# Is the JCJ voting system really coercion-resistant? 

Véronique Cortier, Pierrick Gaudry, Quentin Yang

## To cite this version:

Véronique Cortier, Pierrick Gaudry, Quentin Yang. Is the JCJ voting system really coercion-resistant?. 37th IEEE Computer Security Foundations Symposium (CSF), 2024, Enschede, Netherlands. hal03629587v2

HAL Id: hal-03629587<br>https://inria.hal.science/hal-03629587v2

Submitted on 17 May 2023 (v2), last revised 13 Feb 2024 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers. publics ou privés.

Distributed under a Creative Commons Attribution 4.0 International License

# Is the JCJ voting system really coercion-resistant? 

Véronique Cortier, Pierrick Gaudry, Quentin Yang<br>Université de Lorraine, Inria, CNRS


#### Abstract

Coercion-resistance is a security property of electronic voting, often considered as a must-have for high-stake elections. The JCJ voting scheme, proposed in 2005 by Juels, Catalano and Jakobsson, is still the reference paradigm when designing a coercion-resistant protocol. We highlight a weakness in JCJ that is also present in all the systems following its general structure. This comes from the procedure that precedes the tally, where the trustees remove the ballots that should not be counted. This phase leaks more information than necessary, leading to potential threats for the coerced voters. Fixing this leads to the notion of cleansing-hiding, that we apply to form a variant of JCJ that we call CHide. One reason for the problem not being seen before is the fact that the associated formal definition of coercion-resistance was too weak. We therefore propose a definition that takes into account more behaviors such as revoting or the addition of fake ballots by authorities. We then prove that CHide is coercion-resistant for this definition.


## I. Introduction

Internet voting allows to take part in an election without being physically present at a polling station. It can be used for many reasons, such as providing people with low mobility or expatriates with a way to vote beside postal voting, or as a necessary alternative during a pandemic. As of today, electronic voting has been used for politically-binding elections in several countries (e.g. Australia, Switzerland, Estonia, to name but a few). For such high-stakes contexts, coercion may be an important threat. It occurs when an attacker forces a voter to vote in a specific way, using a threat or a reward. This phenomenon is known to exist in real-world elections, with traditional voting at polling stations. An electronic voting solution which is not designed to tackle coercion could allow the attacker to coerce a larger number of voters, or to gain a more convincing evidence that the coerced voters actually obeyed. Also, since Internet voting is a remote voting process, this introduces new attacks compared to polling station voting. For instance, the coercer can ask the voter to give all the voting materials that they received. The classical verifiability mechanisms will then provide a proof to the coercer that the voter did not cheat.

The JCJ protocol. A famous protocol which aims to counter coercion was proposed in 2005 by Juels, Catalano and Jakobsson [22]. They also provide a formalization of the notion, allowing to give security arguments. This is now called the JCJ protocol and remains the reference for the research on coercion-resistance in electronic voting. The key idea of JCJ is that voters can give fake voting material (a fake credential) to the coercer, and pretend that it is genuine. The coercer, who votes with the provided credential, has no way to detect

[^0]whether the credential is valid or not. In order to guarantee that, during the voting phase, ballots are accepted in the ballot box regardless of their credentials; those which use an invalid or a duplicate one are removed later, during a cleansing phase. The output of this cleansing phase is a set of ballots that is tallied in the usual way. The main security feature is that, given a credential and all the publicly available information, the coercer is unable to tell whether the credential is real or fake. At the same time, for the legitimate voters, verifiability is preserved.

For coercion-resistance, the cleansing phase is critical. In JCJ, some information is leaked, namely the number of ballots sent to the public board and to the tallying procedure, before and after the cleansing phase. It is known that the difference $\Delta$ between those two can reveal some information to the coercer. Therefore it is important to ensure some "noise", coming from revotes or dummy ballots, that would mask the action of a voter resisting a coercer. Still, depending on its exact definition, the cleansing phase could leak more than just $\Delta$. For instance, in [32], the authors present a protocol where the coercer can deduce the number of ballots which pretend to be from each voter, and exploit this additional information (which is not available in the original JCJ). To mitigate this leakage, they propose that the authorities add a random number of dummy ballots for each voter, which mitigates the leakage.

A weakness in JCJ. We reveal that even in the original JCJ protocol, the cleansing step leaks more than the difference $\Delta$ between the sizes of its input and output. Indeed, first, the ballots with the same credentials (i.e. revotes) are handled, keeping only one ballot per credential. Second, the ballots corresponding to invalid credentials are removed. The size of the intermediate ballot box is leaked; moreover, anyone can observe the distribution of the number of revotes. We provide a few examples where this allows to fully break coercionresistance, even if dummy ballots are used. Admittedly, these extreme scenarios are unlikely to occur in practice. Therefore, we explore more realistic scenarios where major revoting may occur, for example due to a technical incident where voters are encouraged to revote. Another case is when a discredit of a candidate happens during the voting phase, yielding voters to flip their votes. In such cases, coercion resistance is not fully broken but the coercer gains some non-negligible advantage due to the extra leakage of JCJ: they know whether it is more likely that the voter obeyed or not. In order to formally measure the loss of coercion-resistance, we follow the approach of Kuesters et al [25] and show that JCJ guarantees a lower level of coercion resistance than a more ideal protocol which only leaks $\Delta$. All the variants and improvements on JCJ
that we know of are also affected by this vulnerability.
A cleansing-hiding protocol. We propose a modification of JCJ, called CHide, that is not subject to this weakness. The key modification is the introduction of a cleansing hiding procedure, that replaces the original cleansing phase. More precisely, while JCJ uses plaintext equivalence tests (PET) as the main cryptographic tool, CHide relies on more complex MPC building blocks. Instead of PETs that return a bit telling whether two ciphertexts represent the same cleartext, we use a primitive Eq that returns an encrypted version of this bit, in order to hide which ballots are removed and for which reason (being a revote or having a fake credential). Then, a few logical gates primitives (e.g. Or, And, etc.) must be operated on these encrypted bits to obtain an encrypted validity bit. Finally, we use a Conditional_Set_Zero primitive that conditionally sets the ballots to (fixed) invalid ones, depending on their encrypted validity bit. Thanks to its cleansing hiding procedure, CHide only reveals the minimal information, namely the number $\Delta$ of ballots that have been removed. Of course, each step comes with a zero-knowledge proof that the expected operations have been performed, so that anyone can check that the result of the election is correct.

A stronger notion of coercion-resistance. The leakage in JCJ was not noticed before because their definition of coercionresistance was too weak. A flaw was already noted in [17], which shows that the definition could not be realized. The fix proposed in [17] repairs this but still does not consider the cases where voters may revote and hence misses the situation where JCJ leaks too much information. These definitions also model the addition of ballots with fake credentials in a contrived way in the sense that each fake ballot must be cast by a voter that sacrifices their vote.

We propose a more generic definition of coercion-resistance that accounts for revoting as well as the addition of fake ballots by authorities. We prove that CHide is coercionresistant according to this definition, and that JCJ is not, thus showing that we indeed capture the weakness of the leakage during the cleansing phase. We also prove that CHide ensures vote privacy (under weaker assumptions than for coercionresistance) as well as universal verifiability.

## Summary of the contributions.

- We discover a vulnerability in the JCJ scheme, which shows that it is not perfectly coercion-resistant.
- We formally measure the loss of coercion-resistance in JCJ in two realistic scenarios.
- We propose CHide, a cleansing-hiding variant of JCJ, that is not subject to this problem.
- We propose a new definition of coercion-resistance, which properly takes revoting and dummy ballots into account.
- We prove that CHide is coercion-resistant for this definition, while JCJ is not, thus showing the definition is precise enough to capture the leakage during cleansing.
Related work. To provide coercion-resistance, many authors used the core idea of JCJ, that is the fake credential paradigm.

Civitas [11] is one of the most notable examples. It is widely considered as an important step towards a practical version of JCJ. Among other things, it introduces the notion of ballot blocks, that mitigates the quadratic cost of the cleansing phase of JCJ. This reduces somehow its coercion-resistance, which can be captured with our approach.

Other attempts were made to improve the efficiency of JCJ. In [32], Spycher et al. claim a linear time cleansing, but this comes with a deterioration of the coercion-resistance, as explained above. Later on, the same authors proposed other schemes with a clear trade-off between efficiency and coercion-resistance, thanks to anonymity sets [33]. Other improvements include [1], where Araújo et al. propose a way to perform the cleansing phase in linear time, and [10], where Clark and Hengartner introduce the idea of over-the-shoulder coercion-resistance. Both schemes would suffer from the same cleansing leakage as JCJ, but it can be observed directly from the public board, when the adversary can still submit ballots. In contrast, the leakage in JCJ can only be observed during the cleansing phase, after the end of the voting phase.

Apart from the fake credential paradigm, a natural approach is to use deniable revoting, where the coerced voter first complies with the coercer but revotes later. The ballot cast under coercion is invalidated by the subsequent revote, in a way that the coercer cannot detect. The Estonian voting system [29] completely relies on revoting to mitigate coercion, and examples of recent academic proposals based on revoting are VoteAgain by Lueks et al. [28] and the scheme of Locher et al. [27]. This approach assumes a weaker adversary which can no longer submit ballots passed a certain point.

All these schemes (including JCJ) do not address the socalled Italian attacks. The latter exist independently of the use of electronic voting and are based on the information given by the tally. They can occur when the ballots are complex enough, so that voters can "sign" them using a specific pattern on the low-stake parts of the answer. When using electronic voting, the typical way to prevent such attacks is tally hiding, i.e. to decrypt only the winners of the election, without decrypting the individual ballots. The challenge is then to preserve verifiability. Therefore tally-hiding usually relies on homomorphic encryption, or more generally on multiparty computation (MPC) techniques [5], [13], [19], [24]. Designing a tally-hiding scheme is out of scope of this paper but interestingly, CHide could be easily coupled with such tally-hiding schemes, right after the cleansing phase.

Regarding formal definition of coercion-resistance, we already mentioned the recent work of Haines and Smyth [17] that attempts to survey and unify the various definitions of the literature, starting with the one of JCJ where the adversary must distinguish between a real and an ideal game modeling the protocol. Another approach to define coercion-resistance is given by Küsters et al. in [25]. The authors define $\delta$-coercionresistance with two conditions: first, the coerced voter must have a strategy to meet their objective with overwhelming probability (i.e. they can actually vote for the desired candidate); second, the adversary cannot distinguish the case when
the voter uses their evasion strategy from the case where the voter forwards all received messages to the adversary, with an advantage greater than $\delta$. Our definition of coercion-resistance can be seen as an instance of [25] where we compare the $\delta$ -coercion-resistance of the real protocol with the $\delta^{\prime}$-coercionresistance of the ideal one, requiring them to be equal up to negligible difference. While [25] provides a very general and abstract definition, we instead consider a fixed objective for the coerced voter (voting for a given candidate), a fixed evasion strategy (the one considered in JCJ) and a formal framework to model the trust assumptions on how messages are delivered. This part is left unspecified in [25] and needs to be shaped for each protocol.

## II. UnVEILING A SHORTCOMING IN JCJ

We provide a high level description of the JCJ protocol and show why coercion-resistance is undermined in case of revoting.

## A. Overview of JCJ

The JCJ voting system consists of the following phases.
Setup. The Election Trustees jointly generate the election public key $\mathrm{pk}_{T}$, that is sent to the public board. Then, the Registrars jointly compute one credential $c$ for each voter. Each credential is sent privately to the voter, possibly with designated zero-knowledge proofs to guarantee voters that their credential is valid [21]. The Registrars send to the public board the list $\boldsymbol{R}=\left\{\operatorname{Enc}\left(c_{1}, \mathrm{pk}_{T}\right), \ldots, \operatorname{Enc}\left(c_{n}, \mathrm{pk}_{T}\right)\right\}$ of encrypted credentials, in some random order.

Voting. To cast a ballot, a voter encrypts their vote $\nu$ with the election public key $\mathrm{pk}_{T}$. They also encrypt their credential $c$ and prove knowledge of $\nu$ and $c$. They prove that the encryptions are valid and linked, yielding a proof $\pi$. The resulting ballot $\mathbf{b}=\left(\operatorname{Enc}\left(\nu, \mathrm{pk}_{T}\right), \operatorname{Enc}\left(c, \mathrm{pk}_{T}\right), \pi\right)$ is sent anonymously to the bulletin board. The voting phase is depicted in Figure 1.

Tallying. The tally phase consists of four steps.

1. Ballots with duplicated credentials are detected using Plaintext Equivalence Tests [20] (PET). At most one ballot (typically, the last) is kept per credential.
2. The trustees mix the remaining ballots.
3. PETs are used again to remove ballots with invalid credential, i.e. whose credential is not encrypted in $\boldsymbol{R}$.
4. Finally, each remaining ballot is decrypted so that the result can be computed.
Each step includes a zero-knowledge proof that the correct operations are performed.
a) Evading coercion.: The JCJ voting system provides an evasion strategy that a voter under coercion must follow to safely disobey a coercer. If Alice is under coercion, she simply provides her coercer with a random (and fake) credential $c^{\prime}$. She can later use her real credential $c$ to cast her vote. Ballots containing $c^{\prime}$ will be removed at Step 3 of the tally phase. Thanks to the mixnet, the coercer is unable to learn that their ballot has been suppressed.

Public board


Fig. 1: Voting phase in JCJ.

## B. Leakage in case of revoting

For a verifiable voting system, it seems unavoidable to leak the number of received ballots in the public board. The number of valid ballots is also leaked unless more sophisticated tally methods are used such as tally-hiding schemes [13], [19], [24].

However, JCJ leaks much more information when revoting is allowed, namely:

- $n_{b}$, the total number of received ballots;
- $n_{v}$, the total number of valid (and counted) ballots;
- $n_{r}$, the total number of revotes;
- the complete distribution of revotes, per (encrypted) credential (hence, for all $k$, the number of credentials used to revote $k$ times).
This can be exploited by a coercer to detect when a coerced voter disobeys. Indeed, there is no reason to assume that revoting is independent from the choice of candidate. On the contrary, revoting is often due to voters changing their mind between candidates, for instance due to some late announcements in the press.
a) An attack against coercion-resistance: We consider an extreme case, with two candidates $A$ and $B$. Suppose that voters voting for $A$ do not revote while those voting for $B$ always revote, exactly once. We denote $r_{A}$ (resp. $r_{B}$ ) the number of votes for $A$ (resp. $B$ ). Due to the considered revoting behaviors, the number of revotes $n_{r}$ corresponds to the number of votes for $B$ sent by the honest voters.

Assume now that Alice wants to vote for $B$ but is instructed by her coercer to vote for $A$ (or abstain).

- If Alice obeys, the coercer will observe $r_{B}=n_{r}$.
- If Alice disobeys and casts one ballot for $B$, the coercer will observe that $r_{B}=n_{r}+1$.
Hence the coercer will detect that Alice has disobeyed, which breaks coercion-resistance.

One could argue that Alice should follow a different evasion strategy and cast one ballot if she votes for $A$ and two if she votes for $B$. This does not work either. Indeed, assume now that Alice wants to vote for $A$, but is instructed to vote for $B$.

- If she obeys, she gives her real credential $c$ to her coercer. The latter then casts exactly one ballot for $B$ using $c$.
- Otherwise, she provides a fake credential $c^{\prime}$, that the coercer uses to vote for $B$. Alice then votes for $A$ using $c$. In the first case, $r_{B}=n_{r}+1$, while in the second case, $r_{B}=n_{r}$. Once again, the coercer is able to detect that Alice disobeyed and coercion-resistance is broken.


## C. Discussion

1) Considering other evasion strategies: One possibility to correct JCJ's flaw is to define other evasion strategies in case of revoting. Indeed, if Alice wants to vote, JCJ's evasion strategy instructs her to do so exactly once. Consequently, if it is usual for everyone to revote several times, the leakage in JCJ allows the coercer to detect that a single person voted once without revoting, and thus that Alice disobeyed. However, it seems very hard to instruct voters to use revoting, according to a certain distribution, when they are under coercion. As illustrated in Section II-B, the natural way to proceed does not work. This is made even harder by the fact that the strategy may evolve depending on new events that could change the revoting distribution for the honest voters.

Hence we propose another option (see Section IV) that consists in reinforcing JCJ in case of revoting, such that there is no leakage besides the total number of ballots and the number of valid ballots. For our proposed protocol, we prove coercion-resistance with the original evasion strategy of JCJ. We acknowledge that the latter is not perfect; in particular, it does not allow a voter under coercion to change their mind and revote. However, modeling a wide variety of behaviors for the coerced voter is too complex and is out of the scope of this article.
2) More noise is needed: A known issue of JCJ is that fake ballots should be randomly added, in order to hide to a coercer that their ballot has been removed. Indeed, if it is usual that absolutely no ballot with a fake credential is removed at Step 3 of the tally, then a coercer, who observes that exactly one ballot is removed, would suspect that the coerced voter has disobeyed.

Hence, it is necessary that an unpredictable number of ballots is removed during the tally. In JCJ, this "noise" comes from honest voters sending dummy ballots, but this source alone may not be sufficient. A natural approach is to have the authorities add a random number of dummies. For instance, [32] uses this to mitigate a leakage during the tally. This noise made of fake ballots should however be calibrated carefully since the computation overhead is important. In a context where revoting is a well spread behavior, it could be judicious to rely on revoting, at least partially, as an additional source of noise. This is not possible in JCJ where a dummy can be distinguished from a revote, but becomes a possibility if our solution from Section IV is used.

## III. The impact on coercion-Resistance

In Section II, we explain the leakage of the JCJ protocol and we illustrate, in an extreme scenario, how this can be exploited to completely break coercion-resistance. In this section, we estimate the impact of the leakage in more realistic scenarios. For this purpose, we use the framework of [25] which allows to quantify the coercion level of a voting protocol.

## A. Quantifying coercion-resistance

We consider $n_{V}$ voters, among which one is under coercion. The others are supposed honest and independent. They choose
a voting option among $C+1$ possibilities, which includes abstention and maybe blank voting. We suppose that the choices follow a probability distribution $\left(P_{0}, \ldots, P_{C}\right)$, where $P_{0}$ is the probability to abstain. Let $\alpha$ be the voting option corresponding to the intention of the coerced voter, and $\beta$ be the one that is the instruction of the coercer. The coerced voter either disobeys, gives a fake credential and votes with option $\alpha$ (the evasion strategy does not imply any revote), or obeys and gives their real credential which the coercer uses to vote with option $\beta$. The coercer must decide whether the voter obeyed or not, given only the result.

The ideal result is $\mathrm{R}^{\text {ldeal }}=\overrightarrow{r e s}=\left(\operatorname{res}_{0}, \cdots, \operatorname{res}_{C}\right)$, the number of voters who opted for each option. In JCJ, however, the real result $\mathrm{R}^{\mathrm{Real}}$ is $\overrightarrow{r e s}$, as well as, for all $k$, the number of voters who revoted $k$ times. In addition, both results should also include the number of invalid ballots. However, to focus on the leakage of JCJ, we assume that a large and unpredictable number of ballots with a fake credential are cast, so that the adversary cannot gain information by observing the number of invalid ballots. This approximation is necessary to use the framework of [25], which does not model the possibility to cast a ballot with an invalid credential in the ideal setting.

We now instantiate [25] in our scenario. To simplify the analysis, we assume that a voter revotes at most once, so that $\mathrm{R}^{\text {Real }}=\left(\overrightarrow{r e s}, n_{R}\right)$, where $n_{R}$ is the total number of revotes. We also assume that all the parties are honest except for the coercer, and that the cryptography is perfect, so that the coercer does not learn any other information than the result. With these assumptions, we define the real and ideal games, where the behavior of the coerced voter is decided at random. The coercer wins the real (resp. ideal) game if they correctly guess the behavior of the coerced voter given the real (resp. ideal) result. For $g \in\{$ Real, Ideal $\}$ and for a pair $(\alpha, \beta)$, we denote $W_{\alpha, \beta}^{g}$ the event when the coercer wins the game, and $\delta_{\alpha, \beta}^{g}=2\left|\operatorname{Pr}\left(W_{\alpha, \beta}^{g}\right)-1 / 2\right|$. Furthermore, we consider the worst case for the voter, and analyze the quantity $\delta^{g}=\max _{\alpha, \beta} \delta_{\alpha, \beta}^{g}$. We call $\delta^{\text {Real }}$ (resp. $\delta^{\text {Ideal }}$ ) the coercion level of the real (resp. ideal) game. Intuitively, it measures how much better the adversary's strategy is compared to a random guess, in a scale from 0 to 1.

We denote by $\operatorname{Pr}\left(\mathrm{R}^{g} \mid \alpha\right)$ (resp. $\operatorname{Pr}\left(\mathrm{R}^{g} \mid \beta\right)$ ), the probability that the result $\mathrm{R}^{g}$ is obtained, assuming the voter votes for $\alpha$ (resp. obeys the coercer and votes for $\beta$ ). According to [25], the best strategy for the coercer is to assume that the voter obeyed if and only if $\operatorname{Pr}\left(\operatorname{R}^{g} \mid \beta\right) \geq \operatorname{Pr}\left(\mathrm{R}^{g} \mid \alpha\right)$. This gives a close formula for the coercion level which can be written as

$$
\begin{equation*}
\delta^{g}=\max _{(\alpha, \beta)} \sum_{\mathrm{R}^{g} \in M_{\alpha, \beta}} \operatorname{Pr}\left(\mathrm{R}^{g} \mid \beta\right)-\operatorname{Pr}\left(\mathrm{R}^{g} \mid \alpha\right) \tag{1}
\end{equation*}
$$

where $M_{\alpha, \beta}$ is the set of all possible results $\mathrm{R}^{g}$ such that $\operatorname{Pr}\left(\mathrm{R}^{g} \mid \beta\right) \geq \operatorname{Pr}\left(\mathrm{R}^{g} \mid \alpha\right)$. In Appendix D , we explain in more details how this formula can be computed. In the remaining of this section, we compare $\delta^{\text {Real }}$ and $\delta^{\text {Ideal }}$ in two scenarios, where external events provoke many revotes. In these scenar-


Fig. 2: Coercion levels as a function of the revote, with 20 voters, 2 candidates, $30 \%$ abstention and a $50 \%-50 \%$ distribution of votes between the candidates.
ios, we assume that there are two candidates $A$ and $B$, no blank vote, but the possibility to abstain or to revote once.

## B. The technical incident scenario

In general, we can expect revoting to be rare. This is something that is not allowed in classical paper-based elections, so that in a context where electronic voting is recent, voters will not be using this possibility. Even in a country such as Estonia, where revoting is available for internet voters since 2005, a recent study revealed a revoting rate of about $2 \%$ [16].

However, much more revotes could occur if an announcement reveals a suspicion of a technical incident, and encourages the voters to revote to be on the safe side. In this case, many voters could be inclined to revote with the same voting option, which would seem harmless if they are not aware of the weaknesses of JCJ. Note that the coercer could be the source of such an announcement, and spread fake news about the necessity to revote.

In Fig. 2, we consider this situation where a proportion of voters (that already voted) revote for the same voting option.

We plot the coercion level in both the real and ideal settings when the proportion $x$ of revotes ranges from 0 to 1 . When $x=0$, both coercion levels are the same since there is no revote. However, when $x=1$, there is no coercion-resistance in JCJ because the coerced voter would be the only one to cast a ballot without revoting. Note that the ideal coercion level remains constant since the overall probability to choose each voting option is unaffected by $x$.

## C. A discredit in the press

We still consider two candidates $A$ and $B$, and in this scenario, we assume that during the period of the voting phase, the candidate $A$ is discredited by an announcement in the press. As a consequence, some proportion of the voters who initially voted for $A$ will change their mind and revote for $B$.

For simplicity, we assume that no revote occurs that is not due to this event.

Such discredits have happened in the past. For instance, Dominique Strauss-Kahn, a former IMF managing director, was highly expected to become the next French president in 2012. However, due to an accusation of sexual assault, his political party chose to support another candidate. This occurred before the time of the election, and no electronic voting was involved. But, we can mention the 2022 Tory leadership election for the succession of Boris Johnson, the members could vote by Internet, and revote was initially authorized (before a security concern forced the organizers to forbid it). The duration was more than a month, which is more than enough for a discredit event to occur (even though it did not occur, for this election).

First, we fix a small number of voters, so that the effect is more visible, and we study the influence of the other parameters.

In Fig. 3, we plot the real and ideal coercion levels as the proportion $x$ of voters who change their mind from $A$ to $B$ ranges from 0 to 1 . When $x=0$, there is no difference since there is no revote. When $x=1$, there is no difference either since nobody votes for $A$ anymore, so that there is no coercionresistance in both the real and ideal games. However, a nonnegligible difference can be observed for the intermediate values of $x$.

