

Election Verifiability: Cryptographic Definitions and an Analysis of Helios, Helios-C, and JCJ

Ben Smyth
University of Luxembourg
Luxembourg
research@bensmyth.com

Steven Frink
IBM Watson
Durham, NC, US
steven.frink@ibm.com

Michael R. Clarkson
Cornell University
Ithaca, NY, US
clarkson@cs.cornell.edu

Abstract—Election verifiability is defined in the computational model of cryptography. The definition formalizes notions of voters verifying their own votes, auditors verifying the tally of votes, and auditors verifying that only eligible voters vote. The Helios (Adida et al., 2009), Helios-C (Cortier et al., 2014) and JCJ (Juels et al., 2010) election schemes are analyzed using the definition. Neither Helios nor Helios-C satisfy the definition because they do not ensure that recorded ballots are tallied in certain cases when the adversary posts malicious material on the bulletin board. A variant of Helios is proposed and shown to satisfy the definition. JCJ similarly does not ensure that recorded ballots are tallied in certain cases. Moreover, JCJ does not ensure that only eligible voters vote, due to a trust assumption it makes. A variant of JCJ is proposed and shown to satisfy a weakened definition that incorporates the trust assumption. Previous definitions of verifiability (Juels et al., 2010; Cortier et al., 2014; Kiayias et al., 2015) and definitions of global verifiability (Küsters et al., 2010; Cortier et al., 2016) are shown to permit election schemes vulnerable to attacks, whereas the new definition prohibits those schemes. And a relationship between the new definition and a variant of global verifiability is shown.

I. INTRODUCTION

Electronic voting systems that have been deployed in real-world, large-scale public elections place extensive trust in software and hardware. Unfortunately, instead of being trustworthy, many systems are vulnerable to attacks that could bring election outcomes into disrepute [29], [30], [84], [142]. So relying solely on trust in voting systems is unwise; verification of election outcomes is essential.¹

Election verifiability enables voters and auditors to ascertain the correctness of election outcomes, regardless of whether the software and hardware of the voting system are trustworthy [1], [2], [39], [85], [111]. Kremer et al. [94] decompose election verifiability into three aspects:

- *Individual verifiability*: voters can check that their own ballots are recorded.
- *Universal verifiability*: anyone can check that the tally of recorded ballots is computed properly.
- *Eligibility verifiability*: anyone can check that each tallied vote was cast by an authorized voter.

We propose new definitions of these three aspects of verifiability in the computational model of cryptography. We show that individual and universal verifiability are orthogonal, and that eligibility verifiability implies individual verifiability. Because some electronic voting systems implement voter

authentication themselves, whereas other systems outsource voter authentication to third parties, we develop two variants of our definitions—one for systems with *internal authentication* and another for systems with *external authentication*.

We employ our definitions to analyze the verifiability of two well-known election schemes, JCJ [87] and Helios [5]. JCJ is an election scheme that achieves *coercion resistance* and has been implemented as Civitas [43]; it implements its own internal authentication. Helios is a web-based voting system that has been deployed in the real-world and outsources authentication. We also analyze the verifiability of Helios-C [47], a variant of Helios that implements internal authentication by digitally signing ballots.

The first implementation of Helios, namely *Helios 2.0*, and the current release, namely *Helios 3.1.4*, are known to have vulnerabilities that can be exploited to violate ballot secrecy and verifiability [23], [32], [52], [53], and the next Helios release [4], henceforth *Helios'12*, is intended to mitigate against those vulnerabilities. Our analysis shows that the mitigations are insufficient to ensure verifiability. In particular, an adversary could record a ballot that causes a voter's ballot to be omitted from tallying. A variant of Helios, henceforth *Helios'16*, is proposed, and shown to satisfy our definition of election verifiability with external authentication. Helios 2.0, Helios 3.1.4 and Helios'12 fail to satisfy our definition.

Our analysis of Helios-C reveals that an adversary could record an ill-formed ballot that causes tallying to abort in a manner that anyone will accept. Yet, our definition of universal verifiability demands that accepted outcomes include the choices used to construct any well-formed ballots. Hence, each voter can be assured that their choice contributed to the outcome. By comparison, Helios-C does not assure this, because ill-formed ballots cause tallying to abort and that abort will be accepted. Thus, Helios-C does not satisfy our definition of universal verifiability. Nevertheless, a straightforward variant of Helios-C that disregards ill-formed ballots should satisfy our definition.

Our analysis of JCJ reveals that an adversary could cause the acceptance of tallies which exclude authorized ballots in favour of unauthorized ballots. Yet, our definition of universal verifiability demands that accepted outcomes include only

1. *Doveriyai, no proveryai* (trust, but verify) says the Russian proverb.

the choices cast by authorized voters. Thus, JCJ does not satisfy our definition of universal verifiability. The JCJ election scheme does not satisfy our definition of eligibility verifiability either, because an adversary who learns the tallier’s private key could cast unauthorized votes. We introduce a weakened definition of eligibility verifiability, incorporating JCJ’s trust assumption that the private key is not known to the adversary, and show that variants of JCJ, henceforth *JCJ’16*, satisfy our weakened definition of election verifiability with internal authentication.

Küsters et al. [97], [98], [100], [101] propose an alternative, holistic notion of verifiability called *global verifiability*, which must be instantiated with a goal. We undertake a formal comparison of election verifiability and global verifiability, when instantiated with a goal proposed by the aforementioned authors and a goal by Cortier et al. [49]. We found that Helios’16 does not satisfy global verifiability with those goals. Nonetheless, we were able to show that Helios’16 satisfies a slightly weaker goal. And, moreover, election verifiability is strictly stronger than global verifiability with that goal.

Our definitions of election verifiability improve upon two previous definitions [47], [87] by detecting a new class of *collusion attacks*, in which the tallying algorithm announces an incorrect tally, and the verification algorithm colludes with the tallying algorithm to accept the incorrect tally. Examples of collusion attacks include vote stuffing, and announcing tallies that are independent of the election. Our definitions also improve upon those previous definitions and a further definition [91] by detecting a new class of *biasing attacks*, in which the verification algorithm rejects some legitimate election outcomes. Examples of biasing attacks include rejecting outcomes in which a particular candidate does not win, and rejecting all election outcomes, even correct outcomes. Moreover, our definitions improve upon global verifiability instantiated with goals by Küsters et al. [101] and Cortier et al. [49] by detecting a new class of *revelation attacks*, in which the verification algorithm accepts incorrect outcomes when coins used to construct some ballots are leaked. Examples of revelation attacks include announcing tallies that exclude or replace some votes.

This paper thus contributes to the security of electronic voting systems by:

- proposing definitions of election verifiability in the computational model;
- showing that individual, universal, and eligibility verifiability are mostly orthogonal properties of voting systems;
- proving that Helios 2.0, Helios 3.1.4, Helios’12, Helios-C and JCJ do not satisfy election verifiability, and that Helios’16 and JCJ’16 do;
- formally comparing election and global verifiability; and
- identifying new classes of attacks on voting systems and demonstrating that they are not detected by earlier works.

Our definitions are sufficient to analyze Helios, Helios-C, and JCJ. They correctly identify Helios 2.0, Helios 3.1.4, Helios’12, Helios-C and JCJ as not satisfying verifiability. And they enable the first proofs that Helios’16 and JCJ’16 satisfy a

definition of verifiability in the computational model. Although some protocols may fall outside the scope of our definitions, they are sufficiently general to be useful.

Structure: Section II defines election verifiability with external authentication. Section III analyzes Helios. Section IV defines election verifiability with internal authentication. Section V analyzes Helios-C. Section VI analyzes JCJ. Section VII presents a comparison between election and global verifiability. Section VIII introduces collusion, biasing and revelation attacks. Section IX reviews related work and Section X concludes. Appendix A defines cryptographic primitives. The remaining appendices explore alternative definitions of verifiability, give the details of Helios and JCJ, and present proofs.

II. EXTERNAL AUTHENTICATION

Some election schemes do not implement authentication themselves, but instead rely on an external authentication mechanism. Helios, for example, supports authentication with Facebook, Google and Twitter credentials.² In essence, the election scheme outsources ballot authentication. We begin by defining election verifiability for that model.

A. Election scheme syntax

We define syntax for an election scheme with external authentication, which henceforth in this section we abbreviate as “election scheme.”³

Definition 1 (Election scheme with external authentication). *An election scheme with external authentication is a tuple (Setup, Vote, Tally, Verify) of probabilistic polynomial-time (PPT) algorithms:*

- **Setup**, denoted⁴ $(PK_{\mathcal{T}}, SK_{\mathcal{T}}, m_B, m_C) \leftarrow \text{Setup}(k)$, is executed by the tallier, who is responsible for tallying ballots.⁵ Setup takes a security parameter k as input and outputs a key pair $(PK_{\mathcal{T}}, SK_{\mathcal{T}})$, a maximum number of ballots m_B , and a maximum number of candidates m_C .⁶
- **Vote**, denoted $b \leftarrow \text{Vote}(PK_{\mathcal{T}}, n_C, \beta, k)$, is executed by voters. A voter makes a choice of candidate from a sequence c_1, \dots, c_{n_C} of candidates. A well-formed choice is an integer β , such that $1 \leq \beta \leq n_C$. Vote takes as input the public key $PK_{\mathcal{T}}$ of the tallier, the number n_C of candidates, the voter’s choice β of candidate,

2. https://github.com/benadida/helios-server/tree/master/helios_auth/auth_systems, accessed 4 Aug 2015.

3. We focus on modeling first-past-the-post voting systems. Smyth shows the syntax is sufficiently versatile to capture ranked-choice voting systems too [125].

4. Let $\text{Alg}(in; r)$ denote the output of probabilistic algorithm Alg on input in and coins r . Let $\text{Alg}(in)$ denote $\text{Alg}(in; r)$, where r is chosen uniformly at random (from the coin space of algorithm Alg). And let \leftarrow denote assignment.

5. Some election schemes (e.g., Helios, Helios-C, and JCJ) permit the tallier’s role to be distributed amongst several talliers. For simplicity, we consider only a single tallier in this paper.

6. The maximum ballots and candidate numbers are used to formalize Correctness. Helios requires that the maximum number of ballots is less than or equal to the size of the underlying encryption scheme’s message space, and JCJ requires that the maximum number of candidates is less than or equal to the size of the underlying encryption scheme’s message space.

and security parameter k . It outputs a ballot b , or error symbol \perp . An error might occur if the candidate choice is not well-formed or for other reasons particular to the election scheme.

- **Tally**, denoted $(\mathbf{X}, P) \leftarrow \text{Tally}(SK_{\mathcal{T}}, BB, n_C, k)$, is executed by the tallier. It involves a public bulletin board BB , which we model as a set.⁷ Tally takes as input the private key $SK_{\mathcal{T}}$ of the tallier, the bulletin board BB , the number of candidates n_C , and security parameter k . It outputs a tally \mathbf{X} and a non-interactive proof P that the tally is correct. A tally is a vector \mathbf{X} of length n_C such that $\mathbf{X}[j]$ indicates the number of votes for candidate c_j .⁸
- **Verify**, denoted $v \leftarrow \text{Verify}(PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, P, k)$, can be executed by anyone to audit the election. Verify takes as input the public key $PK_{\mathcal{T}}$ of the tallier, the bulletin board BB , the number of candidates n_C , a tally \mathbf{X} , a proof P of correct tallying, and security parameter k . It outputs a bit v , which is 1 if the tally successfully verifies and 0 otherwise. We assume that Verify is deterministic.

Election schemes must satisfy Correctness: there exists a negligible function μ , such that for all security parameters k , integers n_B and n_C , and choices $\beta_1, \dots, \beta_{n_B} \in \{1, \dots, n_C\}$, it holds that if \mathbf{Y} is a vector of length n_C whose components are all 0, then

```
Pr[(PKT, SKT, mB, mC) ← Setup(k);
for 1 ≤ i ≤ nB do
  | bi ← Vote(PKT, nC, βi, k);
  | Y[βi] ← Y[βi] + 1;
BB ← {b1, ..., bnB};
(X, P) ← Tally(SKT, BB, nC, k) :
nB ≤ mB ∧ nC ≤ mC ⇒ X = Y] > 1 - μ(k).
```

Correctness asserts that tallies produced by Tally correspond to the choices input to Vote. Note that Correctness does not involve an adversary. Correctness therefore stipulates that, under ideal conditions, an election scheme does indeed produce the correct tally. Correctness is not actually necessary to achieve verifiability: our definition of universal verifiability will ensure that, in the presence of an adversary, Verify detects any errors in the tally. But it is reasonable to rule out election schemes that simply do not work properly under ideal conditions.

Limitations: Our model of election schemes is sufficient to analyze Helios and, after we extend the model to handle internal authentication in Section IV-A, Helios-C and JCJ. These are notable schemes, and formally analyzing their verifiability is a valuable contribution. But there are other notable schemes that fall outside our model:

- Prêt à Voter [39], MarkPledge [107], Scantegrity II [36], and Remotegrity [143] all rely on features implemented with paper, such as scratch-off surfaces and detachable columns.
- Everlasting privacy [105], which requires Vote to output a public ballot and a secret proof, involving temporal information, to the voter.

- Scyt's Pnyx.core ODBP 1.0 [42], which requires the bulletin board to be divided into two parts: a public part visible to all participants, and a secret part visible only to election administrators.

Distributed tallying also falls outside our model. We leave extension of our model to other election schemes and distributed tallying as future work.

B. Election verifiability

Election verifiability comprises three aspects: individual, universal, and eligibility verifiability. We express each as an *experiment*, which is an algorithm that outputs 0 or 1. The adversary *wins* an experiment by causing it to output 1.

1) *Individual verifiability:* In our model of election schemes, all recorded ballots are posted on the bulletin board. So for a voter to verify that their ballot has been recorded, it suffices to enable them to uniquely identify their ballot on the bulletin board.⁹

Individual verifiability experiment $\text{Exp-IV-Ext}(\Pi, \mathcal{A}, k)$, where Π denotes an election scheme, \mathcal{A} denotes the adversary, and k denotes a security parameter, therefore challenges \mathcal{A} to generate a scenario in which the voter cannot uniquely identify their ballot. In essence, Exp-IV-Ext challenges \mathcal{A} to generate a collision from Vote.¹⁰ If \mathcal{A} cannot win, then voters can uniquely identify their ballots on the bulletin board:

```
Exp-IV-Ext(Π, A, k) =
1 (PKT, nC, β, β') ← A(k);
2 b ← Vote(PKT, nC, β, k);
3 b' ← Vote(PKT, nC, β', k);
4 if b = b' ∧ b ≠ ⊥ ∧ b' ≠ ⊥ then
5   | return 1
6 else
7   | return 0
```

Line 1 asks \mathcal{A} to compute two candidate choices β and β' , such that ballots b and b' for those choices, as computed by Vote in lines 2 and 3, are equal.

One way to achieve individual verifiability is to base the election scheme on a probabilistic encryption scheme, such as El Gamal [65]. Intuitively, if Vote encrypts the choice using coins chosen uniformly at random, then it is overwhelmingly unlikely that two votes will result in the same ballot. Our proofs that Helios, Helios-C and JCJ satisfy individual verifiability are based on this idea.

7. Bulletin boards have also been modeled as public broadcast channels [56], [113], [118]. We abstract from the details of channels by employing sets to represent the data sent on them. We favor sets over multisets, because Cortier and Smyth [52], [53] demonstrate attacks against privacy when the bulletin board is modeled as a multiset.

8. Let $\mathbf{X}[i]$ denote component i of vector \mathbf{X} .

9. Section X addresses the complementary issue of whether a recorded ballot corresponds to the candidate choice a voter intended to make.

10. Exp-IV-Ext can be equivalently formulated as an experiment that challenges \mathcal{A} to predict the output of Vote. See Appendix B for details.

Clash attacks: In a *clash attack* [100], the adversary convinces some voters that a single ballot belongs to all of them. Some clash attacks are possible because of vulnerabilities in the design of `Vote`. For example, if `Vote` simply outputs candidate choice β , then a voter has no way to distinguish their vote for β from another voter’s vote for β . Exp-IV-Ext detects clash attacks resulting from vulnerabilities in `Vote`.

Some clash attacks, however, are possible because the adversary subverts the implementation of `Vote`. For example, the adversary might replace some hardware or software, or compromise the random number generator. If any one of these aspects is compromised, then `Vote` has effectively been changed to a different algorithm `Vote'`. The conclusions drawn by a security analyst who uses our definition of individual verifiability to analyze `Vote` would not necessarily be applicable to `Vote'`.

In short, a voter can verify that their ballot has been recorded if and only if they run the correct `Vote` algorithm. We make no guarantees to voters that do not run the correct `Vote` algorithm. One way to make stronger guarantees is to use cut-and-choose protocols to audit ballots [15], [16]. This would require modeling voting as an interactive protocol with the adversary, rather than as an algorithm. We leave this extension as future work.

2) *Universal verifiability:* For an election to be universally verifiable, anyone must be able to check that a tally is correct with respect to recorded ballots—that is, the tally represents the choices used to construct the recorded ballots. Because anyone can execute `Verify`, it suffices that `Verify` accepts if and only if that property holds.

Universal verifiability experiment `Exp-UV-Ext(Π, \mathcal{A}, k)` therefore challenges adversary \mathcal{A} to concoct a scenario in which `Verify` incorrectly accepts, thereby capturing the *only if* requirement:

```

Exp-UV-Ext( $\Pi, \mathcal{A}, k$ ) =
1 ( $PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, P$ )  $\leftarrow$   $\mathcal{A}(k)$ ;
2  $\mathbf{Y} \leftarrow$  correct-tally( $PK_{\mathcal{T}}, BB, n_C, k$ );
3 if  $\mathbf{X} \neq \mathbf{Y} \wedge \text{Verify}(PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, P, k) = 1$  then
4   | return 1
5 else
6   | return 0

```

In line 1, \mathcal{A} is challenged to create a bulletin board BB and purported tally \mathbf{X} of that bulletin board. Line 2 constructs the correct tally \mathbf{Y} of BB (using function *correct-tally*, which we define below), and line 3 checks whether `Verify` accepts an incorrect tally. If \mathcal{A} cannot win `Exp-UV-Ext`, then `Verify` will not accept incorrect tallies. In particular, no ballots can be omitted from the tally, and at most one candidate choice can be included in the tally for each ballot.

Let function *correct-tally* be defined such that for all $PK_{\mathcal{T}}, BB, n_C, k, \ell$, and $\beta \in \{1, \dots, n_C\}$,

$$\begin{aligned} \text{correct-tally}(PK_{\mathcal{T}}, BB, n_C, k)[\beta] &= \ell \\ \iff \exists^{\ell} b \in (BB \setminus \{\perp\}) : & \end{aligned}$$

$$\exists r : b = \text{Vote}(PK_{\mathcal{T}}, n_C, \beta, k; r).$$

The vector produced by *correct-tally* must be of length n_C . Component β of vector *correct-tally*($PK_{\mathcal{T}}, BB, n_C, k$) equals ℓ iff there exist¹¹ ℓ ballots on the bulletin board that are votes for candidate β . It follows that the output of *correct-tally* represents the choices used to construct the recorded ballots. Of course, *correct-tally* cannot be computed by a PPT algorithm for typical cryptographic election schemes. But that does not matter, because *correct-tally* is never actually computed as part of an election scheme—its use is solely in the definition of `Exp-UV-Ext`.¹²

Function *correct-tally* requires that ballots can only be interpreted for one candidate, which can be ensured by *Injectivity*:

Definition 2 (Injectivity). *An election scheme (Setup, Vote, Tally, Verify) satisfies Injectivity, if for all security parameters k , public keys $PK_{\mathcal{T}}$, integers n_C , and choices β and β' , such that $\beta \neq \beta'$, we have*

$$\begin{aligned} \Pr[b \leftarrow \text{Vote}(PK_{\mathcal{T}}, n_C, \beta, k); \\ b' \leftarrow \text{Vote}(PK_{\mathcal{T}}, n_C, \beta', k) : \\ b \neq \perp \wedge b' \neq \perp \Rightarrow b \neq b'] &= 1. \end{aligned}$$

Injectivity ensures that distinct choices are not mapped by `Vote` to the same ballot.¹³ Without Injectivity, an election scheme might produce ballots whose meaning is ambiguous. For example, if $\text{Vote}(PK_{\mathcal{T}}, n_C, \beta, k; r)$ were defined to be $\beta + r$, then a ballot b could be tallied as any well-formed choice β' such that $\beta' = b - r'$ for some r' . But that definition of `Vote` is prohibited by Injectivity. Thus, Injectivity helps to ensure that the choices used to construct ballots can be uniquely tallied.

Security analysts must convince themselves that *correct-tally* is indeed correct. Because of the function’s simplicity, this should be relatively straightforward. By comparison, Tally algorithms for real voting schemes tend to be complicated. For example, compare the complexity of *correct-tally* to Helios’s Tally algorithm, which appears in Definition 24 of Appendix C.

By design, `Exp-UV-Ext` assumes the ballots on bulletin board BB are exactly the ballots that should be tallied. The external authentication mechanism is assumed to prohibit unauthorized ballots from being posted on BB . Helios makes such an assumption about its external authentication mechanism.

11. The definition of *correct-tally* employs a *counting quantifier* [122] denoted \exists^{ℓ} . Predicate $(\exists^{\ell} x : P(x))$ holds exactly when there are ℓ distinct values for x such that $P(x)$ is satisfied. Variable x is bound by the quantifier, whereas ℓ is free.

12. Kiayias et al. [91] use a similar super-polynomial *vote extractor* to recover choices from ballots in an experiment defining verifiability.

13. Individual verifiability resembles Injectivity, but individual verifiability allows choices to be equal and allows adversary \mathcal{A} to choose election parameters.

Election schemes must also satisfy Completeness, which stipulates that tallies produced by Tally will actually be accepted by Verify, capturing the *if* requirement:

Definition 3 (Completeness). *An election scheme (Setup, Vote, Tally, Verify) satisfies Completeness, if for all PPT adversaries \mathcal{A} , there exists a negligible function μ , such that for all security parameters k , it holds that*

$$\begin{aligned} & \Pr[(PK_{\mathcal{T}}, SK_{\mathcal{T}}, m_B, m_C) \leftarrow \text{Setup}(k); \\ & (BB, n_C) \leftarrow \mathcal{A}(PK_{\mathcal{T}}, k); \\ & (\mathbf{X}, P) \leftarrow \text{Tally}(SK_{\mathcal{T}}, BB, n_C, k) : \\ & |BB| \leq m_B \wedge n_C \leq m_C \Rightarrow \\ & \text{Verify}(PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, P, k) = 1] > 1 - \mu(k). \end{aligned}$$

Without Completeness, election schemes might be vulnerable to biasing attacks, as we show in Section VIII-B.

3) *Eligibility verifiability*: For an election to satisfy eligibility verifiability, anyone must be able to check that every tallied vote was cast by an authorized voter—hence, it must be possible to authenticate ballots. In election schemes with external authentication, a trusted third party authenticates ballots. That third party might convince itself that all tallied ballots have been authenticated, but it cannot convince all other parties. Eligibility verifiability, therefore, is not achievable in election schemes with external authentication.

4) *Election verifiability*: With Exp-IV-Ext and Exp-UV-Ext, we define election verifiability with external authentication. Let a PPT adversary’s *success* $\text{Succ}(\text{Exp}(\cdot))$ in an experiment $\text{Exp}(\cdot)$ be the probability that the adversary wins—that is, $\text{Succ}(\text{Exp}(\cdot)) = \Pr[b \leftarrow \text{Exp}(\cdot) : b = 1]$.

Definition 4 (Ver-Ext). *An election scheme Π satisfies election verifiability with external authentication (Ver-Ext) if Completeness and Injectivity are satisfied and for all PPT adversaries \mathcal{A} , there exists a negligible function μ , such that for all security parameters k , it holds that $\text{Succ}(\text{Exp-IV-Ext}(\Pi, \mathcal{A}, k)) + \text{Succ}(\text{Exp-UV-Ext}(\Pi, \mathcal{A}, k)) \leq \mu(k)$.*

An election scheme satisfies individual verifiability if $\text{Succ}(\text{Exp-IV-Ext}(\Pi, \mathcal{A}, k)) \leq \mu(k)$. And universal verifiability is satisfied if the election scheme satisfies Completeness and Injectivity, and $\text{Succ}(\text{Exp-UV-Ext}(\Pi, \mathcal{A}, k)) \leq \mu(k)$.

C. *Example—Toy scheme from nonces*

A toy election scheme satisfying Ver-Ext can be based on nonces. Each voter publishes a nonce paired with their choice of candidate to the bulletin board. This scheme illustrates the essence of election verifiability, even though it does not offer any privacy.

Definition 5. *Election scheme Nonce is defined as follows:*

- $\text{Setup}(k)$ outputs $(\perp, \perp, p_1(k), p_2(k))$, where p_1 and p_2 may be any polynomial functions.
- $\text{Vote}(PK_{\mathcal{T}}, n_C, \beta, k)$ selects a nonce r uniformly at random from \mathbb{Z}_{2^k} and outputs (r, β) .
- $\text{Tally}(SK_{\mathcal{T}}, BB, n_C, k)$ computes a vector \mathbf{X} of length n_C , such that \mathbf{X} is a tally of the votes on BB for which the nonce is in \mathbb{Z}_{2^k} , and outputs (\mathbf{X}, \perp) .

- $\text{Verify}(PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, P, k)$ outputs 1 if $(\mathbf{X}, P) = \text{Tally}(\perp, BB, n_C, k)$, and 0 otherwise.

Proposition 1. *Nonce satisfies Ver-Ext.*

Proof sketch. Nonce satisfies individual verifiability, because voters can use their nonce to check that their own ballot appears on the bulletin board. With overwhelming probability, Vote will select unique nonces for each voter, hence generate distinct ballots. Nonce also satisfies universal verifiability, because plaintext candidate choices are posted on the bulletin board. \square

D. *Orthogonality*

Exp-IV-Ext and Exp-UV-Ext capture orthogonal security properties. A scheme that satisfies individual verifiability but violates universal verifiability can be constructed from Nonce by modifying Verify to always output 1. Voters can still check that their own ballot appears. But an adversary can easily win Exp-UV-Ext, because Verify will accept any tally. A scheme that satisfies universal verifiability but violates individual verifiability can be constructed from Nonce by removing the nonces, leaving just the voter’s choice in the ballots. Call that scheme Choice. Anyone can still verify the tally of the election, but an adversary can easily win Exp-IV-Ext, because two votes for the same candidate will collide.

III. CASE STUDY: HELIOS

Helios [5], [112] is an open-source, web-based electronic voting system,¹⁴ which has been deployed in the real-world. The International Association of Cryptologic Research (IACR) has used Helios annually since 2010 to elect board members [18], [76], the ACM used Helios in an ACM general election [140], the Catholic University of Louvain used Helios to elect the university president [5], and Princeton University has used Helios to elect several student governments [3], [109].

Helios is intended to satisfy verifiability whilst maintaining *ballot secrecy*—i.e., without revealing voters’ votes. For ballot secrecy, voters encrypt candidate choices using a homomorphic encryption scheme, these encrypted choices are homomorphically combined, and the tallier decrypts the homomorphic combination to reveal the tally.¹⁵ For verifiability, encryption and decryption steps are accompanied by zero-knowledge proofs.

Informally, Helios works as follows:

- **Setup.** The tallier generates a key pair for a homomorphic encryption scheme and publishes the public key.
- **Voting.** A voter encrypts their candidate choice with the tallier’s public key, and proves in zero-knowledge that the ciphertext contains a well-formed choice. The voter posts their ballot (i.e., ciphertext and proof) on the

14. <https://vote.heliosvoting.org/>, accessed 16 Nov 2015.

15. Homomorphic combination of ciphertexts is straightforward for two-candidate elections [14], [19], [44], [79], [117], since choices (e.g., “yes” or “no”) can be encoded as 1 or 0. Multi-candidate elections are also possible [19], [59], [78].

bulletin board. (The bulletin board is assumed to correctly authenticate voters during posting.)

