0(y))}^2$ which are defined in such a way that

$$\tilde{r}_l + \tilde{\tau}_{l,\tilde{\sigma}_l^0(z)}^1 + \tilde{\tau}_{l,\tilde{\sigma}_l^1(\tilde{\sigma}_l^0(z))}^2 = r_{\pi_z(l)}^z + \tau_{l,\sigma_l^0(z)}^1 + \tau_{l,\sigma_l^1(\sigma_l^0(z))}^2 \quad (15)$$

$$\tilde{r}_{\pi_y(l)}^y + \tilde{\tau}_{l,\tilde{\sigma}_l^0(y)}^1 + \tilde{\tau}_{l,\tilde{\sigma}_l^1(\tilde{\sigma}_l^0(y))}^2 = r_l + \tau_{l,\sigma_l^0(y)}^1 + \tau_{l,\sigma_l^1(\sigma_l^0(y))}^2. \quad (16)$$

and analogously for $(l + 1)$. Now, one can check that (B7) and (B8) remain the same.

This concludes the description of $h(\omega_2)$. As we have noted, all the parts (B1)–(B9) of the views in both cases—for ω_2 and $h(\omega_2)$ —are exactly the same. What remains to be shown is that h is a bijection from Ω_2 to Ω_2 . To do this, it is enough to prove that ω_2 can be uniquely determined by $\tilde{\omega}_2 = h(\omega_2)$. We only have to deal with those parts of ω_2 that are changed by h . We consider those changed parts of ω_2 case by case:

- It is easy to see that the numbers x_i and x_i^j (for $j \in \{1, \dots, k\}$ and $i \in \{0, \dots, m\}$) are uniquely determined by the numbers \tilde{x}_i and \tilde{x}_i^j .
- $r_{\pi_z(l)}^z$ can be computed from $\tilde{\omega}_2$, as it is uniquely determined by the equality (13) (recall that $x_{\pi_z(l)}^z$, as we already stated, is determined by $\tilde{\omega}_2$). Analogously for $r_{\pi_z(l+1)}^z$, r_l , and r_{l+1} . Apart from these values, r_i^j and r_i coincide with \tilde{r}_i^j and \tilde{r}_i , respectively, and therefore are determined by $\tilde{\omega}_2$.

- The permutations in $\tilde{\omega}_2$ are obtained from the corresponding permutation of ω_2 , by switching some selected positions. It is easy to define the inverse operation.
- if $\tilde{s} = s = 1$, then $\tau_{l,\sigma_l^1}^2(\sigma_l^0(z))$ is uniquely determined by (15), as all the other parts in the equation are determined by $\tilde{\omega}_2$. (Note that $\tau_{l,\sigma_l^0}^1$ is not changed if $s = 1$.) Analogously for $l + 1$ and for the case $s = 2$.

C Proofs for ThreeBallot with Two Candidates

In this section, we prove Theorem 3 and Lemma 4, where for Theorem 3 we only prove the second statement, i.e., the more involved case in which the coercer gets to see the receipts of the honest parties; the proof of the first statement is analogous and simpler. The proof of Lemma 3 is simple (much simpler than the one for Lemma 4) and therefore omitted.

We first introduce some notation. We will assume that the space Ω_1 of all possible combinations of choices made by honest voters determines not only the candidate the voters have chosen, but also the way they vote, i.e. the exact pattern (see Section 6.3 for the definition of a pattern). We define the following random variables on Ω_1 . $F(\omega_1)$ is the number of \times -ballots, $rec(\omega_1)$ is the vector $(r_{\times}, r_{\circ}, r_{\circ})$ of numbers of receipts of honest voters of the corresponding types, $Rec(\omega_1)$ is the vector (r_1, \dots, r_n) , where $r_i \in \{\times, \times, \circ, \circ\}$ is the receipt of the i -th honest voter (without a serial number), $R(\omega_1)$ is the number of votes of honest voters for candidate 1, $S(\omega_1)$ is the set of non-abstaining honest voters, and $N(\omega_1)$ is the number of non-abstaining honest voters (determined by $\omega_1 \in \Omega_1$).

C.1 Proof of Theorem 3

First, 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 coercer, $\vec{r} = (r_{ij})_{i \in \{0, \dots, m\}, j \in \{1, 2, 3\}}$, 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 coerced voter), and π is a permutation applied to the set of ballots before publishing. As usually, by μ_2 we denote the uniform distribution on Ω_2 . (Note that a serial number r_{ij} , $j \in \{1, 2, 3\}$, is not printed, if the i -th voter does not vote.)