In Fig. 4, we plot the real and ideal coercion levels with a fixed value of $x=0.3$ (i.e. $30 \%$ of the voters who voted for $A$ revote for $B$ ) and we let the initial proportion $p$ in favor of $A$ vary from 0 to 1 . When $p=0$, everyone votes for $B$ so that there is no coercion-resistance. When $p$ is large, we get close to the scenario presented in Section II-B, so that there is no coercion-resistance in the real game while the ideal game still offers some coercion-resistance.

In Fig. 5, we plot the real and ideal coercion levels with fixed $x=0.3$ and $p=0.7$ and we let the abstention rate $P_{0}$ range between 0 and 1 . When $P_{0}=0$, there is no coercion-resistance because forced-abstention attacks are trivial; similarly, there is no coercion-resistance when $P_{0}=1$. However, a non-negligible difference can be observed for the intermediate values.

Finally, in Fig. 6, we plot the real and ideal coercion levels with fixed $x=0.3, p=0.7$ and $P_{0}=0.3$, for a number of uncoerced honest voters equals to $16,32,64,128,256,512$ and 1024. This shows that the difference between both coercion levels remains non-negligible even when the number of voters is large. An asymptotic analysis (see e.g. [15]) reveals that the coercion level decreases in $1 / \sqrt{n_{V}}$, so the level of coercion is small anyway.

## IV. CHide: A Cleansing-Hiding protocol

We propose a modification of JCJ that provides full coercion-resistance. During the tally phase, the trustees perform the same tasks of cleansing, mixing and decrypting than in JCJ, but in a hidden way, so that the coercer (or anyone) does not learn how many ballots were deleted because


Fig. 3: Coercion levels as a function of the impact for 20 voters, 2 candidates, $30 \%$ abstention and a $70 \%-30 \%$ distribution between the candidates.
they correspond to revotes or to invalid credentials. For this purpose, we propose a novel cleansing algorithm based on MPC primitives.

## A. Cryptographic primitives

The security of the following primitives relies on the Decisional Diffie-Hellman (DDH) assumption and the Random Oracle Model.

ElGamal encryption. We use the exponential ElGamal encryption scheme on elliptic curves, which is convenient for its efficiency and its homomorphic property. Let $g$ be a public generator and pk the public key, i.e. a group element, the discrete logarithm of which is the corresponding secret key. One encrypts $m$ by choosing a random exponent $r$ and computing $\operatorname{Enc}(m, \mathrm{pk})=\left(g^{r}, g^{m} \mathrm{pk}^{r}\right)$. We impose $m$ to be in a small list of valid messages, so that decryption is feasible (in general, recovering $m$ from $g^{m}$ is hard). An important special case is when $m$ is a bit, because this is the kind of input used by the MPC primitives we mention below. For a generic message $m$, we call bit-wise encryption of $m$ the list of the encryptions of the bits of $m$.

Zero Knowledge Proofs. We use Non-Interactive Zero Knowledge Proofs (ZKP) based on the Fiat-Shamir transformation, mostly for proving relations involving discrete logarithms, à la Chaum-Pedersen.

Distributed key generation / threshold decryption. A DKG is a protocol which allows the participants to generate an ElGamal public key pk in such a way that each participant gets a share of the secret key. The protocol also generates a public commitment to each participant's share. A threshold $t$ is set, such that a collusion of $t$ or less participants can not deduce any information about the secret key, or about the cleartext of any ciphertext. If $t+1$ or more participants collaborate, they can combine their shares to recover the secret key. They can also run a threshold decryption protocol which allows to
jointly decrypt any ciphertext without recovering the secret key nor revealing anything about their shares. The result comes with a ZKP of correct decryption that anyone can verify. See for instance [35] for an instantiation.

Verifiable decryption mixnets. A decryption mixnet is an MPC protocol where the participants take a list of ciphertexts and reveal the corresponding plaintexts, in an order that is unrelated to the initial order, so that it is not possible to tell which one comes from which ciphertext. We will assume that the output plaintexts are sorted in the lexicographic order. The result comes with a corresponding ZKP, for correctness. See for instance [34] for an instantiation.
Logical operations on encrypted bits. We use MPC protocols that allow the owners of the shares to jointly perform logical operations on encrypted bits, based on their threshold decryption protocol. This is done without revealing the cleartexts to anyone, and in a verifiable manner. The main building block we use is the CGate protocol [31], that allows to conditionally set an encrypted value $X$ to (an encryption of) 0 , given an encrypted bit $Y$. In other words, if $x$ and $y$ are the corresponding plaintexts with $y \in\{0,1\}$, CGate $(X, Y)$ is a random encryption of $x y$. By combining this with the homomorphic property of the ElGamal encryption, one can derive a protocol for all the logical operations on bits (e.g. negation, disjunction).
More precisely, we use the And (conjunction) and the Eq (equality test) protocols. The latter is extended to bitwise encrypted data, by computing the conjunction of all the equality tests on encrypted bits. See Appendix A for more details about these MPC protocols.

Sorting encrypted data. Using the above logical operations, it is straightforward to design a comparison test Lt and a conditional swap CSwap. Therefore, it is possible to sort encrypted data, without revealing anything about the data (the conditional swap uses reencryptions, so that it is not possible to determine whether the data were swapped or not). For this purpose, we need a data-oblivious sorting algorithm, that is an algorithm whose control flow does not depend on the result of the comparisons. The popular fast sorting algorithms, such as Quicksort, Mergesort or Heapsort, do not verify this property. Consequently, we use the OddEvenMergeSort by Batcher [3], which has a quasi-linear complexity.

## B. Description of the CHide protocol

Participants. The CHide protocol is similar to JCJ, and the list of participants is the same. We recall them and emphasize the trust assumptions on them, which are the same as in JCJ. For simplicity, we assume that there is a single honest registrar. The literature following the JCJ approach contains techniques to have several registrars and appropriate trust assumptions on them [11]; this is orthogonal to the present discussion.

- The public board is an append-only list of data where all the participants can write. At any time, the content of the board can be read by anyone, and the view is the same. The board is assumed to be honest; see [18] for a possible realization of this.


Fig. 4: Coercion levels as a function of the proportion in favor of A for 20 voters, 2 candidates and $30 \%$ abstention.


Fig. 6: Coercion levels as a function of the number of voters with $30 \%$ abstention, $21 \%$ revotes and a distribution of $70 \%$ $30 \%$ between the candidates.

- The auditors check the consistency of the board, including the validity of all the ZKPs. We assume that there is at least one honest auditor that reveals any detected problem.
- The registrar sends their voting material to the legitimate voters. It is assumed that the registrar is honest.
- The $n_{T}$ election trustees hold the key shares and perform the cleansing and the tally. It is assumed that at most $t$ dishonest trustees can collaborate, where $t<n_{T}$ is the threshold used in the DKG.
- There are $n_{V}$ voters. Some of them may be dishonest and collude with the attacker. Honest voters may be subject to a coercion attempt by the attacker.
Setup phase. During this phase, a security parameter $\lambda$ is


Fig. 5: Coercion levels as a function of the abstention for 20 voters, 2 candidates, $21 \%$ revotes and a distribution of $70 \%$ $30 \%$ between the candidates.
chosen. The election trustees jointly run the DKG protocol, yielding a public key pk for this security level, and a secret share for each trustee. The DKG also produces public data that is sent to the public board for verifiabilty. We denote Setup the corresponding protocol.
Registration phase. For $1 \leq i \leq n_{V}$, the registrar generates $\lambda$ encryptions $\mathcal{R}_{i}=\left(R_{i, 1}, \ldots, R_{i, \lambda}\right)$ of random bits $\left(c_{i, 1}, \ldots, c_{i, \lambda}\right)$. These are called credentials. The registrar sends them to the board as a roster $\boldsymbol{R}=\left(\mathcal{R}_{i}\right)_{1 \leq i \leq n_{V}}$. The registrar shuffles the list of credentials and sends one of them to each voter. Just as in JCJ, we consider that the registration is perfect; but it can be augmented with DVZKPs [21] to allow the voter to verify the validity of their credential with respect to the public roster.

This simplified registration protocol assumes a single honest registrar. For several registrars, the (encrypted) bits of the credentials can be jointly generated, and the credentials can be sent to the voters without revealing them to the registrars, as long as one of them is honest [11]. The registrars can also jointly prove that the credentials are indeed encrypted bit by bit, and thus that the public roster is well-formed (see e.g. [7]).

Voting phase. In order to cast a vote for the option $\nu$ (encoded as a group element), a voter computes $C_{1}=$ $\operatorname{Enc}(\nu, \mathrm{pk})$ and $C_{2}=\left(\operatorname{Enc}\left(c_{1}, \mathrm{pk}\right), \ldots, \operatorname{Enc}\left(c_{\lambda}, \mathrm{pk}\right)\right)$, where $c=\left(c_{1}, \ldots, c_{\lambda}\right)$ is their credential. The neutral element 1 (the encoding $g^{0}$ of the zero bit) should not represent any voting option as it will be used to encode invalid ballots. The voter also produces a ZKP $\pi_{1}$ that proves the knowledge of $\nu$ and $c_{j}$ for all $j$. To ensure a strong Fiat-Shamir transformation [6], the computation of the challenge from the commitment of the $\Sigma$-protocol must include all public informations in the hash, such as $g$, pk, $C_{1}$ and $C_{2}$. Finally, a ZKP $\pi_{2}$ that $\nu$ is a valid voting option and that $c_{1}, \ldots, c_{\lambda}$ are bits must be added to prevent forced-abstention attacks which use write-
ins. The ballot $\left(C_{1}, C_{2}, \pi_{1}, \pi_{2}\right)=\operatorname{Vote}_{\mathrm{pk}}(c, \nu)$ is sent to the public board, using an anonymous channel. The voters check that their ballot is present on the board; this defines the verification step Check. The auditors verify that the ZKPs are valid and that there is no other ballot on the board with the same $\left(C_{1}, C_{2}\right)$; this defines the verification isVal.

Cleansing and tallying phase. Just as in JCJ, the election trustees keep only one ballot per valid credential and remove all the other ones. However, they do so in an oblivious way, by replacing the voting option of an invalid ballot by an encryption of 0 , which represents an invalid voting option. For this purpose, the trustees first compute an encrypted validity boolean for each ballot. Using the toolbox from [13], we could simply propose an MPC variant of JCJ, with a quadratic number of comparisons, as in JCJ. We instead propose a quasi linear approach, which relies on sorting.

The idea is to create a list of pairs of encrypted data $\left(V_{i}, K_{i}\right)_{1 \leq i \leq n_{b}+n_{V}}$, both for ballots posted to the board and credentials from the roster. For ballots, $V_{i}$ contains the voting option (the $C_{1}$ part), and $K_{i}=\left(K_{i}^{\perp}, K_{i}^{\top}\right)$ is formed of the encrypted order of appearance on the public board (i.e., $K_{i}^{\perp}$ is a bitwise-encrypted integer between 0 and $n_{B}-1$ ), and from the encrypted credential (i.e., $K_{i}^{\top}$ is the $C_{2}$ part). This step can be performed by anyone using a fixed randomness for $K_{i}^{\perp}$ (e.g. using the random 0 ) and is publicly verifiable. For entries coming from the roster, $V_{i}$ is the encryption of an invalid option, and $K_{i}=\left(K_{i}^{\perp}, K_{i}^{\top}\right)$ contains the encrypted integer $n_{B}$ in the first part, and the credential in the second part. Again, this step is performed using a fixed randomness for $V_{i}$ and $K_{i}^{\perp}$.

Then the talliers use the toolbox from [13] to sort this list in MPC, in quasi linear time. They use the following order:

- first sort according to $K_{i}^{\top}$ in increasing order. The effect is to group the ballots by credentials;
- for equal $K_{i}^{\top}$ (i.e. for equal credentials), sort in increasing order of $K_{i}^{\perp}$. The effect is to position the entry coming from the roster at the end of the group and the valid ballot (if there is one) just before it.
After this step, an entry is valid iff 1 ) its $K_{i}^{\top}$ part is the same as the one from its successor in the sorted list; and 2) the $K_{i}^{\perp}$ part of its successor encodes $n_{B}$. These tests can be efficiently implemented with the MPC toolbox and we need only a linear number of such tests, yielding an overall procedure in quasi linear time.

The $P_{\text {tally }}$ protocol is more precisely presented as follows. A more detailed description is available in Appendix A.

1. Discard all the ballots marked as invalid by the audit procedure. Let $\left(C_{1}^{i}, C_{2}^{i}\right)_{i=1}^{n_{B}}$ be the remaining ballots, without the ZKPs. We denote $m=\left\lceil\log n_{B}\right\rceil+1$.
2. For all $1 \leq k \leq n_{B}$, set $V_{k}=C_{1}^{k}$ and $K_{k}=U_{k-1} \| C_{2}^{k}$, where $U_{k-1}$ is a bit-wise encryption of $k-1$ over $m$ bits (least significant bit first), using the null randomness.
3. For all $n_{B}+1 \leq k \leq n_{B}+n_{V}$, set $K_{k}=U_{n_{B}} \| \mathcal{R}_{k-n_{B}}$ and $V_{k}$ is the encryption of 0 , using the null randomness.
4. Sort the $\left(V_{k}, K_{k}\right)$ using the keys $K_{k}$, in increasing order. This produces a result $\left(V_{k}^{\prime}, K_{k}^{\prime}\right)_{k=1}^{n_{B}+n_{V}}$ and a transcript $\Pi^{\text {Sort }}$.
5. For all $1 \leq k<n_{B_{\top}}+n_{V}$, compute $D_{k}=$ $\mathrm{Eq}\left(K_{k}^{\prime \top}, K_{k+1}^{\prime \top}\right)$, where $K_{k}^{\prime \top}$ refers to the $\lambda$ most significant bits of $K_{k}^{\prime}$. This produces the transcript $\Pi_{k, 1}^{\mathrm{Eq}}$.
6. For all $1 \leq k<n_{B}+n_{V}$, compute $F_{k}=$ $\mathrm{Eq}\left(K^{\prime} \stackrel{\perp}{k+1}, U_{n_{B}}\right)$, where $K_{k+1}^{\prime \perp}$ refers to the $m$ least significant bits of $K_{k+1}^{\prime}$. This produces the transcript $\Pi_{k, 2}^{\mathrm{Eq}}$.
7. For all $1 \leq k<n_{B}+n_{V}$, replace $V_{k}^{\prime}$ by CGate $\left(V_{k}^{\prime}, \operatorname{And}\left(D_{k}, F_{k}\right)\right)$.
8. Apply the decryption mixnet protocol on the $\left(V_{k}^{\prime}\right)_{k=1}^{n_{B}+n_{V}-1}$. This produces the result of the election as well as a verification transcript $\Pi^{\text {Mixnet }}$.
Each step produces a transcript, published on the board, and verified by the auditors.

Evading coercion. We provide the same evasion strategy as in JCJ, which consists of giving a fake credential to the coercer and to vote (once, if the voter wants to) with the real credential.

Assumptions on the communication channels. As in the original JCJ protocol, the communications between the voters and the registrars must be untapable, and the communication between the voters and the public board must be anonymous. This is required only for the coercion-resistance property. Verifiability and vote secrecy rely on standard communication channels assumptions.

For verifiability, we assume that honest voters check that their ballots appear on the public board and complain as soon as something unexpected happens. For weaker voters that check their ballots only when they vote, the ballot order should be enforced, for instance using a chain of hashes (see e.g. [2]).

## C. Efficiency considerations

In terms of computational and communication costs, CHide is less efficient than JCJ, mainly because the encrypted credentials are now formed by $\lambda$ ciphertexts instead of a single one.

For the talliers, the cleansing phase is more complex but scales better with the number of submitted ballots. While JCJ is quadratic, we propose a quasi-linear tally protocol based on sorting. Another difference is that, due to the MPC toolbox, the number of communication rounds between them is no longer constant, but becomes logarithmic in the number of ballots. Nevertheless, the task is highly parallelizable and remains affordable for medium-size elections.

For the voters, the computational load increases but the total cost for realistic parameters is around a thousand exponentiations, which should be a matter of seconds with a standard implementation in Javascript running within a modern browser.

In Table I, we give estimates of the number of exponentiations and of the transcript size for both JCJ and CHide. For this purpose, we consider a number of $n_{T}=3$ talliers with a threshold $t=2$, a security parameter of $\lambda=128$

TABLE I: Number of exponentiations and transcript size in JCJ and CHide, with $\lambda=128$.

| \# voters | \# exp. |  | estimated CPU time |  | transcript |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| any | JCJ | CHide | JCJ | CHide | JCJ | CHide |
| (Vote) | 27 | 1.4 k | 5.4 ms | 0.28 s | 1.1 kB | 58 kB |
| 10 | JCJ | CHide | JCJ | CHide | JCJ | CHide |
| (Tally) | 4.26 k | 4.1 M | 0.43 s | 6.8 min | 170 kB | 65 MB |
| 100 | JCJ | CHide | JCJ | CHide | JCJ | CHide |
| (Tally) | 380 k | 120 M | 38 s | 3.3 h | 16.0 MB | $1,9 \mathrm{~GB}$ |
| 1000 | JCJ | CHide | JCJ | CHide | JCJ | CHide |
| (Tally) | 37.5 M | 2.4 G | 1.0 h | 2.8 d | 1.59 GB | 39 GB |
| 10000 | JCJ | CHide | JCJ | CHide | JCJ | CHide |
| (Tally) | 3.75 G | 41 G | $4,3 \mathrm{~d}$ | 48 d | 158 GB | 668 GB |
| 100000 | JCJ | CHide | JCJ | CHide | JCJ | CHide |
| (Tally) | 375 G | 658 G | 1.2 y | 2.1 y | 15.8 TB | 10.6 TB |
| 1000000 | JCJ | CHide | JCJ | CHide | JCJ | CHide |
| (Tally) | 37.5 T | 9.4 T | 120 y | 30 y | 1.58 PB | 152 TB |

and a number of $n_{C}=2$ voting options. We also give the corresponding running times, based on an estimate of 5000 exponentiations per second on the client side, and 10000 per second on the server side. For 1000000 voters, CHide starts being more efficient than JCJ. This is due to the fact that JCJ has a quadratic complexity while CHide is quasi linear. While the cost is high, this is not out of reach since the computations are highly parallelizable and hence could be conducted within one day on a (large) cluster.

## V. DEfining coercion-RESISTANCE

One of the reasons why the flaw was not discovered in the original proposition of the JCJ protocol is the fact that the definition of coercion-resistance itself was flawed, in the sense that no scheme can be proved secure w.r.t. this definition. We then introduce our own definition, that is used to analyse the security of CHide.

## A. Voting system

A voting system is a tuple of eight algorithms or protocols $P_{\text {Setup }}$, register, Vote, Check, isVal, Fakecred, $P_{\text {tally }}$, Verify such that:

- $P_{\text {Setup }}\left(\lambda, n_{T}, t\right)$ is a protocol run by $n_{T}$ authorities for the security parameter $\lambda$ and the threshold $t$. It computes the public key pk, as well as the secret and public shares $\left(s_{i}, h_{i}\right)$ for each authority, using a DKG.
- register $\left(\lambda, \mathrm{pk}, n_{V}\right)$ generates a private credential $c_{i}$ for each voter $\in\left[1, n_{V}\right]$. It also returns some public information $\boldsymbol{R}$ that contains the public part of the credentials, and any necessary transcript for proving their validity.
- The algorithm $\operatorname{Vote}_{\mathrm{pk}}(c, \nu)$ takes as input the public key of the election pk , a credential $c$, and a vote $\nu$, and returns a ballot. The public key is often omitted for simplicity.
- The algorithm Check $(B B, \mathrm{pk}, \mathbf{b}, c, \nu)$ takes the bulletin board $B B$, the public key, a ballot $\mathbf{b}$, a private credential $c$ and a voting option $\nu$. It is run by the voters to check that their ballot has been added to the public board.
- The algorithm isVal(pk, $\mathbf{b}, B B)$ takes as input the public key, a ballot and the ballot box. It outputs a bit which states whether the ballot is valid w.r.t. $B B$.
- Fakecred $(c)$ takes a credential and returns a fake one $\tilde{c}$.
- $P_{\text {tally }}\left(B B, \boldsymbol{R}, \mathrm{pk},\left\{h_{i}, s_{i}\right\}, t\right)$ is a protocol run by the authorities that possess the secret shares $\left\{s_{i}\right\}$ with public commitments $\left\{h_{i}\right\}$. It takes as input a list of ballots $B B$, the public roster $\boldsymbol{R}$, the public key pk, the threshold $t$ and returns the result $\boldsymbol{X}$ of the election together with a proof $\Pi$.
- Verify $(B B$, $\Pi$, res) takes as input a ballot box $B B$, a transcript $\Pi$ and a result res. It outputs a bit which states whether the result is valid with respect to $B B$.


## B. The original definition of JCJ

The intuition of the JCJ definition of coercion-resistance is that an adversary must not guess whether a coerced voter obeyed or evaded coercion. When the voter obeys $(b=1$ in the definition), they give their real credential and abstain from doing any other action. Note that a coercer may ask the voter to cast some specific vote or to perform some specific computations, but this is not considered in the definition as the adversary might as well do it themself, with the given credential. When the voter evades ( $b=0$ in the definition), they give a fake credential and cast a single vote for the desired voting option (or abstain, depending on their personal choice).

This yields the game Real ${ }_{J C J}^{C R}$ presented in Algorithm 1. Voting choices are represented as integers between 1 and $n_{C}$, and $\phi$ represents the choice to abstain. During this game, the adversary selects the set of corrupted voters. It is given the corresponding private credentials as well as all the public information $\boldsymbol{R}$ (i.e. the encrypted credentials in the JCJ protocol). It then chooses $(j, \alpha)$, where $j$ denotes the voter under coercion and $\alpha$ their desired voting option. The evasion strategy is modeled in lines 11 and 13: when the voter disobeys, they create a fake credential and cast a vote for $\alpha$ (or abstain if $\alpha=\phi$ ). Otherwise, they give their real credential.