- **Tallying.** The tallier discards any ballots from the bulletin board for which proofs do not hold. The tallier homomorphically combines the ciphertexts in the remaining ballots, decrypts the homomorphic combination, and proves in zero-knowledge that decryption was performed correctly. Finally, the tallier publishes the winning candidate and proof of correct decryption.
- **Verification.** A verifier recomputes the homomorphic combination and checks all the zero-knowledge proofs.

Helios was first implemented as Helios 2.0.^{16,17}

Chang-Fong & Essex [32] have shown that Helios 2.0 does not satisfy universal verifiability. Thus, we would not expect Ver-Ext to hold for Helios 2.0. Indeed, we formalize a generic construction for Helios-like election schemes (Appendix C), which we use to derive a formal description of Helios 2.0 (Appendix D). And using that description, we can prove that Helios 2.0 is not verifiable:

Proposition 2. *Helios 2.0 does not satisfy Ver-Ext.*

Proof sketch. Our proof formalizes the attack by Chang-Fong & Essex in the context of our Completeness definition. □

A proof of Proposition 2 appears in Appendix D. Vulnerabilities can be attributed Helios 2.0 not checking the suitability of cryptographic parameters nor checking that all elements of ballots are constructed using the correct parameters, and the current version of Helios (Helios 3.1.4) is intended to mitigate against those vulnerabilities by performing the necessary checks.¹⁸

Bernhard *et al.* [23] have shown that Helios 3.1.4 does not satisfy universal verifiability. Thus, we would not expect Ver-Ext to hold for Helios 3.1.4 either. Indeed, we use our generic construction to derive a formal description of Helios 3.1.4 (Appendix E). And using that description, we can prove that Helios 3.1.4 is not verifiable:

Proposition 3. *Helios 3.1.4 does not satisfy Ver-Ext.*

Proof sketch. Our proof formalizes the attack by Bernhard *et al.* in the context of our universal verifiability experiment. □

A proof of Proposition 3 appears in Appendix E. Bernhard *et al.* attribute vulnerabilities to application of the Fiat-Shamir transformation without inclusion of statements in hashes (i.e., the weak Fiat-Shamir transformation), and including statements in hashes (i.e., applying the Fiat-Shamir transformation) is postulated as a defense.

Beyond verifiability, Helios 3.1.4 has been shown not to satisfy ballot secrecy,¹⁹ due to tallying meaningfully related ballots,²⁰ and omitting such ballots from the tally (i.e., ballot weeding) is postulated as a defense [21], [22], [52], [53], [124], [131], [132], [134]. The next Helios release (Helios'12) is intended to mitigate against vulnerabilities. In particular, the specification [4] incorporates the Fiat-Shamir transformation (rather than the weak Fiat-Shamir transformation). And there are plans to incorporate ballot weeding.^{21,22} Although ballot

weeding can be sufficient for ballot secrecy (cf. [131, §6] & [126]), we have found that it violates universal verifiability. In particular, an adversary can observe a voter's ballot and cast a related ballot (for a candidate other than the voter's choice), such that the voter's ballot is omitted from tallying. (This could be achieved, for example, by manipulating the bulletin board to ensure that the adversary's ballot is processed before the voter's ballot, since this causes the voter's ballot to be weeded.) Our definition of universal verifiability requires all ballots on the bulletin board to be tallied, thus it is violated by ballot weeding. It follows that Helios'12 does not satisfy Ver-Ext, because that scheme relies upon ballot weeding to defend against ballot secrecy violations.

Remark 4. *Helios'12 does not satisfy Ver-Ext.*

Proof sketch. Helios'12 uses ballot weeding, which violates universal verifiability, as described above. □

An informal proof of Remark 4 follows immediately from our discourse. A formal proof would require a formal description of Helios'12. Such a formal description can be derived as a straightforward variant of Helios 3.1.4 that applies the Fiat-Shamir transformation (rather than the weak Fiat-Shamir transformation) and uses ballot weeding. These details provide little value, so we do not pursue them further.

To ensure universal verifiability, we propose variants of Helios'12. Our variants defend against ballot secrecy violations by incorporating proposals by Smyth *et al.* [135], [136] and Smyth [126] for non-malleable ballots, rather than proposals for ballot weeding. We formalize those variants as a set (Helios'16) of election schemes (Appendix F). Using that formalization, we can prove that Helios'16 is verifiable:²³

16. <https://github.com/benadida/helios/releases/tag/2.0>, released 25 Jul 2009, accessed 16 Nov 2015.

17. Helios 2.0 builds upon Adida's *Helios 1.0* [2]. But, the two systems are rather different. In particular, the Helios 2.0 tallier homomorphically combines encrypted choices and decrypts the homomorphic combination to reveal the tally, whereas the Helios 1.0 tallier mixes encrypted choices and decrypts the ciphertexts output by the mix. Adida has not released an implementation of Helios 1.0. Tsoukalas *et al.* [141] released *Zeus* as a fork of Helios 2.0 spliced with mixnet code to derive an implementation (<https://github.com/gmet/zeus>, accessed 15 Sep 2017) and Yingtong Li released *helios-server-mixnet* as an extension of Zeus with threshold asymmetric encryption and some other minor changes (<https://github.com/RunasSudo/helios-server-mixnet>, accessed 15 Sep 2017). Smyth shows that those implementations do not satisfy universal verifiability and proves that a variant does [128].

18. Cf. <https://github.com/benadida/helios-server/pull/133>, accessed 14 Dec 2016.

19. Eligibility is not satisfied either [104], [137], [138].

20. Meaningfully related ballots can be constructed because Helios ballots are malleable.

21. Cf. <https://github.com/benadida/helios-server/issues/8> and <https://github.com/benadida/helios-server/issues/35>, accessed 9 Aug 2016.

22. Ballot weeding mechanisms have been proposed, e.g., [21], [22], [25], [52], [53], [124], [131], [134], but the specification for Helios'12 does not yet define a particular mechanism. One candidate mechanism would omit any ballot containing a previously observed hash from the tallying procedure.

23. A set of election schemes satisfies Ver-Ext, if every scheme in the set satisfies Ver-Ext.

Theorem 5. *Helios’16 satisfies Ver-Ext.*

Proof sketch. Helios’16 satisfies individual verifiability, because the probabilistic encryption scheme ensures that ballots are unique, with overwhelming probability. And Helios’16 satisfies universal verifiability, because the zero-knowledge proofs can be publicly verified. \square

A formal proof of Theorem 5 appears in Appendix F. The proof assumes the random oracle model [11]. This proof, coupled with the proof of ballot secrecy by Smyth [126], provides strong motivation for future Helios releases being based upon Helios’16, since it is the only variant of Helios which is known to be secure.

IV. INTERNAL AUTHENTICATION

Some election schemes implement their own authentication mechanisms. JCI [85]–[87] and Civitas [43], for example, authenticate ballots based on *credentials* issued to voters by a registration authority. Schemes with this kind of internal authentication enable verification of whether tallied ballots were cast by authorized voters.

A. Election scheme syntax

A *registrar* is responsible for issuing authentication *credentials* to voters.²⁴ Each voter is associated with a credential pair (pk, sk) . The voter uses private credential sk to construct a ballot. Public credential pk is used during tallying and verification. Let L denote the *electoral roll*, which is the set of all public credentials.

We revise our syntax to capture an election scheme with internal authentication, which henceforth in this section we abbreviate as “election scheme.”

Definition 6 (Election scheme with internal authentication). *An election scheme with internal authentication is a tuple (Setup, Register, Vote, Tally, Verify) of PPT algorithms:*

- $(PK_{\mathcal{T}}, SK_{\mathcal{T}}, m_B, m_C) \leftarrow \text{Setup}(k)$
- $(pk, sk) \leftarrow \text{Register}(PK_{\mathcal{T}}, k)$
- $b \leftarrow \text{Vote}(sk, PK_{\mathcal{T}}, n_C, \beta, k)$
- $(\mathbf{X}, P) \leftarrow \text{Tally}(SK_{\mathcal{T}}, BB, L, n_C, k)$
- $v \leftarrow \text{Verify}(PK_{\mathcal{T}}, BB, L, n_C, \mathbf{X}, P, k)$

Election schemes must satisfy Correctness: there exists a negligible function μ , such that for all security parameters k , integers n_B and n_C , and choices $\beta_1, \dots, \beta_{n_B} \in \{1, \dots, n_C\}$, it holds that if \mathbf{Y} is a vector of length n_C whose components are all 0, then

$$\Pr[(PK_{\mathcal{T}}, SK_{\mathcal{T}}, m_B, m_C) \leftarrow \text{Setup}(k);$$

$$\text{for } 1 \leq i \leq n_B \text{ do}$$

$$\quad (pk_i, sk_i) \leftarrow \text{Register}(PK_{\mathcal{T}}, k);$$

$$\quad b_i \leftarrow \text{Vote}(sk_i, PK_{\mathcal{T}}, n_C, \beta_i, k);$$

$$\quad \mathbf{Y}[\beta_i] \leftarrow \mathbf{Y}[\beta_i] + 1;$$

$$L \leftarrow \{pk_1, \dots, pk_{n_B}\};$$

$$BB \leftarrow \{b_1, \dots, b_{n_B}\};$$

$$(\mathbf{X}, P) \leftarrow \text{Tally}(SK_{\mathcal{T}}, BB, L, n_C, k) :$$

$$n_B \leq m_B \wedge n_C \leq m_C \Rightarrow \mathbf{X} = \mathbf{Y} > 1 - \mu(k).$$

Setup is unchanged from election schemes with external authentication (cf. §II-A). The only change to Vote is that it now accepts private credential sk as input. Similarly, the only change to Tally and Verify is that they now accept electoral roll L as input. Register is executed by the registrar. It takes as input the public key $PK_{\mathcal{T}}$ of the tallier and security parameter k , and it outputs a *credential pair* (pk, sk) . After all voters have been registered, the registrar certifies the electoral roll, perhaps by digitally signing and publishing it.²⁵

B. Election verifiability

Secure construction of electoral rolls is not a topic that electronic voting systems usually address—though it seems an important part of any real-world deployment. Indeed, voting systems typically assume the registrar is honest. In our experiments, below, we model an adversary who cannot corrupt the registration process that issues credentials to voters. Hence our definitions will not detect attacks against verifiabilities that result solely from weaknesses in the registration process.²⁶

Recall (from §II-B) that election verifiability is expressed with experiments, and that an adversary wins by causing an experiment to output 1. We henceforth assume that the adversary is *stateful*—that is, information persists across invocations of the adversary in a single experiment. Our experiments in Section II did not need this assumption, because they never invoked the adversary more than once.

1) *Individual verifiability:* The individual verifiability experiment again challenges adversary \mathcal{A} to generate a scenario in which the voter could not uniquely identify their ballot:²⁷

$$\text{Exp-IV-Int}(\Pi, \mathcal{A}, k) =$$

- 1 $(PK_{\mathcal{T}}, n_V) \leftarrow \mathcal{A}(k);$
- 2 **for** $1 \leq i \leq n_V$ **do** $(pk_i, sk_i) \leftarrow \text{Register}(PK_{\mathcal{T}}, k);$
- 3 $L \leftarrow \{pk_1, \dots, pk_{n_V}\};$
- 4 $\text{Crpt} \leftarrow \emptyset;$
- 5 $(n_C, \beta, \beta', i, j) \leftarrow \mathcal{A}^C(L);$
- 6 $b \leftarrow \text{Vote}(sk_i, PK_{\mathcal{T}}, n_C, \beta, k);$
- 7 $b' \leftarrow \text{Vote}(sk_j, PK_{\mathcal{T}}, n_C, \beta', k);$
- 8 **if**
- 8.1 $b = b' \wedge b \neq \perp \wedge b' \neq \perp \wedge i \neq j \wedge sk_i \notin \text{Crpt} \wedge sk_j \notin \text{Crpt}$
- 8.2 **then**
- 9 **return** 1
- 10 **else**
- 11 **return** 0

24. Some election schemes (e.g., Helios-C and JCI) permit the registrar’s role to be distributed among several registrars. For simplicity, we consider only a single registrar in this paper.

25. It might seem surprising that Register does not require the registrar to provide any private keys as input. But in constructions of election schemes with internal authentication, e.g., [43], [87], the registrar does not sign credential pairs with its own private key. Rather, the registrar signs the electoral roll.

26. Küsters and Truderung [96] explore some consequences of permitting adversarial influence during registration.

27. Unlike Exp-IV-Ext, a variant of Exp-IV-Int that challenges \mathcal{A} to predict the output of Vote is strictly stronger. See Appendix B for details.

The main differences from the corresponding experiment for external authentication (§II-B1) are that voters are registered in line 2, and that \mathcal{A} is given access to an oracle C in line 5. The oracle is used to model \mathcal{A} corrupting voters and learning their private credentials: on invocation $C(\ell)$, where $1 \leq \ell \leq n_V$, the oracle records that voter ℓ is corrupted by updating $Crpt$ to be $Crpt \cup \{sk_\ell\}$ and outputs sk_ℓ . In line 5, the voter indices output by \mathcal{A} must be legal with respect to n_V , but we elide that detail from the experiment for simplicity. Line 8 ensures that \mathcal{A} cannot trivially win by corrupting voters.

2) *Universal verifiability*: The universal verifiability experiment again challenges \mathcal{A} to concoct a scenario in which Verify incorrectly accepts:

```

Exp-UV-Int( $\Pi, \mathcal{A}, k$ ) =
1 ( $PK_{\mathcal{T}}, n_V$ )  $\leftarrow \mathcal{A}(k)$ ;
2 for  $1 \leq i \leq n_V$  do ( $pk_i, sk_i$ )  $\leftarrow$  Register( $PK_{\mathcal{T}}, k$ );
3  $L \leftarrow \{pk_1, \dots, pk_{n_V}\}$ ;
4  $M \leftarrow \{(pk_1, sk_1), \dots, (pk_{n_V}, sk_{n_V})\}$ ;
5 ( $BB, n_C, \mathbf{X}, P$ )  $\leftarrow \mathcal{A}(M)$ ;
6  $\mathbf{Y} \leftarrow$  correct-tally( $PK_{\mathcal{T}}, BB, M, n_C, k$ );
7 if  $\mathbf{X} \neq \mathbf{Y} \wedge$  Verify( $PK_{\mathcal{T}}, BB, L, n_C, \mathbf{X}, P, k$ ) = 1 then
8 | return 1
9 else
10 | return 0

```

The main differences from the corresponding experiment for external authentication (§II-B2) are that voters are registered in line 2, and their credential pairs are used in the rest of the experiment.

The tally of recorded ballots should contain at most one vote per voter. Hence, election schemes must handle *revotes*—i.e., multiple ballots submitted by the same voter. Election schemes with external authentication implicitly handle revoting, by assuming a third party ensures that the recorded ballots contain at most one ballot per voter. Election schemes with internal authentication must explicitly handle revoting by tallying only authorized ballots. A ballot is *authorized* if it is constructed with a private credential from M , and that private credential was not used to construct any other ballot on BB .^{28,29}

Function *correct-tally* is now modified to tally only authorized ballots: let function *correct-tally* now be defined such that for all $PK_{\mathcal{T}}, BB, M, n_C, k, \ell$, and $\beta \in \{1, \dots, n_C\}$,

$$\begin{aligned}
& \text{correct-tally}(PK_{\mathcal{T}}, BB, M, n_C, k)[\beta] = \ell \\
& \iff \exists^{\ell} b \in \text{authorized}(PK_{\mathcal{T}}, (BB \setminus \{\perp\}), M, n_C, k) : \\
& \quad \exists sk, r : b = \text{Vote}(sk, PK_{\mathcal{T}}, n_C, \beta, k; r).
\end{aligned}$$

By comparison, the original *correct-tally* function (§II-B2) tallies all the ballots on BB . Function *correct-tally* requires that ballots can only be interpreted for one candidate, which can again be ensured by Injectivity, which we update to include private credentials:

Definition 7 (Injectivity). *An election scheme* (Setup, Register, Vote, Tally, Verify) *satisfies Injectivity, if for all security parameters* k , *public keys* $PK_{\mathcal{T}}$, *integers* n_C , *and*

choices β *and* β' , *such that* $\beta \neq \beta'$, *we have*

$$\begin{aligned}
& \Pr[(pk, sk) \leftarrow \text{Register}(PK_{\mathcal{T}}, k); \\
& \quad (pk', sk') \leftarrow \text{Register}(PK_{\mathcal{T}}, k); \\
& \quad b \leftarrow \text{Vote}(sk, PK_{\mathcal{T}}, n_C, \beta, k); \\
& \quad b' \leftarrow \text{Vote}(sk', PK_{\mathcal{T}}, n_C, \beta', k) : \\
& \quad b \neq \perp \wedge b' \neq \perp \Rightarrow b \neq b'] = 1.
\end{aligned}$$

Let *authorized* be defined as follows:

$$\begin{aligned}
& \text{authorized}(PK_{\mathcal{T}}, BB, M, n_C, k) = \\
& \{b : b \in BB \\
& \quad \wedge \exists pk, sk, \beta, r : b = \text{Vote}(sk, PK_{\mathcal{T}}, n_C, \beta, k; r) \\
& \quad \wedge (pk, sk) \in M \wedge \neg \exists b', \beta', r' : b' \in (BB \setminus \{b\}) \\
& \quad \wedge b' = \text{Vote}(sk, PK_{\mathcal{T}}, n_C, \beta', k; r')\}.
\end{aligned}$$

Function *authorized* discards ballots submitted under the same credential—that is, if there is more than one ballot submitted with a private credential sk , then all ballots submitted under that credential are discarded. Therefore, election schemes that permit revoting cannot be analyzed with this definition of *authorized*. But alternative definitions of *authorized* are possible—for example, if ballots were timestamped, *authorized* could discard all but the most recent ballot submitted under a particular credential.

Election schemes must continue to satisfy Completeness, which we update to include credentials and the electoral roll:

Definition 8 (Completeness). *An election scheme* (Setup, Register, Vote, Tally, Verify) *satisfies Completeness, if for all PPT adversaries* \mathcal{A} , *there exists a negligible function* μ , *such that for all security parameters* k , *it holds that*

$$\begin{aligned}
& \Pr[(PK_{\mathcal{T}}, SK_{\mathcal{T}}, m_B, m_C) \leftarrow \text{Setup}(k); \\
& \quad n_V \leftarrow \mathcal{A}(PK_{\mathcal{T}}, k); \\
& \quad \text{for } 1 \leq i \leq n_V \text{ do } (pk_i, sk_i) \leftarrow \text{Register}(PK_{\mathcal{T}}, k); \\
& \quad L \leftarrow \{pk_1, \dots, pk_{n_V}\}; \\
& \quad M \leftarrow \{(pk_1, sk_1), \dots, (pk_{n_V}, sk_{n_V})\}; \\
& \quad (BB, n_C) \leftarrow \mathcal{A}(M); \\
& \quad (\mathbf{X}, P) \leftarrow \text{Tally}(SK_{\mathcal{T}}, BB, L, n_C, k) : \\
& \quad |BB| \leq m_B \wedge n_C \leq m_C \Rightarrow \\
& \quad \text{Verify}(PK_{\mathcal{T}}, BB, L, n_C, \mathbf{X}, P, k) = 1] > 1 - \mu(k).
\end{aligned}$$

28. Helios-C is claimed to support an alternative definition of *authorized*, whereby only the last ballot cast by a voter is authorized. We found that Helios-C does not support this definition. In particular, an adversary can observe the ballots cast by a voter and replay one of those ballots. The replayed ballot will overwrite the last ballot cast by the voter and will be authorized instead of it.

29. JCJ is claimed to support alternative definitions of *authorized*—e.g., only the last ballot cast by a voter is authorized—using a policy [87, §4.1]. We found that the policy proposed by Juels et al. (namely, “order of postings to [the bulletin board]”) does not support this definition of *authorized*. In particular, an adversary can intercept a voter’s ballot and replay that ballot after observing the voter’s revoke, thus the policy incorrectly defines the first ballot as authorized. This could be prevented by proving knowledge of previously constructed ballots (cf. Clarkson et al. [43]).

3) *Eligibility verifiability*: Recall (from §II-B3) that for an election scheme to satisfy eligibility verifiability, anyone must be able to check that every tallied vote was cast by an authorized voter—hence, it must be possible to authenticate ballots. Because voters are issued credential pairs that can be used to authenticate ballots, it suffices to ensure that knowledge of a private credential is necessary to construct an authentic ballot.

Eligibility verifiability experiment Exp-EV-Int therefore challenges \mathcal{A} to produce a ballot under a private credential that \mathcal{A} does not know:

```

Exp-EV-Int( $\Pi, \mathcal{A}, k$ ) =
1 ( $PK_{\mathcal{T}}, n_V$ )  $\leftarrow \mathcal{A}(k)$ ;
2 for  $1 \leq i \leq n_V$  do ( $pk_i, sk_i$ )  $\leftarrow \text{Register}(PK_{\mathcal{T}}, k)$ ;
3  $L \leftarrow \{pk_1, \dots, pk_{n_V}\}$ ;
4  $Crpt \leftarrow \emptyset$ ;  $Rvld \leftarrow \emptyset$ ;
5 ( $n_C, \beta, i, b$ )  $\leftarrow \mathcal{A}^{C,R}(L)$ ;
6 if  $\exists r : b = \text{Vote}(sk_i, PK_{\mathcal{T}}, n_C, \beta, k; r) \wedge b \neq \perp \wedge b \notin$ 
    $Rvld \wedge sk_i \notin Crpt$  then
7 | return 1
8 else
9 | return 0

```

In line 1, \mathcal{A} chooses the tallier’s public key and the number of voters. Line 2 registers voters. \mathcal{A} is not permitted to influence registration while it is in progress. In particular, \mathcal{A} is not permitted to choose credential pairs, because by doing so \mathcal{A} could trivially win the experiment.

Line 4 initializes two sets: $Crpt$ is a set of voters who have been corrupted, meaning that \mathcal{A} has learned their private credential, and $Rvld$ is a set of ballots that have been revealed to \mathcal{A} . The former set models \mathcal{A} coercing voters to reveal their private credentials. The latter set models \mathcal{A} observing ballots on the bulletin board.

Line 5 challenges \mathcal{A} to produce a ballot b with the help of two oracles. Oracle C is the same oracle as in Exp-IV-Int (cf. §IV-B1); it leaks the private credentials of corrupted voters to \mathcal{A} . Oracle R reveals ballots. On invocation $R(i, \beta, n_C)$, where $1 \leq i \leq n_V$, oracle R does the following:

- Computes a ballot b that represents a vote for candidate β by a voter with private credential sk_i , that is, computes $b \leftarrow \text{Vote}(sk_i, PK_{\mathcal{T}}, n_C, \beta, k)$.
- Records b as being revealed by updating $Rvld$ to be $Rvld \cup \{b\}$.
- Outputs b .

In line 6, \mathcal{A} wins if (i) the ballot is *authentic*, meaning that it is the output of Vote on an authorized credential, and (ii) that credential belongs to a voter that \mathcal{A} did not corrupt, and (iii) that ballot was not revealed. If \mathcal{A} cannot succeed in this experiment, then only authorized votes are tallied.

4) *Election verifiability*: With Exp-IV-Int , Exp-UV-Int , and Exp-EV-Int , we define election verifiability with internal authentication.

Definition 9 (Ver-Int). *An election scheme Π satisfies election verifiability with internal authentication (Ver-Int) if Completeness and Injectivity are satisfied and for all PPT adversaries \mathcal{A} , there exists a negligible function μ , such that for all security parameters k , it holds that $\text{Succ}(\text{Exp-IV-Int}(\Pi, \mathcal{A}, k)) + \text{Succ}(\text{Exp-UV-Int}(\Pi, \mathcal{A}, k)) + \text{Succ}(\text{Exp-EV-Int}(\Pi, \mathcal{A}, k)) \leq \mu(k)$.*

An election scheme satisfies eligibility verifiability if $\text{Succ}(\text{Exp-EV-Int}(\Pi, \mathcal{A}, k)) \leq \mu(k)$, and similarly for individual verifiability. Universal verifiability is satisfied if the election scheme satisfies Completeness and Injectivity, and $\text{Succ}(\text{Exp-UV-Int}(\Pi, \mathcal{A}, k)) \leq \mu(k)$.

C. Example—Toy schemes from digital signatures

A toy election scheme satisfying Ver-Int can be based on a digital signature scheme.³⁰ Each voter publishes their signed candidate choice on the bulletin board.

Definition 10. *Suppose $\Gamma = (\text{Gen}, \text{Sign}, \text{Ver})$ is a digital signature scheme. Let election scheme $\text{Sig}(\Gamma)$ be defined as follows:*

- $\text{Setup}(k)$ outputs $(\perp, \perp, p_1(k), p_2(k))$, where p_1 and p_2 may be any polynomial functions.
- $\text{Register}(PK_{\mathcal{T}}, k)$ outputs a key pair produced by $\text{Gen}(k)$.
- $\text{Vote}(sk, PK_{\mathcal{T}}, n_C, \beta, k)$ computes $\sigma \leftarrow \text{Sign}(sk, \beta)$ and outputs (β, σ) .
- $\text{Tally}(SK_{\mathcal{T}}, BB, L, n_C, k)$ computes a vector \mathbf{X} of length n_C , such that \mathbf{X} is a tally of all the ballots (choice-signature pairs) on BB that are signed by distinct private keys whose corresponding public keys appear in L (formally, signatures can be checked using algorithm Ver), and outputs (\mathbf{X}, \perp) .
- $\text{Verify}(PK_{\mathcal{T}}, BB, L, n_C, \mathbf{X}, P, k)$ outputs 1 if $(\mathbf{X}, P) = \text{Tally}(\perp, \perp, BB, L, n_C, \perp)$ and 0 otherwise.

Let Sig denote $\text{Sig}(\Gamma)$ for an unspecified digital signature scheme Γ satisfying strong unforgeability [7], [27].³¹ The verifiability of Sig follows from the security of the underlying signature scheme:

Proposition 6. *Sig satisfies Ver-Int.*

Proof sketch. Sig satisfies individual verifiability, because voters can verify that their signed choices appear on the bulletin board. Sig satisfies universal verifiability, because signed plaintext choices are posted on BB . Finally, Sig satisfies eligibility verifiability, because anyone can check that the signed choices belong to registered voters. \square

D. Orthogonality

Exp-IV-Int , Exp-UV-Int , and Exp-EV-Int capture mostly orthogonal security properties, as shown in Table I. Individual and universal verifiability are orthogonal, and eligibility verifiability implies individual verifiability.

30. Digital signature schemes are defined in Appendix A.

31. Strong unforgeability is defined in Appendix A.

| Line | IV | UV | EV | Scheme |
|------|----|----|----|-----------------------------------|
| 1 | X | X | X | AlwaysVerify(IgnoreCreds(Choice)) |
| 2 | X | X | ✓ | — |
| 3 | X | ✓ | X | IgnoreCreds(Choice) |
| 4 | X | ✓ | ✓ | — |
| 5 | ✓ | X | X | AlwaysVerify(IgnoreCreds(Nonce)) |
| 6 | ✓ | X | ✓ | AlwaysVerify(Sig) |
| 7 | ✓ | ✓ | X | Malleable Sig |
| 8 | ✓ | ✓ | ✓ | Sig |

TABLE I

ELECTION SCHEMES THAT SATISFY EACH COMBINATION OF INDIVIDUAL, UNIVERSAL AND ELIGIBILITY VERIFIABILITY

Theorem 7. *If an election scheme Π satisfies Exp-EV-Int, then Π also satisfies Exp-IV-Int.*

Proof sketch. If Π satisfies Exp-EV-Int, then no one can construct a ballot that appears to be associated with public credential pk unless they know private credential sk . That means that a voter can uniquely identify their ballot, because no one else knows their private credential. Therefore Π satisfies Exp-IV-Int. \square

A proof of Theorem 7 appears in Appendix G.