A view of the coercer consists of (1) his random coins, (2) the content of the bulletin board, which is a sequence of simple ballots with serial numbers, and (3) the sequence of receipts (where, again, a receipt is a simple ballot with a serial number) associated to the voters. We will use letter η to range over views of the coercer. Recall that ρ is used to denote an essential view of the coercer.

By $\rho(\eta)$ we will denote the essential view determined by η . By $\rho(Z, \omega_1)$, for a pattern Z and $\omega_1 \in \Omega_1$, we denote the essential view obtained when the coerced voter casts simple ballots according to Z and the honest voters casts ballots determined

by ω_1 . By $f(\eta)$ we denote the pattern that the coercer requires the coerced voter to use in a run η , if any; otherwise $f(\eta)$ is undefined. By $Rec(\eta)$ we denote the receipts *without serial numbers* that the honest voters give to the coercer in run η .

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

Let c be a program of the coercer. Let $T = (c \parallel \text{dum} \parallel \mathbf{e}_S)$ and $\tilde{T} = (c \parallel \tilde{v} \parallel \mathbf{e}_S)$.

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 formulated in Lemma 8, with the core stated in the following lemma.

Lemma 7. *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.

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

Let $I = \{1, \dots, n\} \times \{1, 2, 3\}$. For $(i, l) \in I$, by $b_{i,l} \in \{\times, \circ, \times, \circ\}$ and $b'_{i,l} \in \{\times, \circ, \times, \circ\}$ we denote the marking on the l -th ballot cast by the i -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 $b'_{(i,l)} = b_{\sigma(i,l)}$. Moreover, we can assume that σ preserves receipts of honest voters, i.e. if the i -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 $\sigma(i, l') = (i, l)$. Note that, in this case, $b_{(i,l)} = b'_{(i,l')}$.

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

From this, we can conclude.

Lemma 8. *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:*

$$\Pr[T \mapsto \eta] = \Pr_{\omega_1}[\varphi_\eta(\omega_1)] \cdot \Pr_{\omega_2}[T(\omega_1^\eta, \omega_2) \mapsto \eta] \quad (17)$$

$$\Pr[\tilde{T} \mapsto \eta] = \Pr_{\omega_1}[\tilde{\varphi}_\eta(\omega_1)] \cdot \Pr_{\omega_2}[\tilde{T}(\tilde{\omega}_1^\eta, \omega_2) \mapsto \eta] \quad (18)$$

$$\Pr_{\omega_2}[T(\omega_1^\eta, \omega_2) \mapsto \eta] = \Pr_{\omega_2}[\tilde{T}(\tilde{\omega}_1^\eta, \omega_2) \mapsto \eta] \quad (19)$$

Proof. Using Lemma 7 we obtain:

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

This proves (17). One can prove (18) in an analogous way. The equation (19) follows directly from Lemma 7. \square

Now, using Lemma 8, we can link the level of coercion-resistance ThreeBallot provides with the optimal bound $\delta_{TB^+}^i$ stated in Section 6. Clearly we have:

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

and

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

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

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

Now, we are ready to prove that the system S , as defined in Theorem 3 in the case the coercer can see the receipts of honest voters, is δ -coercion resistant w.r.t. γ_i for $\delta = \delta_{TB^+}^i(n, \vec{p})$.

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, and Rec over all possible vectors of receipts. We abbreviate $C(Z, i)$ by $C(Z)$. Finally, let $M_Z^{\rho, Rec} = \{\eta \in M : f(\eta) = Z, Rec(\eta) = Rec, \text{ and } \rho(\eta) = \rho\}$. Let $\omega_1^{Z, \rho, Rec}$ be arbitrary with $\rho(Z, \omega_1^{Z, \rho, Rec}) = \rho$ 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:

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

With $M_Z^+ = M_{Z,i}^+$ as defined in Section 6.3, we get:

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

Next, we use that, 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 , $T(\omega_1^\eta, \omega_2) \mapsto \eta$ implies $f(\omega_2) = Z$. With this, we obtain:

$$\begin{aligned}
\Pr_{\omega_2}[T(\omega_1^\eta, \omega_2) \mapsto \eta] &= \Pr_{\omega_2}[f(\omega_2) = Z] \\
&\quad \cdot \Pr_{\omega_2}[T(\omega_1^\eta, \omega_2) \mapsto \eta \mid f(\omega_2) = Z].
\end{aligned}$$