Honest voters vote according to a distribution which depends on the number of options $n_{C}$ and returns a value that may be:

- any valid vote $\nu \in\left[1, n_{C}\right]$;
- $\phi$, which represents abstention;
- $\lambda$, which represents casting a vote with a fake credential. We extend the Vote function to votes equal to $\lambda$ as follows.

$$
\operatorname{Vote}_{\mathrm{pk}}(c, \lambda)=\operatorname{Vote}_{\mathrm{pk}}(\tilde{c}, \nu)
$$

where $\tilde{c}=\operatorname{Fakecred}(c)$ and $\nu$ is sampled from $\left[1, n_{C}\right]$.
It is worth noting that the advantage of an adversary in game Real ${ }_{J C J}^{C R}$ will always be non negligible since one can compare the result of the tally with the expected result, given the distribution $\mathcal{D}$ of the voting intentions. For example, if the adversary wants to cast a vote for a very unlikely candidate, they may observe cases where the latter does not get a single vote in the result, which is a clear indication that the coerced voter disobeyed. Hence, the JCJ definition compares the advantage of an adversary in game $\mathrm{Real}_{\mathrm{JCJ}}^{\mathrm{CR}}$ with the one in an ideal game Ideal ${ }_{J C J}^{C R}$, where there is no other information

```
```

Algorithm 1: Real ${ }_{\text {JCJ }}^{C R}$

```
```

Algorithm 1: Real ${ }_{\text {JCJ }}^{C R}$
Require: $\mathbb{A}, \lambda, n_{T}, t, n_{V}, n_{A}, n_{C}, \mathcal{D}$
Require: $\mathbb{A}, \lambda, n_{T}, t, n_{V}, n_{A}, n_{C}, \mathcal{D}$
$B B \longleftarrow \varnothing$
$B B \longleftarrow \varnothing$
$\mathrm{pk},\left(s_{i}, h_{i}\right)_{i=1}^{n_{T}} \longleftarrow \operatorname{Setup}^{\mathbb{A}}\left(\lambda, n_{T}, t\right)$
$\mathrm{pk},\left(s_{i}, h_{i}\right)_{i=1}^{n_{T}} \longleftarrow \operatorname{Setup}^{\mathbb{A}}\left(\lambda, n_{T}, t\right)$
$A \longleftarrow \mathbb{A}()$
$A \longleftarrow \mathbb{A}()$
$\left\{c_{i} ; i \in\left[1, n_{V}\right]\right\}, \boldsymbol{R} \longleftarrow \operatorname{register}\left(\lambda, \mathrm{pk}, n_{V}\right)$
$\left\{c_{i} ; i \in\left[1, n_{V}\right]\right\}, \boldsymbol{R} \longleftarrow \operatorname{register}\left(\lambda, \mathrm{pk}, n_{V}\right)$
$\boldsymbol{s}(j, \alpha) \longleftarrow \mathbb{A}\left(\left\{c_{i} ; i \in V\right\}, \boldsymbol{R}\right)$
$\boldsymbol{s}(j, \alpha) \longleftarrow \mathbb{A}\left(\left\{c_{i} ; i \in V\right\}, \boldsymbol{R}\right)$
6 if $|A| \neq n_{A} \vee j \notin\left[1, n_{V}\right] \backslash A \vee \alpha \notin\left[1, n_{C}\right] \bigcup\{\phi\}$ then
6 if $|A| \neq n_{A} \vee j \notin\left[1, n_{V}\right] \backslash A \vee \alpha \notin\left[1, n_{C}\right] \bigcup\{\phi\}$ then
Return 0
Return 0
$b \stackrel{\$}{\leftrightarrows}\{0,1\}$
$b \stackrel{\$}{\leftrightarrows}\{0,1\}$
$\tilde{c} \longleftarrow c_{j}$
$\tilde{c} \longleftarrow c_{j}$
if $b==0$ then
if $b==0$ then
$\tilde{c} \longleftarrow$ Fakecred $\left(c_{j}\right)$
$\tilde{c} \longleftarrow$ Fakecred $\left(c_{j}\right)$
if $\alpha \neq \phi$ then
if $\alpha \neq \phi$ then
$B B \longleftarrow B B \bigcup\left\{\operatorname{Vote}_{\mathrm{pk}}\left(c_{j}, \alpha\right)\right\}$
$B B \longleftarrow B B \bigcup\left\{\operatorname{Vote}_{\mathrm{pk}}\left(c_{j}, \alpha\right)\right\}$
for $i \in\left[1, n_{V}\right] \backslash(A \bigcup\{j\})$ do
for $i \in\left[1, n_{V}\right] \backslash(A \bigcup\{j\})$ do
$\nu_{i} \longleftarrow \mathcal{D}_{n_{C}}()$
$\nu_{i} \longleftarrow \mathcal{D}_{n_{C}}()$
if $\nu_{i} \neq \phi$ then
if $\nu_{i} \neq \phi$ then
$B B \longleftarrow B B \bigcup\left\{\operatorname{Vote}_{\mathrm{pk}}\left(c_{i}, \nu_{i}\right)\right\}$
$B B \longleftarrow B B \bigcup\left\{\operatorname{Vote}_{\mathrm{pk}}\left(c_{i}, \nu_{i}\right)\right\}$
$B B \longleftarrow B B \bigcup \mathbb{A}(\tilde{c}, B B)$
$B B \longleftarrow B B \bigcup \mathbb{A}(\tilde{c}, B B)$
$21 \boldsymbol{X}, \Pi \longleftarrow P_{\text {tally }}^{\mathbb{A}}\left(B B, \boldsymbol{R}, \mathrm{pk},\left\{h_{i}, s_{i}\right\}, t\right)$
$21 \boldsymbol{X}, \Pi \longleftarrow P_{\text {tally }}^{\mathbb{A}}\left(B B, \boldsymbol{R}, \mathrm{pk},\left\{h_{i}, s_{i}\right\}, t\right)$
$b^{\prime} \longleftarrow \mathbb{A}()$
$b^{\prime} \longleftarrow \mathbb{A}()$
Return 1 if $b^{\prime}==b$ else 0

```
    Return 1 if \(b^{\prime}==b\) else 0
```

```
19
20
```

20

```
```

```
Algorithm 2: Ideal \(_{\mathrm{JCJ}}^{\mathrm{CR}}\)
```

```
Algorithm 2: Ideal \(_{\mathrm{JCJ}}^{\mathrm{CR}}\)
    Require: \(\mathbb{A}, \lambda, n_{V}, n_{A}, n_{C}, \mathcal{D}\)
    Require: \(\mathbb{A}, \lambda, n_{V}, n_{A}, n_{C}, \mathcal{D}\)
    \(D \longleftarrow \varnothing\)
    \(D \longleftarrow \varnothing\)
    2
    2
    \(A \longleftarrow \mathbb{A}(\lambda)\)
    \(A \longleftarrow \mathbb{A}(\lambda)\)
4
4
    \((j, \alpha) \longleftarrow \mathbb{A}()\)
    \((j, \alpha) \longleftarrow \mathbb{A}()\)
    if \(|A| \neq n_{A} \vee j \notin\left[1, n_{V}\right] \backslash A \vee \alpha \notin\left[1, n_{C}\right] \bigcup\{\phi\}\) then
    if \(|A| \neq n_{A} \vee j \notin\left[1, n_{V}\right] \backslash A \vee \alpha \notin\left[1, n_{C}\right] \bigcup\{\phi\}\) then
        Return 0
        Return 0
    \(b \stackrel{\$}{\leftarrow}\{0,1\}\)
    \(b \stackrel{\$}{\leftarrow}\{0,1\}\)
    if \(b==0 \wedge \alpha \neq \phi\) then
    if \(b==0 \wedge \alpha \neq \phi\) then
        \(D \longleftarrow D \bigcup\{\alpha\}\)
        \(D \longleftarrow D \bigcup\{\alpha\}\)
12
13
    for \(i \in\left[1, n_{V}\right] \backslash(A \bigcup\{j\})\) do
    for \(i \in\left[1, n_{V}\right] \backslash(A \bigcup\{j\})\) do
        \(\nu_{i} \longleftarrow \mathcal{D}_{n_{C}}()\)
        \(\nu_{i} \longleftarrow \mathcal{D}_{n_{C}}()\)
        if \(\nu_{i} \notin\{\phi, \lambda\}\) then
        if \(\nu_{i} \notin\{\phi, \lambda\}\) then
            \(D \longleftarrow D \bigcup\left\{\nu_{i}\right\}\)
            \(D \longleftarrow D \bigcup\left\{\nu_{i}\right\}\)
    \(\left(\nu_{i}\right)_{i \in A}, \beta \longleftarrow \mathbb{A}()\)
    \(\left(\nu_{i}\right)_{i \in A}, \beta \longleftarrow \mathbb{A}()\)
    if \(b==1 \wedge \beta \in\left[1, n_{C}\right]\) then
    if \(b==1 \wedge \beta \in\left[1, n_{C}\right]\) then
        \(D \longleftarrow D \bigcup\{\beta\}\)
        \(D \longleftarrow D \bigcup\{\beta\}\)
    \(\boldsymbol{X} \longleftarrow \operatorname{result}\left(D \bigcup\left\{\nu_{i} \mid i \in A, \nu_{i} \in\left[1, n_{C}\right]\right\}\right)\)
    \(\boldsymbol{X} \longleftarrow \operatorname{result}\left(D \bigcup\left\{\nu_{i} \mid i \in A, \nu_{i} \in\left[1, n_{C}\right]\right\}\right)\)
    \(b^{\prime} \longleftarrow \mathbb{A}(\boldsymbol{X})\)
    \(b^{\prime} \longleftarrow \mathbb{A}(\boldsymbol{X})\)
    Return 1 if \(b^{\prime}==b\) else 0
```

```
    Return 1 if \(b^{\prime}==b\) else 0
```

```

Fig. 7: JCJ definition of coercion resistance. \(\lambda\) is the security parameter, \(n_{T}\) the number of talliers, \(t\) the threshold, \(n_{V}\) the number of voters, \(n_{A}\) the number of corrupted voters, \(n_{C}\) the number of voting options and \(\mathcal{D}\) the distribution of votes.
than what is unavoidably leaked, that is, the result. The game Ideal \({ }_{J C J}^{C R}\) is presented in Algorithm 2. Compared to the original definition, we present a slightly modified version that reasons on the clear votes only. This simplifies the understanding by focusing on the information given to the adversary. All our claims and remarks hold on the original definition as well.

Definition 1 (adapted from [22]). A voting system is JCJcoercion resistant if for all PPT adversary \(\mathbb{A}\), for all parameters \(n_{T}, t, n_{V}, n_{A}, n_{C}\), and for all distributions \(\mathcal{D}\), there exists a PPT adversary \(\mathbb{B}\) and a negligible function \(\mu\) such that
\[
\begin{aligned}
& \mid \operatorname{Pr}\left(\text { Ideal }_{J C J}^{C R}\left(\mathbb{B}, \lambda, n_{V}, n_{A}, n_{C}, \mathcal{D}\right)=1\right) \\
& -\operatorname{Pr}\left(\operatorname{Real}_{\mathrm{JCJ}}^{C R}\left(\mathbb{A}, \lambda, n_{T}, t, n_{V}, n_{A}, n_{C}, \mathcal{D}\right)=1\right) \mid \leq \mu(\lambda)
\end{aligned}
\]

As noticed in [17], this definition cannot be realized by a scheme which uses a public board. Indeed, in the real game, the adversary observes the length \(n_{B}\) of the board which corresponds to the total number of ballots cast by noncorrupted voters. Then the adversary learns the result and in particular its size \(n_{v}\), that is the number of valid ballots counted. Hence the total number \(\Delta=n_{B}-n_{v}\) of ballots discarded can be deduced, which is not available in the ideal game Ideal \(\mathrm{JCJ}_{\mathrm{CJ}}^{\mathrm{CR}}\). The value of \(\Delta\) can be compared with its
expected number, according to the distribution \(\mathcal{D}\). Since there is an additional ballot discarded (the one of the coercer) when the voter evades coercion, the adversary has a non-negligible advantage in the real game. For instance, if \(\mathcal{D}\) is such that no voters cast a ballot with an invalid credential, either \(n_{B}=n_{v}\), which means that the adversary's ballot has been counted, or \(n_{B}=n_{v}+1\), meaning that the adversary's ballot has been discarded and that the voter has disobeyed. Of course, the same issue applies to the JCJ definition as stated in [22].

The authors of [17] proposed a patch to the issue they discovered: the length of the board should be given to the adversary in the ideal game as well. Intuitively, this corresponds to simply rewriting line 18 of Algorithm 2 as \(\left(\nu_{i}\right)_{i \in V}, \alpha \longleftarrow \mathbb{A}(|D|)\). However, this still does not allow to detect the leakage of the JCJ protocol during the tally. Indeed, the distribution \(\mathcal{D}\) fails to model several aspects:
- First, the addition of a ballot with an invalid credential only happens when a honest voter sacrifices their own vote. This is unlikely in practice, and does not model ballots sent by non-eligible voters (for instance, by the authorities).
- Second, revoting is not considered in \(\mathcal{D}\), which explains why the leakage of the JCJ protocol was not detected.

A final remark about the definition of JCJ concerns its underlying trust assumptions. For clarity, we recall them here. First of all, it is assumed that all the registrars are honest and that the adversary is inactive during the registration phase (or, alternatively, the registration is untapable). Second, the adversary can only corrupt a minority of decryption authorities. Also, ballots are cast through anonymous channels. Finally, the bulletin board is honest.

\section*{C. The quantitative definition of \(K T V\)}

Apart from the definition of \(\mathbf{J C J}\), there are other definitions in the literature (see [17] for a survey). The most prominent one is the quantitative definition of [25] (KTV), where the notion of \(\delta\)-coercion-resistance comes with two conditions: first, the coerced voter must have a strategy to meet their objective with overwhelming probability; second, the adversary cannot decide, with an advantage greater than \(\delta\), whether the voter used this strategy or forwarded all received messages (including their credential).

The KTV definition is abstract. To use it, it is necessary to model the voting protocol, its participants and the evasion strategy. In addition, it does not say much about how ballots sent with an invalid credential should be handled since the honest participants are assumed to vote following a fixed distribution of valid voting options. Finally, it does not tell if a specific \(\delta\) is acceptable or not. Our definition can be seen as an instance of KTV, where \(\delta\) is shown to be minimal, that is, not greater than that of an ideal protocol.

\section*{D. Our definition of coercion-resistance}

If we compare the advantage of the adversary in the real game with its advantage in the ideal one, we need to cover a large family of vote distributions. Otherwise, we may miss security flaws. In particular, we need to cover cases explicitly planned by the protocol such as revote and addition of ballots with fake credentials.

Therefore, given a set \(S\) of unique identifiers and the number \(n_{C}\) of voting options (excluding abstention), we consider a distribution \(\mathcal{B}\left(S, n_{C}\right)\) of sequences of pairs of the form \((j, \nu)\) where \(\nu \in\left[1, n_{c}\right]\) represents a vote and \(j\) represents either a valid voter (when \(j \in S\) ) or a fake voter, with a fake credential. Typically, if \(A\) is the set of corrupted voters, \(S=\left[1, n_{V}\right] \backslash A\). To avoid collisions with identifiers which may be in \(A\), we consider that any \(i \notin S\) holds a negative value. The distribution \(\mathcal{B}\) captures the abstention of a voter \(i\) with the absence of a couple of the form \((i, *)\). It models both revoting and the addition of fake ballots, typically by authorities:
- revoting is reflected in \(\mathcal{B}\) by the fact that a voter may appear several times in the same sequence;
- fake ballots are modeled by pairs \((j, \nu)\) where \(j \notin S\). They may be added by authorities or voters. Note that \(\mathcal{B}\) also models the case of a revote with a fake credential.
For example, in the sequence \((1,1),(2,1),(1,2),(-1,2)\), \((1,1)\) with \(n_{V}=3\), we have three voters \(V_{1}, V_{2}\) and \(V_{3} . V_{1}\) first votes \(1, V_{2}\) votes 1 , then \(V_{1}\) revotes for 2 , then a fake
vote for 2 is added, then \(V_{1}\) changes back her vote to \(1 . V_{3}\) simply chooses to abstain.

Our Real \({ }^{C R}\) game, given in Algorithm 3, is similar to Real \({ }_{\mathrm{JCJ}}^{\mathrm{CR}}\). Votes are drawn according to \(\mathcal{B}\left(\left[1, n_{V}\right] \backslash A, n_{C}\right)\), yielding a sequence \(B\). It typically contains pairs \((i, \nu)\) with \(i<0\), which corresponds to the addition of ballots with fake credentials. For such a pair, we therefore generate a fake credential at lines 11-12. Just as in the definition of JCJ, the adversary must guess a bit \(b\). If \(b=1\), the coerced voter \(j\) obeys, hence any vote from \(j\) is removed from \(B\) and the real credential is provided to the adversary. If \(b=0\), the voter follows the evasion strategy, namely they cast one vote for \(\beta\) (if \(\beta \neq \phi\) ) and provides a fake credential. Hence the votes from \(j\) in \(\mathcal{B}\) are replaced by a single vote for \(\beta\) (if \(\beta \neq \phi\) ). Then ballots are added according to \(\mathcal{B}\). They correspond either to real or fake votes (or revotes). Compared to the original JCJ definition, we also slightly improve the power of the adversary by letting them observe the board after each vote and add ballots if they want to, which better reflects the reality.

Again, the advantage of the adversary in the real game is compared with its advantage in an ideal game Ideal \({ }^{\mathrm{CR}}\), where the adversary can only observe the number of ballots and the result (see Algorithm 4). We assume a function cleanse that removes votes from invalid voters \(j \notin\left[1, n_{V}\right]\) and that takes care of revotes according to the policy (typically, the last vote is kept). The function result, given a set of valid votes, returns the result of the election.

Definition 2. A voting system is coercion resistant if for all PPT adversary \(\mathbb{A}\), for all parameters \(n_{T}, t, n_{V}, n_{A}, n_{C}\), and for all distribution \(\mathcal{B}\), there exists a PPT adversary \(\mathbb{B}\) and a negligible function \(\mu\) such that
\[
\begin{aligned}
& \mid \operatorname{Pr}\left(\text { Ideal }\left.\right|^{\mathrm{CR}}\left(\mathbb{B}, \lambda, n_{V}, n_{A}, n_{C}, \mathcal{B}\right)=1\right) \\
& \quad-\operatorname{Pr}\left(\operatorname{Real}^{C \mathrm{R}}\left(\mathbb{A}, \lambda, n_{T}, t, n_{V}, n_{A}, n_{C}, \mathcal{B}\right)=1\right) \mid \leq \mu(\lambda)
\end{aligned}
\]

The main difference between our definition and the original one is that we consider a larger family of distributions, which allows to analyze a protocol in the context of revotes and fake ballots.

Another difference is that the adversary shall not gain any advantage for any distribution \(\mathcal{B}\), while the JCJ definition defines coercion-resistance with respect to a particular distribution. This is preferable since a protocol should be as secure as the ideal one, whatever the considered distribution. It is counter-intuitive to design a cryptographic protocol that resists only for particular distributions. Of course, it makes sense to analyze the exact advantage in the ideal game for a particular distribution, and devise whether voters are reasonably protected in that case of not. But the cryptographic protocol itself should be as solid as the ideal one nevertheless.

\section*{E. CHide is coercion-resistant}

Thanks to our cleansing procedure, we show that CHide achieves coercion-resistance. This is stated in Theorem 1, proven in Appendix C-B. In this theorem, we suppose secure DKG and Mixnet protocols in the SUC security framework [9],
```

Algorithm 3: Real $^{C R}$
Require: $\mathbb{A}, \lambda, n_{T}, t, n_{V}, n_{A}, n_{C}, \mathcal{B}$
$B B \longleftarrow \varnothing$
$\mathrm{pk},\left(s_{i}, h_{i}\right)_{i=1}^{n_{T}} \longleftarrow P_{\text {Setup }}^{\mathbb{A}}\left(\lambda, n_{T}, t\right)$ (* run the DKG *)
$\left\{c_{i} ; i \in\left[1, n_{V}\right]\right\}, \boldsymbol{R} \longleftarrow$ register $\left(\lambda, \mathrm{pk}, n_{V}\right)$
$A \longleftarrow \mathbb{A}(\boldsymbol{R})(*$ corrupt voters $*)$
$(j, \alpha) \longleftarrow \mathbb{A}\left(\left\{c_{i} ; i \in A\right\}\right)$
(* coerce a voter $j$ who has the intention $\alpha$ *)
if $|A| \neq n_{A} \vee j \notin\left[1, n_{V}\right] \backslash A \vee \alpha \notin\left[1, n_{C}\right] \bigcup\{\phi\}$ then
Return 0
$B \longleftarrow \mathcal{B}\left(\left[1, n_{V}\right] \backslash A, n_{C}\right)$
(* samples a sequence of pairs $\left(i, \nu_{i}\right)$ with
$i \in\left(\left[1, n_{V}\right] \backslash A\right) \bigcup\{n \in \mathbb{Z} \mid n<0\}$ and $\left.\nu_{i} \in\left[1, n_{C}\right] *\right)$
for $(i, *) \in B$, $i \notin\left[1, n_{V}\right]$ do
$c_{i} \longleftarrow$ Fakecred()
(* this captures ballots sent with invalid creds *)
$b \stackrel{\$}{\leftrightarrows}\{0,1\}$
$\tilde{c} \longleftarrow c_{j}$
if $b==1$ then
Remove all $(j, *) \in B$
else
Remove all $(j, *) \in B$ but the last, which is replaced
by $(j, \alpha)$ if $\alpha \neq \phi$ and removed otherwise
$\tilde{c} \longleftarrow$ Fakecred $\left(c_{j}\right)$
$\mathbb{A}(\tilde{c})(* \mathbb{A}$ learns $\tilde{c} *)$
for $\left(i, \nu_{i}\right) \in B$ (in this order) do
$M \longleftarrow \mathbb{A}(B B)(*$ cast ballots $*)$
$B B \longleftarrow B B \bigcup\{m \in \boldsymbol{M} \mid$ isVal $(\mathrm{pk}, m, B B)=1\}$
$B B \leftarrow B B \bigcup\left\{\operatorname{Vote}_{\mathrm{pk}}\left(c_{i}, \nu_{i}\right)\right\}$
$\boldsymbol{M} \longleftarrow \mathbb{A}(B B, " l a s t$ honest ballot sent")
$B B \longleftarrow B B \bigcup\{m \in \boldsymbol{M} \mid$ isVal $(\mathrm{pk}, m, B B)=1\}$
$\boldsymbol{X}, \Pi \leftarrow P_{\text {tally }}^{\mathbb{A}}\left(B B, \boldsymbol{R}, \mathrm{pk},\left\{h_{i}, s_{i}\right\}, t\right)$
$b^{\prime} \longleftarrow \mathbb{A}()$
Return 1 if $b^{\prime}==b$ else 0

```
```