In Table I, `AlwaysVerify(\cdot)` is a function that transforms an election scheme by compromising `Verify` to always return 1. Thus, `AlwaysVerify(Π)` is guaranteed not to satisfy Exp-UV-Int. Similarly, `IgnoreCreds(\cdot)` is a function that accepts as input an election scheme with external authentication and returns as output an election scheme with internal authentication. The resulting scheme, however, simply ignores credentials altogether: `Register` returns (\perp, \perp) , `Vote` ignores sk , and `Tally` and `Verify` ignore L . Thus, `IgnoreCreds(Π)` is guaranteed not to satisfy Exp-EV-Int. Using those functions, we briefly explain each line of the table:

- 1) Recall (from §II-D) that `Choice` is the election scheme in which ballots contain only the plaintext candidate choice. By compromising `Verify` and ignoring credentials, we obtain a scheme that satisfies no properties.
- 2) By Theorem 7, this situation is impossible.
- 3) Compared to line 1 of Table I, this scheme satisfies Exp-UV-Int, because `Verify` is not compromised.
- 4) By Theorem 7, this situation is impossible.
- 5) `Nonce` satisfies Exp-IV-Ext and Exp-UV-Ext. Moreover, `IgnoreCreds(Nonce)` satisfies Exp-IV-Int and Exp-UV-Int. By compromising `Verify`, we obtain a scheme that satisfies only Exp-IV-Int.
- 6) `Sig` satisfies all three properties. By compromising `Verify`, we obtain a scheme that satisfies only Exp-IV-Int and Exp-EV-Int.
- 7) By making `Sig`'s underlying signature scheme malleable,³² we could obtain a scheme that does not satisfy Exp-EV-Int, because the adversary could construct a valid ballot out of a revealed ballot. But the scheme would continue to satisfy Exp-IV-Int and Exp-UV-Int.
- 8) `Sig` satisfies all three properties.

V. CASE STUDY: HELIOS-C

Helios-C [47], [48] is a variant of Helios (cf. §III) for two-candidate elections in which ballots are digitally signed.³³ Informally, Helios-C works as follows [47, §5]:

- **Setup.** As in Section III.
- **Registration.** To register a voter, the registrar generates a key pair for a signature scheme and sends the private key to the voter. After all voters are registered, the registrar publishes electoral roll L .
- **Voting.** A voter generates a ciphertext and proof as in Section III, signs the ciphertext and proof with their private key, and posts their public key, ciphertext, proof, and signature on the bulletin board.
- **Tallying.** The tallier aborts if any ballots on the bulletin board are not signed by distinct private keys whose corresponding public keys appear in L . The tallier also aborts if there exists a proof on the bulletin board that does not hold. The ciphertexts and proofs are processed as in Section III.
- **Verification.** If the tallier aborted, then a verifier immediately accepts. Otherwise, the tallier recomputes the homomorphic combination and checks all the zero-knowledge proofs, as in Section III.

Whilst analyzing Helios-C, we discovered that aborting violates our definition of universal verifiability. In particular, an adversary could post an ill-formed ballot on the bulletin board. (For example, a malicious tallier could secretly tally the recorded ballots while the election is in progress and, if that tally is unfavorable to the tallier's preferred candidate, then the tallier could post an ill-formed ballot on the bulletin board.) That ballot will cause tallying to abort. And verifiers will accept that abort. Yet, our definition of universal verifiability demands that verifiers only accept outcomes representing all the choices used to construct the recorded ballots, which aborting violates. Thus, Helios-C does not satisfy our definition of universal verifiability.³⁴

Remark 8. *Helios-C does not satisfy Ver-Int.*

Proof sketch. Helios-C aborts on errors in a manner that violates universal verifiability, as described above. \square

32. Given a message m and signature σ , a *malleable* signature scheme permits computation of a signature σ' on a related message m' [33]. The malleable signature scheme `Sig` used in line 7 of Table I would need to enable an adversary to transform a signature on a well-formed candidate β into a signature on a distinct, well-formed candidate β' .

33. Helios-C has been implemented (<https://github.com/glondou/helios-server/tree/heliosc>, released c. 2013, accessed 25 Nov 2015), but development has ceased in favour of the *Belenios* variant (<https://github.com/glondou/belenios/releases/tag/1.0>, released 22 Apr 2016, accessed 25 Apr 2016). We analyse Helios-C because a cryptographic definition has been presented in the literature, whereas *Belenios* has not appeared in the literature. (Results for one system do not imply results for the other, because the two systems are rather different. And similarly for a further variant [46] of Helios-C.)

34. Helios 2.0, Helios 3.1.4, Helios'12 and Helios'16 do not abort, so they are not similarly effected.

An informal proof of Remark 8 follows immediately from our discourse and we do not pursue a formal proof. A variant of Helios-C that disregards ill-formed ballots should satisfy our definition of universal verifiability.

Cortier et al. [47] analyzed Helios-C using a different definition of universal verifiability. That definition can be satisfied by schemes in which tallying aborts in a manner that anyone will accept. In particular, the experiment used by that definition cannot be won by an adversary that causes an abort. (As discussed above, this is undesirable, because an adversary might cause an abort when an election is unfavorable for the adversary.) Thus, verifiers accept outcomes that do not include the choices used to construct voters’ ballots. By comparison, our definition demands that verifiers reject such outcomes.

Beyond verifiability, Quaglia & Smyth [115] discovered that Helios-C does not satisfy ballot secrecy (in the presence of an adversary that controls the bulletin board or communication channel). They realised that proving correct signature construction suffices for ballot secrecy and proposed a generic construction for election schemes with internal authentication from schemes with external authentication. Moreover, they proved that their construction produces schemes satisfying ballot secrecy and verifiability.

VI. CASE STUDY: JCJ

JCJ (named for its designers, Juels, Catalano, and Jakobson) [85]–[87] is a *coercion-resistant* election scheme, meaning voters cannot prove whether or how they voted, even if they can interact with the adversary while voting, which protects elections from improper influence by adversaries. JCJ was the first scheme to achieve coercion resistance and has been influential in the design of many subsequent schemes.

To achieve verifiability and coercion resistance, JCJ uses verifiable *mixnets*, which anonymize a set of messages.³⁵ During tallying, all encrypted choices are anonymized by a mixnet, then all choices are decrypted. The tally is computed from the decrypted choices. Informally, JCJ works as follows:

- **Setup.** The tallier generates a key pair for an encryption scheme and publishes the public key.
- **Registration.** To register a voter, the registrar generates a nonce, which is sent to the voter and serves as the private credential. The public credential is computed as an encryption of the private credential with the tallier’s public key. After all voters are registered, the registrar publishes the electoral roll.
- **Voting.** A voter encrypts their candidate choice with the tallier’s public key. They also encrypts their private credential with the tallier’s public key. The voter proves in zero-knowledge that they simultaneously knows both plaintexts, and that their choice is well-formed. The voter posts their ballot (i.e., both ciphertexts and the proof) on the bulletin board.
- **Tallying.** The tallier discards any ballots from the bulletin board for which the zero-knowledge proofs do not verify. All unauthorized ballots are then discarded through a combination of protocols that includes verifiable mixnets

and *plaintext equivalence tests* (PETs) [82]. (A PET enables a proof that two ciphertexts contain the same plaintext without revealing that plaintext.) In particular, the tallier mixes the ciphertexts in the ballots (i.e., the encrypted choices and the encrypted credentials), using the same secret permutation for both mixes, hence, the mixes preserve the relation between encrypted choices and encrypted credentials. The tallier also mixes the public credentials published by the registrar. And discards any mixed encrypted choice if a PET does not hold between the corresponding encrypted credential and a mixed public credential—i.e., ballots cast using ineligible credentials are discarded. Finally, the tallier decrypts the remaining encrypted choices and publishes the corresponding tally, along with a proof that decryption was performed correctly.

- **Verification.** A verifier checks all the proofs included in ballots, and all the proofs published during tallying.

We formalize a generic construction for JCJ-like election schemes (Appendix I), which we instantiate to derive a formal description of JCJ (Appendix J). Whilst analyzing JCJ, we discovered that the mixes are insufficient for universal verifiability, because a verifier cannot distinguish between mixes that preserve the relation between encrypted choices and encrypted credentials, and mixes that do not. In particular, the proofs associated with mixes only prove a mapping between the ciphertexts input and those output. Thus, there is no proof that the relation between encrypted choices and encrypted credentials is maintained during mixing. As such, authorized ballots might be discarded in favour of unauthorized ballots, and the tally will include choices from those unauthorized ballots. Hence, universal verifiability is not satisfied. JCJ does not satisfy eligibility verifiability either, because knowledge of the tallier’s private key suffices to construct ballots that appear authentic: with the private key, any public credential can be decrypted to discover the corresponding private credential. (Note that experiment Exp-EV-Int permits an adversary to choose the tallier’s key pair, so the adversary knows the private key, hence can construct a ballot that suffices to win Exp-EV-Int.)

Proposition 9. *JCJ does not satisfy Ver-Int.*

Proof sketch. As described above, JCJ accepts tallies which exclude authorized ballots in favour of unauthorized ballots. Thus, universal verifiability is not satisfied. Moreover, an adversary can cast unauthorized ballots. Thus, eligibility verifiability is not satisfied. □

A formal proof of Proposition 9 appears in Appendix J. That proof shows that universal verifiability is not satisfied. We have reported these findings to the original authors.³⁶

We can nonetheless prove that JCJ satisfies a variant of eligibility verifiability. Consider the following experiment,

35. Chaum [34] introduced mixnets. Adida [1] surveys verifiable mixnets.
36. Dario Catalano, email communication, 30 November 2016.

which does not permit the adversary to choose the tallier’s key pair:

```

Exp-EV-Int-Weak( $\Pi, \mathcal{A}, k$ ) =
1  $(PK_{\mathcal{T}}, SK_{\mathcal{T}}, m_B, m_C) \leftarrow \text{Setup}(k)$ ;
2  $n_V \leftarrow \mathcal{A}(PK_{\mathcal{T}}, k)$ ;
3 for  $1 \leq i \leq n_V$  do  $(pk_i, sk_i) \leftarrow \text{Register}(PK_{\mathcal{T}}, k)$ ;
4  $L \leftarrow \{pk_1, \dots, pk_{n_V}\}$ ;
5  $Crpt \leftarrow \emptyset$ ;  $Rvld \leftarrow \emptyset$ ;
6  $(n_C, \beta, i, b) \leftarrow \mathcal{A}^{C,R}(L)$ ;
7 if  $\exists r : b = \text{Vote}(sk_i, PK_{\mathcal{T}}, n_C, \beta, k; r) \wedge b \neq \perp \wedge b \notin$ 
    $Rvld \wedge sk_i \notin Crpt$  then
8 | return 1
9 else
10 | return 0

```

Line 1 of Exp-EV-Int has been refactored into lines 1 and 2 of Exp-EV-Int-Weak. In line 1 of Exp-EV-Int-Weak, keys are generated by the experiment. In line 2, \mathcal{A} is given the public key but not the private key.³⁷

We propose a variant of our generic construction for JCJ-like schemes (Appendix K). That variant proves the mixes preserve the relation between encrypted choices and encrypted credentials. Using Exp-EV-Int-Weak, we define a weaker variant of Ver-Int and prove that instantiations of our construction satisfy it.

Definition 11 (Ver-Int-Weak). *An election scheme Π satisfies weak election verifiability with internal authentication (Ver-Int-Weak) if Completeness and Injectivity are satisfied and for all PPT adversaries \mathcal{A} , there exists a negligible function μ , such that for all security parameters k , we have $\text{Succ}(\text{Exp-IV-Int}(\Pi, \mathcal{A}, k)) + \text{Succ}(\text{Exp-UV-Int}(\Pi, \mathcal{A}, k)) + \text{Succ}(\text{Exp-EV-Int-Weak}(\Pi, \mathcal{A}, k)) \leq \mu(k)$.*

An election scheme satisfies *weak eligibility verifiability* if $\text{Succ}(\text{Exp-EV-Int-Weak}(\Pi, \mathcal{A}, k)) \leq \mu(k)$.

Let JCJ’16 be the set of election schemes derived from the variant of our generic construction, assuming cryptographic primitives satisfy certain properties that we identify.³⁸

Theorem 10. *JCJ’16 satisfies Ver-Int-Weak.*

Proof sketch. JCJ’16 satisfies individual verifiability, because the probabilistic encryption scheme ensures that ballots are unique, with overwhelming probability. JCJ’16 satisfies universal verifiability, because the proofs produced throughout tallying can be publicly verified. And JCJ’16 satisfies eligibility verifiability, because \mathcal{A} cannot construct new ballots without knowing a voter’s private credential or the tallier’s private key. \square

A formal proof of Theorem 10 appears in Appendix K. The proof assumes the random oracle model.

The Civitas [43] scheme refines the JCJ scheme. Some refinements relevant to election verifiability are an implementation of a distributed registration protocol, and a mixnet based on randomized partial checking (RPC) [83]. We leave a proof that Civitas satisfies Ver-Int-Weak as future work. In that

proof, it would be necessary to assume the RPC construction satisfies the definition of mixnets given in the appendix. Work by Khazaei and Wikström [89] suggests that actually proving satisfaction is unlikely to be possible. Alternatively, the mixnet could be replaced by one based on zero-knowledge proofs [68], [106].

VII. COMPARISON WITH GLOBAL VERIFIABILITY

Küsters et al. [97], [98], [100] present a definition of global verifiability that can be used with any kind of protocol, not just electronic voting protocols. To analyze the verifiability of a protocol, analysts must define *goals*, which are properties required to hold in runs of the protocol. For example, a goal γ_ℓ is presented in a case study [98, §5.2] of global verifiability applied to voting:

γ_ℓ contains all runs for which there exist choices of the dishonest voters (where a choice is either to abstain or to vote for one of the candidates) such that the result obtained together with the choices made by the honest voters in this run differs only by ℓ votes from the published result (i.e. the result that can be computed from the simple ballots on the bulletin board).

Another goal γ is presented in a case study [100, §6.2] of Helios:

γ is satisfied in a run if the published result exactly reflects the actual votes of the honest voters in this run and votes of dishonest voters are distributed in some way on the candidates, possibly in a different way than how the dishonest voters actually voted.

These informal statements of goals are appealing, but they do not constitute rigorous mathematical definitions. As Kiayias et al. write, “[global verifiability] has the disadvantage that the set γ remains undetermined and thus the level of verifiability that is offered by the definition hinges on the proper definition of γ which may not be simple” [91, p. 476].

In our own work, we found that formal definitions were quite tricky to get right—for example, which ballots should be counted, how to count them, and how to determine whether that count differed from the published tally. So we shared³⁹ and discussed⁴⁰ our results with Küsters. In response, Küsters et al. updated their technical report to propose a formal goal [101, §5.2]. In essence, that goal is satisfied in a run if choices $\beta_1, \dots, \beta_{n_h}$ of honest voters are included in the tally and the tally contains at most $n_h + n_d$ choices, where n_d is the number of dishonest voters. We found that Helios’16 and Nonce do not satisfy global verifiability with that goal, because the goal requires: 1) participation of all voters, 2) ballot posting to always succeed, and 3) bulletin boards not to drop, inject

37. Exp-EV-Int-Weak can be equivalently formulated as an experiment with one registered voter. See Appendix H for details.

38. A set of election schemes satisfies Ver-Int-Weak, if every scheme in the set satisfies Ver-Int-Weak.

39. Ralf Küsters, email communication, 24 June 2014.

40. Ralf Küsters, email communication, October/November 2014.

nor modify ballots. The first and second requirements define availability properties, which an adversary can disrupt. And the third can be disrupted by an adversary that controls the bulletin board. Thus, there exist runs of both Helios’16 and Nonce that cannot satisfy this goal. We defer definitions of global verifiability and the goal by Küsters et al. to Appendix N, and formal results to Appendix O, because the above discussion can be appreciated without the burden of technical details.

Cortier et al. [49, §10.2] propose a variant of the goal by Küsters et al. [101, §5.2]. Their goal is informally claimed to permit some honest voters’ choices to be dropped from the tally, which would intuitively address problems associated with the third requirement. However, this claim is not supported by their formally stated goal, because the goal requires the tally to include $n_h + n_d$ choices, where n_h , respectively n_d , is the number of honest, respectively dishonest, voters. Thus, the goals by Cortier et al. and Küsters et al. have similar drawbacks. We omit recalling further details, because the ideas remain the same. We reported our findings to Cortier et al. and Küsters et al.,⁴¹ but they did not respond. We reported our findings again,⁴² which resulted in confirmation of the error,⁴³ but no fix is yet public.

It is natural to ask whether individual, universal and eligibility verifiability can each be expressed in terms of global verifiability. We believe they can. For instance, they could be expressed, in the informal style of the goals quoted above, as the following goals:

- G_{IV} is satisfied in a run if voters can uniquely identify their ballots on the bulletin board in this run.
- G_{UV} is satisfied in a run if the correct tally of votes cast by authorized voters in this run is the same as the tally that algorithm `Verify` successfully verifies.
- G_{EV} is satisfied in a run if every ballot tallied in this run was created by a voter in possession of a private credential.

Cortier et al. [49], [50] have also expressed goals intended to capture our definitions of individual and universal verifiability. We discuss their work in Section IX.

It is also natural to ask whether election verifiability can be expressed in terms of global verifiability using a single, holistic goal. Indeed, roughly speaking, it can. We introduce a goal δ_{GV} that is satisfied in a run if ballots b_1, \dots, b_n for choices β_1, \dots, β_n appear in the run, such that b_1, \dots, b_n are included on the bulletin board and no further ballots are included, and the run produces a tally for choices β_1, \dots, β_n . We show election verifiability implies global verifiability with that goal. (Hence, Helios’16 and Nonce satisfy global verifiability using goal δ_{GV} .) We also show that global verifiability implies universal verifiability, but not individual verifiability, with that goal. It might seem surprising that individual verifiability is not implied, but this is a consequence of a technical detail. In particular, given a goal defining some properties, global verifiability only requires those properties to hold on runs in which an auditor (or judge) accepts.⁴⁴ Thus, such properties need not hold on runs in which an auditor rejects. Yet, this

does not matter, because auditing suffices to detect problems. To summarise:

- Election verifiability and global verifiability, using goal δ_{GV} , both guarantee that anyone can check whether the tally is properly computed.
- Election verifiability guarantees that collisions can be detected on every run of a protocol, whereas global verifiability using goal δ_{GV} only guarantees that collisions can be detected on runs in which an auditor accepts.

Thus, election verifiability is strictly stronger than global verifiability using goal δ_{GV} . We defer formal results to Appendix P. It is an open problem as to whether election verifiability coincides with global verifiability for some other goal.

One concern that might be raised is whether there still lurk any “gaps” in our decomposition into individual and universal (and eligibility) verifiability. Indeed, there might be. But the definition of global verifiability does not rule out the possibility of gaps, either: any gap in the formal statement of a goal will lead to a vulnerability. That is, if the analyst forgets to include some necessary facet of verifiability when stating the formal goal, then global verifiability will not detect any attacks against that facet. Indeed, Cortier et al. [49, §1] state that some goals have “severe limitations and weaknesses.” Global verifiability does not guarantee a lack of gaps. Although we cannot guarantee the absence of gaps either, we have proved a relationship between election and global verifiability. So, any gap in our definition implies the existence of a gap in the definition of global verifiability using goal δ_{GV} .

VIII. NEW CLASSES OF ATTACK

Our definitions of election verifiability improve upon existing definitions by detecting three previously unidentified classes of attack:

- *Collusion attacks.* An election scheme’s tallying and verification algorithms might be designed such that they collude to accept incorrect tallies.
- *Biasing attacks.* An election scheme’s verification algorithm might be designed to reject some legitimate tallies.
- *Revelation attacks.* An election scheme’s verification algorithm might be designed to accept incorrect tallies when coins used to construct some ballots are leaked.

Although a well-designed election scheme would hopefully not exhibit vulnerabilities to these attacks, it is the job of verifiability definitions to detect malicious schemes, regardless of whether vulnerabilities are due to malice or errors. So definitions of election verifiability should preclude them.

41. Veronique Cortier, David Galindo, Ralf Küsters, Johannes Müller, Tomasz Truderung, & Andreas Vogt, email communication, 18 Oct 2016.

42. Ralf Küsters & Johannes Müller, email communication, 25 Apr 2018.

43. Johannes Müller, email communication, 22 May 2018.

44. In the context of universal verifiability, an auditor accepts when they are satisfied that the tally of recorded ballots is computed properly.

A. Collusion Attacks

Here are two examples of potential collusion attacks:

- **Vote stuffing.** Tally behaves normally, but adds κ votes for candidate β . Verify subtracts κ votes from β , then proceeds with verification as normal. Elections thus verify as normal, except that candidate β receives extra votes.
- **Backdoor tally replacement.** Tally and Verify behave normally, unless a *backdoor* value is posted on the bulletin board BB . For example, if $(SK_{\mathcal{T}}, \mathbf{X}^*)$ appears on BB , then Tally and Verify both ignore the correct tally and instead replace it with tally \mathbf{X}^* . Value $SK_{\mathcal{T}}$ is the backdoor here; it cannot appear on BB (except with negligible probability) unless the tallier is malicious.

Vote stuffing is detected by our definitions of Correctness (§II-A and §IV-A), because these definitions require that the tally produced by Tally corresponds to the choices encapsulated in ballots on the bulletin board. Note that vote stuffing is not a failure of eligibility verifiability, because the stuffed votes do not correspond to any ballots on the bulletin board. Backdoor tally replacement is detected by our definitions of universal verifiability (§II-B2 and §IV-B2), because those definitions require Verify to accept only those tallies that correspond to a correct tally of the bulletin board.

We show, next, that the definition of election verifiability by Juels et al. [87] fails to detect vote stuffing and backdoor tally replacement, and that the definition by Cortier et al. [47] fails to detect backdoor tally replacement.

Juels et al. [87] formalize definitions that we name *JCJ-correctness* and *JCJ-verifiability*. JCJ-correctness is intuitively meant to capture that “ \mathcal{A} cannot pre-empt, alter, or cancel the votes of honest voters [and] that \mathcal{A} cannot cause voters to cast ballots resulting in double voting” [87, p. 45]; it is formalized in terms of whether the adversary can post ballots on the bulletin board that cause the tally to be computed incorrectly. JCJ-verifiability is intuitively “the ability for any player to check whether the tally...has been correctly computed” [87, p. 46]; it is formalized in terms of whether Verify will accept a tally that differs from the output of Tally. We restate the formal definitions in Appendix L.

To show that the JCJ definitions fail to detect collusion attacks, we first formalize the vote stuffing attack. An election scheme $\Pi = (\dots, \text{Tally}, \text{Verify})$ can be modified to derive a vote-stuffing election scheme $\text{Stuff}(\Pi, \beta, \kappa) = (\dots, \text{Tally}_S, \text{Verify}_S)$, which adds κ votes to β , as follows:⁴⁵

- $\text{Tally}_S(\dots)$ computes $(\mathbf{X}, P) \leftarrow \text{Tally}(\dots)$ and outputs $(\text{Add}(\mathbf{X}, \beta, \kappa), P)$.⁴⁶
- $\text{Verify}_S(\dots, \mathbf{X}, P, \dots)$ computes $v \leftarrow \text{Verify}(\dots, \text{Add}(\mathbf{X}, \beta, -\kappa), P, \dots)$ and outputs v .

Proposition 11. *Given an election scheme Π satisfying JCJ-correctness and JCJ-verifiability, and integers β and κ , it holds that $\text{Stuff}(\Pi, \beta, \kappa)$ satisfies JCJ-correctness and JCJ-verifiability.*

A formal proof appears in Appendix M.

Second, we formalize the backdoor tally replacement attack. An election scheme $\Pi = (\dots, \text{Tally}, \text{Verify})$ can be modified to derive an election scheme $\text{Backdoor}(\Pi) = (\dots, \text{Tally}_B, \text{Verify}_B)$, as follows:

- $\text{Tally}_B(\dots, SK_{\mathcal{T}}, BB, \dots)$ checks whether there exists a unique \mathbf{X}^* such that $(SK_{\mathcal{T}}, \mathbf{X}^*) \in BB$. If so, Tally_B outputs (\mathbf{X}^*, \perp) . Otherwise, Tally_B computes $(\mathbf{X}, P) \leftarrow \text{Tally}(\dots, SK_{\mathcal{T}}, BB, \dots)$ and outputs (\mathbf{X}, P) .
- $\text{Verify}_B(PK_{\mathcal{T}}, \dots, BB, \dots, \mathbf{X}, P, \dots)$ checks whether there exists a unique \mathbf{X}^* such that $(SK_{\mathcal{T}}, \mathbf{X}^*) \in BB$.⁴⁷ If so, Verify_B outputs 1. Otherwise, Verify_B computes $v \leftarrow \text{Verify}(PK_{\mathcal{T}}, \dots, BB, \dots, \mathbf{X}, P, \dots)$ and outputs v .

Proposition 12. *Given an election scheme Π satisfying JCJ-correctness and JCJ-verifiability that does not leak the tallier’s private key, it holds that $\text{Backdoor}(\Pi)$ satisfies JCJ-correctness and JCJ-verifiability.*

A formal proof appears in Appendix M, where we also formally define key leakage.

Cortier et al. [47] propose definitions similar to *JCJ-verifiability* and insist that election schemes must satisfy their notions of correctness and partial tallying. Vote stuffing is detected by their correctness property, but backdoor tally replacement is not. The ideas remain the same, so we omit formalized results. We have reported these findings to the original authors.^{48,49}

B. Biasing attacks

Here are three formalizations of biasing attacks, derived from an election scheme $\Pi = (\dots, \text{Verify})$.

- **Reject All.** Let $\text{Reject}(\Pi)$ be (\dots, Verify_R) , where Verify_R always outputs 0. Verify_R therefore always rejects, hence no election can ever be considered valid.
- **Selective Reject.** Let ε be a distinguished value that would not be posted on the bulletin board by honest voters. Let $\text{Selective}(\Pi, \varepsilon)$ be (\dots, Verify_R) , where $\text{Verify}_R(\dots, BB, \dots)$ computes $v \leftarrow \text{Verify}(\dots, BB, \dots)$ and outputs 1 if both $v = 1$ and $\varepsilon \notin BB$. Otherwise, Verify_R outputs 0. Verify_R therefore rejects if ε appears on the bulletin board, hence some elections can be invalidated.
- **Biased Reject.** Suppose Z is a set of tallies. Let $\text{Bias}(\Pi, Z)$ be (\dots, Verify_R) , where $\text{Verify}_R(\dots, \mathbf{X}, \dots)$ computes $v \leftarrow \text{Verify}(\dots, \mathbf{X}, \dots)$ and outputs 1 if both $v = 1$ and $\mathbf{X} \in Z$. Otherwise, Verify_R outputs 0. Verify_R

45. We omit many of the parameters of Tally and Verify here for simplicity; see Appendix M for details.

46. Let $\text{Add}(\mathbf{X}, \beta, \kappa) = (\mathbf{X}[1], \dots, \mathbf{X}[\beta - 1], \mathbf{X}[\beta] + \kappa, \mathbf{X}[\beta + 1], \dots, \mathbf{X}[|\mathbf{X}|])$. And let $|\mathbf{X}|$ denote the length of vector \mathbf{X} .

47. Verify_B also needs to check that $SK_{\mathcal{T}}$ is the private key corresponding to $PK_{\mathcal{T}}$. We omit formalizing this detail, but note that it is straightforward for real-world encryption schemes such as El Gamal and RSA.

48. Véronique Cortier and David Galindo, personal communication, Nancy, France, 13 June 2013.

49. David Galindo and Véronique Cortier, email communication, 19 June 2013 & Summer/Autumn 2014.

therefore only accepts a subset of the tallies accepted by Verify, hence biases tallies toward Z .

These formalizations do not satisfy our definitions of Completeness (§II-B2 and §IV-B2), hence, our definitions of verifiability detect these biasing attacks.