Now we obtain

$$\Phi \leq \sum_Z \Pr_{\omega_2}[f(\omega_2) = Z] \sum_{\rho \in M_Z^+} (A_\rho^Z - A_\rho^{C(Z)}) \tag{20}$$

$$\leq \sum_Z \Pr_{\omega_2}[f(\omega_2) = Z] \cdot \delta_{TB+}^i \tag{21}$$

$$\leq \delta_{TB+}^i = \delta \tag{22}$$

This shows that S is δ -coercion resistant w.r.t. γ_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 $M_{Z^*,i}^+$, where $M_{Z^*,i}^+$ is as defined in Section 6.3, and Z^* is a pattern with

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

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

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

We immediately obtain:

$$\Phi = \sum_{\rho \in M_{Z^*,i}^+} (\Pr[T \mapsto \rho] - \Pr[\tilde{T} \mapsto \rho]) = \sum_{\rho \in M_{Z^*,i}^+} (A_\rho^{Z^*} - A_\rho^{C(Z^*)}) = \delta , \quad (23)$$

which shows that S is δ' -coercion resistant for any $\delta' < \delta$, in case the counter-strategy \tilde{v} is used. To complete the proof of Theorem 3, 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 violate Condition (i) of Definition 1. Conversely, to satisfy Condition (i), the coerced voter only needs to vote if instructed by the coercer. Therefore, to be as indistinguishable from the dummy strategy as possible, it is clear that 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 according to $Z_0 = (\overset{x}{\circ}, \overset{o}{x}, \overset{o}{x})$, assuming $i = 1$; the case $i = 2$ is analogous. By (b) and (c) we know this response must be $(\overset{x}{\circ}, \overset{o}{x}, \overset{x}{\circ})$ or $(\overset{x}{\circ}, \overset{x}{x}, \overset{o}{\circ})$. One of these responses can be chosen randomly, according to some strategy. Recall from Section 6.3 that the response for the counter-strategy \tilde{v} is $C(Z_0) = C(Z_0, 1) = (\overset{x}{\circ}, \overset{o}{x}, \overset{x}{\circ})$. If $\tilde{C}'(Z) = \tilde{C}'(Z, 1)$ denotes the response for a pattern Z in the counter-strategy \tilde{v}' , we know that $\tilde{C}'(Z) = C(Z)$ for every $Z \neq Z_0$. For $Z = Z_0$, as just explained, $\tilde{C}'(Z)$ has two choices which could be chosen randomly, according to some strategy. For simplicity of the argument, we assume that $\tilde{C}'(Z_0)$ always chooses $(\overset{x}{\circ}, \overset{x}{x}, \overset{o}{\circ})$; the case of a randomized choice can be treated similarly. (Note that whenever $\tilde{C}'(Z)$ chooses $(\overset{x}{\circ}, \overset{o}{x}, \overset{x}{\circ})$, then this would coincide with $C(Z_0)$.)

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 $\tilde{M}_{Z^*,i}^+$, where $\tilde{M}_{Z^*,i}^+ = \{\rho : A_\rho^Z \geq A_\rho^{\tilde{C}'(Z^*)}\}$, and Z^* is a pattern such that

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

With this, analogously to (23), we have:

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

Hence it remains to show that

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

Let $Z_1 = (\overset{\circ}{\times}, \overset{\times}{\circ}, \overset{\circ}{\times})$. Then $\tilde{C}'(Z_1)$ is uniquely determined and equal to $C(Z_1)$. As the receipt of the coerced voter is *not* part of the essential view, we have for all ω_1 :

$$\rho(Z_1, \omega_1) = \rho(Z_0, \omega_1) \text{ and } \rho(C(Z_1), \omega_1) = \rho(C(Z_0), \omega_1).$$

It follows:

$$\sum_{\rho \in M_{Z_1,i}^+} (A_\rho^{Z_1} - A_\rho^{C(Z_1)}) = \sum_{\rho \in M_{Z_0,i}^+} (A_\rho^{Z_0} - A_\rho^{C(Z_0)}).$$

Now, we obtain:

$$\begin{aligned} \max_Z \sum_{\rho \in \tilde{M}_{Z,i}^+} (A_\rho^Z - A_\rho^{\tilde{C}'(Z)}) &\geq \max_{Z \neq Z_0} \sum_{\rho \in \tilde{M}_{Z,i}^+} (A_\rho^Z - A_\rho^{\tilde{C}'(Z)}) \\ &= \max_{Z \neq Z_0} \sum_{\rho \in M_{Z,i}^+} (A_\rho^Z - A_\rho^{C(Z)}) \\ &= \max_Z \sum_{\rho \in M_{Z,i}^+} (A_\rho^Z - A_\rho^{C(Z)}). \end{aligned}$$

This concludes the proof of Theorem 3.