Algorithm 4: Ideal $^{C R}$
Require: $\mathbb{A}, \lambda, n_{V}, n_{A}, n_{C}, \mathcal{B}$
1
2
$A \longleftarrow \mathbb{A}(\lambda)$ (* corrupt voters *)
( $j, \alpha) \longleftarrow \mathbb{A}()$
6 (* coerce a voter $j$ who has the intention $\alpha^{*}$ )
if $|A| \neq n_{A} \vee j \notin\left[1, n_{V}\right] \backslash A \vee \alpha \notin\left[1, n_{C}\right] \bigcup\{\phi\}$ then
Return 0
$B \longleftarrow \mathcal{B}\left(\left[1, n_{V}\right] \backslash A, n_{C}\right)$
(* samples a sequence of pairs $\left(i, \nu_{i}\right)$ with
$i \in\left(\left[1, n_{V}\right] \backslash A\right) \bigcup\{n \in \mathbb{Z} \mid n<0\}$ and $\left.\nu_{i} \in\left[1, n_{C}\right] *\right)$
11
12
13
$b \stackrel{\&}{\leftarrow}\{0,1\}$
15
if $b==1$ then
Remove all $(j, *) \in B$
else
Remove all $(j, *) \in B$ but the last, which is
replaced by $(j, \alpha)$ if $\alpha \neq \phi$ and removed
otherwise
$\left(\nu_{i}\right)_{i \in A}, \beta \longleftarrow \mathbb{A}(|B|)$
if $(b==1) \wedge\left(\beta \in\left[1, n_{C}\right]\right)$ then
$B \longleftarrow B \bigcup\{(j, \beta)\}$
$B \longleftarrow B \bigcup\left\{\left(i, \nu_{i}\right) \mid i \in A, \nu_{i} \in\left[1, n_{C}\right]\right\}$
26
27
$28 \boldsymbol{X} \longleftarrow \operatorname{result}(c l e a n s e(B))$
$b^{\prime} \longleftarrow \mathbb{A}(\boldsymbol{X})$
Return 1 if $b^{\prime}==b$ else 0

```

Fig. 8: Definition of coercion-resistance. \(\lambda\) is the security parameter, \(n_{T}\) the number of talliers, \(t\) the threshold, \(n_{V}\) the number of voters, \(n_{A}\) the number of corrupted voters, \(n_{C}\) the number of voting options and \(\mathcal{B}\) the distribution.
which is introduced in Appendix B. Simply speaking, it is a Universally Composable security framework which is suitable for MPC protocols. Examples of UC-secure DKG and decryption mixnet can be found in [34], [35].

We mention that we also prove privacy and verifiability in Appendices C-C and C-D, using a similar approach.

Theorem 1. Under the DDH assumption, assuming SUCsecure DKG and Mixnet protocols and in the Programmable Random Oracle Model, CHide is coercion-resistant.

We can check that JCJ is not coercion-resistant according to our definition. This is due to the leakage during the cleansing phase, as explained in Section II-B.
Proposition 1. The JCJ protocol is not coercion-resistant.

\section*{VI. Discussion}

While JCJ does not satisfy our definition of coercionresistance, it does provide a certain level of security. As future work, we could identify the exact nature of the leakage in JCJ and propose an alternative (weaker) definition of coercionresistance that is achieved by JCJ. Beyond giving a concrete description of the leakage, this would show that there are several shades of coercion-resistance depending on the leakage which is considered acceptable. More generally, for any variant of JCJ in the literature, if the leakage during the cleansing is different from the one in JCJ, our definition of coercionresistance could (in principle) be adapted to capture this. Still, CHide shows that a better notion of coercion-resistance can be achieved.

\section*{REFERENCES}
[1] R. Araújo, S. Foulle, and J. Traoré, "A practical and secure coercionresistant scheme for remote elections," in Frontiers of Electronic Voting. IBFI, 2007.
[2] S. Baloglu, S. Bursuc, S. Mauw, and J. Pang, "Provably Improving Election Verifiability in Belenios," in E-Vote-ID'21. Springer, 2021.
[3] K. E. Batcher, "Sorting Networks and Their Applications," in American Federation of Information Processing Societies, AFIPS'68. Thomson Book Company, Washington D.C., 1968.
[4] M. Bellare and A. Sahai, "Non-malleable Encryption: Equivalence between Two Notions, and an Indistinguishability-Based Characterization," in CRYPTO '99. Springer, 1999.
[5] J. Benaloh, T. Moran, L. Naish, K. Ramchen, and V. Teague, "Shufflesum: coercion-resistant verifiable tallying for STV voting," IEEE Trans. Inf. Forensics Secur., vol. 4, no. 4, pp. 685-698, 2009.
[6] D. Bernhard, O. Pereira, and B. Warinschi, "How Not to Prove Yourself: Pitfalls of the Fiat-Shamir Heuristic and Applications to Helios," in ASIACRYPT'12. Springer, 2012.
[7] D. Boneh, E. Boyle, H. Corrigan-Gibbs, N. Gilboa, and Y. Ishai, "ZeroKnowledge Proofs on Secret-Shared Data via Fully Linear PCPs," in CRYPTO'19. Springer, 2019.
[8] R. Canetti, "Universally Composable Security: A New Paradigm for Cryptographic Protocols," in 42nd Annual Symposium on Foundations of Computer Science, FOCS'01. IEEE, 2001.
[9] R. Canetti, A. Cohen, and Y. Lindell, "A Simpler Variant of Universally Composable Security for Standard Multiparty Computation," in CRYPTO'15. Springer, 2015.
[10] J. Clark and U. Hengartner, "Selections: Internet Voting with Over-theShoulder Coercion-Resistance," in \(F C^{\prime} 11 . \quad\) Springer, 2011.
[11] M. R. Clarkson, S. Chong, and A. C. Myers, "Civitas: Toward a Secure Voting System," in S\&P'08. IEEE, 2008.
[12] V. Cortier, D. Galindo, S. Glondu, and M. Izabachène, "Election Verifiability for Helios under Weaker Trust Assumptions," in ESORICS'14. Springer, 2014.
[13] V. Cortier, P. Gaudry, and Q. Yang, "A Toolbox for Verifiable TallyHiding E-Voting Systems," in ESORICS'22. Springer, 2022.
[14] _, "A toolbox for verifiable tally-hiding e-voting systems," Cryptology ePrint Archive, Paper 2021/491, 2021, https://eprint.iacr.org/2021/ 491.
[15] David Mestel and Johannes Müller and Pascal Reisert, "How Efficient are Replay Attacks against Vote Privacy? A Formal Quantitative Analysis," in CSF'22. IEEE, 2022.
[16] P. Ehin, M. Solvak, J. Willemson, and P. Vinkel, "Internet voting in Estonia 2005-2019: Evidence from eleven elections," Gov. Inf. Q., vol. 39, no. 4, p. 101718, 2022.
[17] T. Haines and B. Smyth, "Surveying definitions of coercion resistance," Cryptology ePrint Archive, Report 2019/822, 2019, https://ia.cr/2019/ 822.
[18] L. Hirschi, L. Schmid, and D. A. Basin, "Fixing the Achilles Heel of E-Voting: The Bulletin Board," in CSF'21. IEEE, 2021.
[19] N. Huber, R. Küsters, T. Krips, J. Liedtke, J. Müller, D. Rausch, P. Reisert, and A. Vogt, "Kryvos: Publicly tally-hiding verifiable evoting," in CCS'22. ACM, 2022.
[20] M. Jakobsson and A. Juels, "Mix and Match: Secure Function Evaluation via Ciphertexts," in ASIACRYPT'00. Springer, 2000.
[21] M. Jakobsson, K. Sako, and R. Impagliazzo, "Designated Verifier Proofs and Their Applications," in EUROCRYPT'96. Springer, 1996.
[22] A. Juels, D. Catalano, and M. Jakobsson, "Coercion-resistant electronic elections," in Proceedings of the 2005 ACM Workshop on Privacy in the Electronic Society, WPES, 2005.
[23] D. Knuth, The Art Of Computer Programming, vol. 3: Sorting And Searching. Addison-Wesley, 1973.
[24] R. Küsters, J. Liedtke, J. Müller, D. Rausch, and A. Vogt, "Ordinos: A Verifiable Tally-Hiding E-Voting System," in IEEE European Symposium on Security and Privacy (EuroS\&P'20). IEEE, 2020.
[25] R. Küsters, T. Truderung, and A. Vogt, "A Game-Based Definition of Coercion-Resistance and Its Applications," in CSF'10. IEEE, 2010.
[26] -, "Verifiability, Privacy, and Coercion-Resistance: New Insights from a Case Study," in \(S \& P\) '11. IEEE, 2011.
[27] P. Locher, R. Haenni, and R. E. Koenig, "Coercion-Resistant Internet Voting with Everlasting Privacy," in FC'16 International Workshops. Springer, 2016.
[28] W. Lueks, I. Querejeta-Azurmendi, and C. Troncoso, "VoteAgain: A scalable coercion-resistant voting system," in USENIX'20. USENIX, 2020.
[29] Ü. Madise and T. Martens, "E-voting in Estonia 2005. The first Practice of Country-wide binding Internet Voting in the World," in 2nd International Workshop in Electronic Voting (EVOTE'06). GI, 2006.
[30] J. B. Nielsen, "On Protocol Security in the Cryptographic Model," Ph.D. dissertation, University of Aarhus, 2003.
[31] B. Schoenmakers and P. Tuyls, "Practical Two-Party Computation Based on the Conditional Gate," in ASIACRYPT'04). Springer, 2004.
[32] O. Spycher, R. E. Koenig, R. Haenni, and M. Schläpfer, "A New Approach towards Coercion-Resistant Remote E-Voting in Linear Time," in \(F C^{\prime} 11\). Springer, 2011.
[33] , "Achieving Meaningful Efficiency in Coercion-Resistant, Verifiable Internet Voting," in 5th International Conference on Electronic Voting, EVOTE'12. GI, 2012.
[34] D. Wikström, "A Universally Composable Mix-Net," in Theory of Cryptography Conference, TCC'04, ser. Lecture Notes in Computer Science. Springer, 2004.
[35] _-, "Universally Composable DKG with Linear Number of Exponentiations," in Security in Communication Networks (SCN)'04. Springer, 2004.

\section*{APPENDIX A \\ FULL specification of \(P_{\text {Tally }}\)}

In this appendix, we give every details that are necessary to understand the tally protocol. First of all, the CGate protocol [31], built upon the exponential ElGamal encryption system and its threshold decryption protocol, is given in Algorithm 5, where we used the adapted version of [14] (which is the full version of [13]), for which there is a security proof in the SUC framework. The CGate protocol enables to perform multiple functionalities, such as computing a logical conjunction And of two encrypted bits and conditionally set an encrypted value to 0 , given an encrypted boolean (we denote this functionality CSZ). Note that in addition to the output \(Z\), a transcript \(\Pi^{\mathrm{CGate}}\) is obtained, which allows to verify that the participants performed the correct operations on \(X\) and \(Y\), and that \(Z\) is indeed an encryption of \(x y\).
```

Algorithm 5: CGate (see [13]).
The total cost is $33 n_{T}$ exponentiations per participants, and a transcript size of $34 n_{T}$ elements of 256 bits.
Require: $G$, a group of prime order $q$ and public generator $g$
pk, an exponential ElGamal public key,
whose shares are distributed among the $n_{T}$ participants
$\tilde{h} \in G$, a public element independent from pk
$X$, an encryption of some $x \in \mathbb{Z}_{q}$
$Y$, an encryption of $y \in\{0,1\}$
Ensure: $Z$, a random encryption of $x y$
$Y_{0} \longleftarrow E_{-1} Y^{2} ; X_{0} \longleftarrow X$
for $i=1$ to $n_{T}$, for the authority $i$, do
$(u, v) \longleftarrow X_{i-1}$
$r_{1}, r_{2} \stackrel{\$}{\leftarrow} \mathbb{Z}_{q} ; s \stackrel{\$}{\leftarrow}\{-1,1\}$
$X_{i} \longleftarrow\left(u^{s} g^{r_{1}}, v^{s} \mathrm{pk}^{r_{1}}\right) ; e \longleftarrow \tilde{h}^{r_{1}}$
$Y_{i} \longleftarrow \operatorname{ReEnc}_{\mathrm{pk}}\left(Y_{i-1}^{s}, r_{2}\right)$
Broadcast $X_{i}, e, Y_{i}$ and a ZKP $\pi_{i}$ that they are well formed
We denote $\Pi=\left(X_{n_{T}}, Y_{n_{T}}, \pi_{n_{T}}\right)\|\cdots\|\left(X_{1}, Y_{1}, \pi_{1}\right)$
Each authority verifies the proof of the other authorities
They collectively rerandomize $X_{n_{T}}$ and $Y_{n_{T}}$ into $X^{\prime}$ and $Y^{\prime}$, which produces a transcript $\Pi^{\text {ReEnc }}$
They collectively decrypt $Y^{\prime}$ into $y_{n_{T}}$, which produces a transcript $\Pi^{\text {Dec }}$.
They output $Z=\left(X X^{\prime y_{n_{T}}}\right)^{\frac{1}{2}}$. The verification transcript is $\Pi^{\mathrm{CGate}}=\left(y_{n_{T}}, \Pi^{\mathrm{Dec}}\right)\left\|\left(X^{\prime}, Y^{\prime}, \Pi^{\mathrm{ReEnc}}\right)\right\| \Pi$

```

As shown in [31], due to the homomorphic property of the ElGamal cryptosystem, we can readily deduce protocols for computing a disjunction (see Algorithm 6), and an equality test (Algorithm 7) between two bits. In addition, the logical negation of an encrypted bit \(B \operatorname{Not}(B)\) can be obtained without interaction by computing Enc(1)/B (for this purpose, a fixed encryption of 1 with a trivial randomness can be used).
```

Algorithm 6: Or
Require: $X, Y$ such that $X$ (resp. $Y$ ) is an encryption
of $x$ (resp. $y$ ), with $x, y \in\{0,1\}$
Ensure: $Z=\operatorname{Enc}(z)$, with $z=1$ if $x=1$ or $y=1 ; 0$
otherwise
Return $X Y / \mathrm{CGate}(X, Y)$ and $\Pi^{\mathrm{CGate}}$

```
```

Algorithm 7: Eq
Require: $X, Y$ such that $X$ (resp. $Y$ ) is an
encryption of $x$ (resp. $y$ ), with $x, y \in\{0,1\}$
Ensure: $Z=\operatorname{Enc}(z)$, with $z=1$ if $x=y ; 0$
otherwise
1 Return Enc(1)CGate $(X, Y)^{2} /(X Y)$ and
$\Pi^{\mathrm{Eq}}=\Pi^{\mathrm{CGate}}$

```

Now, since the logical disjunction (resp. conjunction) operator is associative, we immediately deduce a protocol to extend this binary operator into a \(n\)-ary operator (by abuse of notation, we still denote it Or; resp. And). By using a binary tree structure, the protocol \(\operatorname{Or}\left(B_{1}, \cdots, B_{n}\right)\) (resp. \(\operatorname{And}\left(B_{1}, \cdots, B_{n}\right)\) ); which returns an encryption of 1 if one (resp. all) of the inputs is an encryption of 1 , and 0 otherwise; can be performed in \(O(\log n)\) rounds and a total of \(n-1\) instances of CGate.

We can also extend the Eq protocol to bitwise encrypted integers (see Algorithm 8).
By combining logical operations, it is straightforward to build a comparison protocol Lt (see Algorithm 9 from [13]) which, given two bitwise encryptions of the integers \(x\) and \(y\), returns an encryption of 1 if \(x<y\) and an encryption of 0 otherwise. We also give the conditional swap protocol in Algorithm 10. By combining those two protocols, we can readily adapt the OddEvenMergeSort algorithm in MPC (see Algorithm 11, adapted from [23, Section 5.2.2, Algorithm M]). This sorting algorithm requires approximately \(\frac{1}{4} n \log (n)^{2}\) comparisons and conditional swaps, where \(n\) is the number of elements
```

Algorithm 8: Eq (for bitwise encrypted integers)
Require: $X_{0}, \cdots, X_{m}, Y_{0}, \cdots, Y_{m}$ bit-wise encryptions of $x$ and $y$.
Ensure: $Z$, an encryption of 1 if $x=y, 0$ otherwise.
for all ( in parallel) do
$A_{i} \longleftarrow \mathrm{Eq}\left(X_{i}, Y_{i}\right)$
Return $Z=\operatorname{And}\left(A_{0}, \cdots, A_{m-1}\right)$ and the concatenation of all the verification transcript

```
to be sorted. In the \(P_{\text {tally }}\) protocol, if \(\lambda\) is the security parameter, \(n_{B}\) the number of valid ballots received and \(n_{V}\) the number of eligible voters, then \(n=n_{B}+n_{V}\) and performing a comparison and a conditional swap costs about \(3(\lambda+m)\) CGate in total, where \(m=\left\lceil\log \left(n_{B}\right)\right\rceil+1\). If \(n_{T}\) is the number of talliers, this translates into \(O\left((\lambda+m) n_{T}\left(n_{B}+n_{V}\right) \log \left(n_{B}+n_{V}\right)^{2}\right)\) exponentiations.
```

Algorithm 9: Lt
Require: $\left(X_{0}, \cdots, X_{m-1}\right),\left(Y_{0}, \cdots, Y_{m-1}\right)$, bit-wise
encryptions of $x$ and $y$.
Ensure: $R$, an encryption of 1 if $x<y$ and 0 otherwise.
$A \longleftarrow \operatorname{CGate}\left(X_{0}, Y_{0}\right)$
$Z_{0} \longleftarrow X_{0} Y_{0} / A^{2}\left(* x_{0} \oplus y_{0}{ }^{*}\right)$
$R \longleftarrow Y_{0} / A\left({ }^{*} y_{0} \wedge \neg x_{0}{ }^{*}\right)$
for $k=1$ to $m-1$ do
$A \longleftarrow \operatorname{CGate}\left(Y_{k}, R\right)$
$B \longleftarrow Y_{k} R / A^{2}\left(* y_{k} \oplus r *\right)$
$C \longleftarrow \operatorname{CGate}\left(X_{k}, B\right)$
$Z_{k} \longleftarrow X_{k} B / C^{2}\left({ }^{*} x_{k} \oplus y_{k} \oplus r^{*}\right)$
$R \longleftarrow Y_{k} R /(A C)\left(*\left(y_{k} \wedge r\right) \vee\left[\left(y_{k} \vee r\right) \wedge \neg x_{k}\right] *\right)$

```
    Return \(R\) and the concatenation of all the verification
    transcripts
```

Algorithm 11: OddEvenMergeSort
Require: $\left(V_{i}, K_{i}\right)_{i=0}^{n-1}$, where, for all $i, V_{i}$ is an encryption while $K_{i}$ is a bitwise encryption of an integer $k_{i}$
Ensure: $\left(V_{i}^{\prime}, K_{i}^{\prime}\right)_{i=0}^{n-1}$, reencryptions of the same values, but sorted with increasing $k_{i}$
$t \longleftarrow\lceil\log n\rceil ; p \longleftarrow 2^{t-1} ; \Pi^{\text {Sort }} \longleftarrow \varnothing$
while $p>0$ do
$q \longleftarrow 2^{t-1} ; r \longleftarrow 0 ; d \longleftarrow p$
while $d>0$ do
for $i=0$ to $n-d-1$ (in parallel) do
if BitwiseAnd $(i, p)=r$ then
$B, \Pi^{\mathrm{Lt}} \longleftarrow \operatorname{Lt}\left(K_{i+d}, K_{i}\right) ; \Pi^{\text {Sort }} \longleftarrow \Pi^{\text {Sort }} \| \Pi^{\mathrm{Lt}}$
$V_{i}, V_{i+d}, \Pi^{\text {CSwap }} \longleftarrow \operatorname{CSwap}\left(V_{i}, V_{i+d}, B\right) ; \Pi^{\text {Sort }} \longleftarrow \Pi^{\text {Sort }} \| \Pi^{\text {CSwap }}$
$K_{i}, K_{i+d}, \Pi^{\text {CSwap }} \longleftarrow \operatorname{CSwap}\left(K_{i}, K_{i+d}, B\right) ; \Pi^{\text {Sort }} \longleftarrow \Pi^{\text {Sort }} \| \Pi^{\text {CSwap }}$ (* bit by bit, in parallel *)
$d \longleftarrow q-p ; q \longleftarrow\lfloor q / 2\rfloor ; r \longleftarrow p$
$p \longleftarrow\lfloor p / 2\rfloor$
Return $\left(V_{i}, K_{i}\right)_{i=0}^{n-1}$ and the transcript $\Pi^{\text {Sort }}$

```

With all the above protocols, \(P_{\text {tally }}\) can be written as Algorithm 12, where \(B B\) is the ballot box; \(\boldsymbol{R}\) is the public roster that contains the encrypted credentials.
```

Algorithm 12: $P_{\text {tally }}$
Require: A t-threshold setup phase is assumed to have been run between the $n_{T}$ participants.
Require: $B B, \boldsymbol{R}$
$B=\left\{\left(C_{1}, C_{2}\right) \mid\left(C_{1}, C_{2}, \pi_{1}, \pi_{2}\right) \in B B\right.$,isVal $\left.(\mathrm{pk}, \mathbf{b}, B B)=1\right\}$
We denote $B=\left(C_{1}^{i}, C_{2}^{i}\right)_{i=1}^{n_{B}}$
We denote $\boldsymbol{R}=R_{1}, \cdots, R_{n_{V}}$
$\Pi \longleftarrow \varnothing ; m \longleftarrow\left\lceil\log n_{B}\right\rceil+1$
5 For $k=0$ to $n_{B}$, we denote $U_{k}$ a bitwise encryption of $k$ with $m$ bits (least significant bits first)
6 (In other words, if $\left.k=\sum_{i=0}^{m-1} b_{i, k} 2^{i}, U_{k}=\operatorname{Enc}\left(b_{0, k}\right), \cdots, \operatorname{Enc}\left(b_{m-1, k}\right)\right)$
for $k=1$ to $n_{B}$ do
$V_{k} \longleftarrow C_{1}^{k} ; K_{k} \longleftarrow U_{k-1} \| C_{2}^{k}$
for $k=n_{B}+1$ to $n_{B}+n_{V}$ do
$V_{k} \longleftarrow \operatorname{Enc}(0) ; K_{k} \longleftarrow U_{n_{B}} \| R_{k-n_{B}}$
$\left(V_{k}^{\prime}, K_{k}^{\prime}\right)_{k=1}^{n_{B}+n_{V}}, \Pi^{\text {Sort }} \longleftarrow$ OddEvenMergeSort $\left(\left(V_{k}, K_{k}\right)_{k=1}^{n_{B}+n_{V}}\right)$
$\Pi \longleftarrow \Pi\left|\mid \Pi^{\text {Sort }}\right.$
for $k=1$ to $n_{B}+n_{V}-1$ (in parallel) do
$D_{k}, \Pi_{1}^{\mathrm{Eq}} \longleftarrow \mathrm{Eq}\left(K_{k}^{\prime \top}, K_{k+1}^{\prime \top}\right)(*$ compare the $\lambda$ most significant bits $*)$
$F_{k}, \Pi_{2}^{\mathrm{Eq}} \longleftarrow \mathrm{Eq}\left(K^{\prime} \stackrel{\perp}{k+1}, U_{n_{B}}\right)(*$ compare the $m$ least significant bits *)
$\Pi \longleftarrow \Pi\left|\left|\Pi_{1}^{\mathrm{Eq}}\right|\right| \Pi_{2}^{\mathrm{Eq}}$
$G_{k}, \Pi^{\text {And }} \longleftarrow \operatorname{And}\left(D_{k}, F_{k}\right) ; \Pi \longleftarrow \Pi| | \Pi^{\text {And }}$
$V_{k}^{\prime}, \Pi^{\mathrm{CGate}} \longleftarrow \mathrm{CGate}\left(V_{k}^{\prime}, G_{k}\right) ; \Pi \longleftarrow \Pi| | \Pi^{\mathrm{CGate}}$
$\boldsymbol{X}, \Pi^{\text {DecMixnet }} \longleftarrow \operatorname{DecMixnet}\left(V^{\prime}{ }_{1}, \cdots, V^{\prime}{ }_{n_{B}+n_{V}-1}\right)$
$\Pi \longleftarrow \Pi\left|\mid \Pi^{\text {DecMixnet }}\right.$
Return $\boldsymbol{X}, \Pi$

```

\section*{Appendix B \\ The SUC security of \(P_{\text {Tally }}\)}

In our security proofs, we use the SUC framework [9] which is a simpler variant of the Universally Composable (UC) framework [8]. In this appendix, we give a brief introduction to the SUC framework and prove that \(P_{\text {tally }}\) is SUC-secure.

\section*{A. Definitions, notations and composition}

In the SUC framework, a protocol is a collection of participants, which are modeled as Interactive Turing Machines (ITM). As the name implies, the ITMs communicate with each others using common tapes. An execution of a protocol is called a process, during which the participants are executed concurrently. In turn, a process can invoke sub-processes, but with the same participants.

Some of the participants may be impersonated by an adversary \(\mathbb{A}\) which also has a full control over the communication network: it can read, block and deliver the messages as it sees fit, which includes changing the order of two messages. However, the adversary cannot change (or forge) the content nor the recipient of a message, nor can it replay a message which was already delivered. In a sense, this is the same as considering strongly authenticated (but insecure) channels with enforced integrity. In addition to \(\mathbb{A}\), there is a second adversarial ITM \(\mathcal{Z}\) which represents the environment. It may only interact with \(\mathbb{A}\), and can output either 0 or 1 .

The security of a process is assessed by a comparison with an ideal process, where a trusted party computes the output of the process with the inputs of the participants. Intuitively, a protocol is SUC-secure if, for all adversary \(\mathbb{A}\) in the real process, there exists a simulator \(\mathcal{S}\) in the ideal process such that for all environment \(\mathcal{Z}\), the latter cannot tell whether it interacts with \(\mathbb{A}\) in the real process or \(\mathcal{S}\) in the ideal process.
1) The real and ideal processes: The real process is mostly determined by the adversary and the protocol, which is the specification of every participants. Conversely, the ideal process is mostly determined by the simulator and the trusted party. If \(\mathcal{F}\) is a functionality, the trusted party which computes \(\mathcal{F}\) from the inputs of the participants is denoted \(T_{\mathcal{F}}\). In the ideal process, it is assumed that the simulator (but not the environment) can interact with \(T_{\mathcal{F}}\), which may react in various manners depending on the functionality. Typically, if the adversary corrupts so many participants that it becomes impossible for the remaining honest ones to run the protocol successfully, it is reasonable to expect that the protocol aborts and output \(\perp\). In this case, in the ideal process, the simulator is given the possibility to have the trusted party abort and output \(\perp\). Another quirk of the ideal process is that the simulator and the trusted parties know which participants are corrupted.

In both the ideal and real processes, the environment is given an arbitrary auxiliary input \(z \in\{0,1\}^{*}\) and chooses the inputs of all the participants. When the process terminates, \(\mathcal{Z}\) can also read the output of all the participants (which is the output of the trusted party in the ideal process).
2) The hybrid process: In the hybrid process, we consider an ideal functionality \(\mathcal{G}\) and the corresponding trusted party \(T_{\mathcal{G}}\). In addition to regular operations, we suppose that the participants can send (secure) queries to \(T_{\mathcal{G}}\) which answers according to its specification. (For instance, computing the mean of several values can only be done when all values are known, so that \(T_{\mathcal{G}}\) may have to wait until sufficiently many inputs are sent.) Contrary to the other communication channels, the one which connects the participants to \(T_{\mathcal{G}}\) in the hybrid model is secure, meaning that the adversary cannot read the queries nor their answers (except if the participant communicating with \(T_{\mathcal{G}}\) is corrupted). However, it can still block or delay the messages.

In this article, we denote \(\operatorname{Rea}_{P, \mathbb{A}, \mathcal{Z}}^{S U C, \mathcal{G}}(\lambda, z)\) (resp. Ideal \(\left.\right|_{T_{\mathcal{F}}, \mathcal{S}, \mathcal{Z}} ^{S U C}(\lambda, z)\) ) the probability that \(\mathcal{Z}\) outputs 1 , given the auxiliary input \(z\), when interacting with \(\mathbb{A}\) in the \(\mathcal{G}\)-hybrid process (resp. with \(\mathcal{S}\) which interacts with \(T_{\mathcal{F}}\) in the ideal process), where \(\lambda\) is the security parameter. We can now give the definition of SUC-security.

Definition 3 ( [9]). Let \(P\) be a protocol and \(T_{\mathcal{F}}\) and \(T_{\mathcal{G}}\) two trusted parties, which computes the functionalities \(\mathcal{F}\) and \(\mathcal{G}\) respectively. We say that \(P\) SUC-securely computes \(\mathcal{F}\) in the \(\mathcal{G}\)-hybrid model if, for all PPT real adversary \(\mathbb{A}\), there exists a PPT ideal simulator \(\mathcal{S}\) such that for every PPT environment \(\mathcal{Z}\) and every constant \(d \in \mathbb{N}\), there exists a negligible function \(\mu\) such that every security parameter \(\lambda \in \mathbb{N}\) and every auxiliary input \(z \in\{0,1\}^{\lambda^{d}}\),
\[
\left|\operatorname{Real}_{P, \mathbb{A}, \mathcal{Z}}^{\mathrm{SUC}, \mathcal{G}}(\lambda, z)-\operatorname{Idea|}\right|_{T_{\mathcal{F}}, \mathcal{S}, \mathcal{Z}}^{\mathrm{SUC}}(\lambda, z) \mid \leq \mu(\lambda)
\]
3) The composition theorem: UC frameworks are extremely powerful due to their composition theorem, which allow to deduce the security of a protocol from the security of its sub-protocols. More precisely, suppose that \(P\) is a protocol to compute \(\mathcal{F}\) in the \(\mathcal{G}\)-hybrid model. Suppose, in addition, that there exists a protocol \(Q\) to compute \(\mathcal{G}\) in the \(\mathcal{H}\)-hybrid model. Then we can consider the protocol \(P^{Q}\) in the \(\mathcal{H}\)-hybrid model, where every query to \(T_{\mathcal{G}}\) is replaced by an instance of \(Q\). Given this transformation, the composition theorem can be written as follows.

Theorem 2 ( [9]). Let \(P\) be a protocol that SUC-securely computes \(\mathcal{F}\) in the \(\mathcal{G}\)-hybrid model and \(Q\) a protocol that SUCsecurely computes \(T_{\mathcal{G}}\) in the \(\mathcal{H}\)-hybrid model. Then \(P^{Q}\), as defined above, \(S U C\)-securely computes \(\mathcal{F}\) in the \(\mathcal{H}\)-hybrid model.

\section*{B. Application to CHide}

To assert the security of CHide, we use the SUC framework. For this purpose, we consider a static and malicious adversary which can fully impersonate the corrupted parties, but can only corrupt a fixed set of participants, which is decided at the beginning of the protocol. (We also consider that this set's cardinal may not exceed the threshold.) We consider a fixed number of participants with a fixed threshold (participants cannot leave or arrive and the threshold may not change). We use the so-called programmable random oracle model which is modeled as an ideal functionality \(\mathcal{F}_{\mathcal{R O}}\), and we also suppose that the DKG and the decryption mixnet are SUC -secure. We denote \(\mathcal{F}_{\mathrm{DKG}}\) and \(\mathcal{F}_{\text {DecMixnet }}\) the corresponding ideal functionalities. For more details, including a construction to obtain a UC-secure protocol for those, see for instance the work of Wikström [34], [35].

Hence, we place ourselves in the \(\mathcal{G}\)-hybrid model, where \(\mathcal{G}=\left(\mathcal{F}_{\mathcal{R O}}, \mathcal{F}_{\text {DecMixnet }}\right)\).
Following [13], we restrict ourselves to environment which only gives as input "valid" bulletin boards; i.e. the ballots are of the form \(\left(C_{1}, C_{2}, \pi_{1}, \pi_{2}\right)\) with valid ZKP \(\pi_{1}, \pi_{2}\) as defined in Section IV, where \(C_{1}\) is a valid ElGamal encryption (that is, any pair of group elements) and \(C_{2} \lambda\) valid ElGamal encryptions of bits. In his PhD thesis [30, Section 3.5], Nielsen proves that the composition theorem still holds if the inputs chosen by the environment are restricted to some decidable language. This allows to use [13]'s result, where the CGate protocol is shown to SUC-securely computing \(T_{\mathrm{CSZ}}\) in the \(\mathcal{G}\)-hybrid model, where \(\operatorname{CSZ}\) is the Conditional Set Zero functionality (i.e. \(\operatorname{CSZ}(A, B)\), for any ciphertext \(A\) and any encryption \(B\) of 0 or 1 , returns a random reencryption of \(A\) if \(B\) is an encryption of 1 and a random encryption of 0 otherwise). We can now state Theorem 3, which is a direct consequence of the composition theorem. In this theorem, we consider \(T_{\text {tally }}\), the trusted party that computes the result of the election protocol from the public board and the inputs of all the parties, and returns the result as well as the output of all the conditional gate protocols.

Theorem 3. Under the DDH assumption, \(P_{\text {tally }} S U C\)-securely computes \(T_{\text {tally }}\) in the \(\mathcal{G}\)-hybrid model.
Proof sketch. As shown in [13], CGate SUC-securely computes \(T_{\mathrm{CSz}}\) in the \(\mathcal{G}\)-hybrid model. Consequently, by the composition theorem, it is sufficient to show that \(Q=P_{\text {tally }}^{T_{\text {Csz }}, T_{\text {DecMixnet }}} \mathrm{SUC}\)-securely computes \(T\) tally in the \(\left(\mathcal{G}, \mathcal{F}_{\mathrm{CSz}}\right)\)-hybrid model, where \(Q\) is \(P_{\text {tally }}\) where all the instances of CGate are replaced by a query to \(T_{\mathrm{CSZ}}\) while the instance of the decryption mixnet protocol is replaced by a query to \(T_{\text {DecMixnet }}\). In other words, it is sufficient to show that, given the ideal result computed by \(T_{\mathrm{tally}}\), the simulator can simulate the output of all the trusted parties, which comes immediately since it can use the result for the output of the decryption mixnet, and that it has the output of the conditional gate protocols from \(T_{\text {tally }}\).

\section*{Appendix C \\ SECURITY PRoofs FOR CHide}

\section*{A. Preliminary: on the IND-PAO game}

To prove coercion-resistance, we use a reduction to the IND-PA0 security of the encryption scheme. This security notion is defined, for instance, in [4], where it is shown to be equivalent to NM-CPA. In this section, we recall the definition of IND-PA0 and formalize the voting protocol as an IND-PA0 encryption scheme.

First, using the notations of Bellare and Sahai, let \((\mathcal{K}, \mathcal{E}, \mathcal{D})\) be an encryption scheme. We say that it is IND-PA0 if, for all (stateful) PPT adversary \(\mathbb{A}\), its advantage in the game depicted in Algorithm 13 is negligible in the security parameter \(\lambda\). Note that at line 5, the adversary can send a list of several ciphertexts, and gets the element-wise decryptions of those ciphertexts at the following line.

Now, let (Gen, Enc, Dec) be the ElGamal encryption scheme, which is known to be IND-CPA secure under DDH assumption. For simplicity, we assume that any voting option can be encoded as a single group element. For any \(\lambda\)-bits credential \(c=\) \(c_{1}, \cdots, c_{\lambda}\) and voting option \(\nu\), recall that \(\operatorname{Vote}_{\mathrm{pk}}(c, \nu)\) is obtained by computing the encryptions \(C_{1}=\operatorname{Enc}_{\mathrm{pk}}(\nu)\) and \(C_{2}=\left(\operatorname{Enc}_{\mathrm{pk}}\left(g^{c_{1}}\right), \cdots, \operatorname{Enc}_{\mathrm{pk}}\left(g^{c_{\lambda}}\right)\right)\), where \(g\) is a public group element. Then, a NIZKPoK \(\pi_{1}\) of the plaintexts \((c, \nu)\) is computed using the Fiat-Shamir heuristic. To make sure that the ballot remains non-malleable, all the public informations such as the public encryption key, \(g, C_{1}\) and \(C_{2}\) are included into the hash to compute the challenge from the commitment of the \(\Sigma\)-protocol. Finally, a ZKP \(\pi_{2}\) that \((c, \nu)\) is valid is computed and the ballot \(\left(C_{1}, C_{2}, \pi_{1}, \pi_{2}\right)\) is formed. Clearly, Vote \(\mathrm{ek}(c, \nu)\) can be seen as an encryption of the tuple \((c, \nu)\) by the encryption scheme (Gen, Vote, Extract), where the Extract algorithm first checks the ZKPs and outputs \(\perp\) if any of the proof is invalid. Otherwise, it decrypts all the ciphertexts with the secret key and returns the tuple \((c, \nu)\). By a result of [6], this defines an NM-CPA encryption scheme, which is therefore IND-PA0 secure. Note that \(\pi_{2}\) is not necessary for the IND-PA9 security but does not deteriorate it either. It is useful for other properties of the protocol. In the following section, whenever we use the IND-PA0 security, we refer to the security of this encryption scheme whose plaintext (resp. ciphertext) space consists of \((\lambda+1)\)-tuples of plaintexts (resp. ciphertexts, along with the corresponding ZKP).
```