The definition of verifiability by Juels et al. [87] fails to detect all three of the above attacks, because that definition has no notion of Completeness. For example, it is vulnerable to Biased Reject attacks:

Proposition 13. *Given an election scheme Π satisfying JCJ-correctness and JCJ-verifiability, and given a multiset Z , it holds that $\text{Bias}(\Pi, Z)$ satisfies JCJ-correctness and JCJ-verifiability.*

A formal proof appears in Appendix M.

The definition of verifiability by Kiayias et al. [91] fails to detect Selective Reject attacks, because (like JCJ) the definition has no notion of Completeness. Their notion of Correctness does rule out Reject All and Biased Reject attacks.

Similarly, the definition of verifiability by Cortier et al. [47] detects Biased Reject and Reject All attacks, but fails to detect Selective Reject attacks, because that definition’s notion of Completeness does not quantify over all bulletin boards.

C. Revelation attacks

Here are two formalizations of revelation attacks, derived from an election scheme $\Pi = (\dots, \text{Verify})$ with ballots that do not leak coins.

- **Replace choices.** Let $\text{Replace}(\Pi)$ be (\dots, Verify_R) , where $\text{Verify}_R(PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, P, k)$ proceeds as follows. The algorithm checks whether $BB = \{b_1, \dots, b_\ell, (\beta_1, \beta'_1, r_1), \dots, (\beta_k, \beta'_k, r_k)\}$ such that $\bigwedge_{1 \leq i \leq k} b_i = \text{Vote}(PK_{\mathcal{T}}, n_C, \beta_i, k; r_i) \wedge 1 \leq \beta_i, \beta'_i \leq n_C$. If so, the algorithm computes $v \leftarrow \text{Verify}(PK_{\mathcal{T}}, BB, n_C, \mathbf{X}^*, P, k)$, where tally \mathbf{X}^* is derived from \mathbf{X} by replacing choices $\beta'_1, \dots, \beta'_k$ with β_1, \dots, β_k . Otherwise, the algorithm computes $v \leftarrow \text{Verify}(PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, P, k)$. Finally, the algorithm outputs v .
- **Drop choices.** Let Drop be a variant of Replace that derives tally \mathbf{X}^* from \mathbf{X} by adding choices β_1, \dots, β_k .

These revelation attacks do not satisfy our definitions of universal verifiability (§II-B2 and §IV-B2), because the adversary constructs the ballots posted on the bulletin board, hence, can also post the coins used to construct those ballots. Similarly, these attacks do not satisfy global verifiability instantiated with goal δ_{GV} .

Global verifiability fails to detect the above attacks when instantiated with the goal by Küsters et al. [101], because coins are implicitly assumed never to leak, even when the software, hardware, voter, etc., that selected those coins has the ability to leak them. Consequently, voters may verify that their correctly constructed ballot has been recorded, yet their vote can be excluded from the tally. We defer a formal result to Appendix O, where we also formally define coin leakage. Global verifiability fails similarly when instantiated with the goal by Cortier et al. [49].

IX. RELATED WORK

Kiayias [90] & Schoenmakers [121] present overviews of security properties for election schemes. Many election schemes in the literature state properties called correctness, accuracy, or (universal) verifiability without formally defining those terms.

In the computational model, Juels et al. [85]–[87] and Cortier et al. [47] give game-based definitions of verifiability. Those definitions fail to detect biasing and collusion attacks (cf. §VIII). Definitions of universal verifiability (which is just one aspect of election verifiability) in the computational model seem to originate with Benaloh and Tuinstra [17], who define a *correctness* property that says every participant is convinced that the tally is accurate with respect to the votes cast, and with Cohen and Fischer [44], who define *verifiability* to mean that there exists a *check* function that returns *good* iff the announced tally of the election corresponds to the cast votes.

Kiayias et al. [91] define a property they name *E2E verifiability* (E2E abbreviates “end-to-end”). This property combines our intuitive notions of individual and universal verifiability into a single definition. Their definition fails to detect Selective Reject attacks (cf. §VIII). Their definitions, like ours, do not address voter intent—that is, verification by humans that ballots correctly encode candidate choices—as we discuss in Section X.

Cortier et al. [49], [50] survey definitions of verifiability and cast them into the context of global verifiability. In particular, they express goals intended to capture definitions of verifiability by Cohen and Fischer [14], [44], Kiayias et al. [91], and Cortier et al. [47]. They also express goals intended to capture our definitions of individual and universal verifiability. Using these goals, Cortier et al. compare different notions of verifiability.

Cortier et al. [49, §8.5 & §10.1] claim that our definition of election verifiability admits an election scheme which it should not: the election scheme in which “Vote always [outputs error symbol \perp] for some dishonestly generated public key [and Tally behaves normally].” We believe our definition *should* admit this scheme, because it *is* verifiable. Indeed, ballot construction will result in an error, alerting voters to malice. Cortier et al. [49, §10.1] also claim that we trust the bulletin board and assume all voters will run the correct Vote algorithm, we do not (cf. §II-B1 and §II-B2).

Küsters & Müller claim “it is often believed that individual [verifiability] together with universal verifiability implies [global] verifiability...However, [we] have demonstrated that individual and universal verifiability are neither sufficient nor necessary for [global verifiability].” They state their claim shortly after an explicit reference to our definitions of individual and universal verifiability [95, §2.2]. Yet, those definitions are proven to be strictly stronger than global verifiability, which seemingly contradicts their claim. We contacted Küsters & Müller for clarification.⁵⁰ They stated that their claim only holds for the goal by Küsters et al. [101, §5.2] and the goal

50. Ralf Küsters & Johannes Müller, email communication, 25 April 2018.

they proposed in collaboration with Cortier et al. [49, §10.2].⁵¹ But, those goals are uninteresting, since they omit attacks (§VIII-C).

Also in the computational model, Groth [74], and Moran and Naor [105], state definitions of verifiability in terms of *universal composability* [31]. These definitions involve defining an *ideal functionality*; part of that is similar to our *correct-tally* function. Groth’s definition does not guarantee universal verifiability [74, p. 2], but Moran and Naor’s does [105, p. 386].

In the symbolic model, Smyth et al. [139] define the first definition of election verifiability. This definition is amenable to automated reasoning, but is stronger than necessary and cannot be satisfied by many election schemes, including Helios and Civitas. Kremer et al. [94] overcome this limitation with a weaker definition that sacrifices amenability to automated reasoning, and Smyth [123, §3] extends this definition. Additionally, the scope of automated reasoning, using the definition by Smyth et al., is limited by analysis tools (e.g., ProVerif [26]), because the function symbols and equational theory used to model cryptographic primitives might not be suitable for automated analysis (cf. [8], [61], [110], [130]). Cortier et al. [45] overcome this limitation with an alternative definition based on refinement type systems.

Also in the symbolic model, Kremer and Ryan [93] and Backes et al. [9] formalize definitions of *eligibility*. These definitions are not intended to provide assurances if the election authorities are dishonest (cf. [104, §1]). For example, the definition of Kremer and Ryan does not detect whether corrupt election authorities insert votes [93, §5.2]. Likewise, the definition of Backes et al. assumes that election authorities are honest [9, §3].

Our definition of election verifiability has been adapted to auction schemes by Quaglia & Smyth [116]. And the definition of election verifiability by Kremer et al. [94] has been adapted to auction [63] and examination [62], [64] schemes. Moreover, McCarthy et al. [103] have shown that auction schemes can be constructed from Helios and JCJ. Thus, our results are applicable beyond voting.

Our definition of election verifiability follows Smyth et al. [94], [123], [139] by deconstructing it into individual, universal, and eligibility verifiability. Other deconstructions of election verifiability are possible. For example, Adida and Neff [6, §2] identify four aspects of verifiability:

- *Cast as intended*: the ballot is cast at the polling station as the voter intended.
- *Recorded as cast*: cast ballots are preserved with integrity through the ballot collection process.
- *Counted as recorded*: recorded ballots are counted correctly.
- *Eligible voter verification*: only eligible voters can cast a ballot in the first place.

Those definitions are not mathematical, so we cannot attempt a precise comparison. Nonetheless, eligibility verifiability and eligible voter verification seem to be addressing similar concerns. Likewise, individual and universal verifiability together

seem to be addressing concerns similar to that of recorded as cast and counted as recorded together. We postpone a discussion of cast as intended to Section X.

Privacy properties [60], [87], [98], [99], [126], [131], [133]—such as ballot secrecy,⁵² receipt freeness, and coercion resistance—complement verifiability. Chevallier-Mames et al. [40], [41] and Hosp and Vora [80], [81] show an incompatibility result: election schemes cannot unconditionally satisfy privacy and universal verifiability. But weaker versions of these properties can hold simultaneously, as can be witnessed from Theorems 5 and 10 coupled with existing privacy results such as the ballot secrecy proofs for Helios’12 [23, Theorem 3], [20, Theorem 6.12], and the coercion resistance proof for JCJ [87, §5].

Cortier & Lallemand claim privacy implies individual verifiability [51]. But, they assume a trusted tallier. For privacy, this assumption is necessary to ensure ballots cannot be tallied individually, which would reveal votes. By comparison, the assumption is counter-intuitive for individual verifiability, because attacks by malicious talliers must be detected. Our definition of individual verifiability detects such attacks and Smyth proves it is not implied by privacy [129, Appendix C].

In an analysis of Helios, Küsters et al. [100] use goal γ to conclude that global verifiability is satisfied. Yet Bernhard et al. [23] and Chang-Fong & Essex [32] demonstrate vulnerabilities against verifiability, and in Appendix E we show that Ver-Ext detects these vulnerabilities. This seeming discrepancy arises because the analysis in [100] does not formalize all the cryptographic primitives used by Helios, hence the vulnerabilities go unnoticed. So another contribution of our own work is to correctly distinguish between unverifiable and verifiable variants of Helios by rigorously analyzing the cryptography used in Helios.

X. CONCLUDING REMARKS

When we began this work, we were studying the Juels et al. [87] definition of election verifiability. We discovered that the definition fails to detect biasing and collusion attacks. While attempting to improve the Juels et al. definition to rule out those attacks, we discovered that factoring it into individual, universal, and eligibility verifiability led to an elegant decomposition of (mostly) orthogonal properties. We later sought to apply our new definitions to existing electronic voting systems, and Helios [5] and JCJ [87] were natural choices. But they treat authentication differently—Helios out-sources authentication, whereas JCJ does not—so we were led to separate our definitions into variants for external and internal authentication. We were at first surprised to discover that JCJ does not satisfy the strong definition of eligibility verifiability. But upon reflection, it became apparent that an adversary who knows the tallier’s private key can easily forge

51. Ralf Küsters & Johannes Müller, email communication, 22 May 2018.

52. Quaglia & Smyth [114] and Smyth [127] provide overviews of ballot-secrecy definitions and provide comparisons between definitions. Smyth [126, §7] and Bernhard et al. [21], [22] provide more detailed comparisons.

ballots that appear to be from eligible voters. Helios-C [47], however, avoids this problem by employing digital signatures.

Our definitions of verifiability have not addressed the issue of voter intent—that is, verification by a human that the ballot submitted by a voter corresponds to the candidate choice the voter intended to make. Adida and Neff call this property “cast as intended” [6]. Many election schemes (e.g., [67], [79], [87], [91]) do not satisfy cast as intended, because the schemes implicitly or explicitly assume that voters can themselves verify the cryptographic operations required to construct ballots. Nevertheless, schemes by Chaum [35], Neff [107], and Benaloh [15], [16] introduce cryptographic mechanisms to verify voter intent. It would be natural to explore strengthening our definitions to address voter intent.

The goal of this research is to enable verifiability of the voting systems we use in real-life, rather than merely trusting them. Research on verifiability can generalize beyond voting to other systems that must guarantee strong forms of integrity. Verifiable voting systems thus have the potential to contribute to the science of security, to democracy, and to broader society.

ACKNOWLEDGMENTS

We thank David Bernhard, Dario Catalano, Jeremy Clark, Véronique Cortier, Aleksander Essex, David Galindo, Stéphane Glondu, Markus Jakobsson, Steve Kremer, Ralf Küsters, Elizabeth Quaglia, Mark Ryan, Susan Thomson, and Poorvi Vora for insightful discussions that have influenced this paper. This work is partly supported by the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013) / ERC project *CRYSP* (259639), by AFOSR grants FA9550-12-1-0334 and FA9550-14-1-0334, by NSF grant 1421373, by the National Security Agency, and the Luxembourg National Research Fund (FNR) under the FNR-INTER-VoteVerif project. This work was performed in part at George Washington University and at INRIA.

DEDICATION⁵³

Ben Smyth dedicates his contribution to the loving memory of Anne Konishi, 1971 – 2015. What matters most of all is the dash. We had a great time.

He writes for Christina Mai Konishi. Smile like your mother, for good fortune seeks those who smile (*warau kado niwa fuku kitaru*, says the Japanese proverb).

APPENDIX A CRYPTOGRAPHIC PRIMITIVES

A. Basic definitions

Definition 12 (Negligible function [70]). *A function $\mu : \mathbb{N} \rightarrow \mathbb{R}$ is negligible if for every positive polynomial function $p(\cdot)$, there exists an N , such that for all $n > N$,*

$$\mu(n) < \frac{1}{p(n)}.$$

An event $E(k)$, where k is a security parameter, occurs with *negligible probability* if $\Pr[E(k)] \leq \mu(k)$ for some negligible

function μ . The event occurs with *overwhelming probability* if the complement of the event occurs with negligible probability.

Definition 13 (Asymmetric encryption scheme [88]). *An asymmetric encryption scheme is a tuple of PPT algorithms $(\text{Gen}, \text{Enc}, \text{Dec})$ such that:*

- **Gen**, denoted $(pk, sk, \mathfrak{m}) \leftarrow \text{Gen}(k)$, takes a security parameter k as input and outputs a key pair (pk, sk) and message space \mathfrak{m} .
- **Enc**, denoted $c \leftarrow \text{Enc}(pk, m)$, takes a public key pk and message $m \in \mathfrak{m}$ as input, and outputs a ciphertext c .
- **Dec**, denoted $m \leftarrow \text{Dec}(sk, c)$, takes a private key sk , and ciphertext c as input, and outputs a message m or error symbol \perp . We assume $\perp \notin \mathfrak{m}$ and Dec is deterministic.

Moreover, the scheme must be correct: there exists a negligible function μ , such that for all security parameters k and messages m , we have $\Pr[(pk, sk, \mathfrak{m}) \leftarrow \text{Gen}(k); c \leftarrow \text{Enc}(pk, m) : m \in \mathfrak{m} \Rightarrow \text{Dec}(sk, c) = m] > 1 - \mu(k)$.

Our definition of asymmetric encryption schemes differs from Katz and Lindell’s definition [88, Definition 10.1] in that we formally state the plaintext space.

Definition 14 (Homomorphic encryption). *An asymmetric encryption scheme $\Gamma = (\text{Gen}, \text{Enc}, \text{Dec})$ is homomorphic, with respect to ternary operators \odot , \oplus , and \otimes ,⁵⁴ if there exists a negligible function μ , such that for all security parameters k , we have the following.⁵⁵ First, for all messages m_1 and m_2 we have $\Pr[(pk, sk, \mathfrak{m}) \leftarrow \text{Gen}(k); c_1 \leftarrow \text{Enc}(pk, m_1); c_2 \leftarrow \text{Enc}(pk, m_2) : m_1, m_2 \in \mathfrak{m} \Rightarrow \text{Dec}(sk, c_1 \otimes_{pk} c_2) = \text{Dec}(sk, c_1) \odot_{pk} \text{Dec}(sk, c_2)] > 1 - \mu(k)$. Secondly, for all messages m_1 and m_2 , and coins r_1 and r_2 , we have $\Pr[(pk, sk, \mathfrak{m}) \leftarrow \text{Gen}(k) : m_1, m_2 \in \mathfrak{m} \Rightarrow \text{Enc}(pk, m_1; r_1) \otimes_{pk} \text{Enc}(pk, m_2; r_2) = \text{Enc}(pk, m_1 \odot_{pk} m_2; r_1 \oplus_{pk} r_2)] > 1 - \mu(k)$.*

We say Γ is additively homomorphic, respectively multiplicatively homomorphic, if for all security parameters k , key pairs pk, sk , and message spaces \mathfrak{m} , such that there exists coins r and $(pk, sk, \mathfrak{m}) = \text{Gen}(k; r)$, we have \odot_{pk} is the addition operator, respectively multiplication operator, in group $(\mathfrak{m}, \odot_{pk})$.

Indistinguishability under chosen-plaintext attack (IND-CPA) [10], [12], [13], [71], [72] is a standard definition of security for encryption schemes. Intuitively, if an encryption scheme satisfies IND-CPA, then an adversary without access to a decryption oracle is unable to distinguish ciphertexts. A variant (IND-PA0) allows the adversary a *parallel decryption query*—i.e., it requests the decryption of a vector of ciphertexts.

53. The dedication references Linda Ellis (1996) *The Dash*.

54. We shall implicitly bind ternary operators occasionally—i.e., we write Γ is a homomorphic asymmetric encryption scheme as opposed to the more verbose Γ is a homomorphic asymmetric encryption scheme, with respect to ternary operators \odot , \oplus , and \otimes .

55. We write $X \circ_{pk} Y$ for the application of ternary operator \circ to inputs X , Y , and pk . We occasionally abbreviate $X \circ_{pk} Y$ as $X \circ Y$, when pk is clear from the context.

Definition 15 (IND-PA0 [12]). An asymmetric encryption scheme satisfies IND-PA0, if for all probabilistic polynomial-time adversaries \mathcal{A} , there exists a negligible function μ , such that for all security parameters k , we have $\text{Succ}(\text{IND-PA0}(\Pi, \mathcal{A}, k)) \leq 1/2 + \mu(k)$, where experiment IND-PA0 is defined as follows:⁵⁶

```

IND-PA0( $\Pi, \mathcal{A}, k$ ) =
1  $(pk, sk, m) \leftarrow \text{Gen}(k)$ ;
2  $(m_0, m_1) \leftarrow \mathcal{A}(pk, m, k)$ ;
3  $\beta \leftarrow_R \{0, 1\}$ ;
4  $c \leftarrow \text{Enc}(pk, m_\beta)$ ;
5  $\mathbf{c} \leftarrow \mathcal{A}(c)$ ;
6  $\mathbf{m} \leftarrow (\text{Dec}(sk, \mathbf{c}[1]), \dots, \text{Dec}(sk, \mathbf{c}[|\mathbf{c}|]))$ ;
7  $g \leftarrow \mathcal{A}(\mathbf{m})$ ;
8 return  $g = \beta \wedge \bigwedge_{1 \leq i \leq |\mathbf{c}|} c \neq \mathbf{c}[i] \wedge m_0, m_1 \in \mathbf{m} \wedge |m_0| = |m_1|$ ;

```

Definition 16 (Signature scheme [88]). A signature scheme is a tuple $(\text{Gen}, \text{Sign}, \text{Ver})$ of PPT algorithms such that:

- **Gen**, denoted $(pk, sk) \leftarrow \text{Gen}(k)$, takes a security parameter k as input and outputs a key pair (pk, sk) .
- **Sign**, denoted $\sigma \leftarrow \text{Sign}(sk, m)$, takes a private key sk and message m as input, and outputs a signature σ .
- **Verify**, denoted $v \leftarrow \text{Ver}(pk, m, \sigma)$, takes a public key pk , message m , and signature σ as input, and outputs a bit v , which is 1 if the signature successfully verifies and 0 otherwise. We assume Ver is deterministic.

Moreover, the scheme must be correct: there exists a negligible function μ , such that for all security parameters k and messages m , we have $\Pr[(pk, sk) \leftarrow \text{Gen}(k); \sigma \leftarrow \text{Sign}(sk, m); \text{Ver}(pk, m, \sigma) = 1] > 1 - \mu(k)$.

Definition 17. A signature scheme $\Gamma = (\text{Gen}, \text{Sign}, \text{Ver})$ satisfies strong unforgeability if for all PPT adversaries \mathcal{A} , there exists a negligible function μ , such that for all security parameters k , we have $\text{Succ}(\text{Exp-StrongSign}(\Gamma, \mathcal{A}, k)) \leq \mu(k)$, where experiment Exp-StrongSign is defined as follows:

```

Exp-StrongSign( $\Gamma, \mathcal{A}, k$ ) =
1  $(pk, sk) \leftarrow \text{Gen}(k)$ ;
2  $Msg \leftarrow \emptyset$ ;
3  $(m, \sigma) \leftarrow \mathcal{A}^\mathcal{O}(pk, k)$ ;
4 if  $\text{Ver}(pk, m, \sigma) = 1 \wedge (m, \sigma) \notin Msg$  then
5 |   return 1
6 else
7 |   return 0

```

The experiment defines an oracle \mathcal{O} .⁵⁷ On invocation $\mathcal{O}(m)$, oracle \mathcal{O} computes a signature $\sigma \leftarrow \text{Sign}(sk, m)$, records the request and response (m, σ) by updating Msg to be $Msg \cup \{(m, \sigma)\}$, and outputs σ .

B. Proof systems

A proof system (originally known as an interactive proof system [73]) is a two-party protocol between a prover and a verifier. The prover convinces the verifier that a string x is in

a language L . Here, we assume that there is a witness relation R , such that $s \in L$ iff there exists a witness w , such that $(s, w) \in R$. For any $(s, w) \in R$, it must also hold that the length of w is at most polynomial in the length of s . Proof systems ensure that a prover can convince a verifier of any valid claim (completeness), and that a verifier cannot be fooled into accepting a false claim (soundness).

A sigma protocol [28], [58], [77], [120] is a proof system with a particular three-move structure: commit, challenge, respond.

Definition 18 (Sigma protocol). A sigma protocol for a relation R is a tuple $(\text{Comm}, \text{Chal}, \text{Resp}, \text{Verify})$ of PPT algorithms such that:

- **Comm**, denoted $(\text{comm}, t) \leftarrow \text{Comm}(s, w, k)$, is executed by a prover. Comm takes a statement s , witness w and security parameter k as input, and outputs a commitment comm and some state information t .
- **Chal**, denoted $\text{chal} \leftarrow \text{Chal}(s, \text{comm}, k)$, is executed by a verifier. Chal takes a statement s , a commitment comm and a security parameter k as input, and outputs a string chal .
- **Resp**, denoted $\text{resp} \leftarrow \text{Resp}(\text{chal}, t, k)$, is executed by a prover. Resp takes a challenge chal , state information t and security parameter k as input, and outputs a response resp .
- **Verify**, denoted $v \leftarrow \text{Verify}(s, (\text{comm}, \text{chal}, \text{resp}), k)$ is executed by a verifier. Verify takes a statement s , a transcript $(\text{comm}, \text{chal}, \text{resp})$ and a security parameter k as input, and outputs a bit v , which is 1 if the transcript successfully verifies and 0 otherwise. We assume Verify is deterministic.

Moreover, the sigma protocol must be complete: there exists a negligible function μ , such that for all statements and witnesses $(s, w) \in R$ and security parameters k , we have $\Pr[(\text{comm}, t) \leftarrow \text{Comm}(s, w, k); \text{chal} \leftarrow \text{Chal}(s, \text{comm}, k); \text{resp} \leftarrow \text{Resp}(\text{chal}, t, k) : \text{Verify}(s, (\text{comm}, \text{chal}, \text{resp}), k) = 1] > 1 - \mu(k)$.

Some sigma protocols ensure special soundness and special honest-verifier zero-knowledge. We will make use of a result by Bernhard et al. that requires these properties, but we will not need the details of those definitions in our proofs, so we omit them here; see Bernhard et al. [23] for a formalization.

C. Non-interactive proof systems

A proof system is non-interactive if a single message is sent from the prover to the verifier.

Definition 19 (Non-interactive proof system). A non-interactive proof system for a relation R is a tuple of PPT algorithms $(\text{Prove}, \text{Verify})$ such that:

⁵⁶ Let $x \leftarrow_R S$ denote assignment to x of an element chosen uniformly at random from set S .

⁵⁷ The oracle in experiment Exp-Sign may access parameter sk . Henceforth, we continue to allow oracles to access experiment parameters without explicitly mentioning them.

- **Prove**, denoted $\sigma \leftarrow \text{Prove}(s, w, k)$, is executed by a prover to prove $(s, w) \in R$.
- **Verify**, denoted $v \leftarrow \text{Verify}(s, \sigma, k)$, is executed by anyone to check the validity of a proof. We assume Verify is deterministic.

Moreover, the system must be complete: there exists a negligible function μ , such that for all statement and witnesses $(s, w) \in R$ and security parameters k , we have $\Pr[\sigma \leftarrow \text{Prove}(s, w, k) : \text{Verify}(s, \sigma, k) = 1] > 1 - \mu(k)$.

We can derive non-interactive proof systems from sigma protocols using the Fiat-Shamir transformation [66], which replaces the verifier’s challenge with a hash of the prover’s commitment, concatenated with the prover’s statement.

Definition 20 (Fiat-Shamir transformation [66]). Given a sigma protocol $\Sigma = (\text{Comm}, \text{Chal}, \text{Resp}, \text{Verify}_\Sigma)$ for relation R and a hash function \mathcal{H} , the Fiat-Shamir transformation, denoted $\text{FS}(\Sigma, \mathcal{H})$, is the tuple $(\text{Prove}, \text{Verify})$ of algorithms, defined as follows:

$\text{Prove}(s, w, k) =$
 1 $(\text{comm}, t) \leftarrow \text{Comm}(s, w, k);$
 2 $\text{chal} \leftarrow \mathcal{H}(\text{comm}, s);$
 3 $\text{resp} \leftarrow \text{Resp}(\text{chal}, t, k);$
 4 **return** $(\text{comm}, \text{resp})$

$\text{Verify}(s, (\text{comm}, \text{resp}), k) =$
 1 $\text{chal} \leftarrow \mathcal{H}(\text{comm}, s);$
 2 **return** $\text{Verify}_\Sigma(s, (\text{comm}, \text{chal}, \text{resp}), k)$

It is straightforward to check that FS produces non-interactive proof systems. In particular, given sigma protocol Σ for relation R , and a hash function \mathcal{H} , we have $\text{FS}(\Sigma, \mathcal{H})$ is a non-interactive proof system for relation R .

Some applications of the Fiat-Shamir transformation produce non-interactive proof systems satisfying *zero-knowledge*: anything a verifier can derive about a witness can be derived without interaction with a prover—that is, the prover can be simulated by a PPT algorithm called a *simulator*. We will not need the details of zero-knowledge in our proofs, so we omit them here; see Bernhard et al. [23] or Quaglia & Smyth [116] for formalizations.

In addition, some applications of the Fiat-Shamir transformation produce non-interactive proof systems satisfying *simulation sound extractability*: an *extractor* can recover witnesses from proofs by *rewinding* the prover, as discussed below. (We use extractors in our proofs of theorems, to obtain witnesses from proofs.) We define simulation sound extractability in the *random oracle model* [11]. A random oracle can be *programmed* or *patched*. We will not need the details of how patching works in our proofs, so we omit them here; see Bernhard et al. [23] for a formalization.