C.2 Proof of Lemma 4

In the proof of Lemma 4, we will use the following easy to prove facts (see [9] for similar results).

Lemma 9. *Consider honest, non-abstaining voters.*

1. *The probability that a voter takes receipt $\overset{\times}{\times}$ is $\frac{1}{9}$.*
2. *The probability that a voter takes receipt $\overset{\times}{\circ}$ and the probability that she takes receipt $\overset{\circ}{\times}$ is $\frac{2}{9}$.*
3. *The probability that a voter who does not abstain votes for candidate 1 (or candidate 2) is independent of the receipt she gets and is $\frac{p_1}{p_1+p_2}$ (or $\frac{p_2}{p_1+p_2}$, respectively).*
4. *The probability that a voter produces a $\overset{\times}{\times}$ -ballot is $\frac{1}{2}$ in either of the following cases: (a) if we assume that she votes for candidate 1 and takes $\overset{\circ}{\times}$ as a receipt, and (b) if she votes for candidate 2 and takes $\overset{\times}{\circ}$ as a receipt.*

We introduce two new random variables: $\tau_1(\omega_1)$, indicating the number of voters that vote for 2 and take $\overset{\times}{\times}$ as receipt; $\tau_2(\omega_1)$, indicating the number of voters that vote for 1 and take $\overset{\circ}{\times}$ as receipt.

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

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

Moreover,

$$\begin{aligned} & \Pr_{\omega_1}[\rho(\omega_1) = \rho, \tau_1(\omega_1) = \tau_1, \tau_2(\omega_1) = \tau_2] = \\ &= \Pr_{\omega_1}[F(\omega_1) = n_x^x, R(\omega_1) = R, \tau_1(\omega_1) = \tau_1, \tau_2(\omega_1) = \tau_2, \text{rec}(\omega_1) = (r_x^x, r_x^o, r_x^o), N(\omega_1) = N] = \\ &= \Pr_{\omega_1}[F(\omega_1) = n_x^x \mid R(\omega_1) = R, \tau_1(\omega_1) = \tau_1, \tau_2(\omega_2) = \tau_2, \text{rec}(\omega_1) = (r_x^x, r_x^o, r_x^o), N(\omega_1) = N] \cdot \end{aligned} \quad (24)$$

$$\cdot \Pr_{\omega_1}[R(\omega_1) = R \mid \tau_1(\omega_1) = \tau_1, \tau_2(\omega_2) = \tau_2, \text{rec}(\omega_1) = (r_x^x, r_x^o, r_x^o), N(\omega_1) = N] \cdot \quad (25)$$

$$\cdot \Pr_{\omega_1}[\tau_1(\omega_1) = \tau_1, \tau_2(\omega_2) = \tau_2 \mid \text{rec}(\omega_1) = (r_x^x, r_x^o, r_x^o), N(\omega_1) = N] \cdot \quad (26)$$

$$\cdot \Pr_{\omega_1}[\text{rec}(\omega_1) = (r_x^x, r_x^o, r_x^o) \mid N(\omega_1) = N] \cdot \quad (27)$$

$$\cdot \Pr_{\omega_1}[N(\omega_1) = N] \cdot \quad (28)$$

For (28), we have

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

For (27), we have to distribute (independently) the receipts to the N non-abstaining voters. With Lemma 9 we obtain:

$$\begin{aligned} & \Pr_{\omega_1}[\text{rec}(\omega_1) = (r_x^x, r_x^o, r_x^o) \mid N(\omega_1) = N] \\ &= \binom{N}{r_x^x, r_x^o, r_x^o} \left(\frac{1}{9}\right)^{r_x^x} \left(\frac{2}{9}\right)^{r_x^o + r_x^o} \left(\frac{4}{9}\right)^{N - r_x^x - r_x^o - r_x^o}. \end{aligned}$$