Algorithm 13: IND-PA0 game
$(\mathrm{pk}, \mathrm{sk}) \longleftarrow \mathcal{K}(\lambda)$
$2\left(x_{0}, x_{1}\right) \longleftarrow \mathbb{A}(\mathrm{pk})$
$3 b \stackrel{\mathscr{\&}}{\leftarrow}\{0,1\}$
$4 \longleftarrow \mathcal{E}_{\mathrm{pk}}\left(x_{b}\right)$
$5 c \longleftarrow \mathbb{A}(y)$
${ }^{\boldsymbol{f}} \boldsymbol{p} \longleftarrow \mathcal{D}_{\text {sk }}(\boldsymbol{c})$
$b^{\prime} \longleftarrow \mathbb{A}(\boldsymbol{p})$
Return 1 if and only if $(y \notin c) \wedge\left(b^{\prime}=b\right)$

```

\section*{B. Proof of coercion-resistance}

The definition of coercion-resistance is given in Definition 2. In this section, we prove Theorem 1, that we restate below. For this purpose, we give a succession of games such that Game 0 is the real game and Game 10 is the ideal game. We consider a PPT \(\mathbb{A}_{0}\) for Game 0 . For Game \(i\), we construct a PPT adversary \(\mathbb{A}_{i}\) for this game and we denote \(S_{i}\) the probability that \(\mathbb{A}_{i}\) wins this game. (To ease the notation, we drop the dependency in \(\lambda\) when the context is clear.) For all \(i\), we show that \(\left|S_{i+1}-S_{i}\right|\) is negligible, which proves that \(\left|S_{0}-S_{9}\right|\) is also negligible.
Theorem 1. Under the DDH assumption, assuming SUC-secure DKG and Mixnet protocols and in the Programmable Random Oracle Model, CHide is coercion-resistant.

Proof. Game 1: In this game, the adversary no longer takes part into the tally process at line 28 . Instead, it is given the output of \(T_{\text {tally }}\).

By Theorem 3, there exists a simulator \(\mathcal{S}\) such that, for all environment \(\mathcal{Z},|\operatorname{Real}|_{P_{\text {tally }}, \mathbb{A}_{0}, \mathcal{Z}}^{\operatorname{SUC}, \mathcal{G}}(\lambda, z)(\lambda, 0)-\operatorname{Ideal} \operatorname{lic}_{T_{\text {tally }}, \mathcal{S}, \mathcal{Z}}^{S U C}(\lambda, 0) \mid\) is negligible. In particular, we consider the environment which defines Game 0 , so that Real \(P_{P_{\text {tally }}, \mathbb{A}_{0}, \mathcal{Z}}^{\mathrm{SUC}, \mathcal{A}}(\lambda, z)(\lambda, 0)=S_{0}\). Then, \(\mathbb{A}_{1}\) can interact with \(\mathbb{A}_{0}\) by simulating the real game using \(\mathcal{S}\), so that Ideal \(T_{T_{\text {taly }}, \mathcal{S}, \mathcal{Z}}^{\mathrm{SUC}}(\lambda, 0)=S_{1}\). Hence, \(\left|S_{1}-S_{0}\right|\) is negligible.

Game 2: In this game, we remove line 2. Instead, we generate a perfectly random public encryption key pk and give it to the adversary at line 4 . Just as in the previous transition, assuming an ideal DKG protocol, there exists a negligible function \(\mu_{2}\) and an adversary \(\mathbb{A}_{2}\) for Game 2 such that \(\left|S_{2}-S_{1}\right| \leq \mu_{2}\).

Game 3: In this game, the adversary is no longer given the output of the conditional gates, that we denote \(\Pi^{Z}\).
We construct \(\mathbb{A}_{3}\) that interacts with \(\mathbb{A}_{2}\) by simulating \(\Pi^{Z}\). For this purpose, \(\mathbb{A}_{3}\) uses uniformly random ciphertexts.
To argue the validity of this transition, we construct an adversary \(\mathbb{B}\) for DDH as follows. First, \(\mathbb{B}\) gets the challenge tuple \(\left(g, g_{2}, g_{3}, g_{4}\right)\) from the DDH game and sets \(\mathrm{pk}=g_{2}\). To run the setup, \(\mathbb{B}\) recovers the set \(S\) of the corrupted participants from \(\mathbb{A}_{2}\), and picks \(s_{i} \in \mathbb{Z}_{q}\) at random for all \(i \in S\). It the completes \(S\) into \(I\) by picking some additional \(s_{i} \in \mathbb{Z}_{q}\) at random for all \(i \in I \backslash S\), where \(I \subset\left[1, n_{T}\right]\) is a set of size \(t\) that contains \(S\), and \(n_{T}\) is the number of talliers. For \(i \in I\), it computes \(h_{i}=g_{1}^{s_{i}}\) and, for \(i \in\left[1, n_{T}\right] \backslash I\), it deduces \(h_{i}\) with Lagrange interpolation.

It then runs the remaining of Game 3honestly, but each time \(\mathbb{A}_{2}\) casts a ballot, \(\mathbb{B}\) extracts the corresponding voting option from \(\mathbb{A}_{2}\) 's proof of knowledge. In the ROM, this is possible in polynomial time, as a consequence of the forking lemma (see for instance Theorem [6, Theorem 1]). This way, \(\mathbb{B}\) can compute the result \(r\) of the tally without knowing the secret key sk. Finally, since \(\mathbb{B}\) knows the cleartexts of the ballots to tally, \(\mathbb{B}\) can run the tally protocol "on the cleartexts", i.e. it can compute the cleartext of each of the outputs of each conditional gate, as the product of two cleartexts. To simulate the output of a conditional gate, \(\mathbb{B}\) "encrypts" the corresponding cleartext \(z\) by choosing two random \(\rho_{1}, \rho_{2} \in \mathbb{Z}_{q}\) and computing \(Z=\left(g^{\rho_{1}} g_{3}^{\rho_{2}}, g^{z} g_{2}^{\rho_{1}} g_{4}^{\rho_{2}}\right)\). Finally, if \(\mathbb{A}_{2}\) wins the game, \(\mathbb{B}\) states that the challenge was a DDH tuple; otherwise, it states that is was a random tuple. Remark that if \(\left(g, g_{2}, g_{3}, g_{4}\right)\) is a DDH tuple, then \(\mathbb{B}\) played a perfect simulation of Game 2 to \(\mathbb{A}_{2}\) and hence wins with probability \(S_{2}\). On the other hand, if the challenge tuple is a random tuple, \(\mathbb{B}\) played \(\mathbb{A}_{3}\) 's simulation of Game 2and wins with probability \(1-S_{3}\). Yet, under the DDH assumption, \(\mathbb{B}\) 's advantage in the DDH game must be negligible, hence \(\left|S_{2}-S_{3}\right|\) is negligible.

Game 4: In this game, we modify the sequence \(B\) so that the honest voters no longer revote. Instead, for all honest voter \(x\), we replace all but the last occurrence of the form \((x, \nu)\) in \(B\) by an occurrence of the form \((\tilde{x}, \nu)\) which uses a fresh and unique \(\tilde{x}<0\). This way, the last vote remains the same but the previous votes are replaced by a vote with a fresh, random (and fake) credential. Note that this modification happens on the sequence \(B\), before the voters actually cast their votes according to this sequence.

Let \(n_{r}\) be the number of revotes. We set \(\mathbb{A}_{4}=\mathbb{A}_{3}\) and to argue that \(\left|S_{4}-S_{3}\right|\) is negligible, we give another succession of hops \(H_{0}, \cdots, H_{n_{r}}\) such that in game \(H_{i}\), we replace the last \(i\) revotes as described above. This way, \(H_{0}\) is Game 3 and \(H_{n_{r}}\) is Game 4. For each of these games, we denote \(W_{i}\) the probability that \(\mathbb{A}_{4}\) wins this game.

Now, let \(i\) be some index. We construct an adversary \(\mathbb{B}\) for IND-PA0 (see Section \(C\) - \(A\) ) as follows. First, \(\mathbb{B}\) receives pk from the IND-PA0 game. It simulates game \(H_{i}\) by giving this pk to \(\mathbb{A}_{4}\) and generating the credentials at random. It then gets \((j, \alpha)\) from \(H_{i}\) and chooses \(b\) at random. Afterwards, \(\mathbb{B}\) generates \(B\) from the distribution \(\mathcal{B},(j, \alpha)\) and \(b\). It then continues the simulation of \(H_{i}\) by replacing the \(i\) last revotes as necessary. However, for the next remaining revote, it looks up for the previous vote \((x, \nu)\) with the same \(x\) and generates a random, fresh credential \(c\). It plays \(\left(c_{x}, \nu\right),(c, \nu)\) in the IND-PA0 game and gets an encrypted ballot \(C\) which it plays in the simulation of \(H_{i}\) instead of a honestly generated ballot for \(\left(c_{x}, \nu\right)\) (therefore, \(C\) is added on the board). Finally, to compute the tally, \(\mathbb{B}\) simply plays the concatenation of all the valid ballots sent by \(\mathbb{A}_{4}\) in the IND-PA0 game, which returns a decryption of these ballots. Remark that due to the nature of isVal which uses a form of weeding, a valid ballot is not already on the board. Therefore, none of the valid ballots sent by \(\mathbb{A}_{4}\) can be equal to \(C\), so that the IND-PA0 will indeed accept to decrypt them. \(\mathbb{B}\) computes the result of the tally from the plaintexts and gives it to \(\mathbb{A}_{4}\) which answers with some bit \(b^{\prime}\). This is possible because every valid ballot of the board has a valid ZKP \(\pi_{2}\), whose soundness guarantees that the corresponding plaintext \((c, \nu)\) is such that \(c\) is a \(\lambda\)-bit credential and \(\nu\) a valid voting option. If \(b=b^{\prime}, \mathbb{B}\) states that the IND-PA0 game encrypted \(\left(c_{x}, \nu\right)\), and \((c, \nu)\) otherwise.

Clearly, when the IND-PA0 game encrypts \(\left(c_{x}, \nu\right)\) (resp. \((c, \nu)\) ), \(\mathbb{B}\) plays a perfect simulation of game \(H_{i}\) (resp. \(H_{i+1}\) ) so that \(b=b^{\prime}\) with probability \(W_{i}\) (resp. \(W_{i+1}\) ). Hence, \(\mathbb{B}\) 's advantage in the IND-PA0 game is \(\left|\frac{1}{2}\left(W_{i}+1-W_{i+1}\right)-\frac{1}{2}\right| \leq \varepsilon_{\text {PAo }}\). By the triangular inequality,
\[
\left|S_{4}-S_{3}\right| \leq 2 n_{r} \varepsilon_{\mathrm{PAO}}
\]

Game 5: In this game, the adversary no longer has access to the roster \(\boldsymbol{R}\) at line 4 . We construct \(\mathbb{A}_{5}\) which interacts with \(\mathbb{A}_{4}\) by simulating Game 4 . For this purpose, it generates \(n_{V} \lambda\) ElGamal encryptions of random bits and uses it to simulate \(\boldsymbol{R}\). For the remaining of the game, it can play a perfect simulation since both games are identical. To argue that \(\left|S_{5}-S_{4}\right|\) is negligible, we construct a succession of hop \(H_{0}, \cdots, H_{n_{V}}\) such that in game \(H_{i}\), the last \(i\) elements of \(\boldsymbol{R}\) are replaced by a random encryption. This way, \(H_{0}\) is Game 4 and \(H_{n_{V}}\) is Game 5. For each of these games, we denote \(W_{i}\) the probability that \(\mathbb{A}_{4}\) wins this game.

Now, let \(i\) be some index. We construct an adversary \(\mathbb{B}\) for IND-PA0 as follows. \(\mathbb{B}\) gets pk from the IND-PA0 game and plays this pk to \(\mathbb{A}_{4}\) to simulate \(H_{i}\). It generates the credentials and replaces the last \(i\) elements of the roster by encryptions of random bits just as in \(H_{i}\). However, for the \(\left(n_{V}-i\right)\) th element of \(\boldsymbol{R}\), it generates a second random credential \(c\) and plays the pair \(\left(c_{n_{V}-i}, 1\right),(c, 1)\) in the IND-PA0 game which answers with some encryption \(C\). \(\mathbb{B}\) retrieves an ElGamal bitwise encryption of the credential from \(C\) and uses it as the \(\left(n_{V}-i\right)\) th element of \(\boldsymbol{R}\) instead of an honest bitwise encryption of \(c_{n_{V}-i}\). Afterwards, \(\mathbb{B}\) continues the simulation of \(H_{i}\) (see the transition to Game 4 for more details) and outputs 1 if and only if \(\mathbb{A}_{4}\) wins the game.

Clearly, when the IND-PA0 encrypts \(\left(c_{n_{V}-i}, 1\right)\) (resp. \((c, 1)\) ), \(\mathbb{B}\) plays a perfect simulation of game \(H_{i}\) (resp. \(H_{i+1}\) ) to \(\mathbb{A}_{4}\), so that the advantage of \(\mathbb{B}\) in the IND-PA0 game is \(\left|\frac{1}{2}\left(W_{i}+1-W_{i+1}\right)-\frac{1}{2}\right| \leq \varepsilon_{\text {PAO }}\). By the triangular inequality,
\[
\left|S_{5}-S_{4}\right| \leq 2 n_{V} \varepsilon_{\mathrm{PAO}}
\]

Game 6: In this game, before computing the tally, we decrypt every valid ballot sent by the adversary at lines 23 and 26. If one of these ballots uses the same credential as a ballot sent by an honest voter (i.e. a ballot added to the board at line 25 for some \(\left(i, \nu_{i}\right)\) with \(i \in\left[1, n_{V}\right] \backslash(A \bigcup\{j\})\), we abort the game and output a random bit.

Now, we set \(\mathbb{A}_{6}=\mathbb{A}_{5}\) and, to argue that \(\left|S_{6}-S_{5}\right|\) is negligible, we remark that \(\left|S_{6}-S_{5}\right|=\varepsilon / 2\), where \(\varepsilon\) is the probability that we abort in Game 6 . Let \(E\) be the event of an abortion. We construct an adversary \(\mathbb{B}\) for IND-PA0 which wins with a non-negligible advantage whenever \(E\) occurs and wins with probability \(1 / 2\) otherwise, which shows that \(\varepsilon\) is negligible.

First, \(\mathbb{B}\) gets \(p k\) from the IND-PA0 game and forwards it to \(\mathbb{A}_{5}\). It simulate Game 5 as in the transition to Game 4. However, it chooses a random honest voter \(x\) that would send a ballot (if no honest voter votes, \(E\) cannot happen) and, for this voter, generates a fresh, second random credential \(c\). When this voter votes (which happens at most once due to the transition to Game 4\(), \mathbb{B}\) plays the pair \(\left(c_{x}, \nu\right),(c, \nu)\) in the IND-PA0 game, where \(\nu\) is the voting option chosen by \(x\). It gets back an encrypted ballot \(C\), which it uses in the simulation instead of the ballot from \(x\). Afterwards, \(\mathbb{B}\) gets the list of all valid ballots sent by \(\mathbb{A}_{5}\) in the simulation and plays them in the IND-PA0 game to get their decryption. If there is a ballot which uses the credential \(c_{x}\) (resp. \(c\) ), \(\mathbb{B}\) states that the IND-PA0 encrypted \(\left(c_{x}, \nu\right)\) (resp. \((c, \nu)\) ). If there is no such ballot, a ballot which uses \(c\) and a ballot which uses \(c_{x}\) or if no honest voter votes, \(\mathbb{B}\) guesses at random.

Now, suppose that \(b=0\) (resp. 1) in the IND-PA0 game; in other words, that \(C\) is an encryption of \(\left(c_{x}, \nu\right)\) (resp. \((c, \nu)\) ). Let \(q\) be the number of valid ballots sent by \(\mathbb{A}_{5}\). With probability \(\varepsilon, \mathbb{A}_{5}\) managed to produce a ballot which uses the same credential as a ballot sent by some honest voter. In this case, with probability at least \(1 / n_{H}\), one of the concerned honest voter is \(x\). Then, when \(\mathbb{B}\) gets the decryption from the IND-PA0 game, there is a ballot of the form \(\left(c_{x}, \gamma\right)\) (resp. ( \(\left.c, \gamma\right)\) ). In addition, \(\mathbb{A}_{5}\) has no information about \(c\) (resp. \(c_{x}\) ) so that with probability at least \(1-q / 2^{\lambda}\), there is no ballot of the form ( \(c, \gamma\) ) (resp. \(\left(c_{x}, \gamma\right)\) ). Hence \(\mathbb{B}\) wins with probability at least \(1-q / 2^{\lambda+1}\). Otherwise, no ballot uses the credential \(c_{x}\) (resp. \(c\) ) and since the
adversary has no information about \(c\) (resp. \(c_{x}\) ), the probability that a ballot uses the credential \(c\) (resp. \(c_{x}\) ) is at most \(q / 2^{\lambda}\). Therefore, the probability that \(\mathbb{B}\) wins the IND-PA0 game is at least \(\left(1-q / 2^{\lambda}\right) / 2\). Overall, \(\mathbb{B}\) 's probability to win is at least
\[
\frac{\varepsilon}{n_{H}}\left(1-q / 2^{\lambda+1}\right)+\left(1-\frac{\varepsilon}{n_{H}}\right)\left(1-q / 2^{\lambda}\right) / 2=\frac{1}{2}+\frac{\varepsilon}{2 n_{H}}-\frac{q}{2^{\lambda+1}}
\]

Therefore, we have
\[
\frac{\varepsilon}{2 n_{H}}-\frac{q}{2^{\lambda+1}} \leq \operatorname{Adv}_{\mathbb{B}}^{\mathrm{IND}-\mathrm{PAO}} \leq \varepsilon_{\mathrm{PA} 0}
\]
hence \(\left|S_{6}-S_{5}\right|=\varepsilon / 2 \leq n_{H} \varepsilon_{\text {PA0 }}+\frac{q n_{H}}{2^{\lambda+1}}\), where \(q\) is the number of valid ballots sent by the adversary.
Game 7: In this game, we remove lines 23 and 24 so that the adversary can no longer insert its own ballots between two honest ballots. In other words, the adversary must send all its ballots at the end, after every honest voter has voted. We construct \(\mathbb{A}_{7}\) which interacts with \(\mathbb{A}_{6}\) by simulating Game 6 . For this purpose, \(\mathbb{A}_{7}\) gets \(B B=\left(B_{1}, \cdots, B_{n}\right)\) at line 26 and creates a fake empty ballot box \(B B^{\prime}\). Then, in the \(k\) th iteration of the for loop, it simply appends to \(B B^{\prime}\) the valid ballots output by \(\mathbb{A}_{6}\) and then \(B_{k}\). The remaining of the simulation is similar to that of the transition to Game 4.

Clearly, \(\mathbb{A}_{7}\) plays a perfect simulation of Game 6 if the result of the tally is the same. Besides, the latter can only differ if the credential of a ballot sent by \(\mathbb{A}_{6}\) is the same as the credential of a ballot sent by some honest voter. In this case, both games abort with a random output and \(\mathbb{A}_{7}\) 's probability to win is the same as \(\mathbb{A}_{6}\) 's in Game 6 . Consequently, \(S_{7}=S_{6}\).

Game 8: In this game, the adversary has no longer access to the ballot box \(B B\) at line 26 but instead has \(|B|\) (which is equal to \(|B B|\) ). With a similar argument as in Game 5 , we construct \(\mathbb{A}_{8}\) and we have
\[
\left|S_{8}-S_{7}\right| \leq|B B| \varepsilon_{\mathrm{PAO}}
\]

Game 9: In this game, the lines 22 to 27 are replaced by lines 22 to 25 of the ideal game. We construct an adversary \(\mathbb{A}_{9}\) which interacts with \(\mathbb{A}_{8}\) by simulating Game 8 . For this purpose, \(\mathbb{A}_{9}\) creates a hashmap with keys \(\left\{c_{i} ; i \in A\right\}\) and \(\tilde{c}\) and values \(\left(\nu_{i}\right)_{i \in A}\) and \(\beta\) which are initially \(\phi\). For each ballot \(C \in \boldsymbol{M}, \mathbb{A}_{9}\) extracts \((c, \nu)\) from the ZKPoK and, if \(c\) is a key of the hashmap, it changes the corresponding value to \(\nu\). It plays the obtained values in Game 9 and receives the result of the tally which it forwards to \(\mathbb{A}_{8}\). Finally, it outputs \(\mathbb{A}_{8}\) 's output. Clearly, \(S_{9}=S_{8}\).

Game 10: The final game is the ideal game, where the credentials and the public encryption key have been removed. Note that they did not play any role in Game 9 so that the adversary might as well generate them itself. Clearly, \(S_{10}=S_{9}\).

Conclusion. With all the above transitions, we showed that for all adversary \(\mathbb{A}\) for the real game, there exists an adversary \(\mathbb{B}\) for the ideal game which wins with the same probability (up to a negligible difference. By definition, this shows that CHide is coercion-resistant.

\section*{C. Privacy}
1) Our definition of privacy: In [26], a quantitative definition of privacy is proposed. In this definition, the adversary is an observer who can corrupt some voters and talliers. The honest voters vote according to a fixed distribution s in Section III-A and a voter under observation either chooses option \(\nu_{0}\) or \(\nu_{1}\), where \(\left(\nu_{0}, \nu_{1}\right)\) can be chosen by the adversary. The latter must decide which was chosen, and do so with an advantage \(\operatorname{Adv}(\mathbb{A})\). A Voting system is then said \(\delta\)-private if, for all adversary, \(2 \operatorname{Adv}(\mathbb{A}) \leq \delta+\mu\), where \(\mu\) is some negligible function.

This definition can be turned into a qualitative one when \(\delta\) is shown to be minimal, in a sense that an ideal protocol achieves \(\delta^{\prime}\)-privacy with a negligible \(\left|\delta-\delta^{\prime}\right|\). Hence, a natural definition of privacy is to compare the probability of success of the adversary in a real process to that in an ideal one, and to show that the difference is negligible. Just as in [26], we consider a definition where the adversary tries to guess the vote of a single voter. Note that, as opposed to the coercion-resistance case, we do not assume that ballots are sent anonymously. Consequently, the adversary can always determine whether a voter abstained or voted. Finally, we consider that the registrar is honest and that up to \(t\) talliers can be corrupted by the adversary. To model this, we consider a fixed set \(\left[1, n_{C}\right]\) of valid voting options (excluding abstention) and we give Algorithms 14 and 15.

The real process can be read as follows. First, the public encryption key and the shares are generated with \(P_{\text {Setup }}\), in which the adversary can impersonate up to \(t\) corrupted authorities. Then the voters are registered and the roster \(\boldsymbol{R}\) is published. The adversary can corrupt voters and chooses a voter under observation \(j\). Then the honest voters vote according to a distribution \(\mathcal{B}\) and, as in Definition 2, we let the adversary freely insert any number of ballots between two ballots cast by an honest voter. To model the fact that the ballots are not sent anonymously, we give away the identity of the voter to the adversary at line 13 . Also, to take into account the fact that the latter could revote a certain number of times we include \(j\) as a possible identity for the honest voters at line 9 . Then the adversary must choose the voting options \(\left(\nu_{0}, \nu_{1}\right)\) and the voter under observation (re)votes with one of those two options. The adversary is given a last opportunity to cast ballots; after this, the tally is computed with \(P_{\text {tally }}\), during which the adversary can impersonate the corrupted authorities (we only consider static corruption). Finally, the adversary must guess whether \(\nu_{0}\) or \(\nu_{1}\) was chosen.
```

```
Algorithm 14: Real \({ }^{\text {Priv }}\)
```

```
Algorithm 14: Real \({ }^{\text {Priv }}\)
    Require: \(\mathbb{A}, \lambda, n_{T}, t, n_{V}, n_{A}, n_{C}, \mathcal{B}\)
    Require: \(\mathbb{A}, \lambda, n_{T}, t, n_{V}, n_{A}, n_{C}, \mathcal{B}\)
    \(B B \longleftarrow \varnothing\)
    \(B B \longleftarrow \varnothing\)
    pk, \(\left(s_{i}, h_{i}\right)_{i=1}^{n_{T}} \longleftarrow P_{\text {Setup }}^{\mathbb{A}}\left(\lambda, n_{T}, t\right)\)
    pk, \(\left(s_{i}, h_{i}\right)_{i=1}^{n_{T}} \longleftarrow P_{\text {Setup }}^{\mathbb{A}}\left(\lambda, n_{T}, t\right)\)
    \(\left\{c_{i} ; i \in\left[1, n_{V}\right]\right\}, \boldsymbol{R} \longleftarrow\) register \(\left(\lambda, \mathrm{pk}, n_{V}\right)\)
    \(\left\{c_{i} ; i \in\left[1, n_{V}\right]\right\}, \boldsymbol{R} \longleftarrow\) register \(\left(\lambda, \mathrm{pk}, n_{V}\right)\)
    \(A \longleftarrow \mathbb{A}(\boldsymbol{R})\)
    \(A \longleftarrow \mathbb{A}(\boldsymbol{R})\)
    \(j, \nu_{0}, \nu_{1} \longleftarrow \mathbb{A}\left(\left\{c_{i} ; i \in A\right\}\right)\)
    \(j, \nu_{0}, \nu_{1} \longleftarrow \mathbb{A}\left(\left\{c_{i} ; i \in A\right\}\right)\)
    6 (* \(\mathbb{A}\) chooses the voter under observation \(*)\)
    6 (* \(\mathbb{A}\) chooses the voter under observation \(*)\)
    if \(|A| \neq n_{A} \vee j \notin\left[1, n_{V}\right] \backslash A\) then
    if \(|A| \neq n_{A} \vee j \notin\left[1, n_{V}\right] \backslash A\) then
        Return 0
        Return 0
    \(B \longleftarrow \mathcal{B}\left(\left[1, n_{V}\right] \backslash A, n_{C}\right)\)
    \(B \longleftarrow \mathcal{B}\left(\left[1, n_{V}\right] \backslash A, n_{C}\right)\)
    for \(\left(i, \nu_{i}\right) \in B\) do
    for \(\left(i, \nu_{i}\right) \in B\) do
        \(\boldsymbol{M} \longleftarrow \mathbb{A}(B B)\)
        \(\boldsymbol{M} \longleftarrow \mathbb{A}(B B)\)
        \(B B \leftarrow B B \bigcup\{m \in M \mid\) isVal \((\mathrm{pk}, m, B B)=1\}\)
        \(B B \leftarrow B B \bigcup\{m \in M \mid\) isVal \((\mathrm{pk}, m, B B)=1\}\)
        \(B B \longleftarrow B B \bigcup\left\{\operatorname{Vote}_{\mathrm{pk}}\left(c_{i}, \nu_{i}\right)\right\} ; \mathbb{A}(i)\)
        \(B B \longleftarrow B B \bigcup\left\{\operatorname{Vote}_{\mathrm{pk}}\left(c_{i}, \nu_{i}\right)\right\} ; \mathbb{A}(i)\)
    \(\boldsymbol{M} \longleftarrow \mathbb{A}(B B\), "last honest ballot sent")
    \(\boldsymbol{M} \longleftarrow \mathbb{A}(B B\), "last honest ballot sent")
    \(B B \longleftarrow B B \bigcup\{m \in \boldsymbol{M} \mid\) isVal \((\mathrm{pk}, m, B B)=1\}\)
    \(B B \longleftarrow B B \bigcup\{m \in \boldsymbol{M} \mid\) isVal \((\mathrm{pk}, m, B B)=1\}\)
    \(b \stackrel{\$}{\leftarrow}\{0,1\}\)
    \(b \stackrel{\$}{\leftarrow}\{0,1\}\)
    \(B B \longleftarrow B B \bigcup\left\{\operatorname{Vote}_{\mathrm{pk}}\left(c_{j}, \nu_{b}\right)\right\}\)
    \(B B \longleftarrow B B \bigcup\left\{\operatorname{Vote}_{\mathrm{pk}}\left(c_{j}, \nu_{b}\right)\right\}\)
    \(\boldsymbol{M} \longleftarrow \mathbb{A}(B B)\)
    \(\boldsymbol{M} \longleftarrow \mathbb{A}(B B)\)
    \(B B \longleftarrow B B \bigcup\{m \in \boldsymbol{M} \mid \operatorname{isVal}(\mathrm{pk}, m, B B)=1\}\)
    \(B B \longleftarrow B B \bigcup\{m \in \boldsymbol{M} \mid \operatorname{isVal}(\mathrm{pk}, m, B B)=1\}\)
    \(\boldsymbol{X}, \Pi \longleftarrow P_{\text {tally }}^{\mathbb{A}}\left(B B, \boldsymbol{R}, \mathrm{pk},\left\{h_{i}, s_{i}\right\}, t\right)\)
    \(\boldsymbol{X}, \Pi \longleftarrow P_{\text {tally }}^{\mathbb{A}}\left(B B, \boldsymbol{R}, \mathrm{pk},\left\{h_{i}, s_{i}\right\}, t\right)\)
    \(b^{\prime} \longleftarrow \mathbb{A}()\)
    \(b^{\prime} \longleftarrow \mathbb{A}()\)
    2 Return 1 if \(\nu_{0}, \nu_{1} \in\left[1, n_{C}\right] \wedge b^{\prime}==b\) else 0
```

```
    2 Return 1 if \(\nu_{0}, \nu_{1} \in\left[1, n_{C}\right] \wedge b^{\prime}==b\) else 0
```

```
```