Definition 21 (Simulation sound extractability [23], [75]). Suppose Σ is a sigma protocol for relation R , \mathcal{H} is a random oracle, and $(\text{Prove}, \text{Verify})$ is a non-interactive proof system, such that $\text{FS}(\Sigma, \mathcal{H}) = (\text{Prove}, \text{Verify})$. Further suppose \mathcal{S}

is a simulator for $(\text{Prove}, \text{Verify})$ and \mathcal{H} can be patched by \mathcal{S} . Proof system $(\text{Prove}, \text{Verify})$ satisfies simulation sound extractability if there exists a PPT algorithm \mathcal{K} , such that for all PPT adversaries \mathcal{A} and coins r , there exists a negligible function μ , such that for all security parameters k , we have:⁵⁸

$$\Pr[\mathbf{P} \leftarrow (); \mathbf{Q} \leftarrow \mathcal{A}^{\mathcal{H}, \mathcal{P}}(-; r); \mathbf{W} \leftarrow \mathcal{K}^{\mathcal{A}'}(\mathbf{H}, \mathbf{P}, \mathbf{Q}) : |\mathbf{Q}| \neq |\mathbf{W}| \vee \exists j \in \{1, \dots, |\mathbf{Q}|\} \cdot (\mathbf{Q}[j][1], \mathbf{W}[j]) \notin R \wedge \forall (s, \sigma) \in \mathbf{Q}, (t, \tau) \in \mathbf{P} \cdot \text{Verify}(s, \sigma, k) = 1 \wedge \sigma \neq \tau] \leq \mu(k)$$

where $\mathcal{A}(-; r)$ denotes running adversary \mathcal{A} with an empty input and coins r , where \mathbf{H} is a transcript of the random oracle’s input and output, and where oracles \mathcal{A}' and \mathcal{P} are defined below:

- $\mathcal{A}'()$. Computes $\mathbf{Q}' \leftarrow \mathcal{A}(-; r)$, forwarding any of \mathcal{A}' ’s oracle calls to \mathcal{K} , and outputs \mathbf{Q}' . By running $\mathcal{A}(-; r)$, \mathcal{K} is rewinding the adversary.
- $\mathcal{P}(s)$. Computes $\sigma \leftarrow \mathcal{S}(s, k); \mathbf{P} \leftarrow (\mathbf{P}[1], \dots, \mathbf{P}[|\mathbf{P}|], (s, \sigma))$ and outputs σ .

Algorithm \mathcal{K} is an extractor for $(\text{Prove}, \text{Verify})$.

Our definition of simulation sound extractability in the random oracle model is an analogue of Groth’s definition in the common reference string model [75, §2]. (See Bernhard et al. [23, §1] for a detailed comparison.) Our presentation of simulation sound extractability differs from the presentation by Bernhard et al. [23] by formalizing some of the details.

Bernhard et al. [23] show that non-interactive proof systems derived using the Fiat-Shamir transformation satisfy zero-knowledge and simulation sound extractability:

Theorem 14 (from [23]). Let Σ be a sigma protocol for relation R , and let \mathcal{H} be a random oracle. If Σ satisfies special soundness and special honest verifier zero-knowledge, then $\text{FS}(\Sigma, \mathcal{H})$ satisfies zero-knowledge and simulation sound extractability.

The Fiat-Shamir transformation can be generalized to include an optional string m in the hashes produced by functions Prove and Verify. We write $\text{Prove}(s, w, m, k)$ and $\text{Verify}(s, (\text{comm}, \text{resp}), m, k)$ for invocations of Prove and Verify which include an optional string. When m is provided, it is included in the hashes in both algorithms. That is, given $\text{FS}(\Sigma, \mathcal{H}) = (\text{Prove}, \text{Verify})$, the hashes are computed as follows in both algorithms: $\text{chal} \leftarrow \mathcal{H}(\text{comm}, s, m)$. Theorem 14 can be extended to this generalization.

APPENDIX B VARIANTS OF Exp-IV

Our individual verifiability experiment with external authentication (§II-B1) can be equivalently formulated as an experiment that challenges \mathcal{A} to predict the output of Vote:

⁵⁸ We extend set membership notation to vectors: we write $x \in \mathbf{x}$ if x is an element of the set $\{\mathbf{x}[i] : 1 \leq i \leq |\mathbf{x}|\}$.

```

Exp-IV-Ext'(Π, A, k) =
1 (PKT, nC, β, b) ← A(k);
2 b' ← Vote(PKT, nC, β, k);
3 if b = b' ∧ b' ≠ ⊥ then
4   | return 1
5 else
6   | return 0

```

Proposition 15. *Given an election scheme Π, we have*

$$\forall \mathcal{A} \exists \mu \forall k . \text{Succ}(\text{Exp-IV-Ext}(\Pi, \mathcal{A}, k)) \leq \mu(k)$$

$$\Leftrightarrow \forall \mathcal{A}' \exists \mu' \forall k' . \text{Succ}(\text{Exp-IV-Ext}'(\Pi, \mathcal{A}', k')) \leq \mu'(k'),$$

where \mathcal{A} and \mathcal{A}' are PPT adversaries, μ and μ' are negligible functions, and k and k' are security parameters.

Intuitively, if \mathcal{A} can predict the output of Vote, then \mathcal{A} can use that prediction to generate a collision. And if \mathcal{A} can generate collisions, then \mathcal{A} can use them to predict outputs.

Proof. For the forward implication, suppose \mathcal{A}' is a PPT adversary such that $\text{Succ}(\text{Exp-IV-Ext}'(\Pi, \mathcal{A}', k')) > \frac{1}{p(k')}$ for some polynomial function p and security parameter k' . We construct an adversary \mathcal{A} against Exp-IV-Ext. On input k' , adversary \mathcal{A} computes $(PK_{\mathcal{T}}, n_C, \beta, b) \leftarrow \mathcal{A}'(k')$ and outputs $(PK_{\mathcal{T}}, n_C, \beta, \beta)$. Since \mathcal{A}' wins Exp-IV-Ext' with non-negligible probability, we have

$$\Pr[b' \leftarrow \text{Vote}(PK_{\mathcal{T}}, n_C, \beta, k') : b = b' \wedge b \neq \perp] > \frac{1}{p(k')}.$$

Moreover, since calls to algorithm Vote are independent, we have

$$\Pr[b_1 \leftarrow \text{Vote}(PK_{\mathcal{T}}, n_C, \beta, k');$$

$$b_2 \leftarrow \text{Vote}(PK_{\mathcal{T}}, n_C, \beta, k')$$

$$: b_1 = b \wedge b_2 = b \wedge b_1 \neq \perp \wedge b_2 \neq \perp] > \frac{1}{p(k')^2}.$$

It follows that $\text{Succ}(\text{Exp-IV-Ext}(\Pi, \mathcal{A}, k')) > \frac{1}{p(k')^2}$.

For the reverse implication, suppose \mathcal{A} is a PPT adversary such that $\text{Succ}(\text{Exp-IV-Ext}(\Pi, \mathcal{A}, k)) > \frac{1}{p(k)}$ for some polynomial function p and security parameter k . We construct an adversary \mathcal{A}' against Exp-IV-Ext'. On input k , adversary \mathcal{A}' computes $(PK_{\mathcal{T}}, n_C, \beta_1, \beta_2) \leftarrow \mathcal{A}(k)$; $b_1 \leftarrow \text{Vote}(PK_{\mathcal{T}}, n_C, \beta_1, k)$ and outputs $(PK_{\mathcal{T}}, n_C, \beta_2, b_1)$. Since \mathcal{A} wins Exp-IV-Ext with probability no less than $\frac{1}{p(k)}$, we have

$$\Pr[b_2 \leftarrow \text{Vote}(PK_{\mathcal{T}}, n_C, \beta_2, k) : b_1 = b_2 \wedge b_1 \neq \perp] > \frac{1}{p(k)}.$$

It follows that $\text{Succ}(\text{Exp-IV-Int}'(\Pi, \mathcal{A}', k)) > \frac{1}{p(k)}$. \square

Our individual verifiability experiment with internal authentication (§IV-B1) can also be reformulated as an experiment that challenges \mathcal{A} to predict the output of Vote algorithms:

```

Exp-IV-Int'(Π, A, k) =
1 (PKT, nV) ← A(k);
2 for 1 ≤ i ≤ nV do (pki, ski) ← Register(PKT, k);
3 L ← {pk1, ..., pknV};
4 Crpt ← ∅;
5 (nC, β, i, b) ← AC(L);
6 b' ← Vote(ski, PKT, nC, β, k);
7 if b = b' ∧ b' ≠ ⊥ ∧ ski ∉ Crpt then
8   | return 1
9 else
10  | return 0

```

Similarly to Section IV-B1, the adversary is given access to oracle C and the voter index output on line 5 must be legal with respect to n_V .

Experiment Exp-IV-Int' is strictly stronger than our original experiment Exp-IV-Int, since predicting the output of Vote does not imply the existence of collisions, whereas collisions can be used to predict the output of Vote. For instance, consider the following variant of Nonce (Definition 5):

Definition 22. *Election scheme Nonce' is defined as follows:*

- Setup(k) outputs $(\perp, \perp, \infty, \infty)$.
- Register($PK_{\mathcal{T}}, k$) computes $r \in \mathbb{Z}_{2^k}$ and outputs (r, r) .
- Vote($r, PK_{\mathcal{T}}, n_C, \beta, k$) outputs (r, β) .
- Tally($SK_{\mathcal{T}}, BB, L, n_C, k$) computes a vector \mathbf{X} of length n_C , such that \mathbf{X} is a tally of the votes on BB for which the nonce is in L , and outputs (\mathbf{X}, \perp) .
- Verify($PK_{\mathcal{T}}, BB, L, n_C, \mathbf{X}, P, k$) outputs 1 if $(\mathbf{X}, P) = \text{Tally}(\perp, \perp, BB, L, n_C, k)$ and 0 otherwise.

Intuitively, an adversary can predict the output of Vote, because the algorithm is deterministic and the electoral roll lists private credentials. However, the Register algorithm ensures that voters' credentials are distinct with overwhelming probability, hence, instantiations of the Vote algorithm with distinct voter credentials will never collide.

Proposition 16. *Given an election scheme Π, PPT adversary A, negligible function μ, and security parameter k, if $\text{Succ}(\text{Exp-IV-Int}'(\Pi, \mathcal{A}, k)) \leq \mu(k)$, then there exists a PPT adversary B such that $\text{Succ}(\text{Exp-IV-Int}(\Pi, \mathcal{B}, k)) \leq \mu(k)$.*

The proof of Proposition 16 is similar to the reverse implication proof of Proposition 15.

APPENDIX C GENERALIZED HELIOS SCHEME

We formalize a generic construction for Helios-like election schemes (Definition 24). Our construction is parameterized on the choice of homomorphic encryption scheme and sigma protocols for the relations introduced in the following definition.

Definition 23. *Let (Gen, Enc, Dec) be a homomorphic asymmetric encryption scheme and Σ be a sigma protocol for a*

binary relation R .⁵⁹

- Σ proves correct key construction if $((k, pk, m), (sk, s)) \in R \Leftrightarrow (pk, sk, m) = \text{Gen}(k; s)$.

Suppose $(pk, sk, m) = \text{Gen}(k; s)$, for some security parameter k and coins s .

- Σ proves plaintext knowledge in a subspace if $((pk, c, m'), (m, r)) \in R \Leftrightarrow c = \text{Enc}(pk, m; r) \wedge m \in m' \wedge m' \subseteq m$.
- Σ proves correct decryption if $((pk, c, m), sk) \in R \Leftrightarrow m = \text{Dec}(sk, c)$.

Definition 24 (Generalized Helios). Suppose $\Gamma = (\text{Gen}, \text{Enc}, \text{Dec})$ is an additively homomorphic asymmetric encryption scheme, Σ_1 proves correct key construction, Σ_2 proves plaintext knowledge in a subspace, Σ_3 proves correct decryption, and \mathcal{H} is a hash function. Let $\text{FS}(\Sigma_1, \mathcal{H}) = (\text{ProveKey}, \text{VerKey})$, $\text{FS}(\Sigma_2, \mathcal{H}) = (\text{ProveCiph}, \text{VerCiph})$, and $\text{FS}(\Sigma_3, \mathcal{H}) = (\text{ProveDec}, \text{VerDec})$. We define generalized Helios as $\text{Helios}(\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \mathcal{H}) = (\text{Setup}, \text{Vote}, \text{Tally}, \text{Verify})$:

- **Setup**(k). Select coins s uniformly at random, compute $(pk, sk, m) \leftarrow \text{Gen}(k; s)$; $\rho \leftarrow \text{ProveKey}((k, pk, m), (sk, s), k)$; $PK_{\mathcal{T}} \leftarrow (pk, m, \rho)$; $SK_{\mathcal{T}} \leftarrow (pk, sk)$, let m be the largest integer such that $\{0, \dots, m\} \subseteq \{0\} \cup m$, and output $(PK_{\mathcal{T}}, SK_{\mathcal{T}}, m, m)$.
- **Vote**($PK_{\mathcal{T}}, n_C, \beta, k$). Parse $PK_{\mathcal{T}}$ as a vector (pk, m, ρ) . Output \perp if parsing fails or $\text{VerKey}((k, pk, m), \rho, k) \neq 1 \vee \beta \notin \{1, \dots, n_C\}$. Select coins r_1, \dots, r_{n_C-1} uniformly at random and compute:

for $1 \leq j \leq n_C - 1$ **do**
 if $j = \beta$ **then** $m_j \leftarrow 1$; **else** $m_j \leftarrow 0$;
 $c_j \leftarrow \text{Enc}(pk, m_j; r_j)$;
 $\sigma_j \leftarrow \text{ProveCiph}((pk, c_j, \{0, 1\}), (m_j, r_j), j, k)$;
 $c \leftarrow c_1 \otimes \dots \otimes c_{n_C-1}$;
 $m \leftarrow m_1 \odot \dots \odot m_{n_C-1}$;
 $r \leftarrow r_1 \oplus \dots \oplus r_{n_C-1}$;
 $\sigma_{n_C} \leftarrow \text{ProveCiph}((pk, c, \{0, 1\}), (m, r), n_C, k)$;

Output ballot $(c_1, \dots, c_{n_C-1}, \sigma_1, \dots, \sigma_{n_C})$.

- **Tally**($SK_{\mathcal{T}}, BB, n_C, k$). Initialize vectors \mathbf{X} of length n_C and \mathbf{P} of length $n_C - 1$. Compute **for** $1 \leq j \leq n_C$ **do** $\mathbf{X}[j] \leftarrow 0$. Parse $SK_{\mathcal{T}}$ as a vector (pk, sk) . Output (\mathbf{X}, \mathbf{P}) if parsing fails. Let $\{b_1, \dots, b_\ell\}$ be the largest subset of BB such that $b_1 < \dots < b_\ell$ and for all $1 \leq i \leq \ell$ we have b_i is a vector of length $2 \cdot n_C - 1$ and $\bigwedge_{j=1}^{n_C-1} \text{VerCiph}((pk, b_i[j], \{0, 1\}), b_i[j+n_C-1], j, k) = 1 \wedge \text{VerCiph}((pk, b_i[1] \otimes \dots \otimes b_i[n_C-1], \{0, 1\}), b_i[2 \cdot n_C - 1], n_C, k) = 1$. If $\{b_1, \dots, b_\ell\} = \emptyset$, then output (\mathbf{X}, \mathbf{P}) , otherwise, compute:

for $1 \leq j \leq n_C - 1$ **do**
 $c \leftarrow b_1[j] \otimes \dots \otimes b_\ell[j]$;
 $\mathbf{X}[j] \leftarrow \text{Dec}(sk, c)$;
 $\mathbf{P}[j] \leftarrow \text{ProveDec}((pk, c, \mathbf{X}[j]), sk, k)$;
 $\mathbf{X}[n_C] \leftarrow \ell - \sum_{j=1}^{n_C-1} \mathbf{X}[j]$;

Output (\mathbf{X}, \mathbf{P}) .

- **Verify**($PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, \mathbf{P}, k$). Parse \mathbf{X} as a vector of

length n_C , parse \mathbf{P} as a vector of length $n_C - 1$, parse $PK_{\mathcal{T}}$ as a vector (pk, m, ρ) . Output 0 if parsing fails or $\text{VerKey}((k, pk, m), \rho, k) \neq 1$. Let $\{b_1, \dots, b_\ell\}$ be the largest subset of BB satisfying the conditions given by the tally algorithm and let m_B be the largest integer such that $\{0, \dots, m_B\} \subseteq m$. If $\{b_1, \dots, b_\ell\} = \emptyset \wedge \bigwedge_{j=1}^{n_C} \mathbf{X}[j] = 0$ or $\bigwedge_{j=1}^{n_C-1} \text{VerDec}((pk, b_1[j] \otimes \dots \otimes b_\ell[j], \mathbf{X}[j]), \mathbf{P}[j], k) = 1 \wedge \mathbf{X}[n_C] = \ell - \sum_{j=1}^{n_C-1} \mathbf{X}[j] \wedge 1 \leq \ell \leq m_B$, then output 1, otherwise, output 0.

The above algorithms assume $n_C > 1$ and we define special cases of Vote, Tally and Verify when $n_C = 1$:

- **Vote**($PK_{\mathcal{T}}, n_C, \beta, k$). Parse $PK_{\mathcal{T}}$ as a vector (pk, m, ρ) . Output \perp if parsing fails or $\text{VerKey}((k, pk, m), \rho, k) \neq 1 \vee \beta \neq 1$. Select coins r uniformly at random, compute $m \leftarrow 1$; $c \leftarrow \text{Enc}(pk, m; r)$; $\sigma \leftarrow \text{ProveCiph}((pk, c, \{0, 1\}), (m, r), k)$, and output ballot (c, σ) .
- **Tally**($SK_{\mathcal{T}}, BB, n_C, k$). Initialize \mathbf{X} and \mathbf{P} as vectors of length 1. Compute $\mathbf{X}[1] \leftarrow 0$. Parse $SK_{\mathcal{T}}$ as a vector (pk, sk) . Output (\mathbf{X}, \mathbf{P}) if parsing fails. Let $\{b_1, \dots, b_\ell\}$ be the largest subset of BB such that for all $1 \leq i \leq \ell$ we have b_i is a vector of length 2 and $\text{VerCiph}((pk, b_i[1], \{0, 1\}), b_i[2], k) = 1$. If $\{b_1, \dots, b_\ell\} = \emptyset$, then output (\mathbf{X}, \mathbf{P}) . Otherwise, compute $c \leftarrow b_1[1] \otimes \dots \otimes b_\ell[1]$; $\mathbf{X}[1] \leftarrow \text{Dec}(sk, c)$; $\mathbf{P}[1] \leftarrow \text{ProveDec}((pk, c, \mathbf{X}[1]), sk, k)$ and output (\mathbf{X}, \mathbf{P}) .
- **Verify**($PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, \mathbf{P}, k$). Parse \mathbf{X} and \mathbf{P} as vectors of length 1, and parse $PK_{\mathcal{T}}$ as a vector (pk, m, ρ) . Output 0 if parsing fails or $\text{VerKey}((k, pk, m), \rho, k) \neq 1$. Let $\{b_1, \dots, b_\ell\}$ be the largest subset of BB satisfying the conditions given by the tally algorithm and let m_B be the largest integer such that $\{0, \dots, m_B\} \subseteq m$. If $\{b_1, \dots, b_\ell\} = \emptyset \wedge \mathbf{X}[1] = 0$ or $\text{VerDec}((pk, b_1[1] \otimes \dots \otimes b_\ell[1], \mathbf{X}[1]), \mathbf{P}[1], k) = 1 \wedge 1 \leq \ell \leq m_B$, then output 1, otherwise, output 0.

Generalized Helios works as follows. Setup generates the tallier's key pair. The public key includes a non-interactive proof demonstrating that the key pair is correctly constructed. Vote takes a choice $\beta \in \{1, \dots, n_C\}$ and outputs ciphertexts c_1, \dots, c_{n_C-1} such that if $\beta < n_C$, then ciphertext c_β contains plaintext 1 and the remaining ciphertexts contain plaintext 0, otherwise, all ciphertexts contain plaintext 0. Vote also outputs proofs $\sigma_1, \dots, \sigma_{n_C}$ so that this can be verified. In particular, proof σ_j demonstrates ciphertext c_j contains 0 or 1, for all $1 \leq j \leq n_C - 1$. And proof σ_{n_C} demonstrates that the homomorphic combination of ciphertexts $c_1 \otimes \dots \otimes c_{n_C-1}$ contains 0 or 1. (It follows that the voter's ballot contains a vote for exactly one candidate.) Tally homomorphically combines ciphertexts representing votes for a particular candidate and decrypts the homomorphic combinations. The number of votes for a candidate $\beta \in \{1, \dots, n_C - 1\}$ is simply the homomorphic combination of ciphertexts representing votes

59. Given a binary relation R , we write $((s_1, \dots, s_l), (w_1, \dots, w_k)) \in R \Leftrightarrow P(s_1, \dots, s_l, w_1, \dots, w_k)$ for $(s, w) \in R \Leftrightarrow P(s_1, \dots, s_l, w_1, \dots, w_k) \wedge s = (s_1, \dots, s_l) \wedge w = (w_1, \dots, w_k)$, hence, R is only defined over pairs of vectors of lengths l and k .

for that candidate. The number of votes for candidate n_C is equal to the number of votes for all other candidates subtracted from the total number of valid ballots on the bulletin board. Verify checks that each of the above steps has been performed correctly.

Lemma 17 demonstrates that generalized Helios is a construction for election schemes.

Lemma 17. Helios($\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \mathcal{H}$) satisfies Correctness, where $\Gamma, \Sigma_1, \Sigma_2, \Sigma_3$ and \mathcal{H} satisfy the preconditions of Definition 24.

APPENDIX D

PROOF: HELIOS 2.0 IS NOT VERIFIABLE

Chang-Fong & Essex [32] demonstrate that Helios 2.0 is not verifiable and we prove that Helios 2.0 does not satisfy Ver-Ext. Our proof formalizes the attack by Chang-Fong & Essex [32, §4.1] in the context of our Completeness definition using the adversary we define in Figure 1. Intuitively, that adversary computes a ciphertext with a masked term (Line 1) and falsifies a proof of correct construction in a manner that hides malice (Lines 2–12). In particular, the proof ensures $c_1 \not\equiv 0 \pmod{2}$, which causes cancellation of the mask during verification. A ballot is constructed from that ciphertext and proof, and added to a bulletin board (Line 14). The ballot is valid, hence, it will be decrypted during tallying, yet correct decryption cannot be proved, due to the masked ciphertext, thus, verification will fail and Completeness is not satisfied.

Definition 25 (Weak Fiat-Shamir transformation [23]). *The weak Fiat-Shamir transformation is a function wFS that is identical to FS, except that it excludes statement s in the hashes computed by Prove and Verify, as follows: $\text{chal} \leftarrow \mathcal{H}(\text{comm})$.*

Definition 26 (Helios 2.0). *Let $\widehat{\text{Helios}}$ be Helios after replacing all instances of the Fiat-Shamir transformation with the weak Fiat-Shamir transformation and excluding the (optional) messages input to ProveCiph—i.e., ProveCiph should be used as a ternary function. Helios 2.0 is $\widehat{\text{Helios}}(\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \mathcal{H})$, where Γ is additively homomorphic El Gamal [56, §2], Σ_1 is the sigma protocol for proving knowledge of discrete logarithms by Chaum et al. [37, Protocol 2], Σ_2 is the sigma protocol for proving knowledge of disjunctive equality between discrete logarithms by Cramer et al. [55, Figure 1], Σ_3 is the sigma protocol for proving knowledge of equality between discrete logarithms by Chaum and Pedersen [38, §3.2], and \mathcal{H} is SHA-256 [108].*

We assume the sigma protocols used by Helios 2.0 satisfy the preconditions of generalized Helios—that is, [37, Protocol 2] is a sigma protocol for proving correct key construction, [55, Figure 1] is a sigma protocol for proving plaintext knowledge in a subspace, and [38, §3.2] is a sigma protocol for proving decryption. We leave formally proving this assumption as future work. Under this assumption, Lemma 17 demonstrates that Helios 2.0 is an election scheme.

Fig. 1 Adversary against Helios 2.0

Given a public key $PK_{\mathcal{T}}$ and security parameter k as input, adversary \mathcal{A} parses $PK_{\mathcal{T}}$ as a vector (pk, m, ρ) and pk as (p, q, g, h) , computes a generator g' of a sub-group of order 2 such that $g' \mid p - 1$, selects coins r , and computes:

```

1  $e \leftarrow (g' \cdot g^r \pmod{p}, h^r \cdot g \pmod{p});$ 
2 do
3    $(c_0, f_0) \leftarrow_R \mathbb{Z}_q^2;$ 
4    $A_0 \leftarrow g^{f_0} \cdot e[1]^{-c_0} \pmod{p};$ 
5    $B_0 \leftarrow h^{f_0} \cdot e[2]^{-c_0} \pmod{p};$ 
6    $w \leftarrow_R \mathbb{Z}_q;$ 
7    $A_1 \leftarrow g^w \pmod{p};$ 
8    $B_1 \leftarrow h^w \pmod{p};$ 
9    $c_1 \leftarrow \mathcal{H}(A_0, B_0, A_1, B_1) - c_0 \pmod{q};$ 
10 while  $c_1 \not\equiv 0 \pmod{2};$ 
11  $f_1 \leftarrow w + c_1 \cdot r \pmod{q};$ 
12  $\sigma \leftarrow (A_0, B_0, c_0, f_0, A_1, B_1, c_1, f_1);$ 
13  $n_C \leftarrow 2;$ 
14  $BB \leftarrow \{(e, \sigma)\};$ 
15 return  $(n_C, BB)$ 
```

Proof of Proposition 2. Let Setup, Tally and Verify be the setup, tallying and verification algorithms defined by Helios 2.0. Moreover, let $\Gamma = (\text{Gen}, \text{Enc}, \text{Dec})$, $wFS(\Sigma_1, \mathcal{H}) = (\text{ProveKey}, \text{VerKey})$, and $wFS(\Sigma_3, \mathcal{H}) = (\text{ProveDec}, \text{VerDec})$. We construct an adversary \mathcal{A} (Figure 1) against the Completeness experiment.

Suppose k is a security parameter, $(PK_{\mathcal{T}}, SK_{\mathcal{T}}, m_B, m_C)$ is an output of Setup(k), and (BB, n_C) is an output of $\mathcal{A}(PK_{\mathcal{T}}, k)$, such that $|BB| \leq m_B \wedge n_C \leq m_C$. By definition of Setup, we have $PK_{\mathcal{T}}$ parses as (pk, m, ρ) and $SK_{\mathcal{T}}$ parses as (pk, sk) , such that $(pk, sk, m) = \text{Gen}(k; s)$ and ρ is an output of ProveKey($(k, pk, m), (sk, s), k$) for some coins s chosen uniformly at random by Setup. By definition of Gen, we have pk parses as (p, q, g, h) . And by definition of \mathcal{A} , we have $n_C = 2$ and $BB = \{(e, \sigma)\}$, where e and σ are computed by the adversary. Further suppose (\mathbf{X}, \mathbf{P}) is an output of Tally($SK_{\mathcal{T}}, BB, n_C, k$).

Let us recall the definition of VerCiph (cf. [55, Figure 1], Definition 25, and Helios 2.0 source code) and consider whether $\text{VerCiph}((pk, e, \{0, 1\}), \sigma, k) = 1$:

- $\text{VerCiph}((pk, e, \{0, 1\}), \sigma, k)$. Parses pk as (p, q, g, h) , e as (R, S) , and σ as $(A_0, B_0, c_0, f_0, A_1, B_1, c_1, f_1)$, outputting 0 if parsing fails. If $g^{f_0} \equiv A_0 \cdot R^{c_0} \pmod{p} \wedge h^{f_0} \equiv B_0 \cdot S^{c_0} \pmod{p} \wedge g^{f_1} \equiv A_1 \cdot R^{c_1} \pmod{p} \wedge h^{f_1} \equiv B_1 \cdot (S/g)^{c_1} \pmod{p} \wedge \mathcal{H}(A_0, B_0, A_1, B_1) \equiv c_0 + c_1 \pmod{q}$, then output 1, otherwise, output 0.