For (26), we have to distribute $r_x^o - \tau_1$ votes resp. τ_2 votes for cand_1 in the set of those voters that get $\overset{x}{\circ}$ resp. $\overset{o}{\circ}$ as receipt. Using Lemma 9, we obtain

$$\begin{aligned} & \Pr[\tau_1(\omega_1) = \tau_1, \tau_2(\omega_1) = \tau_2 \mid \text{rec}(\omega_1) = (r_x^x, r_x^o, r_x^o), N(\omega_1) = N] \\ &= \binom{r_x^o}{\tau_1} q^{r_x^o - \tau_1} (1 - q)^{\tau_1} \cdot \binom{r_x^o}{\tau_2} q^{\tau_2} (1 - q)^{r_x^o - \tau_2}. \end{aligned}$$

where $q = \frac{p_1}{p_1 + p_2}$.

For (25), we have to distribute the rest of the votes for cand_1 (i.e. $R - (r_x^o - \tau_1) - \tau_2$) to those non-abstaining voters that do not get $\overset{x}{\circ}$ or $\overset{o}{\circ}$ as receipt. With Lemma 9 we have that the probability that a non-abstaining voter votes for cand_1 is q , regardless

of the receipt. Hence we have

$$\begin{aligned} \Pr_{\omega_1} [R(\omega_1) = R \mid \tau_1(\omega_1) = \tau_1, \tau_2(\omega_1) = \tau_2, \text{rec}(\omega_1) = (r_x, r_o, r_x), N(\omega_1) = N] = \\ = \binom{N - r_x - r_o}{R - (r_x - \tau_1) - \tau_2} \cdot q^{R - (r_x - \tau_1) - \tau_2} \cdot (1 - q)^{N - r_o - R - \tau_1 + \tau_2}. \end{aligned}$$

For (24), we have to spot n_x voters that submit an $\overset{x}{\times}$ -ballot among all voters. Clearly, every voter that takes $\overset{x}{\times}$ or $\overset{o}{\circ}$ as receipt, submits an $\overset{x}{\times}$ -ballot. Also, the voters who vote according to τ_1 or τ_2 do not submit a $\overset{x}{\times}$ -ballot (that was the reason for introducing τ_1, τ_2). Hence we have to distribute $n_x - r_x - (N - r_x - r_o - r_x) = n_x - N + r_x + r_o$ among $N - r_x - (N - r_x - r_o - r_x) - \tau_1 - \tau_2$ voters. Note that any of those voters either votes for cand_1 with receipt $\overset{x}{\circ}$ or for cand_2 with receipt $\overset{o}{\times}$. The probability that such a voter submits a $\overset{x}{\times}$ -ballot is $\frac{1}{2}$, according to Lemma 9. Hence we have

$$\begin{aligned} \Pr [F(\omega_1) = n_1 \mid R(\omega_1) = R, \tau_1(\omega_1) = \tau_1, \tau_2(\omega_1) = \tau_2, \text{rec}(\omega_1) = (r_x, r_o, r_x), N(\omega_1) = N] = \\ = \binom{r_x + r_o - \tau_1 - \tau_2}{n_x - N + r_x + r_o} \left(\frac{1}{2}\right)^{r_x + r_o - \tau_1 - \tau_2}. \end{aligned}$$

By putting everything together and rewriting the formula, we obtain the formula in Lemma 4.

D Proof of Theorem 4

Let us consider the program $c \in \mathcal{C}$ which:

- instructs the coerced voter to vote for the candidate j for which the sum from equation (6) achieves its maximum. (Note that, by the definition of \mathcal{C} , the exact pattern the coerced voter is supposed to use is determined.)
- accepts a run if and only the receipt given by the voter is as required and the restricted view ρ in this run is in $M_{i,j}$.

Let v^* be the counter-strategy as defined in Section 6. As we argued Section 6, this strategy is optimal for \mathcal{C} and therefore for c . Hence, to prove Theorem 4, it suffices to show that

$$\Phi = \Pr[T \mapsto 1] - \Pr[\tilde{T} \mapsto 1] \geq \delta,$$

where $T = (c \parallel \text{dum} \parallel e_S)$, $\tilde{T} = (c \parallel v^* \parallel e_S)$, and $\delta = \delta_i(n, k, \vec{p})$. We show, in fact, that $\Phi = \delta$:

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

where we use the equalities

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

These hold true, because the events $T \mapsto \rho$ and $\tilde{T} \mapsto \rho$ depend only on the choices made by honest voters and the coerced voter. This concludes the proof of Theorem 4.