Algorithm 15: Ideal $^{\text {Priv }}$
Require: $\mathbb{A}, \lambda, n_{T}, t, n_{V}, n_{A}, n_{C}, \mathcal{B}$
1
2
$A \longleftarrow \mathbb{A}(\lambda)$
$j, \nu_{0}, \nu_{1} \longleftarrow \mathbb{A}()$
6
if $|A| \neq n_{A} \vee j \notin\left[1, n_{V}\right] \backslash A$ then
Return 0
$B \longleftarrow \mathcal{B}\left(\left[1, n_{H}\right] \backslash(A \cup\{j\}), n_{C}\right)$
$I \leftarrow(i)_{(i, \nu) \in B}$
11
12
13
$14\left(\alpha_{i}\right)_{i \in A} \longleftarrow \mathbb{A}(I)$
$B \longleftarrow B \bigcup\left\{\left(i, \alpha_{i}\right) \mid i \in A, \alpha_{i} \in\left[1, n_{C}\right]\right\}$
$b \stackrel{\$}{\leftarrow}\{0,1\}$
$B \longleftarrow B \bigcup\left\{\left(j, \nu_{b}\right)\right\}$
8

19
$20 r$ tally (cleanse $(B))$
$b^{\prime} \longleftarrow \mathbb{A}(r)$
Return 1 if $\nu_{0}, \nu_{1} \in\left[1, n_{C}\right] \wedge b^{\prime}==b$ else 0

```

Fig. 9: Definition of privacy, \(\lambda\) is the security parameter, \(n_{T}\) the number of talliers, \(t\) the threshold, \(n_{V}\) the number of voters, \(n_{A}\) the number of corrupted voters, \(n_{C}\) the number of voting options (excluding abstention) and \(\mathcal{B}\) the distribution.

The ideal process is simpler. The adversary still corrupts \(n_{A}\) voters and chooses \(\left(\nu_{0}, \nu_{1}\right)\), but must guess which option was chosen given only the result \(r\) of the election, and the list \(I\) that indicates which voter (re)votes at which moment. The latter is given in the ideal game to model that the ballots are not sent anonymously.
Definition 4 (Vote privacy). We say that a voting protocol ( \(P_{\text {Setup }}\), register, Vote, Check, isVal,Fakecred, \(P_{\text {tally }}\), Verify) guarantees vote privacy w.r.t. a result function tally if, for all parameters \(n_{T}, t, n_{V}, n_{A}, n_{C}\) with \(t<a\) and \(n_{A} \leq n_{V}\), for all distribution \(\mathcal{B}\) and for all PPT adversary \(\mathbb{A}\), there exists a PPT adversary \(\mathbb{B}\) and a negligible function \(\mu\) such that
\[
\mid \operatorname{Pr}\left(\text { Real }^{\text {Priv }}\left(\mathbb{A}, \lambda, n_{T}, t, n_{V}, n_{A}, n_{C}, \mathcal{B}\right)=1\right)-\operatorname{Pr}\left(\text { Ideal }^{\operatorname{Priv}}\left(\mathbb{B}, \lambda, n_{T}, t, n_{V}, n_{A}, n_{C}, \mathcal{B}\right)=1\right) \mid \leq \mu(\lambda)
\]
2) Proof that CHide is private: We now prove that CHide is private, as of Definition 4 . The proof is extremely similar to that of coercion-resistance in Section C-B and uses roughly the same transitions. We reproduce the transitions, with a few little tweaks to take into account the differences between Real \({ }^{C R}\) and Real \({ }^{\text {Priv }}\). Just as for coercion-resistance, we make several assumptions; the main difference is that we no longer assume anonymous channels for voting.

Theorem 4. Under the DDH assumption and in the ROM, CHide has privacy as of Definition 4.
Proof. We give a succession of games such that Game 0 is the real game and Game 9 is the ideal game. We consider a PPT \(\mathbb{A}_{0}\) for Game 0 . For Game \(i\), we construct a PPT adversary \(\mathbb{A}_{i}\) for this game and we denote \(S_{i}\) the probability that \(\mathbb{A}_{i}\) wins this game. (To ease the notations, we drop the dependency in \(\lambda\) when the context is clear.) For all \(i\), we show that \(\left|S_{i+1}-S_{i}\right|\) is negligible, which proves that \(\left|S_{0}-S_{9}\right|\) is also negligible.

Game 1: In this game, the adversary no longer takes part into the tally process at line 28. Instead, it is given the output of \(T_{\text {tally }}\).

Game 2: In this game, we remove line 2. Instead, we generate a perfectly random public encryption key pk and give it to the adversary at line 4 .

Game 3: In this game, the adversary is no longer given the output of the conditional gates, that we denote \(\Pi^{Z}\).
Game 4: In this game, we modify the sequence \(B\) so that the honest voters no longer revote. Instead, for all honest voter \(x\), we replace all but the last occurrence of the form \((x, \nu)\) in \(B\) by an occurrence of the form \((x \rightarrow \tilde{x}, \nu)\) where \(\tilde{x}\) is a fresh
and unique negative number. Then, at line 13 , when \(i\) is of the form \(x \rightarrow \tilde{x}\), we give \(x\) to the adversary but we add a ballot of the form \(\operatorname{Vot} \mathrm{e}_{\mathrm{pk}}(c, \nu)\) with a fresh random (fake) credential \(c\).

Although this transition is slightly different from the transition to Game 4 in the proof of coercion resistance, the reduction argument is the same.

Game 5: In this game, the adversary no longer has access to the roster \(\boldsymbol{R}\) at line 4.
Game 6: In this game, before computing the tally, we decrypt every valid ballot sent by the adversary at lines 11,14 and 18. If one of these ballots uses the same credential as a ballot sent by an honest voter (i.e. a ballot added to the board at line 13 for some \(\left(i, \nu_{i}\right)\) with \(i \in\left[1, n_{V}\right] \backslash A\), we abort the game and output a random bit.

Game 7: In this game, we remove lines \(11,12,14\) and 15 so that the adversary can no longer insert its own ballots between two honest ballots (including those sent by the voter under observation). In other words, the adversary must send all its ballots at the end, after every honest voter has voted.

Game 8: In this game, the adversary has no longer access to the ballot box \(B B\) at lines 18 and 5 but instead is given \(I\).
Compared to the transition to Game 8 in the proof of coercion resistance, the adversary know has access to \(I\) instead of \(|B|\). This extra information is used to simulate line 13 .

Game 9: The final game is the ideal game.
We construct an adversary \(\mathbb{A}_{9}\) which interacts with \(\mathbb{A}_{8}\) by simulating Game 8 . For this purpose, \(\mathbb{A}_{9}\) runs the setup and the registration honestly, by generating the secret key and the credentials. This allows \(\mathbb{A}_{9}\) to get \(j, \nu_{0}, \nu_{1}\) from \(\mathbb{A}_{8}\), that it plays in the ideal game. Then, when given \(I\) in the ideal game, it forwards it to \(\mathbb{A}_{8}\) which answers with \(\boldsymbol{M}\). To deduce the corresponding voting options \(\left(\nu_{i}\right)_{i \in A}, \mathbb{A}_{9}\) creates a hashmap with the keys \(\left\{c_{i} ; i \in A\right\}\), and values \(\left(\nu_{i}\right)_{i \in A}\) (initially \(\phi\), for abstention). For each valid ballot cast by \(\mathbb{A}_{8}, \mathbb{A}_{9}\) decrypts the ballot using the secret key and deduces \((\nu, c)\). Since the ballot is valid, by the soundness of the ZKP, \(c\) consists of \(\lambda\) bits and \(\nu\) is a valid voting option. If \(c\) is a key of the hashmap, \(\mathbb{A}_{9}\) changes the corresponding value to \(\nu\). (Otherwise, it ignores the ballot.) It plays the obtained values in Game 9 and receives the result of the tally which it forwards to \(\mathbb{A}_{8}\). Finally, it outputs \(\mathbb{A}_{8}\) 's output. Remark that \(\mathbb{A}_{9}\) played a perfect simulation of Game 8, so that \(S_{9}=S_{8}\).

\section*{D. Verifiability}
1) Our definition of verifiability: Verifiability is often split into several sub-properties, namely cast-as-intended, where the voters can verify that their voting device produced a ballot which indeed contains the desired voting option; recorded-as-cast, where the voters can verify that the ballot produced by their voting device is indeed added to the ballot box (typically, a public bulletin board); tally-as-recorded, where anyone can verify that the result obtained with the tally protocol indeed corresponds to that obtained by opening every (valid) ballot of the ballot box and then applying the counting function on the cleartexts; and eligibility, where anyone can verify that the ballot box indeed contains ballot that comes from an eligible voter, and that each voter can have at most one ballot counted. The first two properties form the individual verifiability while the two last form the universal verifiability.

In this paper, we focus on the universal verifiability. We do not consider cast-as-intended verification and assume that the voting device is not compromised. Achieving both cast-as-intended verification and coercion-resistance is a non-trivial task that is out of the scope of this article. In addition, we assume that each time a honest voter casts a ballot, they verify that the ballot appears on the board. Since the bulletin board is considered honest, this immediately gives individual verifiability.

To formalize the notion of universal privacy, we use the approach of [12] which gives a game-based definition. Originally, their definition mixes the universal and individual verifiability and divides the honest voters into two subsets, depending on whether they perform the verification or not. For those who do not check, they require that their ballots can be dropped but that their content cannot be modified. This, however, comes with a difficulty when revoting is allowed. Indeed, if a voter's votes are, in order, \(\left(\nu_{1}, \nu_{2}, \nu_{3}\right)\) and if the voter does not verify that the corresponding ballots appear on the board, then the adversary might drop any of those ballots and have the voter abstain, or have their vote counted as \(\nu_{1}, \nu_{2}\) or \(\nu_{3}\) as it sees fit. A stronger and more desirable definition would require the adversary to either have the voter abstain or have their last choice \(\nu_{3}\) be counted. This, however, requires specific measures which are yet to appear in the literature. In this paper, we consider that those considerations are out of scope and thus consider a weaker definition where every voter systematically verifies that their ballot appears on the board.

This gives Definition 5, where the adversary first takes part in the setup and corrupts a subset \(A\) of the voters. Then, given the roster and the credentials of the corrupted voters, the adversary must generate a valid bulletin board. For this purpose, it has access to the oracle \(\mathcal{O}_{\text {cast }}\) which takes as input a ballot \(\mathbf{b}\) (see Algorithm 17). If the ballot is valid w.r.t. \(B B, \mathcal{O}_{\text {cast }}\) adds \(\mathbf{b}\) to \(B B\). The adversary also has access to \(\mathcal{O}_{\text {vote }}\), which models the ballots sent by the honest voters (see Algorithm 16). However, \(\mathcal{O}_{\text {vote }}\) modifies two inner tables H and HV . The first one model the fact that a honest voter is supposed to check that their ballot appears on the board. If a honest voter vote but does not check or if their check fails, they become unhappy. The second tabular represent the state of a voter, and contains their ballot, their credential and their chosen voting option. The adversary
```