By definition of \mathcal{A} , we have $\sigma = (A_0, B_0, c_0, f_0, A_1, B_1, c_1, f_1)$. Moreover, we have $e[1] \equiv g' \cdot g^r \pmod{p}$, $e[2] \equiv h^r \cdot g \pmod{p}$, $A_0 \equiv g^{f_0} \cdot e[1]^{-c_0} \pmod{p}$, and $B_0 \equiv h^{f_0} \cdot e[2]^{-c_0} \pmod{p}$, where g' is a generator of a sub-group of order 2 such that $g' \mid p - 1$ and c_0, f_0 and r are coins. Hence, we trivially have

$$g^{f_0} \equiv g^{f_0} \cdot e[1]^{-c_0} \cdot e[1]^{c_0} \equiv A_0 \cdot e[1]^{c_0} \pmod{p}$$

$$h^{f_0} \equiv h^{f_0} \cdot e[2]^{-c_0} \cdot e[2]^{c_0} \equiv B_0 \cdot e[2]^{c_0} \pmod{p}$$

By definition of \mathcal{A} , we also have $A_1 \equiv g^w \pmod{p}$, $B_1 \equiv h^w \pmod{p}$, $c_1 \equiv \mathcal{H}(A_0, B_0, A_1, B_1) - c_0 \pmod{q}$, and $f_1 \equiv w + c_1 \cdot r \pmod{q}$, such that $c_1 \equiv 0 \pmod{2}$, where w are coins. Hence, we have

$$g^{f_1} \equiv g^w \cdot g^{c_1 \cdot r} \pmod{p}$$

and, since $c_1 \equiv 0 \pmod{2}$, we have $g^{c_1} \equiv 1 \pmod{p}$, thus,

$$\begin{aligned} &\equiv g^w \cdot g^{c_1} \cdot g^{c_1 \cdot r} \pmod{p} \\ &\equiv g^w \cdot (g' \cdot g^r)^{c_1} \pmod{p} \\ &\equiv g^w \cdot e[1]^{c_1} \pmod{p} \end{aligned}$$

Moreover, we trivially have

$$h^{f_1} \equiv h^w \cdot h^{c_1 \cdot r} \equiv h^w \cdot (h^r \cdot g/g)^{c_1} \equiv B_1 \cdot (e[2]/g)^{c_1} \pmod{p}$$

Furthermore, we have $\mathcal{H}(A_0, B_0, A_1, B_1) \equiv c_0 + c_1 \pmod{q}$. Hence, $\text{VerCiph}((pk, e, \{0, 1\}), \sigma, k) = 1$. It follows that BB is the largest subset of BB satisfying the conditions defined by algorithm Tally. Thus, $\mathbf{X} = (\text{Dec}(sk, e), 1 - \text{Dec}(sk, e))$ and \mathbf{P} is an output of $\text{ProveDec}((pk, e, \mathbf{X}[1]), sk, k)$. It remains to show $\text{Verify}(PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, P, k) \neq 1$ with non-negligible probability. By definition of Verify, it suffices to show $\text{VerDec}((pk, e, \mathbf{X}[1]), \mathbf{P}[1], k) \neq 1$.

Let us recall definitions of ProveDec and VerDec (cf. [38, §3.2], Definition 25, and Helios 2.0 source code):

- $\text{ProveDec}((pk, e, m), sk, k)$. Parses pk as (p, q, g, h) , outputting 0 if parsing fails. Computes $w \leftarrow_R \mathbb{Z}_q$; $A \leftarrow g^w \pmod{p}$; $B \leftarrow e[1]^w \pmod{p}$; $c \leftarrow \mathcal{H}(A, B) \pmod{q}$; $f \leftarrow w + c \cdot sk \pmod{q}$. And outputs (A, B, f)
- $\text{VerDec}((pk, e, m), \tau, k)$. Parses pk as (p, q, g, h) and τ as (A, B, f) , outputting 0 if parsing fails. If $g^f \equiv A \cdot h^c \pmod{p}$ and $e[1]^f \equiv B \cdot (e[2]/g^m)^c \pmod{p}$, then output 1, otherwise, output 0, where $c \equiv \mathcal{H}(A, B) \pmod{q}$.

Hence, we have $\mathbf{P} = (A, B, f)$ such that $B \equiv e[1]^w \pmod{p}$ and $f \equiv w + c \cdot sk \pmod{q}$, where $c \equiv \mathcal{H}(A, B) \pmod{q}$ and coins w were selected by ProveDec. Thus, $e[1]^f \not\equiv B \cdot (e[2]/g^{\mathbf{X}[1]})^c \pmod{p}$, concluding our proof. \square

APPENDIX E

PROOF: HELIOS 3.1.4 IS NOT VERIFIABLE

Helios 2.0 is vulnerable to attacks because it does not check the suitability of cryptographic parameters, nor does it check that all elements of ballots are constructed using the correct parameters. Chang-Fong & Essex [32] address these vulnerabilities by performing the necessary checks.

Definition 27 (Helios 3.1.4). *Election scheme Helios 3.1.4 is Helios 2.0 after modifying the sigma protocols to perform the checks proposed by Chang-Fong & Essex [32, §4].*

Bernhard et al. [23] demonstrate that Helios 2.0 is not verifiable and we prove that Helios 3.1.4 does not satisfy Ver-Ext. Our proof formalizes the attack by Bernhard et al. [23, §3]

in the context of our universal verifiability experiment using the adversary we define in Figure 2. That adversary computes the challenge hash (Line 9) before computing a ciphertext. (This is possible because weak Fiat-Shamir does not include statements in hashes, hence, ciphertexts are not included in hashes.) Moreover, the adversary computes: a private key as a function of that hash (Line 11), challenges as functions of the hash and the private key (Lines 13 & 14), and responses as functions of the challenges and some coins (Lines 18 & 19). Furthermore, the adversary computes a public key from the private key (Line 23) and a proof of correct key generation (Line 25). That proof is valid, because the private key could have been correctly computed. The adversary encrypts a plaintext m (such that $m > 1$) using the aforementioned coins (Line 27) and proves correct decryption of that ciphertext (Line 33). That proof is valid, because the ciphertext is well-formed. Finally, the adversary claims $(m, m-1)$ is the election outcome corresponding to the ballot containing the ciphertext and falsified proof of correct construction. The verification procedure will accept that outcome, because all proofs hold, yet the election outcome is clearly invalid, hence, universal verifiability is not satisfied.

Proof of Proposition 3. Let Vote and Tally be the vote and tallying algorithms defined by Helios 3.1.4. Moreover, let $\text{wFS}(\Sigma_1, \mathcal{H}) = (\text{ProveKey}, \text{VerKey})$, $\text{wFS}(\Sigma_2, \mathcal{H}) = (\text{ProveCiph}, \text{VerCiph})$ and $\text{wFS}(\Sigma_3, \mathcal{H}) = (\text{ProveDec}, \text{VerDec})$. We construct an adversary \mathcal{A} (Figure 2) against the universal verifiability experiment.

Suppose an execution of Exp-UV-Ext computes

$$\begin{aligned} &(PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, P) \leftarrow \mathcal{A}(k); \\ &\mathbf{Y} \leftarrow \text{correct-tally}(pk, BB, n_C, k) \end{aligned}$$

Since $m > 1$, there is no choice $\beta \in \{1, 2\}$ nor coins r such that $\text{Vote}(PK_{\mathcal{T}}, n_C, \beta, k; r) \in BB$. By definition of function *correct-tally*, we have $\mathbf{Y} = (0, 0)$. Moreover, since $\mathbf{X} = (m, 1 - m)$, we have $\mathbf{X} \neq \mathbf{Y}$ and $\mathbf{X}[2] = 1 - \mathbf{X}[1]$. Let us show that $\text{Verify}(PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, P, k) = 1$.

By definition of \mathcal{A} , we have $PK_{\mathcal{T}}$ is a vector (pk, m, ρ) . Moreover, by the completeness of $(\text{ProveKey}, \text{VerKey})$ and $(\text{ProveDec}, \text{VerDec})$, we have $\text{VerKey}((k, pk, m), \rho, k) = 1$ and $\text{VerDec}((pk, e, \mathbf{X}[1]), \mathbf{P}[1], k) = 1$. It remains to show that BB is the largest subset of BB satisfying the conditions given by the Tally algorithm. Since $BB = \{(e, \sigma, \sigma)\}$ and (e, σ, σ) is a vector of length $2 \cdot n_C - 1$, it suffices to show that $\text{VerCiph}((pk, e, \{0, 1\}), \sigma, k) = 1$. Let us recall the definition of VerCiph (cf. [55, Figure 1], Definition 25, and Helios source code) with the additional checks proposed by Chang-Fong & Essex [32, §4]:

- $\text{VerCiph}((pk, e, \{0, 1\}), \sigma, k)$. Parses pk as (p, q, g, h) , e as (R, S) , and σ as $(A_0, B_0, c_0, f_0, A_1, B_1, c_1, f_1)$, outputting 0 if parsing fails or R, S, A_0, B_0, A_1 or B_1 belong to the wrong group. If $g^{f_0} \equiv A_0 \cdot R^{c_0} \pmod{p} \wedge h^{f_0} \equiv B_0 \cdot S^{c_0} \pmod{p} \wedge g^{f_1} \equiv A_1 \cdot R^{c_1} \pmod{p} \wedge h^{f_1} \equiv B_1 \cdot (S/g)^{c_1} \pmod{p} \wedge \mathcal{H}(A_0, B_0, A_1, B_1) \equiv c_0 + c_1 \pmod{q}$, then output 1, otherwise, output 0.

Fig. 2 Adversary against Helios 3.1.4

Given a security parameter k as input, adversary \mathcal{A} computes primes p and q such that $p = 2 \cdot q + 1$ and q is of length k , and also computes a generator g of the multiplicative group \mathbb{Z}_p^* . Let $n_C \leftarrow 2$ and $\mathfrak{m} \leftarrow \mathbb{N}_{q-1}$, moreover, let $m > 1$ be an element of \mathfrak{m} . The adversary proceeds as follows:

```

1 %coins
2  $(a_0, b_0, a_1, b_1) \leftarrow_R \mathbb{Z}_q^4$ ;
3 %witnesses
4  $A_0 \leftarrow g^{a_0} \pmod{p}$ ;
5  $B_0 \leftarrow g^{b_0} \pmod{p}$ ;
6  $A_1 \leftarrow g^{a_1} \pmod{p}$ ;
7  $B_1 \leftarrow g^{b_1} \pmod{p}$ ;
8 %challenge hash
9  $c \leftarrow \mathcal{H}(A_0, B_0, A_1, B_1) \pmod{q}$ ;
10 %private key
11  $x \leftarrow \frac{(b_0+c \cdot m) \cdot (1-m) - b_1 \cdot m}{a_0 \cdot (1-m) - a_1 \cdot m} \pmod{q}$ ;
12 %challenges
13  $c_1 \leftarrow \frac{b_1 - a_1 \cdot x}{1-m} \pmod{q}$ ;
14  $c_0 \leftarrow c - c_1 \pmod{q}$ ;
15 %coins
16  $r \leftarrow_R \mathbb{Z}_q$ ;
17 %responses
18  $f_0 \leftarrow a_0 + c_0 \cdot r \pmod{q}$ ;
19  $f_1 \leftarrow a_1 + c_1 \cdot r \pmod{q}$ ;
20 %proof of plaintext knowledge
21  $\sigma \leftarrow (A_0, B_0, c_0, f_0, A_1, B_1, c_1, f_1)$ ;
22 %public key
23  $h \leftarrow g^x \pmod{p}$ ;  $pk \leftarrow (p, q, g, h)$ ;
24 %proof of correct key construction
25  $\rho \leftarrow \text{ProveKey}((k, pk, \mathfrak{m}), (x, r'), k)$ ;
26 %ciphertext
27  $e \leftarrow (g^r \pmod{p}, h^r \cdot g^m \pmod{p})$ ;
28 %bulletin board
29  $BB \leftarrow \{(e, \sigma, \rho)\}$ ;
30 %tally
31  $\mathbf{X} \leftarrow (m, 1 - m)$ ;
32 %proof of decryption
33  $\mathbf{P} \leftarrow (\text{ProveDec}((pk, e, m), x, k))$ ;
34 return  $((pk, \mathfrak{m}, \rho), BB, n_C, \mathbf{X}, \mathbf{P})$ 

```

where r' is computed such that $(pk, x, \mathfrak{m}) = \text{Gen}(k; r')$.

By definition of \mathcal{A} , we have R, S, A_0, B_0, A_1 and B_1 belong to the right group. And we have

$$\begin{aligned} g^{f_0} &\equiv g^{a_0+c_0 \cdot r} \equiv g^{a_0} \cdot (g^r)^{c_0} \equiv A_0 \cdot R^{c_0} \pmod{p} \\ g^{f_1} &\equiv g^{a_1+c_1 \cdot r} \equiv g^{a_1} \cdot (g^r)^{c_1} \equiv A_1 \cdot R^{c_1} \pmod{p} \end{aligned}$$

Moreover, we have $h^{f_0} \equiv g^{x(a_0+c_0 \cdot r)} \pmod{p}$ and $B_0 \cdot S^{c_0} \equiv g^{b_0+c_0(x \cdot r+m)}$ (mod p), hence, to show $h^{f_0} \equiv B_0 \cdot S^{c_0}$

(mod p), it is sufficient to show $(b_0+c_0 \cdot m) \equiv x \cdot a_0 \pmod{q}$:

$$\begin{aligned} &b_0 + c_0 \cdot m \\ &\equiv b_0 + c \cdot m - m \cdot c_1 \\ &\equiv b_0 + c \cdot m - \frac{b_1 \cdot m - a_1 \cdot m \cdot x}{1-m} \\ &\equiv \frac{(b_0+c \cdot m)(1-m) - b_1 \cdot m + a_1 \cdot m \cdot x}{1-m} \\ &\equiv \frac{(b_0+c \cdot m)(1-m) - b_1 \cdot m + \frac{a_1 \cdot m \cdot ((b_0+c \cdot m)(1-m) - b_1 \cdot m)}{a_0(1-m) - a_1 \cdot m}}{1-m} \\ &\equiv \frac{(a_0(1-m) - a_1 \cdot m)((b_0+c \cdot m)(1-m) - b_1 \cdot m)}{1-m} \\ &\quad + \frac{a_1 \cdot m \cdot ((b_0+c \cdot m)(1-m) - b_1 \cdot m)}{(1-m)(a_0(1-m) - a_1 \cdot m)} \\ &\equiv \frac{a_0(1-m)((b_0+c \cdot m)(1-m) - b_1 \cdot m)}{(1-m)(a_0(1-m) - a_1 \cdot m)} \\ &\equiv \frac{a_0 \cdot ((b_0+c \cdot m)(1-m) - b_1 \cdot m)}{a_0(1-m) - a_1 \cdot m} \\ &\equiv x \cdot a_0 \pmod{q} \end{aligned}$$

Similarly, $h^{f_1} \equiv g^{x(a_1+c_1 \cdot r)} \pmod{p}$ and $B_1 \cdot (S/g)^{c_1} \equiv g^{b_1+c_1(x \cdot r+m-1)}$ (mod p), hence, to show $h^{f_1} \equiv B_1 \cdot (S/g)^{c_1}$ (mod p), it is sufficient to show $b_1 + c_1(m-1) \equiv a_1 \cdot x$ (mod q):

$$\begin{aligned} &b_1 + c_1(m-1) \\ &\equiv b_1 + \frac{(m-1)(b_1 - a_1 \cdot x)}{1-m} \\ &\equiv \frac{b_1(1-m) + (m-1)(b_1 - a_1 \cdot x)}{1-m} \\ &\equiv \frac{a_1 \cdot x(1-m)}{1-m} \\ &\equiv a_1 \cdot x \pmod{q} \end{aligned}$$

Furthermore, we have

$$\begin{aligned} \mathcal{H}(A_0, B_0, A_1, B_1) &\equiv c_0 + c_1 \equiv c - c_1 + c_1 \\ &\equiv \mathcal{H}(A_0, B_0, A_1, B_1) - c_1 + c_1 \pmod{q} \end{aligned}$$

It follows that $\text{VerCiph}((pk, e, \{0, 1\}), \sigma, k) = 1$, concluding our proof. \square

APPENDIX F

PROOF: HELIOS'16 IS VERIFIABLE

Elections schemes constructed from generalized Helios satisfy individual (§F-A) and universal (§F-B) verifiability, assuming cryptographic primitives satisfy certain properties that we identify. It follows that Helios'16 satisfies election verifiability with external authentication (§F-C).

A. Individual verifiability

Definition 28 (Collision-free). *Suppose $\Gamma = (\text{Gen}, \text{Enc}, \text{Dec})$ is an asymmetric encryption scheme, Σ_1 proves correct key construction, \mathcal{H} is a hash function, and \mathfrak{m} and \mathfrak{m}' are message spaces such that $\mathfrak{m} \subseteq \mathfrak{m}'$. Let $\text{FS}(\Sigma_1, \mathcal{H}) = (\text{ProveKey}, \text{VerKey})$. If for all security parameters k , public keys pk , proofs ρ , messages $m_1, m_2 \in \mathfrak{m}$, and coins r_1 and r_2 , we have*

$$\begin{aligned} \text{VerKey}((k, pk, \mathfrak{m}'), \rho, k) &= 1 \wedge (m_1 \neq m_2 \vee r_1 \neq r_2) \\ &\Rightarrow \text{Enc}(pk, m_1; r_1) \neq \text{Enc}(pk, m_2; r_2) \end{aligned}$$

Then we say Γ is collision-free for \mathfrak{m} .

Proposition 18. *Suppose $\Gamma, \Sigma_1, \Sigma_2, \Sigma_3$ and \mathcal{H} satisfy the preconditions of Definition 24. Further suppose that Γ is collision-free for $\{0, 1\}$. We have $\text{Helios}(\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \mathcal{H})$ satisfies individual verifiability.*

Proof. Let $\text{Helios}(\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \mathcal{H}) = (\text{Setup}, \text{Vote}, \text{Tally}, \text{Verify})$, $\Gamma = (\text{Gen}, \text{Enc}, \text{Dec})$, and $\text{FS}(\Sigma_1, \mathcal{H}) = (\text{ProveKey}, \text{VerKey})$. Suppose k is a security parameter, $PK_{\mathcal{T}}$ is a public key, n_C is an integer, and β and β' are choices. Further suppose b is an output of $\text{Vote}(PK_{\mathcal{T}}, n_C, \beta, k)$ and b' is an output of $\text{Vote}(PK_{\mathcal{T}}, n_C, \beta', k)$ such that $b \neq \perp$ and $b' \neq \perp$. By definition of Vote , we have $PK_{\mathcal{T}}$ parses as a vector (pk, m, ρ) and $\text{VerKey}((k, pk, m), \rho, k) = 1$. Moreover, $b[1]$ is an output of $\text{Enc}(pk, m)$, and $b'[1]$ is an output of $\text{Enc}(pk, m')$, where $m, m' \in \{0, 1\}$. Furthermore, the ciphertexts are constructed using coins chosen uniformly at random—i.e., the coins used by $b[1]$ and $b'[1]$ will be distinct with overwhelming probability. Since Γ is collision-free for $\{0, 1\}$, we have $b[1] \neq b'[1]$ and $b \neq b'$ with overwhelming probability, concluding our proof. \square

B. Universal verifiability

Lemma 19. *Suppose $\Gamma, \Sigma_1, \Sigma_2, \Sigma_3$ and \mathcal{H} satisfy the preconditions of Definition 24. Further suppose Γ is collision-free for $\{0, 1\}$. We have $\text{Helios}(\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \mathcal{H})$ satisfies Injectivity.*

The proof of Lemma 19 is similar to the proof of Proposition 18.

Proof. Let $\text{Helios}(\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \mathcal{H}) = (\text{Setup}, \text{Vote}, \text{Tally}, \text{Verify})$, $\Gamma = (\text{Gen}, \text{Enc}, \text{Dec})$, and $\text{FS}(\Sigma_1, \mathcal{H}) = (\text{ProveKey}, \text{VerKey})$. Suppose k is a security parameter, $PK_{\mathcal{T}}$ is a public key, n_C is an integer, and β and β' are choices such that $\beta \neq \beta'$. Further suppose b is an output of $\text{Vote}(PK_{\mathcal{T}}, n_C, \beta, k)$ and b' is an output of $\text{Vote}(PK_{\mathcal{T}}, n_C, \beta', k)$ such that $b \neq \perp$ and $b' \neq \perp$. By definition of Vote , we have $PK_{\mathcal{T}}$ is a vector (pk, m, ρ) and $\text{VerKey}((k, pk, m), \rho, k) = 1$. Moreover, there exist coins r and r' such that

$$b[1] = \text{Enc}(pk, m; r), \text{ where } m = \begin{cases} 1 & \text{if } \beta = 1 \\ 0 & \text{otherwise} \end{cases}$$

and

$$b'[1] = \text{Enc}(pk, m'; r'), \text{ where } m' = \begin{cases} 1 & \text{if } \beta' = 1 \\ 0 & \text{otherwise} \end{cases}$$

Since $\beta \neq \beta'$, we have $m \neq m'$. And, since Γ is collision-free for $\{0, 1\}$, we have $b[1] \neq b'[1]$ and, therefore, $b \neq b'$, concluding our proof. \square

Proposition 20. *Suppose $\Gamma, \Sigma_1, \Sigma_2, \Sigma_3$ and \mathcal{H} satisfy the preconditions of Definition 24. Further suppose Γ is perfectly correct, perfectly homomorphic, and collision-free for $\{0, 1\}$, Σ_1, Σ_2 and Σ_3 satisfy special soundness and special honest verifier zero-knowledge, and \mathcal{H} is a random oracle. We have $\text{Helios}(\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \mathcal{H})$ satisfies universal verifiability.*

Proof. Let $\Pi = \text{Helios}(\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \mathcal{H}) = (\text{Setup}, \text{Vote}, \text{Tally}, \text{Verify})$, $\text{FS}(\Sigma_1, \mathcal{H}) = (\text{ProveKey}, \text{VerKey})$, $\text{FS}(\Sigma_2, \mathcal{H}) = (\text{ProveCiph}, \text{VerCiph})$, and $\text{FS}(\Sigma_3, \mathcal{H}) = (\text{ProveDec}, \text{VerDec})$. By Theorem 14, each of the non-interactive proof systems satisfies simulation sound extractability.

Suppose k is a security parameter and \mathcal{A} is a PPT adversary. Further suppose that an execution of $\text{Exp-UV-Ext}(\Pi, \mathcal{A}, k)$ computes

$$\begin{aligned} (PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, P) &\leftarrow \mathcal{A}(k); \\ \mathbf{Y} &\leftarrow \text{correct-tally}(PK_{\mathcal{T}}, BB, n_C, k) \end{aligned}$$

such that $\text{Verify}(PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, P, k) = 1$. (If $\text{Verify}(PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, P, k) \neq 1$, then we can conclude immediately.) We focus on the case $n_C > 1$; the case $n_C = 1$ is similar.

By definition of the verification algorithm, vector \mathbf{X} is of length n_C and P is a vector of length $n_C - 1$. Moreover, $PK_{\mathcal{T}}$ is a vector (pk, m, ρ) . Let $\{b_1, \dots, b_\ell\}$ be the largest subset of BB such that for all $1 \leq i \leq \ell$ we have b_i is a vector of length $2 \cdot n_C - 1$ and $\bigwedge_{j=1}^{n_C-1} \text{VerCiph}((pk, b_i[j], \{0, 1\}), b_i[j + n_C - 1], j, k) = 1 \wedge \text{VerCiph}((pk, b_i[1] \otimes \dots \otimes b_i[n_C - 1], \{0, 1\}), b_i[2 \cdot n_C - 1], n_C, k) = 1$.

We have for all choices $\beta \in \{1, \dots, n_C\}$, coins r and ballots $b = \text{Vote}(PK_{\mathcal{T}}, n_C, \beta, k; r)$ that $b \notin BB \setminus \{b_1, \dots, b_\ell\}$ with overwhelming probability, since such an occurrence would imply a contradiction: $\{b_1, \dots, b_\ell\}$ is not the largest subset of BB satisfying the conditions given by the tally algorithm, because b is a vector of length $2 \cdot n_C - 1$ such that $\bigwedge_{j=1}^{n_C-1} \text{VerCiph}((pk, b[j], \{0, 1\}), b[j + n_C - 1], j, k) = 1 \wedge \text{VerCiph}((pk, b[1] \otimes \dots \otimes b[n_C - 1], \{0, 1\}), b[2 \cdot n_C - 1], n_C, k) = 1$ with overwhelming probability, but $b \notin \{b_1, \dots, b_\ell\}$. It follows that:

$$\begin{aligned} \text{correct-tally}(PK_{\mathcal{T}}, BB, n_C, k) \\ = \text{correct-tally}(PK_{\mathcal{T}}, \{b_1, \dots, b_\ell\}, n_C, k) \end{aligned} \quad (1)$$

A proof of (1) follows from the definition of function correct-tally . If $\{b_1, \dots, b_\ell\} = \emptyset$, then \mathbf{Y} is a vector of length n_C such that $\bigwedge_{j=1}^{n_C} \mathbf{Y}[j] = 0$ by definition of function correct-tally and (1), and, since $\bigwedge_{i=j}^{n_C} \mathbf{X}[j] = 0$, we have $\mathbf{X} = \mathbf{Y}$ by definition of the verification algorithm, hence, $\text{Exp-UV-Ext}(\Pi, \mathcal{A}, k)$ outputs 0 with overwhelming probability and $\text{Succ}(\text{Exp-UV-Ext}(\Pi, \mathcal{A}, k))$ is negligible, concluding our proof. Otherwise ($\{b_1, \dots, b_\ell\} \neq \emptyset$), we proceed as follows.

By definition of the verification algorithm, we have $\text{VerKey}((k, pk, m), \rho, k) = 1$. Moreover, by simulation sound extractability, we are assured that pk is an output of Gen with overwhelming probability—i.e., there exists s and sk such that $(pk, sk, m) = \text{Gen}(k; s)$.

By simulation sound extractability, with overwhelming probability, for all $1 \leq i \leq \ell$ there exists messages $m_{i,1}, \dots, m_{i,n_C-1} \in \{0, 1\}$ and coins $r_{i,1}, \dots, r_{i,2 \cdot n_C - 2}$ such that for all $1 \leq j \leq n_C - 1$ we have

$$\begin{aligned} b_i[j + n_C - 1] &= \text{ProveCiph}((pk, b_i[j], \{0, 1\}), \\ &\quad (m_{i,j}, r_{i,j}), j, k; r_{i,j+n_C-1}) \end{aligned}$$

and

$$b_i[j] = \text{Enc}(pk, m_{i,j}; r_{i,j}).$$

Moreover, for all $1 \leq i \leq \ell$ we have $\sum_{j=1}^{n_C-1} m_{i,j} \in \{0,1\}$ and there exist coins $r_{i,2 \cdot n_C - 1}$ such that

$$b_i[2 \cdot n_C - 1] = \text{ProveCiph}(pk, c, \{0,1\}), \\ (m, r), n_C, k; r_{i,2 \cdot n_C - 1})$$

with overwhelming probability, where $c \leftarrow b_i[1] \otimes \cdots \otimes b_i[n_C - 1]$, $m \leftarrow m_{i,1} \odot \cdots \odot m_{i,n_C-1}$, and $r \leftarrow r_{i,1} \oplus \cdots \oplus r_{i,n_C-1}$.