Algorithm 19: Ver
Require: $\mathbb{A}, \lambda, n_{T}, t, n_{V}, n_{A}, n_{C}$
$1 \mathrm{pk},\left(s_{i}, h_{i}\right)_{i=1}^{n_{T}} \longleftarrow P_{\text {Setup }}^{\mathbb{A}}\left(\lambda, n_{T}, t\right)$
$\left\{c_{i} ; i \in\left[1, n_{V}\right]\right\}, \boldsymbol{R} \longleftarrow$ register $\left(\lambda, \mathrm{pk}, n_{V}\right)$
$3 \longleftarrow \mathbb{A}(\boldsymbol{R})$
$B B \longleftarrow \varnothing$; for $i \in\left[1, n_{V}\right] \backslash A$ do $\mathrm{HV}_{i} \longleftarrow \perp ; \mathrm{H}_{i} \longleftarrow 1$
$\mathbb{A}^{\mathcal{O}_{\text {vote }}}, \mathcal{O}_{\text {cast }}, \mathcal{O}_{\text {check }}\left(\left\{c_{i} ; i \in A\right\}\right)$
6 if $|A| \neq n_{A} \vee \exists i \in\left[1, n_{V}\right] \backslash A \mid \mathrm{H}_{i}<1$ then Return 0
$7 \boldsymbol{X}, \Pi \longleftarrow P_{\text {tally }}^{\mathbb{A}}\left(B B, \boldsymbol{R}, \mathrm{pk},\left\{h_{i}, s_{i}\right\}, t\right)$
if Verify $(B B, \Pi, \boldsymbol{X})=0$ then Return 0
9 if $\exists\left(\operatorname{cor}_{i}, \alpha_{i}\right)_{i=1}^{n} \in\left(A \times\left[1, n_{C}\right]\right) \mid$
$\boldsymbol{X}=$ tally $\left(\operatorname{cleanse}\left(\left\{\left(i, \mathrm{HV}_{i}[4]\right) \mid i \in\left[1, n_{V}\right] \backslash A, \mathrm{HV}_{i} \neq \perp\right\} \bigcup\left\{\left(\operatorname{cor}_{i}, \alpha_{i}\right) \mid i \in[1, n]\right\}\right)\right)$ then Return 0 else
Return 1

```

Fig. 10: Definition of verifiability, \(\lambda\) is the security parameter, \(n_{T}\) the number of talliers, \(t\) the threshold, \(n_{V}\) the number of voters, \(n_{A}\) the number of corrupted voters and \(n_{C}\) the number of voting options (excluding abstention)
can have a voter \(i\) check their vote by calling \(\mathcal{O}_{\text {check }}(i)\) (see Algorithm 18) which calls the Check procedure with the current state \(\mathrm{HV}_{i}\) of the voter and the current bulletin board. The adversary can call \(\mathcal{O}_{\text {cast }}, \mathcal{O}_{\text {vote }}\) and \(\mathcal{O}_{\text {check }}\) any (polynomial) number of times, in any order. It then takes part into the tally to produce the result \(\boldsymbol{X}\) as well as the transcript \(\Pi\). The goal of the adversary is to produce a valid \(\boldsymbol{X}, \Pi\) which is different from any result obtained with the sequence \(\left(i, \mathrm{HV}_{i}\right)_{i}\) concatenated with another sequence \(\left(\operatorname{cor}_{i}, \alpha_{i}\right)_{i}\), where the \(\operatorname{cor}_{i}\) 's are corrupted voters. In addition, all honest voter who have voted must be happy.

We acknowledge that, in our definition, we assume that the registrar is honest, and that up to \(t\) decryption trustees can be corrupted. Usually, one would want verifiability even if all the talliers and the registrar are corrupted. In [12], for instance, it is assumed that either the registrar or the bulletin board is honest. However, in JCJ-like voting systems, the registrar knows the credential of the voter and can therefore break eligibility by voting instead of the abstaining voters (ballot stuffing), or even change the choice of any voter by revoting with their credential afterwards. Similarly, since an encryption of the credentials is published in the roster \(\boldsymbol{R}\), if more than a threshold \(t\) of talliers is corrupted, they can decrypt the credentials and perform the same attacks. Therefore, it is necessary to assume that the registrar is honest and that up to a threshold of talliers can be corrupted. Those shortcomings were already present in JCJ.

Definition 5 (Verifiability). We say that the a voting protocol ( \(P_{\text {Setup }}\), register,Vote, Check, isVal,Fakecred, \(P_{\text {tally }}\), Verify ) is verifiable if, for all adversary \(\mathbb{A}\), for all parameters \(\left(n_{T}, t, n_{V}, n_{C}\right)\), there exists a negligible function \(\mu\) such that \(\operatorname{Pr}\left(\operatorname{Ver}\left(\mathbb{A}, \lambda, n_{T}, t, n_{V}, n_{A}, n_{C}\right)=1\right) \leq \mu(\lambda)\).
```

Algorithm 16: $\mathcal{O}_{\text {vote }}$ (updates the tabulars HV and H)
Require: $i, \nu\left(n_{V}\right.$ is the number of voters, $A$ the
corrupted voters, $n_{C}$ the number of voting
options, pk the public key and $c_{i}$ the
credential of voter $i$ )
if $i \in\left[1, n_{V}\right] \backslash A \wedge \nu \in\left[1, n_{C}\right]$ then
$\mathrm{H}_{i} \longleftarrow \mathrm{H}_{i}-1 ;$
$\mathbf{b} \longleftarrow \operatorname{Vote}_{\mathrm{pk}}\left(c_{i}, \nu\right)$;
$\mathrm{HV}_{i} \longleftarrow\left(\mathrm{pk}, \mathbf{b}, c_{i}, \nu\right) ;$
Return b
Algorithm 17: $\mathcal{O}_{\text {cast }}$ (updates the bulletin board $B B$ )
Require: b
$\mathbf{1}$ if isVal $(\mathrm{pk}, B B, \mathbf{b})=1$ then $B B \longleftarrow B B \bigcup\{\mathbf{b}\}$

```
Algorithm 18: \(\mathcal{O}_{\text {check }}\) (updates HV and H )
```