By inspection of Vote, for all $1 \leq i \leq \ell$ there exists β_i, r_i such that

$$b_i = \text{Vote}(PK_{\mathcal{T}}, n_C, \beta_i, k; r_i)$$

and either $\beta_i = n_C \wedge \bigwedge_{j=1}^{n_C-1} m_{i,j} = 0$ or $\beta_i \in \{1, \dots, n_C - 1\} \wedge m_{i,\beta_i} = 1 \wedge \bigwedge_{j \in \{1, \dots, \beta_i-1, \beta_i+1, \dots, n_C-1\}} m_{i,j} = 0$. It follows for all $1 \leq i \leq \ell$ and $1 \leq j \leq n_C - 1$ that:

$$m_{i,j} = 0 \iff \beta_i = n_C \vee \beta_i \neq j \quad (2)$$

$$m_{i,j} = 1 \iff \beta_i = j \quad (3)$$

Moreover, for all $1 \leq i \leq \ell$ we have:

$$\sum_{j=1}^{n_C-1} m_{i,j} = 0 \iff \beta_i = n_C \quad (4)$$

Furthermore, we have the following facts:

Fact 1. For all integers β and n such that $1 \leq \beta \leq n_C$, we have:

$$\exists^{=n} b \in (\{b_1, \dots, b_\ell\} \setminus \{\perp\}) : \\ \exists r : b = \text{Vote}(PK_{\mathcal{T}}, n_C, \beta, k; r) \\ \iff \exists^{=n} i \in \{1, \dots, \ell\} : \beta = \beta_i$$

Fact 2. For all integers j and n such that $1 \leq j \leq n_C - 1$, we have:

$$\exists^{=n} i \in \{1, \dots, \ell\} : \beta_i = j \iff n = \sum_{i=1}^{\ell} m_{i,j}$$

Proof of Fact 2. For the forward implication, suppose j, n are integers such that $1 \leq j \leq n_C - 1$ and $\exists^{=n} i \in \{1, \dots, \ell\} : \beta_i = j$. We proceed by induction on ℓ . In the base case ($\ell = 0$), we have $n = 0$, hence, $n = \sum_{i=1}^{\ell} m_{i,j}$. In the inductive case, we distinguish two cases. Case I: $\exists^{=n} i \in \{1, \dots, \ell - 1\} : \beta_i = j$ holds. We have $\beta_\ell \neq j$ by definition of the counting quantifier and, hence, $m_{i,j} = 0$ by (2). By our induction hypothesis, we derive $n = \sum_{i=1}^{\ell-1} m_{i,j} = \sum_{i=1}^{\ell} m_{i,j}$. Case II: $\exists^{=n} i \in \{1, \dots, \ell - 1\} : \beta_i = j$ does not hold. We have $\beta_\ell = j$ by definition of the counting quantifier and, hence, $m_{i,j} = 1$ by (3). Moreover, we have $\exists^{=n-1} i \in \{1, \dots, \ell - 1\} : \beta_i = j$ holds. By our induction hypothesis, we derive $n - 1 = \sum_{i=1}^{\ell-1} m_{i,j}$, that is, $n = \sum_{i=1}^{\ell} m_{i,j}$.

For the reverse implication, suppose j, n are integers such that $1 \leq j \leq n_C - 1$ and $n = \sum_{i=1}^{\ell} m_{i,j}$. We proceed by induction on ℓ . In the base case ($\ell = 0$), we have $n = 0$, hence, $\exists^{=n} i \in \{1, \dots, \ell\} : \beta_i = j$. In the inductive case, we distinguish two cases. Case I: $n = \sum_{i=1}^{\ell-1} m_{i,j}$. We have $m_{\ell,j} = 0$, hence, $\beta_\ell \neq j$ by (2). By our induction hypothesis, we have $\exists^{=n} i \in \{1, \dots, \ell - 1\} : \beta_i = j$. Since $\beta_\ell \neq j$, the

result follows. Case II: $n \neq \sum_{i=1}^{\ell-1} m_{i,j}$. Since $m_{\ell,j} \in \{0,1\}$, we have $m_{\ell,j} = 1$, hence, $\beta_\ell = j$ by (3). Moreover, we have $n - 1 = \sum_{i=1}^{\ell-1} m_{i,j}$. By our induction hypothesis, we derive $\exists^{=n-1} i \in \{1, \dots, \ell - 1\} : \beta_i = j$. The result follows.

Fact 3. For all integers n , we have

$$\exists^{=n} i \in \{1, \dots, \ell\} : \beta_i = n_C \iff n = \ell - \sum_{j=1}^{n_C-1} \sum_{i=1}^{\ell} m_{i,j}$$

Proof of Fact 3. For the forward implication, suppose $\exists^{=n} i \in \{1, \dots, \ell\} : \beta_i = n_C$. We proceed by induction on ℓ . In the base case ($\ell = 0$), we have $n = 0$, hence, $n = \ell - \sum_{j=1}^{n_C-1} \sum_{i=1}^{\ell} m_{i,j}$. In the inductive case, we distinguish two cases. Case I: $\exists^{=n} i \in \{1, \dots, \ell - 1\} : \beta_i = n_C$ holds. We have $\beta_\ell \neq n_C$ by definition of the counting quantifier and we derive $\sum_{j=1}^{n_C-1} m_{\ell,j} \neq 0$ by (4). Moreover, since $\sum_{j=1}^{n_C-1} m_{\ell,j} \in \{0,1\}$, we have $\sum_{j=1}^{n_C-1} m_{\ell,j} = 1$. By our induction hypothesis, we derive $n = \ell - 1 - \sum_{j=1}^{n_C-1} \sum_{i=1}^{\ell-1} m_{i,j} = \ell - \sum_{j=1}^{n_C-1} \sum_{i=1}^{\ell} m_{i,j}$. Case II: $\exists^{=n} i \in \{1, \dots, \ell - 1\} : \beta_i = n_C$ does not hold. We have $\beta_\ell = n_C$ by definition of the counting quantifier and we derive $\sum_{j=1}^{n_C-1} m_{i,j} = 0$ by (4). Moreover, we have $\exists^{=n-1} i \in \{1, \dots, \ell - 1\} : \beta_i = n_C$ holds. By our induction hypothesis, we derive $n - 1 = \ell - 1 - \sum_{j=1}^{n_C-1} \sum_{i=1}^{\ell-1} m_{i,j}$, that is, $n = \ell - \sum_{j=1}^{n_C-1} \sum_{i=1}^{\ell} m_{i,j} = \ell - \sum_{j=1}^{n_C-1} \sum_{i=1}^{\ell} m_{i,j}$.

For the reverse implication, suppose $n = \ell - \sum_{j=1}^{n_C-1} \sum_{i=1}^{\ell} m_{i,j}$. We proceed by induction on ℓ . In the base case ($\ell = 0$), we have $n = 0$, hence, $\exists^{=n} i \in \{1, \dots, \ell\} : \beta_i = n_C$. In the inductive case, we distinguish two cases. Case I: $n = \ell - 1 - \sum_{j=1}^{n_C-1} \sum_{i=1}^{\ell-1} m_{i,j}$. We have $\sum_{j=1}^{n_C-1} m_{\ell,j} = 1$. Since $m_{\ell,1}, \dots, m_{\ell,n_C-1} \in \{0,1\}$, there exists j such that $1 \leq j \leq n_C - 1$ and $m_{\ell,j} = 1$, moreover, $\beta_\ell = j$ by (3), hence, $\beta_\ell \neq n_C$. By our induction hypothesis, we derive $\exists^{=n} i \in \{1, \dots, \ell - 1\} : \beta_i = n_C$. The result follows. Case II: $n \neq \ell - 1 - \sum_{j=1}^{n_C-1} \sum_{i=1}^{\ell-1} m_{i,j}$. Since $\sum_{j=1}^{n_C-1} m_{\ell,j} \in \{0,1\}$, we have $\sum_{j=1}^{n_C-1} m_{\ell,j} = 0$, and we derive $\beta_i = n_C$ by (4). Moreover, we have $n - 1 = \ell - 1 - \sum_{j=1}^{n_C-1} \sum_{i=1}^{\ell-1} m_{i,j}$. By our induction hypothesis, we derive $\exists^{=n-1} i \in \{1, \dots, \ell - 1\} : \beta_i = n_C$. The result follows.

We proceed the proof of Proposition 20 using the above facts.

By definition of the verification algorithm, we have $\bigwedge_{j=1}^{n_C-1} \text{VerDec}((pk, b_1[j] \otimes \cdots \otimes b_\ell[j], \mathbf{X}[j]), P[j], k) = 1 \wedge \mathbf{X}[n_C] = \ell - \sum_{j=1}^{n_C-1} \mathbf{X}[j]$. By simulation sound extractability, we have for all $1 \leq j \leq n_C - 1$ that $\mathbf{X}[j] = \text{Dec}(sk, b_1[j] \otimes \cdots \otimes b_\ell[j])$ with overwhelming probability. Although, public key pk may not have been constructed using coins chosen uniformly at random, we nevertheless have for all $1 \leq j \leq n_C - 1$ that $b_1[j] \otimes \cdots \otimes b_\ell[j]$ is a ciphertext with overwhelming probability, because Γ is perfectly homomorphic. Similarly, for all $1 \leq j \leq n_C - 1$, although ciphertext $b_1[j] \otimes \cdots \otimes b_\ell[j]$ may not have been constructed using coins chosen uniformly at random nor using a public key that was constructed using coins chosen uniformly, and although private key sk may not have been constructed using coins chosen uniformly, we have $\text{Dec}(\mathbf{X}[j], k) = 1$.

$sk, b_1[j] \otimes \cdots \otimes b_\ell[j]) = m_{1,j} \odot \cdots \odot m_{\ell,j}$ with overwhelming probability, because Γ is perfectly correct. Let m_B be the largest integer such that $\{0, \dots, m_B\} \subseteq \mathfrak{m}$. By definition of the verification algorithm, we have $\ell \leq m_B$. It follows that $m_{1,j} \odot \cdots \odot m_{\ell,j} = \sum_{i=1}^{\ell} m_{i,j}$, hence,

$$\mathbf{X}[j] = \sum_{i=1}^{\ell} m_{i,j} r$$

with overwhelming probability. By definition of function *correct-tally*, (1) and Fact 1, we have \mathbf{Y} is a vector of length n_C such that for all $1 \leq \beta \leq n_C$ we have

$$\mathbf{Y}[\beta] = n \text{ if } \exists \beta_i \in \{1, \dots, \ell\} : \beta = \beta_i$$

It follows by Facts 2 and 3 that for all $1 \leq \beta \leq n_C$ we have $\mathbf{X}[\beta] = \mathbf{Y}[\beta]$ with overwhelming probability, hence, $\mathbf{X} = \mathbf{Y}$ with overwhelming probability, therefore, $\text{Exp-UV-Ext}(\Pi, \mathcal{A}, k)$ outputs 0 with overwhelming probability and $\text{Succ}(\text{Exp-UV-Ext}(\Pi, \mathcal{A}, k))$ is negligible, concluding our proof. \square

Proposition 21. *Suppose $\Gamma, \Sigma_1, \Sigma_2, \Sigma_3$ and \mathcal{H} satisfy the preconditions of Definition 24. Further suppose Σ_2 satisfies special soundness and special honest verifier zero-knowledge, and \mathcal{H} is a random oracle. We have $\text{Helios}(\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \mathcal{H})$ satisfies Completeness.*

Proof. Let $\text{Helios}(\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \mathcal{H}) = (\text{Setup}, \text{Vote}, \text{Tally}, \text{Verify})$, $\text{FS}(\Sigma_1, \mathcal{H}) = (\text{ProveKey}, \text{VerKey})$, $\text{FS}(\Sigma_2, \mathcal{H}) = (\text{ProveCiph}, \text{VerCiph})$, and $\text{FS}(\Sigma_3, \mathcal{H}) = (\text{ProveDec}, \text{VerDec})$. Suppose k is a security parameter and \mathcal{A} is a PPT adversary. Further suppose $(PK_{\mathcal{T}}, SK_{\mathcal{T}}, m_B, m_C)$ is an output of $\text{Setup}(k)$, (BB, n_C) is an output of $\mathcal{A}(PK_{\mathcal{T}}, k)$, and (\mathbf{X}, P) is an output of $\text{Tally}(SK_{\mathcal{T}}, BB, n_C, k)$. Moreover, suppose $|BB| \leq m_B$. We focus on the case $n_C > 1$; the case $n_C = 1$ is similar. By definition of Setup , there exist coins s such that $(pk, sk, \mathfrak{m}) = \text{Gen}(k; s)$, $PK_{\mathcal{T}} = (pk, \mathfrak{m}, \rho)$, $SK_{\mathcal{T}} = (pk, sk)$ and m_B is the largest integer such that $\{0, \dots, m_B\} \subseteq \{0\} \cup \mathfrak{m}$, where ρ is an output of $\text{ProveKey}((k, pk, \mathfrak{m}), (sk, s), k)$. By definition of Tally , we have \mathbf{X} is a vector of length n_C and P is a vector of length $n_C - 1$. It follows that Verify can successfully parse \mathbf{X}, P , and $PK_{\mathcal{T}}$. Moreover, by the completeness of $(\text{ProveKey}, \text{VerKey})$, we have $\text{VerKey}((k, pk, \mathfrak{m}), \rho, k) = 1$ with overwhelming probability. Let $\{b_1, \dots, b_\ell\}$ be the largest subset of BB satisfying the conditions given by the tally algorithm. If $\{b_1, \dots, b_\ell\} = \emptyset$, then \mathbf{X} is a zero-filled vector and Verify outputs 1, concluding our proof, otherwise, we proceed as follows. Since $\{b_1, \dots, b_\ell\}$ is a subset of BB , we have $\ell \leq m_B$. By definition of Tally , we have for all $1 \leq i \leq \ell$ that $\bigwedge_{j=1}^{n_C-1} \text{VerCiph}((pk, b_i[j], \{0, 1\}), b_i[j + n_C - 1], j, k) = 1$. By Theorem 14, we have $(\text{ProveCiph}, \text{VerCiph})$ satisfies simulation sound extractability, hence, for all $1 \leq i \leq \ell$ and all $1 \leq j \leq n_C - 1$ we have $b_i[j]$ is a ciphertext with overwhelming probability. And, because Γ is homomorphic, we have $b_1[j] \otimes \cdots \otimes b_\ell[j]$ is also a ciphertext with overwhelming probability. By definition of Tally and completeness

of $(\text{ProveDec}, \text{VerDec})$, we have $\bigwedge_{j=1}^{n_C-1} \text{VerDec}((pk, b_1[j] \otimes \cdots \otimes b_\ell[j], \mathbf{X}[j]), P[j], k) = 1 \wedge \mathbf{X}[n_C] = \ell - \sum_{j=1}^{n_C-1} \mathbf{X}[j]$ with overwhelming probability, hence, Verify outputs 1 with overwhelming probability, concluding our proof. \square

C. Proof: Theorem 5

By Propositions 18, 20 & 21 and Lemma 19, election schemes constructed from generalized Helios satisfy election verifiability with external authentication:

Corollary 22. *Suppose $\Gamma, \Sigma_1, \Sigma_2, \Sigma_3$ and \mathcal{H} satisfy the preconditions of Definition 24. Further suppose that Γ is perfectly correct, perfectly homomorphic and collision-free for $\{0, 1\}$, Σ_1, Σ_2 and Σ_3 satisfy special soundness and special honest verifier zero-knowledge, and \mathcal{H} is a random oracle. We have $\text{Helios}(\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \mathcal{H})$ satisfies election verifiability with external authentication.*

Proof of Theorem 5. Let $\text{Helios}'16$ be the set of election schemes derived from $\text{Helios}(\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \mathcal{H})$, where primitives $\Gamma, \Sigma_1, \Sigma_2, \Sigma_3$ and \mathcal{H} satisfy the conditions identified in Corollary 22. Hence, Theorem 5 is an immediate consequence of Corollary 22 \square

A non-interactive proof system $(\text{ProveKey}, \text{VerKey})$ derived from a sigma protocol for proving correct key construction is sufficient to ensure that additively homomorphic El Gamal [56, §2] is collision-free (Lemma 23), assuming algorithm VerKey guarantees that public keys are constructed from suitable parameters: if $\text{VerKey}((k, pk, \mathfrak{m}), \rho, k) = 1$, then there exists p, q, g and h such that $pk = (p, q, g, h)$ and (p, q, g) are *cryptographic parameters*—i.e., $p = 2 \cdot q + 1$, $|q| = k$, and g is a generator of \mathbb{Z}_p^* of order q . Thus, since El Gamal is perfectly correct and perfectly homomorphic, we have additively homomorphic El Gamal is a suitable asymmetric encryption scheme to instantiate Helios'16.

Lemma 23. *Suppose Σ_1 is a sigma protocol that proves correct key construction and \mathcal{H} is a hash function. Let $\text{FS}(\Sigma_1, \mathcal{H}) = (\text{ProveKey}, \text{VerKey})$. Further suppose for all security parameters k , public keys pk , message spaces \mathfrak{m} , and proofs ρ , we have $\text{VerKey}((k, pk, \mathfrak{m}), \rho, k) = 1$ implies $h \neq 0$ and there exists p, q, g and h such that $pk = (p, q, g, h)$ and (p, q, g) are cryptographic parameters. It follows that additively homomorphic El Gamal is collision-free for $\{0, 1\}$.*

Proof. Suppose k is a security parameter, pk is a public key, ρ is a proof, $m_1, m_2 \in \{0, 1\}$ are messages and r_1 and r_2 are coins such that $\text{VerKey}((k, pk, \mathfrak{m}), \rho, k) = 1$, $m_1 \neq m_2 \vee r_1 \neq r_2$, $pk = (p, q, g, h)$ and (p, q, g) are cryptographic parameters, for some p, q, g and h . Further suppose that c_1 and c_2 are ciphertexts such that $c_1 = \text{Enc}(pk, m_1; r_1)$, $c_2 = \text{Enc}(pk, m_2; r_2)$, and Enc is El Gamal's encryption algorithm. If $r_1 \neq r_2$, then we proceed as follows. By definition of Enc , we have $c_1[1] = g^{r_1} \pmod{p}$ and $c_2[1] = g^{r_2} \pmod{p}$. Since r_1 and r_2 are distinct, we have $g^{r_1} \not\equiv g^{r_2} \pmod{p}$. (We implicitly assume that coins r_1 and r_2 are selected from the coin space \mathbb{Z}_q^* , hence, $g^{r_1} = g^{r_1} \pmod{p}$)

and $g^{r_2} = g^{r_2} \pmod p$.) It follows that $c_1 \neq c_2$. Otherwise ($r_1 = r_2$), we have $m_1 \neq m_2$ and we proceed as follows. By definition of Enc, we have $c_1[2] = h^{r_1} \cdot g_1^{m_1} \pmod p$ and $c_2[2] = h^{r_2} \cdot g_2^{m_2} \pmod p$. Since (p, q, g) are cryptographic parameters and $h \neq 0$, we have $h^{r_1} \not\equiv h^{r_1} \cdot g \pmod p$, which is sufficient to conclude, because $m_1, m_2 \in \{0, 1\}$. \square

The sigma protocol for proving knowledge of discrete logarithms by Chaum et al. [37, Protocol 2] does not explicitly require the suitability of cryptographic parameters to be checked, hence, Lemma 23 is not immediately applicable. Nonetheless, we can trivially make the necessary checks explicit and, hence, the non-interactive proof system derived from the sigma protocol for proving knowledge of discrete logarithms by Chaum et al. is sufficient to ensure that El Gamal is collision-free for $\{0, 1\}$. We can also trivially include the checks proposed by Chang-Fong & Essex [32, §4]. These modifications should suffice to ensure special soundness and special honest verifier zero-knowledge. Similarly, it should be possible to modify the sigma protocols for proving knowledge of disjunctive equality between discrete logarithms by Cramer et al. [55, Figure 1] and for proving knowledge of equality between discrete logarithms by Chaum and Pedersen [38, §3.2] to ensure that they satisfy special soundness and special honest verifier zero-knowledge. Thus, the modified sigma protocols should be suitable to instantiate Helios'16.

APPENDIX G

PROOF: Exp-EV-Int \Rightarrow Exp-IV-Int

Our eligibility verifiability experiment (§IV-B3) asserts that no one can construct a ballot that appears to be associated with public credential pk unless they know private credential sk . It follows that a voter can uniquely identify their ballot on the bulletin board, because no one else knows their private credential. Eligibility verifiability therefore implies individual verifiability (Theorem 7).

Our proof of Theorem 7 is reliant on distinct credentials, which is a consequence of eligibility verifiability:

Lemma 24. *If an election scheme Π satisfies strong eligibility verifiability, then there exists a negligible function μ , such that for all security parameters k , we have*

$$\begin{aligned} &Pr[(PK_{\mathcal{T}}, SK_{\mathcal{T}}, m_B, m_C) \leftarrow \text{Setup}(k); \\ &\quad (pk_0, sk_0) \leftarrow \text{Register}(PK_{\mathcal{T}}, k); \\ &\quad (pk_1, sk_1) \leftarrow \text{Register}(PK_{\mathcal{T}}, k) : \\ &\quad\quad sk_0 = sk_1] \leq \mu(k) \end{aligned}$$

Proof. Suppose an election scheme Π satisfies Exp-EV-Int, but

$$\begin{aligned} &Pr[(PK_{\mathcal{T}}, SK_{\mathcal{T}}, m_B, m_C) \leftarrow \text{Setup}(k); \\ &\quad (pk_0, sk_0) \leftarrow \text{Register}(PK_{\mathcal{T}}, k); \\ &\quad (pk_1, sk_1) \leftarrow \text{Register}(PK_{\mathcal{T}}, k) : \\ &\quad\quad sk_0 = sk_1] \geq \frac{1}{p(k)} \end{aligned}$$

for some polynomial function p and security parameter k . Then we can construct an adversary \mathcal{A} that wins Exp-EV-Int as follows. Adversary \mathcal{A} is given input k and runs Setup to obtain a key pair $(PK_{\mathcal{T}}, SK_{\mathcal{T}})$, chooses some positive integer n_V , and outputs $(PK_{\mathcal{T}}, n_V)$. The challenger then generates n_V key pairs and gives the set L of public keys to \mathcal{A} . Now \mathcal{A} simply runs Register($PK_{\mathcal{T}}, k$) to get a key pair (pk, sk) , chooses some positive integers n_C and β such that $1 \leq \beta \leq n_C$, computes $b \leftarrow \text{Vote}(sk, PK_{\mathcal{T}}, n_C, \beta, k)$, and outputs (n_C, b) . We know that secret keys generated by Register collide with probability at least $\frac{1}{p(k)}$, so Register must generate a particular secret key sk' with probability $\frac{1}{p(k)}$. Therefore, this sk' will correspond to one of the public keys in L with probability $\frac{n_V}{p(k)}$. Furthermore, the key sk generated by the adversary will be sk' with probability $\frac{1}{p(k)}$. Therefore, b will be a vote constructed under a voter's secret key with probability $\frac{n_V}{p(k)^2}$, so \mathcal{A} wins the experiment with non-negligible probability. \square

A. Proof: Theorem 7

Suppose there exists an adversary \mathcal{A}' that wins Exp-IV-Int(Π, \mathcal{A}', k) with probability $\frac{1}{p(k)}$ for some polynomial function p . Then we can construct an adversary \mathcal{A} that wins Exp-EV-Int(Π, \mathcal{A}, k) with non-negligible probability. Adversary \mathcal{A} is given k as input, which it passes to \mathcal{A}' . Adversary \mathcal{A}' may ask for secret keys from its oracle C , in which case \mathcal{A} forwards these queries to its own, identical oracle. Adversary \mathcal{A} then forwards the oracle's response back to \mathcal{A}' . Adversary \mathcal{A}' then outputs $(PK_{\mathcal{T}}, n_V)$, which is then output by \mathcal{A} . Next, \mathcal{A} is given the public keys (pk_1, \dots, pk_{n_V}) . Adversary \mathcal{A} passes these keys to \mathcal{A}' , which returns $(n_C, \beta, \beta', i, j)$. Any oracle queries made by \mathcal{A}' are handled exactly as before. Now \mathcal{A} queries its oracle C on i . The oracle returns sk_i . Adversary \mathcal{A} computes $b = \text{Vote}(sk_i, PK_{\mathcal{T}}, n_C, \beta)$ and outputs (n_C, β', j, b) . Adversary \mathcal{A}' wins Exp-IV-Int(Π, \mathcal{A}, k) with non-negligible probability, so with non-negligible probability $b = \text{Vote}(sk_j, PK_{\mathcal{T}}, n_C, \beta')$ and \mathcal{A}' (and therefore \mathcal{A}) did not query the oracle on input j . Adversary \mathcal{A} only makes one additional oracle query on input i , so again, \mathcal{A} does not query the oracle on j . Furthermore, by Lemma 24, $sk_i = sk_j$ with only negligible probability. Therefore \mathcal{A} wins Exp-EV-Int(Π, \mathcal{A}, k) with probability $\frac{1}{p(k)} - \text{negl}(k)$. \square

APPENDIX H

VARIANT OF Exp-EV-Int-Weak

Our weak election verifiability experiment with internal authentication (§VI) can be equivalently formulated as an experiment with just one voter:

$$\text{Exp-EV-Int-Weak}'(\Pi, \mathcal{A}, k) =$$

```

1  $(PK_{\mathcal{T}}, SK_{\mathcal{T}}, m_B, m_C) \leftarrow \text{Setup}(k)$ ;
2  $(pk, sk) \leftarrow \text{Register}(PK_{\mathcal{T}}, k)$ ;
3  $Rvld \leftarrow \emptyset$ ;
4  $(n_C, \beta, b) \leftarrow \mathcal{A}^{R'}(PK_{\mathcal{T}}, pk, k)$ ;
5 if  $\exists r : b = \text{Vote}(sk, PK_{\mathcal{T}}, n_C, \beta, k; r) \wedge b \neq \perp \wedge$ 
    $b \notin Rvld$  then
6   | return 1
7 else
8   | return 0

```

Oracle R' is similar to oracle R in Exp-EV-Int-Weak. On invocation $R'(\beta, n_C)$, oracle R' computes $b \leftarrow \text{Vote}(sk, PK_{\mathcal{T}}, n_C, \beta, k)$; $Rvld \leftarrow Rvld \cup \{b\}$ and outputs b .

Lemma 25. *Given an election scheme Π , we have*

$$\forall \mathcal{A} \exists \mu \forall k . \text{Succ}(\text{Exp-EV-Int-Weak}(\Pi, \mathcal{A}, k)) \leq \mu(k)$$

$$\Leftrightarrow \forall \mathcal{A}' \exists \mu' \forall k' . \text{Succ}(\text{Exp-EV-Int-Weak}'(\Pi, \mathcal{A}', k')) \leq \mu'(k'),$$

where \mathcal{A} and \mathcal{A}' are PPT adversaries, μ and μ' are negligible functions, and k and k' are security parameters.

A proof of the forward implication is straightforward, so we omit formalizing a proof. The reverse implication is formally proved below.

Proof. Suppose there exists an adversary \mathcal{A} that wins Exp-EV-Int-Weak with non-negligible probability. Let us construct an adversary \mathcal{B} against Exp-EV-Int-Weak'.

- $\mathcal{B}(PK_{\mathcal{T}}, pk, k)$ computes


```

 $n_V \leftarrow \mathcal{A}(PK_{\mathcal{T}}, k)$ ;
 $i^* \leftarrow_R \{1, \dots, n_V\}$ ;
 $pk_{i^*} \leftarrow pk$ ;
for  $i \in \{1, \dots, i^* - 1, i^* + 1, \dots, n_V\}$  do
  |  $(pk_i, sk_i) \leftarrow \text{Register}(PK_{\mathcal{T}}, k)$ ;
 $L \leftarrow \{pk_1, \dots, pk_{n_V}\}$ ;
 $(n_C, \beta, i, b) \leftarrow \mathcal{A}(L)$ ;
return  $(n_C, \beta, b)$ 

```

responding to \mathcal{A} 's oracle calls $R(i, \beta, n_C)$ by computing **if** $i = i^*$ **then** $b \leftarrow R'(\beta, n_C)$ **else** $b \leftarrow \text{Vote}(sk_i, PK_{\mathcal{T}}, n_C, \beta, k)$ and returning b , and oracle calls $C(i)$ by returning sk_i if $i \neq i^*$ and aborting otherwise.

We prove that \mathcal{B} wins Exp-EV-Int-Weak' with non-negligible probability.