Algorithm 18: $\mathcal{O}_{\text {check }}$ (updates HV and H )

```
Algorithm 18: \(\mathcal{O}_{\text {check }}\) (updates HV and H )
    Require: \(i\)
    Require: \(i\)
    Require: \(i\)
    \(\mathrm{pk}, \mathbf{b}, c_{i}, \nu \longleftarrow \mathrm{HV}_{i} ; \mathrm{HV}_{i}[2] \longleftarrow \perp\)
    \(\mathrm{pk}, \mathbf{b}, c_{i}, \nu \longleftarrow \mathrm{HV}_{i} ; \mathrm{HV}_{i}[2] \longleftarrow \perp\)
    \(\mathrm{pk}, \mathbf{b}, c_{i}, \nu \longleftarrow \mathrm{HV}_{i} ; \mathrm{HV}_{i}[2] \longleftarrow \perp\)
    if \(\operatorname{Check}\left(B B, \mathrm{pk}, \mathbf{b}, c_{i}, \nu\right)=1\) then \(\mathrm{H}_{i} \longleftarrow \mathrm{H}_{i}+1\)
```

```
    if \(\operatorname{Check}\left(B B, \mathrm{pk}, \mathbf{b}, c_{i}, \nu\right)=1\) then \(\mathrm{H}_{i} \longleftarrow \mathrm{H}_{i}+1\)
```

```
    if \(\operatorname{Check}\left(B B, \mathrm{pk}, \mathbf{b}, c_{i}, \nu\right)=1\) then \(\mathrm{H}_{i} \longleftarrow \mathrm{H}_{i}+1\)
```

```
2) Proof that CHide is verifiable: We now prove that CHide is verifiable, as of Definition 5. Just as for coercion-resistance, we make several assumptions; the main difference is that the adversary has a complete control over the vote of the honest voters.

\section*{Theorem 5. Under the DDH assumption and in the ROM, CHide is verifiable as of Definition 5.}

Proof. Despite the difference between coercion-resistance and verifiability, there are arguments in common in the two proofs and the first transitions are the same.

Game 1: In this game, the adversary no longer takes part into the tally process at line 28. Instead, it is given the output of \(T_{\text {tally }}\).

Game 2: In this game, we remove line 2 . Instead, we generate a perfectly random public encryption key pk and give it to the adversary at line 4 .

Game 3: In this game, the adversary is no longer given the output of the conditional gates, that we denote \(\Pi^{Z}\).
Game 4: In this game, we create a secret internal state \(\mathrm{Cred}_{i}\) for each voter. It contains a sequence of credentials which is initially \(\left[c_{i}\right]\), where \(c_{i}\) is \(i\) 's credential. In addition, we also modify \(\mathcal{O}_{\text {vote }}\) so that it uses the last element of Cred \({ }_{i}\) instead of \(c_{i}\). It also generates a fresh random credential \(c\) which is added at the end of \(\mathrm{Cred}_{i}\), so that the next instance of \(\mathcal{O}_{\text {vote }}\) with the same voter will use this new credential instead of the old one. (In case if \(c\) is a collision with another voter, we pick another random credential.) Finally, we change the way the tally \(\boldsymbol{X}\) is computed. Now, we decrypt each ballot of \(B B\) into a couple \((c, \nu)\). For each ballot, we look for a voter \(i\) such that \(c\) appears in the sequence \(\mathrm{Cred}_{i}\) and we set \(i\) 's vote as \(\nu\) (since there is no collision, there may be up to one such voter). Finally, we define \(\boldsymbol{X}\) as the voters' votes.

We take \(\mathbb{A}_{4}=\mathbb{A}_{3}\) and we show that \(\left|S_{4}-S_{3}\right|\) is negligible. For this purpose, for all honest voter \(j\), we consider another succession of games \(H_{0}^{j}, \cdots, H_{n_{j}}^{j}\), where \(n_{j}\) is the number of times that \(\mathbb{A}_{3}\) calls \(\mathcal{O}_{\text {vote }}\) with the voter \(j\) and a valid \(\nu\). (For the sake of notations, we set \(n_{j}=0\) when \(j\) is a corrupted voter.) Game \(H_{k}^{j}\) is like Game 4 , except that the first \(n_{j}-k\) times \(\mathcal{O}_{\text {vote }}\) is called with voter \(j\), we do not add the new credential \(c\) in Cred \({ }_{j}\), so that the next instance of \(\mathcal{O}_{\text {vote }}\) with voter \(j\) will still use the credential \(c_{j}\). As for all voter \(i>j\), we never update their \(\mathrm{Cred}_{i}\), so that they are treated as in Game 3 . This way, \(H_{0}^{0}\) is Game \(3, H_{n_{j}}^{j}\) is \(H_{0}^{j+1}\) for all \(j\) and \(H_{n_{n_{V}}}^{n_{V}}\) is Game 4. For each of these games, we denote \(W_{k}^{j}\) the probability that \(\mathbb{A}_{3}\) wins the game.

Now, let \(k\) be some index and \(j\) be some voter. We construct \(\mathbb{B}\) for IND-PA0 as follows. \(\mathbb{B}\) gets pk from the IND-PA0 game and uses this value to play a simulation of \(H_{k}^{j}\) to \(\mathbb{A}_{3}\). For this purpose, it creates and updates \(B B, \mathrm{HV}, \mathrm{H}\) and Cred, and simulates the oracles. It also simulates the register procedure by generating a random credential for each voter and computing \(\boldsymbol{R}\) by encrypting the bits of those credentials, using pk. However, the \(\left(n_{j}-k-1\right)\) th time \(\mathcal{O}_{\text {vote }}\) is called with voter \(j, \mathbb{B}\) generates a new \(c\) in \(\mathrm{Cred}_{j}\) as in game \(H_{k+1}^{j}\). In addition, the next time \(\mathcal{O}_{\text {vote }}\) is called with voter \(j\) (let \(\nu\) be the corresponding voting option), instead of using the last credential of \(\mathrm{Cred}_{j}\) (which is \(c\) ), \(\mathbb{B}\) plays \(\left(c_{i}, \nu\right),(c, \nu)\) in the IND-PA0 game and uses the result as the output \(\mathbf{b}\). At some point, \(\mathbb{A}_{3}\) terminates and \(\mathbb{B}\) must decide whether the IND-PA0 game encrypted \(\left(c_{i}, \nu\right)\) or \((c, \nu)\). For this purpose, \(\mathbb{B}\) uses \(B B\), the list of the ballots cast by \(\mathbb{B}\). If \(\mathbf{b} \in B B\), \(\mathbb{B}\) removes it from \(B B\) but keeps in memory its index. For the remaining ballots, \(\mathbb{B}\) decrypts them using the decryption oracle of the IND-PA0 game. With all these informations, \(\mathbb{B}\) can compute the tally as in Game 4. If the condition of line 9 is not verified (i.e. \(\mathbb{A}_{3}\) won the simulation), \(\mathbb{B}\) states that \(\left(c_{i}, \nu\right)\) was encrypted; otherwise, it states that ( \(c, \nu\) was encrypted.

Clearly, when \(\left(c_{i}, \nu\right)\) is encrypted, \(\mathbb{B}\) plays a perfect simulation of \(H_{k}^{j}\) while, when \((c, \nu)\) is encrypted, \(\mathbb{B}\) plays a perfect simulation of \(H_{k+1}^{j}\). Therefore, \(\mathbb{B}\) 's advantage in the IND-PA0 game is \(\left|1 / 2\left(W_{k}^{j}+1-W_{k+1}^{j}\right)-1 / 2\right| \leq \varepsilon_{\text {PA0 }}\). By triangular inequality, \(\left|S_{4}-S_{3}\right| \leq \sum_{j=1}^{n_{V}} 2 n_{j} \varepsilon_{\text {PA0 }}=2 q \varepsilon_{\text {PAO }}\), where \(q\) is the number of calls to \(\mathcal{O}_{\text {vote }}\).

Game 5: In this game, the adversary no longer has access to the roster \(\boldsymbol{R}\) at line 3 . We construct \(\mathbb{A}_{5}\) which interacts with \(\mathbb{A}_{4}\) by simulating Game 4 . For this purpose, it generates \(n_{V} \lambda\) ElGamal encryptions of random bits and uses it to simulate \(\boldsymbol{R}\). For the remaining of the game, it can play a perfect simulation since both games are identical. To argue that \(\left|S_{5}-S_{4}\right|\) is negligible, we construct a succession of hop \(H_{0}, \cdots, H_{n_{V}}\) such that in game \(H_{i}\), the last \(i\) elements of \(\boldsymbol{R}\) are replaced by a random encryption. This way, \(H_{0}\) is Game 4 and \(H_{n_{V}}\) is Game 5. For each of these games, we denote \(W_{i}\) the probability that \(\mathbb{A}_{4}\) wins this game.

Now, let \(i\) be some index. We construct an adversary \(\mathbb{B}\) for IND-PA0 as follows. \(\mathbb{B}\) gets pk from the IND-PA0 game and plays this pk to \(\mathbb{A}_{4}\) to simulate \(H_{i}\). It generates the credentials and replaces the last \(i\) elements of the roster by encryptions of random bits just as in \(H_{i}\). However, for the \(\left(n_{V}-i\right)\) th element of \(\boldsymbol{R}\), it generates a second random credential \(c\) and plays the pair \(\left(c_{n_{V}-i}, 1\right),(c, 1)\) in the IND-PA0 game which answers with some encryption \(C\). \(\mathbb{B}\) retrieves an ElGamal bitwise encryption of the credential from \(C\) and uses it as the \(\left(n_{V}-i\right)\) th element of \(\boldsymbol{R}\) instead of an honest bitwise encryption of \(c_{n_{V}-i}\). Afterwards, \(\mathbb{B}\) continues the simulation of \(H_{i}\) (see the transition to Game 4 for more details) and outputs 1 if and only if \(\mathbb{A}_{4}\) wins the game.

Clearly, when the IND-PA0 encrypts \(\left(c_{n_{V}-i}, 1\right)\) (resp. \((c, 1)\) ), \(\mathbb{B}\) plays a perfect simulation of game \(H_{i}\) (resp. \(\left.H_{i+1}\right)\) to \(\mathbb{A}_{4}\), so that the advantage of \(\mathbb{B}\) in the IND-PA0 game is \(\left|\frac{1}{2}\left(W_{i}+1-W_{i+1}\right)-\frac{1}{2}\right| \leq \varepsilon_{\text {PA0 }}\). By the triangular inequality,
\[
\left|S_{5}-S_{4}\right| \leq 2 n_{V} \varepsilon_{\mathrm{PAO}}
\]

Game 6: In this game, before computing the result, we decrypt every valid ballot of \(B B\) which is not the output of some \(\mathcal{O}_{\text {vote }}\). If one of these ballots uses the same credential as a ballot output by \(\mathcal{O}_{\text {vote }}\), we abort the game and output a random bit.

Now, we set \(\mathbb{A}_{6}=\mathbb{A}_{5}\) and, to argue that \(\left|S_{6}-S_{5}\right|\) is negligible, we remark that \(\left|S_{6}-S_{5}\right|=\varepsilon / 2\), where \(\varepsilon\) is the probability that we abort in Game 6 . Let \(E\) be the event of an abortion. We construct an adversary \(\mathbb{B}\) for IND-PA0 which wins with a non-negligible advantage whenever \(E\) occurs and wins with probability \(1 / 2\) otherwise, which shows that \(\varepsilon\) is negligible.

First, \(\mathbb{B}\) gets pk from the IND-PA0 game and forwards it to \(\mathbb{A}_{5}\). It simulate Game 5 as in the transition to Game 4. However, it chooses a random instance of \(\mathcal{O}_{\text {vote }}\) (let \((x, \nu)\) be its inputs) and, for this instance, generates a fresh, second random credential
\(\tilde{c} . \mathbb{B}\) plays the pair \((c, \nu),(\tilde{c}, \nu)\) in the IND-PA0 game, where \(c\) is the last credential in \(C_{r e d}\). It gets back an encrypted ballot \(\mathbf{b}\), which it uses in the simulation as the output of \(\mathcal{O}_{\text {vote }}\). Afterwards, \(\mathbb{B}\) removes from \(B B\) any ballot output by \(\mathcal{O}_{\text {vote }}\) (including \(\mathbf{b}\) ) and plays \(B B\) in the IND-PA0 game to get the decryption of the remaining ballots. If one uses the credential \(c\) (resp. \(\tilde{c}\) ), \(\mathbb{B}\) states that the \(\operatorname{IND}-\mathrm{PA} 0\) encrypted \((c, \nu)\) (resp. \((\tilde{c}, \nu)\) ). If there is no such ballot or if there is a ballot which uses \(c\) and a ballot which uses \(\tilde{c}, \mathbb{B}\) guesses at random.

Now, suppose that \(b=0\) (resp. 1) in the IND-PA0 game; in other words, that \(\mathbf{b}\) is an encryption of \((c, \nu)\) (resp. \((\tilde{c}, \nu)\) ). Let \(q_{v}\) be the number of (valid) calls to \(\mathcal{O}_{\text {vote }}\) and \(q_{c}\) be the number of ballots from \(B B\) which are not an output of \(\mathcal{O}_{\text {vote }}\). With probability \(\varepsilon, \mathbb{A}_{5}\) managed to produce a ballot which uses the same credential as a ballot output by \(\mathcal{O}_{\text {vote. }}\). In this case, with probability at least \(1 / q_{v}\), one of the concerned ballot is \(\mathbf{b}\). Then, when \(\mathbb{B}\) gets the decryption from the IND-PA0 game, there is a ballot of the form \((c, \gamma)\) (resp. \((\tilde{c}, \gamma)\) ). In addition, \(\mathbb{A}_{5}\) has no information about \(\tilde{c}\) (resp. \(c\) ) so that with probability at least \(1-q_{c} / 2^{\lambda}\), there is no ballot of the form \((\tilde{c}, \gamma)\) (resp. \((c, \gamma)\) ). Hence \(\mathbb{B}\) wins with probability at least \(1-q_{c} / 2^{\lambda+1}\). Otherwise, no ballot uses the credential \(c\) (resp. \(\tilde{c}\) ) and since the adversary has no information about \(\tilde{c}\) (resp. \(c\) ), the probability that a ballot uses the credential \(\tilde{c}\) (resp. \(c\) ) is at most \(q_{c} / 2^{\lambda}\). Therefore, the probability that \(\mathbb{B}\) wins the IND-PA0 game is at least \(\left(1-q_{c} / 2^{\lambda}\right) / 2\). Overall, \(\mathbb{B}\) 's probability to win is at least
\[
\frac{\varepsilon}{q_{v}}\left(1-q_{c} / 2^{\lambda+1}\right)+\left(1-\frac{\varepsilon}{q_{v}}\right)\left(1-q_{c} / 2^{\lambda}\right) / 2=\frac{1}{2}+\frac{\varepsilon}{2 q_{v}}-\frac{q_{c}}{2^{\lambda+1}} .
\]

Therefore, we have
\[
\frac{\varepsilon}{2 q_{v}}-\frac{q_{c}}{2^{\lambda+1}} \leq \operatorname{Adv}_{\mathbb{B}}^{\mathrm{IND}-\mathrm{PAO}} \leq \varepsilon_{\mathrm{PAO}}
\]
hence \(\left|S_{6}-S_{5}\right|=\varepsilon / 2 \leq q_{v} \varepsilon_{\text {PAO }}+\frac{q_{c} q_{v}}{2^{\lambda+1}}\), where \(q_{v}\) is the number of calls to \(\mathcal{O}_{\text {vote }}\) and \(q_{c}\) the number of ballots in \(B B\) which are not an output of \(\mathcal{O}_{\text {vote }}\).

Conclusion. Now, remark that due to the nature of \(\mathcal{O}_{\text {vote }}\) which overwrite HV , it is clear that the adversary must call \(\mathcal{O}_{\text {check }}(i)\) between each call of \(\mathcal{O}_{\text {vote }}\) with the same voter \(i\), otherwise \(\mathrm{H}_{i}<1\) at line 6 . Since \(\mathcal{O}_{\text {check }}\) erases \(\mathrm{HV}_{i}\), it is also necessary for the adversary to call \(\mathcal{O}_{\text {cast }}(\mathbf{b})\) before \(\mathcal{O}_{\text {check }}(i)\), where \(\mathbf{b}\) is the output of \(\mathcal{O}_{\text {vote }}(i, \nu)\). Note that due to the randomness involved in Vote, except with negligible probability \(\mu\), the adversary cannot call \(\mathcal{O}_{\text {cast }}(\mathbf{b})\) before \(\mathcal{O}_{\text {vote }}(i, \nu)\) (at which point it has no information about \(\mathbf{b}\) ). Therefore, the order of the revotes of each voter is enforced in \(B B\) and, because of the last transition and the nature of the tally, we readily have that the condition at line 9 always result in a 0 output, which concludes the proof.

\section*{Appendix D}

\section*{QUANTIFYING COERCION-RESISTANCE}

As explained in Section III, the coercion level can be evaluated thanks to Eq. (1). This assumes that the cryptography is perfect, that a large and unpredictable number of ballots is removed during the tally phase and that the adversary is able to compare \(\operatorname{Pr}\left(\mathrm{R}^{g} \mid \alpha\right)\) and \(\operatorname{Pr}\left(\mathrm{R}^{g} \mid \beta\right)\), given \(\mathrm{R}^{g}\) with \(g \in\left\{\operatorname{Real}{ }^{\mathrm{CR}}\right.\), Ideal \(\left.{ }^{\mathrm{CR}}\right\}\). In this appendix, we give more details about how this comparison can be done as well as some efficient ways to evaluate the formula from Eq. (1). In particular, this allows to understand how the figures from Section III were obtained.

\section*{A. The coercion level in the ideal game}

In the ideal game, \(\mathrm{R}^{\text {ldeal }}\) is the vector \(\overrightarrow{r e s}\), and computing \(\operatorname{Pr}(\overrightarrow{r e s} \mid \alpha)\) for any \(\alpha\) can be done thanks to a formula given in the following result from [25].
Theorem 6 ( [25]). Let \((\alpha, \beta)\) be two options, \(n_{H}\) the number of honest voters and a pure result \(\overrightarrow{r e s}\) such that \(\sum_{i=0}^{C} \operatorname{res}_{i}=\) \(n_{H}+1\). Let \(\vec{P}\) be the probability distribution for the honest voters. Assuming \(P_{\alpha} P_{\beta} \neq 0\), we have \(\operatorname{Pr}(\overrightarrow{r e s} \mid \beta) \geq \operatorname{Pr}(\overrightarrow{r e s} \mid \alpha)\) if and only if \(\operatorname{res}_{\beta} P_{\alpha} \geq \operatorname{res}_{\alpha} P_{\beta}\).

In addition, we have
\[
\operatorname{Pr}(\overrightarrow{r e s} \mid \alpha)=\frac{n_{H}!}{\prod_{k=0}^{C} \mathrm{res}_{k}!}\left(\prod_{k=0}^{C} P_{k}^{\mathrm{res}_{k}}\right) \frac{\operatorname{res}_{\alpha}}{P_{\alpha}}
\]

Finally, with \(M_{\alpha, \beta}=\{\overrightarrow{r e s} \mid \operatorname{Pr}(\overrightarrow{r e s} \mid \beta) \geq \operatorname{Pr}(\overrightarrow{r e s} \mid \alpha)\}\), the optimal value of \(\delta_{\alpha, \beta}^{\text {Ideal }}\) is
\[
\delta_{\alpha, \beta}^{\text {ldeal }^{C R}}=\sum_{\overrightarrow{r e s} \in M_{\alpha, \beta}}(\operatorname{Pr}(\overrightarrow{r e s} \mid \beta)-\operatorname{Pr}(\overrightarrow{r e s} \mid \alpha))=\sum_{\overrightarrow{r e s} \in M_{\alpha, \beta}} \frac{n_{H}!}{\prod_{k=0}^{C} \operatorname{res}_{k}!}\left(\prod_{k=0}^{C} P_{k}^{\mathrm{res}_{k}}\right)\left(\frac{\mathrm{res}_{\beta}}{P_{\beta}}-\frac{\mathrm{res}_{\alpha}}{P_{\alpha}}\right) .
\]

Thanks to this result, it appears that the adversary can compare \(\operatorname{Pr}\left(\mathrm{R}^{\text {ldeal }} \mid \alpha\right)\) and \(\operatorname{Pr}\left(\mathrm{R}^{\text {ldeal }} \mid \beta\right)\) in \(O(1)\) floating operations, as long as it has access to \(P_{\alpha}\) and \(P_{\beta}\) (note that it does not need the whole distribution), and where the asymptotic analysis is
done for \(n_{H}\) growing to infinity and other parameters are fixed. We can further analyze this expression to give a representation which is easier to compute. For this purpose, we give Lemma 1, which allows to compute \(\delta^{\text {Ideal }}\) in \(O\left(n_{H}\right)\) floating operations.
Lemma 1. Let \((\alpha, \beta)\) be two options, \(n_{H}\) the number of honest voters and a pure result \(\overrightarrow{r e s}\) such that \(\sum_{i=0}^{C} \operatorname{res}_{i}=n_{H}+1\). Let \(\vec{P}\) be the probability distribution for the honest voters. Assuming \(P_{\alpha} P_{\beta} \neq 0\), we have
\[
\delta_{\alpha, \beta}^{\left.\mathrm{Idea}\right|^{\mathrm{RR}}}=n_{H}!\sum_{x=0}^{n_{H}} \frac{T_{x} P_{\beta}^{T_{x}-1} P_{\alpha}^{N_{x}-T_{x}}}{x!T_{x}!\left(N_{x}-T_{x}\right)!}\left(1-P_{\beta}-P_{\alpha}\right)^{x}, \text { where } N_{x}=n_{H}-x+1 \text { and } T_{x}=\left\lceil\frac{P_{\beta}}{P_{\beta}+P_{\alpha}} N_{x}\right\rceil
\]

Proof. First, we partition \(M_{\alpha, \beta}\) into \(\bigcup_{x=0}^{n_{H}} M_{\alpha, \beta}^{x}\), where \(M_{\alpha, \beta}^{x}\) is the subset of all results in \(M_{\alpha, \beta}\) where the options other than \(\alpha\) and \(\beta\) received exactly \(x\) votes. Then, we further partition \(M_{\alpha, \beta}^{x}\) into subsets where res \({ }_{\beta}=y\) (and, thus, res \({ }_{\alpha}=N_{x}-y\) ). Note that due to the condition from Theorem 6 , \(\operatorname{res}_{\beta} P_{\alpha} \geq \operatorname{res}_{\alpha} P_{\beta}\) hence \(y\) ranges from \(T_{x}\) to \(N_{x}\). To express the formula from Theorem 6 using this partition, we assume that \(\alpha\) and \(\beta\) are the last two voting options (otherwise we reorder them) and denote \(\tilde{M}^{x}=\left\{\overrightarrow{r e s} \in \mathbb{N}^{C-2} \mid \sum_{k=0}^{C-2}\right.\) res \(\left._{k}=x\right\}\). With these notations, we have
\[
\begin{aligned}
\delta_{\alpha, \beta}^{\mathrm{Idea}}{ }^{\mathrm{CR}} & =\sum_{x=0}^{n_{H}} \sum_{\underset{r e s \in \tilde{M}^{x}}{ } \sum_{y=T_{x}}^{N_{x}} \frac{n_{H}!}{\prod_{k=0}^{C-2} \operatorname{res}_{k}!}\left(\prod_{k=0}^{C-2} P_{k}^{\mathrm{res}_{k}}\right) \frac{P_{\beta}^{y} P_{\alpha}^{N_{x}-y}}{y!\left(N_{x}-y\right)!}\left(\frac{y}{P_{\beta}}-\frac{N_{x}-y}{P_{\alpha}}\right)} \\
& =\sum_{x=0}^{n_{H}} \sum_{\frac{r e s}{} \in \tilde{M}^{x}} \frac{n_{H}!}{\prod_{k=0}^{C-2} \operatorname{res}_{k}!}\left(\prod_{k=0}^{C-2} P_{k}^{\mathrm{res}_{k}}\right) \sum_{y=T_{x}}^{N_{x}} \frac{P_{\beta}^{y} P_{\alpha}^{N_{x}-y}}{y!\left(N_{x}-y\right)!}\left(\frac{y}{P_{\beta}}-\frac{N_{x}-y}{P_{\alpha}}\right)
\end{aligned}
\]

Now, we show that, for all \(0 \leq T \leq N, \sum_{y=T}^{N} \frac{P_{\beta}^{y} P_{\alpha}^{N-y}}{y!(N-y)!}\left(\frac{y}{P_{\beta}}-\frac{N-y}{P_{\alpha}}\right)=\frac{T P_{\beta}^{T-1} P_{\alpha}^{N-T}}{T!(N-T)!}\).
For this purpose, we first fix some \(N \geq 0\) and we proceed by backward iteration over \(T\). First, it is true when \(T=N\). Now, suppose that it is true for some \(0<T \leq N\); we show that it is also true for \(T-1\) :
\[
\begin{aligned}
\sum_{y=T-1}^{N} \frac{P_{\beta}^{y} P_{\alpha}^{N-y}}{y!(N-y)!}\left(\frac{y}{P_{\beta}}-\frac{N-y}{P_{\alpha}}\right) & =\frac{P_{\beta}^{T-1} P_{\alpha}^{N-T+1}}{(T-1)!(N-T+1)!}\left(\frac{T-1}{P_{\beta}}-\frac{N-T+1}{P_{\alpha}}\right)+\sum_{y=T}^{N} \frac{P_{\beta}^{y} P_{\alpha}^{N-y}}{y!(N-y)!}\left(\frac{y}{P_{\beta}}-\frac{N-y}{P_{\alpha}}\right) \\
& =\frac{P_{\beta}^{T-1} P_{\alpha}^{N-T+1}}{(T-1)!(N-T+1)!}\left(\frac{T-1}{P_{\beta}}-\frac{N-T+1}{P_{\alpha}}\right)+\frac{T P_{\beta}^{T-1} P_{\alpha}^{N-T}}{T!(N-T)!} \\
& =\frac{(T-1) P_{\beta}^{T-2} P_{\alpha}^{N-T+1}}{(T-1)!(N-T+1)!}
\end{aligned}
\]

Hence, with \(T=T_{x}\) and \(N=N_{x}\), we deduce that
\[
\begin{aligned}
& \delta_{\alpha, \beta}^{\mathrm{Idea}{ }^{\mathrm{CR}}}=\sum_{x=0}^{n_{H}} \sum_{\frac{\mathrm{res} \in \tilde{M}^{x}}{}} \frac{n_{H}!}{\prod_{k=0}^{C-2} \mathrm{res}_{k}!}\left(\prod_{k=0}^{C-2} P_{k}^{\mathrm{res}_{k}}\right) \frac{T_{x} P_{\beta}^{T_{x}-1} P_{\alpha}^{N_{x}-T_{x}}}{T_{x}!\left(N_{x}-T_{x}\right)!} \\
& =n_{H}!\sum_{x=0}^{n_{H}} \frac{T_{x} P_{\beta}^{T_{x}-1} P_{\alpha}^{N_{x}-T_{x}}}{x!T_{x}!\left(N_{x}-T_{x}\right)!}\left(1-P_{\beta}-P_{\alpha}\right)^{x} \sum_{\frac{r e s \in \tilde{M}^{x}}{}} \frac{x!}{\prod_{k=0}^{C-2} \text { res }_{k}!} \prod_{k=0}^{C-2}\left(\frac{P_{k}}{1-P_{\beta}-P_{\alpha}}\right)^{\mathrm{res}_{k}} \\
& =n_{H}!\sum_{x=0}^{n_{H}} \frac{T_{x} P_{\beta}^{T_{x}-1} P_{\alpha}^{N_{x}-T_{x}}}{x!T_{x}!\left(N_{x}-T_{x}\right)!}\left(1-P_{\beta}-P_{\alpha}\right)^{x} .
\end{aligned}
\]

Indeed, the last summation is the sum over all possibilities of a multinomial distribution, which is 1.

\section*{B. Modeling the real game}

For the real game where revoting is possible, we propose a model in which each honest voter does the following
- Abstain with probability \(P_{0}\),
- Otherwise, vote for option \(l>0\) with probability \(p_{l}\),
- Revote with probability \(r_{l}\),
- In this case, choose option \(k>0\) with probability \(r_{l, k}\).

In this model, a voter may revote at most once. Also, a voter who abstains does not "revote". If a voter who initially wanted to abstain changes their mind, it will count as voting once. This approximation is made because the adversary has access to the number of revotes, and not the number of voters who changed their mind. Also, the probability to revote depends on the


Fig. 11: Tree structure for two candidates.
initial choice. This is to capture the fact that an announcement on the press can make some voter revote if their candidate has been compromised. Similarly, the probability distribution when revoting depends on the first choice. See Fig. 11 for an illustration. In this figure, the label on the left of an edge denotes the nature of the transition and the label on the right denotes the probability of the transition. On leafs, we used labels of the form \((n, P)\) when \(n\) denotes the total number of honest voters who choose this path and \(P\) the probability to choose this path. For instance, \(P_{1,2}\) is the probability to first vote for option 1 but to finally revote for option 2 while \(P_{2, \varnothing}\) is the probability to vote for option 2 and not to revote. Since abstention plays a specific role, we denote it as choice 0 , while the other choices actually imply to send a ballot.

Now, we must decide on the information the adversary has access to. The most conservative approach is that it may have access to anything which is not a secret; here, all the probability transitions. With this assumption, if \(\overrightarrow{r e s}=\left(\operatorname{res}_{0}, \cdots\right.\), res \(\left._{C}\right)\) is the number of votes for each option and \(n_{r}\) is the number of revotes, then \(\operatorname{Pr}\left(\left(\overrightarrow{r e s}, n_{r}\right) \mid \alpha\right)\) is the probability that the honest voters votes res \({ }_{k}-\mathbb{1}_{k=\alpha}\) times for each candidate \(k\) and revote \(n_{r}\) times. Given the above tree, for any possible outcome \(\left(\overrightarrow{r e s}, n_{r}\right)\) with \(n_{H}\) honest voters, we have
\[
\begin{equation*}
\operatorname{Pr}\left(\overrightarrow{r e s}, n_{r}\right)=\sum_{H_{\overrightarrow{r e s}}^{n_{r}}} \frac{n_{H}!}{\operatorname{res}_{0}!\prod_{k=1}^{C} n_{k, \varnothing}!\prod_{l=1}^{C} n_{l, k}!} P_{0}^{\mathrm{res}_{0}} \prod_{k=1}^{C} P_{k, \varnothing}^{n_{k, \varnothing}} \prod_{l=1}^{C} P_{l, k}^{n_{l, k}}, \tag{2}
\end{equation*}
\]
where \(H_{\overrightarrow{r e s}}^{n_{r}}=\left\{\left(n_{k, \varnothing}\right),\left(n_{l, k}\right) \mid \sum_{l, k} n_{l, k}=n_{r}, \forall k>0, n_{k, \varnothing}+\sum_{k=1}^{C} n_{l, k}=\operatorname{res}_{k}\right\}\) is the set of all possible atomic outcomes compatible with \(\overrightarrow{r e s}\) and \(n_{r}\).

\section*{C. Quantifying the coercion level in some specific cases}

Because of the complexity of Eq. (2), we only consider the specific cases where there are two candidates and the possibility to abstain (i.e. \(C=2\) ). In these simpler cases, we give Lemma 2 which allows to compute the coercion level in \(O\left(n_{H}{ }^{4}\right)\) floating operations. Interestingly, \(\delta^{\text {Real }}\) does not depend on the \(P_{l, k}\) directly but only on the sums \(P_{r 1}\) and \(P_{r 2}\), which correspond to the probability to revote for 1 or 2 respectively. This is because the adversary can only observe the number of revotes and has no information about the initial intention of the voter. Hence, they cannot exploit the dependency between the final choice and the first choice. In practice, we can imagine a scenario where the probability to revote and the probability distributions when revoting and when not revoting are known to the adversary, for instance thanks to exit poll or social media, where the voters tell whether they revote and what was their final choice. In this more realistic scenario when the adversary does not know all the probabilities on the tree, the latter can still compare \(\operatorname{Pr}\left(\mathrm{R}^{\text {Real }} \mid \alpha\right)\) and \(\operatorname{Pr}\left(\mathrm{R}^{\text {Real }} \mid \beta\right)\), which takes \(O\left(n_{H}\right)\) floating operations.

Lemma 2. Let \(n_{H}\) be the number of honest voters, \(\left(\operatorname{res}_{0}, \mathrm{res}_{1}, \mathrm{res}_{2}\right)\) such that \(\mathrm{res}_{0}+\mathrm{res}_{1}+\operatorname{res}_{2}=n_{H}+1\) and \(n_{r} \leq n_{H}-\operatorname{res}_{0}\). Let \(\beta \in\{0,1,2\}\) be a voting option. Then the probability \(\operatorname{Pr}\left(\left(\mathrm{res}_{0}, \mathrm{res}_{1}, \mathrm{res}_{2}\right), n_{r} \mid \beta\right)\) that there are \(\mathrm{res}_{0}\) abstentions, res \({ }_{1}\) votes for option 1, res 2 votes for option 2 and \(n_{r}\) revotes when the coerced voter chooses the voting option \(\beta\) is given by the following algorithm.
- First, set \(\overrightarrow{r e s}=\left(\right.\) res \(_{0}\), res \(_{1}\), res \(\left._{2}\right)\).
- If \(\operatorname{res}_{\beta}=0\), return 0 , otherwise, do \(\operatorname{res}_{\beta} \leftarrow \operatorname{res}_{\beta}-1\).
- Set \(N=n_{H}-\operatorname{res}_{0}-n_{r}, P_{r 1}=P_{1,1}+P_{2,1}\) and \(P_{r 2}=P_{1,2}+P_{2,2}\). Return
\[
\operatorname{Pr}\left(\overrightarrow{r e s}, n_{r}\right)=\frac{n_{H}!P_{0}^{\mathrm{res}_{0}}}{\operatorname{res}_{0}!} \sum_{n_{1, \varnothing}} \frac{P_{1, \varnothing}^{n_{1, \varnothing}} P_{2, \varnothing}^{N-n_{1, \varnothing}}}{n_{1, \varnothing}!\left(N-n_{1, \varnothing}\right)!} \frac{P_{r 1}^{\mathrm{res}_{1}-n_{1, \varnothing}}}{\left(\mathrm{res}_{1}-n_{1, \varnothing}\right)!} \frac{P_{r 2}^{n_{r}+n_{1, \varnothing}-\mathrm{res}_{1}}}{\left(n_{r}+n_{1, \varnothing}-\mathrm{res}_{1}\right)!},
\]
where \(\max \left(0\right.\), res \(\left._{1}-n_{r}\right) \leq n_{1, \varnothing} \leq \min \left(\right.\) res \(\left._{1}, n_{H}-\operatorname{res}_{0}-n_{r}\right)\).
Proof. As shown in Section D-B, \(\operatorname{Pr}\left(\left(\operatorname{res}_{0}, \operatorname{res}_{1}, \operatorname{res}_{2}\right), n_{r} \mid \beta\right)\) is given by Eq. (2). We explicit the sum over \(H_{r e s}^{n_{r}}\). First, \(0 \leq n_{2, \varnothing} \leq\) res \(_{2}\) and \(n_{1, \varnothing}+n_{2, \varnothing}=n_{H}-\operatorname{res}_{0}-n_{r}\) and res \({ }_{0}+\) res \(_{1}+\) res \(_{2}=n_{H}\), therefore
\[
\begin{aligned}
0 & \leq n_{1, \varnothing}
\end{aligned} \leq \text { res }_{1}, ~=n_{H}-\text { res }_{0}-n_{r} .
\]

In addition, \(n_{1, \varnothing}+n_{1,1}+n_{2,1}=\operatorname{res}_{1}\) and \(0 \leq n_{2,1} \leq n_{r}\) so that
\[
\begin{aligned}
& 0 \leq n_{1,1} \leq n_{r} \\
& \operatorname{res}_{1}-n_{1, \varnothing}-n_{r} \leq n_{1,1} \leq \operatorname{res}_{1}-n_{1, \varnothing} .
\end{aligned}
\]

Finally, \(n_{1,2}+n_{2,2}=n_{r}-n_{1,1}-n_{2,1}=n_{r}+n_{1, \varnothing}-\) res \(_{1}\), thus
\[
0 \leq n_{1,2} \leq \operatorname{res}_{2}-n_{2, \varnothing}=n_{r}+n_{1, \varnothing}-\operatorname{res}_{1}
\]

Let \(N=n_{H}-\operatorname{res}_{0}-n_{r}, N_{n_{1, \varnothing}}=\operatorname{res}_{1}-n_{1, \varnothing}\) and \(N_{n_{1, \varnothing}}^{\prime}=n_{r}+n_{1, \varnothing}-\operatorname{res}_{1}\). Considering all these inequalities, we have
\[
\begin{aligned}
\operatorname{Pr}\left(\overrightarrow{r e s}, n_{r}\right) & =\frac{n_{H}!P_{0}^{\mathrm{res}_{0}}}{\operatorname{res}_{0}!} \sum_{n_{1, \varnothing}} \frac{P_{1, \varnothing}^{n_{1, \varnothing}} P_{2, \varnothing}^{N-n_{1, \varnothing}}}{n_{1, \varnothing}!\left(N-n_{1, \varnothing}\right)!} \sum_{n_{1,1}} \frac{P_{1,1}^{n_{1,1}} P_{2,1}^{N_{n_{1, \varnothing}}-n_{1,1}}}{n_{1,1}!\left(N_{n_{1, \varnothing}}-n_{1,1}\right)!} \sum_{n_{1,2}} \frac{P_{1,2}^{n_{1,2}} P_{2,2}^{N_{n_{1, \varnothing}}^{\prime}-n_{1,2}}}{n_{1,2}!\left(N_{n_{1, \varnothing}}^{\prime}-n_{1,2}\right)!} \\
& =\frac{n_{H}!P_{0}^{\mathrm{res}_{0}}}{\operatorname{res}_{0}!} \sum_{n_{1, \varnothing}} \frac{P_{1, \varnothing}^{n_{1, \varnothing}} P_{2, \varnothing}^{N-n_{1, \varnothing}}}{n_{1, \varnothing}!\left(N-n_{1, \varnothing}\right)!} \sum_{n_{1,1}} \frac{P_{1,1}^{n_{1,1}} P_{2,1}^{N_{n_{1, \varnothing}}-n_{1,1}}}{n_{1,1}!\left(N_{n_{1, \varnothing}}-n_{1,1}\right)!} \frac{\left(P_{1,2}+P_{2,2}\right)^{N_{n_{1, \varnothing}}^{\prime}}}{N_{n_{1, \varnothing}}^{\prime}!}
\end{aligned}
\]

Now, res \(_{1}-n_{1, \varnothing}\) is equal to the number of revotes for 1 , so that res \(_{1}-n_{1, \varnothing} \leq n_{r}\). Therefore, the innermost summation is also on the full domain for \(n_{1,2}\). Hence,
\[
\begin{aligned}
\operatorname{Pr}\left(\overrightarrow{r e s}, n_{r}\right) & =\frac{n_{H}!P_{0}^{\mathrm{res}}}{\operatorname{res}_{0}!} \sum_{n_{1, \varnothing}} \frac{P_{1, \varnothing}^{n_{1, \varnothing}} P_{2, \varnothing}^{N-n_{1, \varnothing}}}{n_{1, \varnothing}!\left(N-n_{1, \varnothing}\right)!} \frac{\left(P_{1,1}+P_{2,1}\right)^{N_{n_{1, \varnothing}}}}{N_{n_{1, \varnothing}}!} \frac{\left(P_{1,2}+P_{2,2}\right)^{N_{n_{1, \varnothing}}^{\prime}}}{N_{n_{1, \varnothing}}^{\prime}!} \\
& =\frac{n_{H}!P_{0}^{\mathrm{res}}}{\operatorname{res}_{0}!} \sum_{n_{1, \varnothing}} \frac{P_{1, \varnothing}^{n_{1, \varnothing}} P_{2, \varnothing}^{N-n_{1, \varnothing}}}{n_{1, \varnothing}!\left(N-n_{1, \varnothing}\right)!} \frac{P_{r 1}^{\mathrm{res}_{1}-n_{1, \varnothing}}}{\left(\mathrm{res}_{1}-n_{1, \varnothing}\right)!} \frac{P_{r 2}^{n_{r}+n_{1, \varnothing}-\mathrm{res}_{1}}}{\left(n_{r}+n_{1, \varnothing}-\mathrm{res}_{1}\right)!},
\end{aligned}
\]
where \(P_{r 1}=P_{1,1}+P_{2,1}\) and \(P_{r 2}=P_{1,2}+P_{2,2}\).
Note that this can be computed in \(O\left(n_{H}\right)\) floating operations. Since we must range over all possibilities for res \({ }_{0}\), res \({ }_{1}\), res \({ }_{2}, n_{r}\), the overall \(\delta^{\text {Real }}\) can be computed in \(O\left(n_{H}{ }^{4}\right)\) floating operations.

In Section III, we used the formulas from Lemma 2 and 1 to evaluate the coercion level. The only exceptions are Figures 16 and 6 , where the real coercion level for 513 and 1025 voters were obtained with a Monte-Carlo estimates with 1000000 iterations, which gives a sufficiently small confidence interval for our purpose. The reason why is that evaluating the formula from Lemma 2 would be too long for such values of \(n_{V}\).


Fig. 12: Simplified tree structure for two candidates.


Fig. 13: Coercion levels as a function of the probability of revoting; with 20 voters, 2 candidates and \(30 \%\) abstention, with three different distributions between the candidates.

\section*{D. The impact of the parameters}

In Lemma 2, we saw that \(\delta^{\text {Real }}\) only depends on \(P_{0}\), the probability of abstention, \(P_{r 1}\) (resp. \(P_{r 2}\) ), the probability to revote for option 1 (resp. 2) and \(P_{1, \varnothing}\) (resp. \(P_{2, \varnothing}\) ), the probability to vote for option 1 (resp. 2) without revoting. To give a better representation of these parameters, we define
\[
r=\left(P_{r 1}+P_{r 2}\right) /\left(P_{r 1}+P_{r 2}+P_{1, \varnothing}+P_{2, \varnothing}\right),
\]
which is the probability to revote (after voting),
\[
p_{1}=P_{1, \varnothing} /\left(P_{1, \varnothing}+P_{2, \varnothing}\right) \text { and } p_{2}=P_{2, \varnothing} /\left(P_{1, \varnothing}+P_{2, \varnothing}\right),
\]
which are the probabilities to vote for option 1 and option 2 when voting exactly once and
\[
r_{1}=P_{r 1} /\left(P_{r 1}+P_{r 2}\right) \text { and } r_{2}=P_{r 2} /\left(P_{r 1}+P_{r 2}\right),
\]
the probabilities to vote for option 1 and 2 when revoting (see Fig. 12 for an illustration).
A first important parameter is \(r\), the probability of revoting. In Fig. 13, we let it vary from 0 to 1, and we plot the coercion level in JCJ and in the ideal protocol for various values of \(\left(p_{1}, p_{2}\right)\) and \(\left(r_{1}, r_{2}\right)\). As expected, since the leakage we detected comes from the revotes, both are the same when there is no revote. Note that when everyone revotes, there is no coercion resistance anymore. This is due to the considered evasion strategy, which instructs the voter to vote once (which is detected if they are the only one to do so).

Now, an important point to notice is that, if \(\left(p_{1}, p_{2}\right)\) and \(\left(r_{1}, r_{2}\right)\) are close, the difference between the ideal and the real coercion level is small, unless there are a lot of revotes. On the other hand, if both distributions are opposite, the difference is


Fig. 14: Coercion levels as a function of the bias when revoting; with 20 voters, 2 candidates and \(30 \%\) abstention.


Fig. 15: Coercion level as a function of the abstention rate; with 20 voters and 2 candidates.
noticeable, even with a reasonable proportion of revotes. Therefore, another interesting parameter to consider is the distance between \(\left(p_{1}, p_{2}\right)\) and \(\left(r_{1}, r_{2}\right)\), defined as \(2\left|p_{1}-r_{1}\right|\), which we call the bias. In Fig. 14 , we let the bias vary from 0 to 2 and we plot the coercion level in the real and ideal protocols, for various probability of revoting. As expected, when the bias is maximal, there is no coercion resistance at all (this corresponds to the scenario of Section II-B). Note, however, that the leakage is non-zero when there is no bias. This, once again, is due to the fact that the adversary can count the number of revotes, and therefore have a non-negligible advantage to decide whether the coerced voter voted or abstained.

Another natural parameter to consider is the abstention rate. Until now, it was fixed at 0.3 in our experimentations. However, it can variate a lot in real-life elections. In Fig. 15, we plot the real and ideal coercion levels as a function of the abstention rate, with various probabilities of revoting and bias. As expected, when there is no abstention, the attacker can trivially break coercion-resistance with a forced-abstention attack. Similarly, when everyone but the coerced voter abstains, there is no coercion-resistance (both situations are captured by the ideal model).

Finally, the last parameter of interest is the number of honest voters. In Fig. 16, we plot the real and ideal coercion levels for \(16,32,64,128,256,512\) and 1024 honest uncoerced voters, with various probabilities of revoting and bias. As expected, the coercion level decreases as \(n_{H}\) grows for both the ideal and real game. Interestingly, the difference remains noticeable as long as there is enough revotes.


Fig. 16: Coercion level as a function of the number of voters; with 2 candidates and \(30 \%\) abstention.```


[^0]:    This work was partly supported by the PEPR ANR-22-PECY-0006 grant.