Suppose $(PK_{\mathcal{T}}, SK_{\mathcal{T}}, m_B, m_C)$ is an output of Setup(k) and (pk, sk) is an output of Register($PK_{\mathcal{T}}, k$). Further suppose we compute $\mathcal{B}(PK_{\mathcal{T}}, pk, k)$. If \mathcal{B} does not abort, then it is trivial to see that \mathcal{B} wins Exp-EV-Int-Weak' with non-negligible probability, because \mathcal{B} simulates \mathcal{A} 's challenger and oracles to \mathcal{A} . Hence, it suffices to prove that \mathcal{B} does not abort with non-negligible probability. Suppose n_V is an output of $\mathcal{A}(PK_{\mathcal{T}}, k)$. If $n_V = 1$, then \mathcal{B} aborts with negligible probability, otherwise, \mathcal{B} aborts with probability less than $\frac{1}{n_V}$. Thus, \mathcal{B} does not abort with non-negligible probability, concluding our proof. \square

We formalize a generic construction for JCJ-like election schemes (Definition 30). Our construction is parameterized on the choice of homomorphic encryption scheme and sigma protocols, using the relations introduced in the following definition.⁶⁰

Definition 29. *Let $(\text{Gen}, \text{Enc}, \text{Dec})$ be a homomorphic asymmetric encryption scheme and Σ be a sigma protocol for a binary relation R . Suppose $(pk, sk, m) = \text{Gen}(k; r)$, for some security parameter k and coins r .*

- Σ proves conjunctive plaintext knowledge if $((pk, c_1, \dots, c_k), (m_1, r_1, \dots, m_k, r_k)) \in R \Leftrightarrow \bigwedge_{1 \leq i \leq k} c_i = \text{Enc}(pk, m_i; r_i) \wedge m_i \in m$.
- Σ is a plaintext equivalence test (PET) if $((pk, c, c', i), sk) \in R \Leftrightarrow ((i = 0 \wedge \text{Dec}(sk, c) \neq \text{Dec}(sk, c')) \vee (i = 1 \wedge \text{Dec}(sk, c) = \text{Dec}(sk, c'))) \wedge \text{Dec}(sk, c) \neq \perp \wedge \text{Dec}(sk, c') \neq \perp$.
- Σ is a mixnet if $((pk, \mathbf{c}, \mathbf{c}'), (\mathbf{r}, \chi)) \in R \Leftrightarrow \bigwedge_{1 \leq i \leq |\mathbf{c}|} \mathbf{c}'[i] = \mathbf{c}[\chi(i)] \otimes \text{Enc}(pk, \mathbf{e}; \mathbf{r}[i]) \wedge |\mathbf{c}| = |\mathbf{c}'| = |\mathbf{r}|$, where \mathbf{r} is a vector of coins, χ is a permutation on $\{1, \dots, |\mathbf{c}|\}$, and \mathbf{e} is an identity element of the encryption scheme's message space with respect to \odot .

Definition 30 (Generalized JCJ). *Suppose $\Gamma = (\text{Gen}, \text{Enc}, \text{Dec})$ is a multiplicatively homomorphic asymmetric encryption scheme with a message space over \mathbb{Z}_m^* for some integer m that is super-polynomial in the security parameter, \mathbf{e} is an identity element of Γ 's message space with respect to \odot , Σ_1 proves correct key construction, Σ_2 proves plaintext knowledge in a subspace, Σ_3 proves conjunctive plaintext knowledge, Σ_4 proves correct decryption, Σ_5 is a PET, Σ_6 is a mixnet, and \mathcal{H} is a hash function. Let $\text{FS}(\Sigma_1, \mathcal{H}) = (\text{ProveKey}, \text{VerKey})$, $\text{FS}(\Sigma_2, \mathcal{H}) = (\text{ProveCiph}, \text{VerCiph})$, $\text{FS}(\Sigma_3, \mathcal{H}) = (\text{ProveBind}, \text{VerBind})$, $\text{FS}(\Sigma_4, \mathcal{H}) = (\text{ProveDec}, \text{VerDec})$, $\text{FS}(\Sigma_5, \mathcal{H}) = (\text{ProvePET}, \text{VerPET})$, and $\text{FS}(\Sigma_6, \mathcal{H}) = (\text{ProveMix}, \text{VerMix})$. We define generalized JCJ as $\text{JCJ}(\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5, \Sigma_6, \mathcal{H}) = (\text{Setup}, \text{Register}, \text{Vote}, \text{Tally}, \text{Verify})$:*

- Setup(k). *Select coins r uniformly at random, compute $(pk_T, sk_T, m) \leftarrow \text{Gen}(k; r)$; $\rho \leftarrow \text{ProveKey}((k, pk_T, m), (sk_T, r), k)$; $PK_{\mathcal{T}} \leftarrow (pk_T, m, \rho)$; $SK_{\mathcal{T}} \leftarrow (pk_T, sk_T)$; $m_C \leftarrow |m|$, and output $(PK_{\mathcal{T}}, SK_{\mathcal{T}}, \text{poly}(k), m_C)$.*
- Register($PK_{\mathcal{T}}, k$). *Parse $PK_{\mathcal{T}}$ as (pk_T, m, ρ) , outputting (\perp, \perp) if parsing fails or $\text{VerKey}((k, pk_T, m), \rho, k) \neq 1$. Compute $d \leftarrow_R m$; $pd \leftarrow \text{Enc}(pk_T, d)$ and output (pd, d) .*
- Vote($d, PK_{\mathcal{T}}, n_C, \beta, k$). *Parse $PK_{\mathcal{T}}$ as a vector (pk_T, m, ρ) , outputting \perp if parsing fails or $\text{VerKey}((k, pk_T, m), \rho, k) \neq 1 \vee \beta \notin \{1, \dots, n_C\} \vee \{1, \dots, n_C\} \not\subseteq m$.*

⁶⁰ For brevity, the encryption scheme's message space m is assumed to be $\{1, \dots, |m|\}$.

Select coins r_1 and r_2 uniformly at random, and compute
 $c_1 \leftarrow \text{Enc}(pk_T, \beta; r_1)$;
 $c_2 \leftarrow \text{Enc}(pk_T, d; r_2)$;
 $\sigma \leftarrow \text{ProveCiph}((pk_T, c_1, \{1, \dots, n_C\}), (\beta, r_1), k)$;
 $\tau \leftarrow \text{ProveBind}((pk_T, c_1, c_2), (\beta, r_1, d, r_2), k)$;

Output ballot (c_1, c_2, σ, τ) .

- Tally(SK_T, BB, L, n_C, k). Parse SK_T as (pk_T, sk_T) . Initialize \mathbf{X} as a zero-filled vector of length n_C , and \mathbf{P} as a vector of length 9. Proceed as follows.

1) Remove invalid ballots: Let $\{b_1, \dots, b_\ell\}$ be the largest subset of BB such that $b_1 < \dots < b_\ell$ and for all $1 \leq i \leq \ell$ we have b_i is a vector of length 4 and $\text{VerCiph}((pk_T, b_i[1], \{1, \dots, n_C\}), b_i[3], k) = 1 \wedge \text{VerBind}((pk_T, b_i[1], b_i[2]), b_i[4], k) = 1$. If $\{b_1, \dots, b_\ell\} = \emptyset$, then output (\mathbf{X}, \mathbf{P}) .

2) Eliminating duplicates: Initialize \mathbf{P}_{dupl} as a vector of length ℓ . For each $1 \leq i \leq \ell$, if there exists $j \in \{1, \dots, i-1, i+1, \dots, \ell\}$ such that $\text{VerPET}((pk_T, b_i[2], b_j[2], 1), \sigma, k) = 1$ for some output σ of $\text{ProvePET}((pk_T, b_i[2], b_j[2], 1), sk_T, k)$, then assign $\mathbf{P}_{\text{dupl}}[i] \leftarrow (j, \sigma)$, otherwise, compute $\sigma_j \leftarrow \text{ProvePET}((pk_T, b_i[2], b_j[2], 0), sk_T, k)$ for each $j \in \{1, \dots, i-1, i+1, \dots, \ell\}$ and assign $\mathbf{P}_{\text{dupl}}[i] \leftarrow (0, \sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_\ell)$. Initialize \mathbf{BB} as the empty vector and compute **for** $1 \leq i \leq \ell \wedge \mathbf{P}_{\text{dupl}}[i][1] = 0$ **do** $\mathbf{BB} \leftarrow \mathbf{BB} \parallel (b_i)$, where $\mathbf{BB} \parallel (b_i)$ denotes the concatenation of vectors \mathbf{BB} and (b_i) —i.e., $\mathbf{BB} \parallel (b_i) = (\mathbf{BB}[1], \dots, \mathbf{BB}[\mathbf{BB}], b_i)$.

3) Mixing: Suppose $\mathbf{BB} = (b'_1, \dots, b'_{|\mathbf{BB}|})$, select a permutation χ on $\{1, \dots, |\mathbf{BB}|\}$ uniformly at random, initialize $\mathbf{C}_1, \mathbf{C}_2, \mathbf{r}_1$ and \mathbf{r}_2 as vectors of length $|\mathbf{BB}|$, and fill \mathbf{r}_1 and \mathbf{r}_2 with coins chosen uniformly at random. Compute

```

for  $1 \leq i \leq |\mathbf{BB}|$  do
   $\mathbf{C}_1[i] \leftarrow b'_{\chi(i)}[1] \otimes \text{Enc}(pk_T, \mathbf{e}; \mathbf{r}_1[i])$ ;
   $\mathbf{C}_2[i] \leftarrow b'_{\chi(i)}[2] \otimes \text{Enc}(pk_T, \mathbf{e}; \mathbf{r}_2[i])$ ;
 $\mathbf{BB}_1 \leftarrow (b'_1[1], \dots, b'_{|\mathbf{BB}|}[1])$ ;
 $\mathbf{BB}_2 \leftarrow (b'_1[2], \dots, b'_{|\mathbf{BB}|}[2])$ ;
 $P_{\text{mix},1} \leftarrow \text{ProveMix}((pk_T, \mathbf{BB}_1, \mathbf{C}_1), (\mathbf{r}_1, \chi), k)$ ;
 $P_{\text{mix},2} \leftarrow \text{ProveMix}((pk_T, \mathbf{BB}_2, \mathbf{C}_2), (\mathbf{r}_2, \chi), k)$ ;

```

Similarly, suppose $L = \{pd_1, \dots, pd_{|L|}\}$ such that $pd_1 < \dots < pd_{|L|}$, select a permutation χ' on $\{1, \dots, |L|\}$ uniformly at random, initialize \mathbf{C}_3 and \mathbf{r}_3 as vectors of length $|L|$, fill \mathbf{r}_3 with coins chosen uniformly at random, and compute

```

for  $1 \leq i \leq |L|$  do
   $\mathbf{C}_3[i] \leftarrow pd_{\chi'(i)} \otimes \text{Enc}(pk_T, \mathbf{e}; \mathbf{r}_3[i])$ ;
 $\mathbf{pd} \leftarrow (pd_1, \dots, pd_{|L|})$ ;
 $P_{\text{mix},3} \leftarrow \text{ProveMix}((pk_T, \mathbf{pd}, \mathbf{C}_3), (\mathbf{r}_3, \chi'), k)$ ;

```

4) Remove ineligible ballots: Initialize $\mathbf{P}_{\text{inelig}}$ as a vector of length $|\mathbf{C}_2|$. For each $1 \leq i \leq |\mathbf{C}_2|$, if there exists $j \in \{1, \dots, |\mathbf{C}_3|\}$ such that $\text{VerPET}((pk_T,$

$\mathbf{C}_2[i], \mathbf{C}_3[j], 1), \sigma, k) = 1$ for some output σ of $\text{ProvePET}((pk_T, \mathbf{C}_2[i], \mathbf{C}_3[j], 1), sk_T, k)$, then compute $\mathbf{P}_{\text{inelig}}[i] \leftarrow (j, \sigma)$, otherwise, compute $\sigma_j \leftarrow \text{ProvePET}((pk_T, \mathbf{C}_2[i], \mathbf{C}_3[j], 0), sk_T, k)$ for each $j \in \{1, \dots, |\mathbf{C}_3|\}$ and assign $\mathbf{P}_{\text{inelig}}[i] \leftarrow (0, \sigma_1, \dots, \sigma_{|\mathbf{C}_3|})$. Initialize \mathbf{C}'_1 as the empty vector and compute **for** $1 \leq i \leq \ell \wedge \mathbf{P}_{\text{inelig}}[i][1] \neq 0$ **do** $\mathbf{C}'_1 \leftarrow \mathbf{C}'_1 \parallel (\mathbf{C}_1[i])$.

5) Decrypting: Initialize \mathbf{P}_{dec} as the empty vector. Compute

```

for  $1 \leq i \leq |\mathbf{C}'_1|$  do
   $\beta \leftarrow \text{Dec}(sk_T, \mathbf{C}'_1[i])$ ;
   $\sigma \leftarrow \text{ProveDec}((pk_T, \mathbf{C}'_1[i], \beta), sk_T, k)$ ;
   $\mathbf{X}[\beta] \leftarrow \mathbf{X}[\beta] + 1$ ;
   $\mathbf{P}_{\text{dec}} \leftarrow \mathbf{P}_{\text{dec}} \parallel (\beta, \sigma)$ ;

```

Assign $\mathbf{P} \leftarrow (\mathbf{P}_{\text{dupl}}, \mathbf{C}_1, P_{\text{mix},1}, \mathbf{C}_2, P_{\text{mix},2}, \mathbf{C}_3, P_{\text{mix},3}, \mathbf{P}_{\text{inelig}}, \mathbf{P}_{\text{dec}})$ and output (\mathbf{X}, \mathbf{P}) .

- Verify($PK_T, BB, L, n_C, \mathbf{X}, \mathbf{P}, k$). Parse PK_T as a vector (pk_T, \mathbf{m}, ρ) , \mathbf{X} as a vector of length n_C , and \mathbf{P} as a vector $(\mathbf{P}_{\text{dupl}}, \mathbf{C}_1, P_{\text{mix},1}, \mathbf{C}_2, P_{\text{mix},2}, \mathbf{C}_3, P_{\text{mix},3}, \mathbf{P}_{\text{inelig}}, \mathbf{P}_{\text{dec}})$, outputting 0 if parsing fails, $\text{VerKey}((k, pk_T, \mathbf{m}), \rho, k) \neq 1$, or $|\mathbf{m}| < n_C$. Perform the following checks and output 0 if any check does not hold.

1) Check removal of invalid ballots: Compute $\{b_1, \dots, b_\ell\}$ as per Step 1 of the tallying algorithm. Check that $\{b_1, \dots, b_\ell\} = \emptyset$ implies \mathbf{X} is a zero-filled vector.

2) Check duplicate elimination: Check that \mathbf{P}_{dupl} is a vector of length ℓ and that for all $1 \leq i \leq \ell$, either: i) $\mathbf{P}_{\text{dupl}}[i]$ parses as a vector (j, σ) , $\text{VerPET}((pk_T, b_i[2], b_j[2], 1), \sigma, k) = 1$, and $j \in \{1, \dots, i-1, i+1, \dots, \ell\}$, or ii) $\mathbf{P}_{\text{dupl}}[i]$ parses as a vector $(0, \sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_\ell)$ and for all $j \in \{1, \dots, i-1, i+1, \dots, \ell\}$ we have $\text{VerPET}((pk_T, b_i[2], b_j[2], 0), \sigma_j, k) = 1$.

3) Check mixing: Compute \mathbf{BB} as per Step 2 of the tallying algorithm. Suppose $\mathbf{BB} = (b'_1, \dots, b'_{|\mathbf{BB}|})$ and $L = \{pd_1, \dots, pd_{|L|}\}$ such that $pd_1 < \dots < pd_{|L|}$. Check $\text{VerMix}((pk_T, (b'_1[1], \dots, b'_{|\mathbf{BB}|}[1]), \mathbf{C}_1), P_{\text{mix},1}, k) = 1 \wedge \text{VerMix}((pk_T, (b'_1[2], \dots, b'_{|\mathbf{BB}|}[2]), \mathbf{C}_2), P_{\text{mix},2}, k) = 1 \wedge \text{VerMix}((pk_T, (pd_1, \dots, pd_{|L|}), \mathbf{C}_3), P_{\text{mix},3}, k) = 1$.

4) Check removal of ineligible ballots: Check that $\mathbf{P}_{\text{inelig}}$ is a vector of length $|\mathbf{C}_2|$ and that for all $1 \leq i \leq |\mathbf{C}_2|$, either: i) $\mathbf{P}_{\text{inelig}}[i]$ parses as a vector (j, σ) , $\text{VerPET}((pk_T, \mathbf{C}_2[i], \mathbf{C}_3[j], 1), \sigma, k) = 1$, and $j \in \{1, \dots, |\mathbf{C}_3|\}$, or ii) $\mathbf{P}_{\text{inelig}}[i]$ parses as a vector $(0, \sigma_1, \dots, \sigma_{|\mathbf{C}_3|})$ and for all $1 \leq j \leq |\mathbf{C}_3|$ we have $\text{VerPET}((pk_T, \mathbf{C}_2[i], \mathbf{C}_3[j], 0), \sigma_j, k) = 1$.

5) Check decryption: Compute \mathbf{C}'_1 as per Step 4 of the tallying algorithm. Check that \mathbf{P}_{dec} parses as a vector $((\beta_1, \sigma_1), \dots, (\beta_{|\mathbf{C}'_1|}, \sigma_{|\mathbf{C}'_1|}))$ such that for all $1 \leq i \leq |\mathbf{C}'_1|$ we have $\text{VerDec}((pk_T, \mathbf{C}'_1[i], \beta_i), \sigma_i, k) = 1$ and for all $1 \leq \beta \leq n_C$ we have $\exists = \mathbf{X}[\beta] j \in \{1, \dots, |\mathbf{C}'_1|\} : \beta = \beta_j$.

Output 1 if all the above checks hold.

The specification of algorithms Setup, Register and Vote follow from our informal descriptions (§VI). The tallying algorithm performs the following steps:

- 1) *Remove invalid ballots*: The tallier discards any ballots from the bulletin board for which proofs do not hold.
- 2) *Eliminating duplicates*: The tallier performs pairwise PETs on the encrypted credentials and discard any ballots for which a test holds, that is, ballots using the same credential are discarded.⁶¹
- 3) *Mixing*: The tallier mixes the ciphertexts in the ballots (i.e., the encrypted choices and the encrypted credentials), using the same secret permutation for both mixes, hence, the mix preserves the relation between encrypted choices and credentials. Let \mathbf{C}_1 and \mathbf{C}_2 be the vectors output by these mixes. The tallier also mixes the public credentials published by the registrar. Let \mathbf{C}_3 be the vector output by this mix.
- 4) *Remove ineligible ballots*: The tallier discards ciphertexts $\mathbf{C}_1[i]$ from \mathbf{C}_1 if there is no ciphertext c in \mathbf{C}_3 such that a PET holds for c and $\mathbf{C}_2[i]$, that is, ballots cast using ineligible credentials are discarded.
- 5) *Decrypting*: The tallier decrypts the remaining encrypted choices in \mathbf{C}_1 and proves that decryption was performed correctly. The tallier identifies the winning candidate from the decrypted choices.

The Verify algorithm checks that each of the above steps has been performed correctly.

Lemma 26 demonstrates that generalized JCJ is a construction for election schemes.

Lemma 26. *Suppose Γ , Σ_1 , Σ_2 , Σ_3 , Σ_4 , Σ_5 , Σ_6 and \mathcal{H} satisfy the preconditions of Definition 30. We have $\text{JCJ}(\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5, \Sigma_6, \mathcal{H})$ satisfies Correctness.*

Proof. Let $\text{JCJ}(\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5, \Sigma_6, \mathcal{H}) = (\text{Setup}, \text{Register}, \text{Vote}, \text{Tally}, \text{Verify})$, $\Gamma = (\text{Gen}, \text{Enc}, \text{Dec})$, $\text{FS}(\Sigma_1, \mathcal{H}) = (\text{ProveKey}, \text{VerKey})$, $\text{FS}(\Sigma_2, \mathcal{H}) = (\text{ProveCiph}, \text{VerCiph})$, and $\text{FS}(\Sigma_3, \mathcal{H}) = (\text{ProveBind}, \text{VerBind})$.

Suppose k is a security parameter, n_B and n_C are integers, and $\beta_1, \dots, \beta_{n_B} \in \{1, \dots, n_C\}$ are choices. Further suppose $(PK_{\mathcal{T}}, SK_{\mathcal{T}}, m_B, m_C)$ is an output of $\text{Setup}(k)$. Moreover, for all $1 \leq i \leq n_B$ suppose (pd_i, d_i) is an output of $\text{Register}(PK_{\mathcal{T}}, k)$ and b_i is an output of $\text{Vote}(d_i, PK_{\mathcal{T}}, n_C, \beta_i, k)$. Further suppose \mathbf{Y} is derived by initializing \mathbf{Y} as a zero-filled vector of length n_C and computing **for** $1 \leq i \leq n_B$ **do** $\mathbf{Y}[\beta_i] \leftarrow \mathbf{Y}[\beta_i] + 1$. If $n_B \not\leq m_B \vee n_C \not\leq m_C$, then Correctness is trivially satisfied, otherwise ($n_B \leq m_B \wedge n_C \leq m_C$), we proceed as follows.

By definition of Setup, we have $PK_{\mathcal{T}} = (pk_T, \mathbf{m}, \rho)$, $SK_{\mathcal{T}} = (pk_T, sk_T)$, $m_B = \text{poly}(k)$, and $m_C = |\mathbf{m}|$, where $(pk_T, sk_T, \mathbf{m}) = \text{Gen}(k; r)$ and ρ is an output of $\text{ProveKey}((k, pk_T, \mathbf{m}), (sk_T, r), k)$ for some coins r chosen uniformly at random by Setup. By completeness of (ProveKey, VerKey), we have $\text{VerKey}((k, pk_T, \mathbf{m}), \rho, k) = 1$. And, since Γ has a message space over \mathbb{Z}_m^* for some integer m and since

$n_C \leq |\mathbf{m}|$, we have $\{1, \dots, n_C\} \subseteq \mathbf{m}$. Therefore, by definition of Vote, we have for all $1 \leq i \leq n_B$ that $b_i[1] = \text{Enc}(pk_T, \beta_i; r_{i,1})$, $b_i[2] = \text{Enc}(pk_T, d_i; r_{i,2})$, $b_i[3]$ is an output of $\text{ProveCiph}((pk_T, b_i[1], \{1, \dots, n_C\}), (\beta_i, r_{i,1}), k)$, and $b_i[4]$ is an output of $\text{ProveBind}((pk_T, b_i[1], b_i[2]), (\beta_i, r_{i,1}, d_i, r_{i,2}), k)$, where $r_{i,1}$ and $r_{i,2}$ are coins chosen uniformly at random by Vote. Let us consider the computation of (\mathbf{X}, P) by $\text{Tally}(SK_{\mathcal{T}}, \{b_1, \dots, b_{n_B}\}, \{pd_1, \dots, pd_{n_B}\}, n_C, k)$.

Suppose a subset of $\{b_1, \dots, b_{n_B}\}$ is computed as per Step 1 of algorithm Tally. By completeness of (ProveCiph, VerCiph) and (ProveBind, VerBind), that subset is $\{b_{\pi(1)}, \dots, b_{\pi(n_B)}\}$, where π is a permutation on $\{1, \dots, n_B\}$ such that $b_{\pi(1)} < \dots < b_{\pi(n_B)}$. If $n_B = 0$, then \mathbf{X} and \mathbf{Y} are both zero-filled vectors of length n_C , and we conclude immediately, otherwise, we proceed as follows.

Suppose \mathbf{BB} is computed as per Step 2 of algorithm Tally. By definition of Register, we have d_1, \dots, d_{n_B} are chosen uniformly at random from \mathbf{m} , where $n_B \leq \text{poly}(k)$ and $|\mathbf{m}|$ is super-polynomial in the security parameter. Thus, for all distinct integers $i, j \in \{1, \dots, n_B\}$ we have $d_i \neq d_j$, with overwhelming probability. It follows for all $1 \leq i \leq \ell$, all $j \in \{1, \dots, i-1, i+1, \dots, \ell\}$, and outputs σ of $\text{ProvePET}((pk_T, b_i[2], b_j[2], 1), sk_T, k)$ that $\text{VerPET}((pk_T, b_i[2], b_j[2], 1), \sigma, k) \neq 1$, with overwhelming probability. Thus, $\mathbf{BB} = (b_{\pi(1)}, \dots, b_{\pi(n_B)})$.

Suppose \mathbf{C}_1 , \mathbf{C}_2 and \mathbf{C}_3 are computed as per Step 3 of algorithm Tally. We have for all $1 \leq i \leq n_B$ that $\mathbf{C}_1[i] = b'_{\chi(\pi(i))}[1] \otimes \text{Enc}(pk_T, \mathbf{e}; \mathbf{r}_1[i])$ and $\mathbf{C}_2[i] = b'_{\chi(\pi(i))}[2] \otimes \text{Enc}(pk_T, \mathbf{e}; \mathbf{r}_2[i])$. Moreover, since Γ is a homomorphic asymmetric encryption scheme and \mathbf{e} is an identity element, we have for all $1 \leq i \leq n_B$ that

$$\begin{aligned} \mathbf{C}_1[i] &= \text{Enc}(pk_T, \beta_{\chi(\pi(i))}; r_{\chi(\pi(i)),1} \oplus \mathbf{r}_1[i]) \\ \mathbf{C}_2[i] &= \text{Enc}(pk_T, d_{\chi(\pi(i))}; r_{\chi(\pi(i)),2} \oplus \mathbf{r}_2[i]) \end{aligned}$$

Similarly, we have for all $1 \leq i \leq n_B$ that

$$\mathbf{C}_3[i] = \text{Enc}(pk_T, d_{\chi'(\pi'(i))}; r_{\chi'(\pi'(i))} \oplus \mathbf{r}_3[i])$$

where coins r_1, \dots, r_{n_B} were used to construct pd_1, \dots, pd_{n_B} and π' is a permutation on $\{1, \dots, n_B\}$ such that $pd_{\pi'(1)} < \dots < pd_{\pi'(n_B)}$.

Suppose \mathbf{C}'_1 is computed as per Step 4 of algorithm Tally. We have for all $1 \leq i \leq n_B$ that there exists $j \in \{1, \dots, n_B\}$ such that $\text{VerPET}((pk_T, \mathbf{C}_2[i], \mathbf{C}_3[j], 1), \sigma, k) = 1$ for some output σ of $\text{ProvePET}((pk_T, \mathbf{C}_2[i], \mathbf{C}_3[j], 1), sk_T, k)$, because \mathbf{C}_2 , respectively \mathbf{C}_3 , is a vector of ciphertexts on plaintexts $d_{\chi(\pi(1))}, \dots, d_{\chi(\pi(n_B))}$, respectively $d_{\chi'(\pi'(1))}, \dots, d_{\chi'(\pi'(n_B))}$, that is, \mathbf{C}_2 and \mathbf{C}_3 contain ciphertexts on the same plaintexts. Thus, $\mathbf{C}'_1 = (\mathbf{C}_1[1], \dots, \mathbf{C}_1[n_B])$.

61. JCJ defines discarding ballots in accordance with a revoting policy [87, §4.1]. However, we have shown that JCJ fails to satisfy universal verifiability when the policy proposed by Juels et al. is adopted (§IV-B2). So, we consider a policy that discards ballots using the same credential—i.e., choices by voters that cast multiple ballots will be discarded.

Suppose \mathbf{X} is computed as per Step 5 of algorithm Tally, namely, **for** $1 \leq i \leq n_B$ **do** $\beta \leftarrow \text{Dec}(sk_T, \mathbf{C}'_1[i]); \mathbf{X}[\beta] \leftarrow \mathbf{X}[\beta] + 1$. By correctness of Γ , we have for all $1 \leq i \leq n_B$ that $\text{Dec}(sk_T, \mathbf{C}'_1[i]) = \beta_{\chi(\pi(i))}$. Hence, \mathbf{X} can be equivalently computed as **for** $1 \leq i \leq n_B$ **do** $\mathbf{X}[\beta_{\chi(\pi(i))}] \leftarrow \mathbf{X}[\beta_{\chi(\pi(i))}] + 1$. And, since \mathbf{Y} is derived by initializing \mathbf{Y} as a zero-filled vector of length n_C and computing **for** $1 \leq i \leq n_B$ **do** $\mathbf{Y}[\beta_i] \leftarrow \mathbf{Y}[\beta_i] + 1$, we have $\mathbf{X} = \mathbf{Y}$, concluding our proof. \square

APPENDIX J

PROOF: JCJ IS NOT VERIFIABLE

Generalized JCJ can be instantiate to derive JCJ:

Definition 31 (JCJ [87]). JCJ is $\text{JCJ}(\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5, \Sigma_6, \mathcal{H})$, where Γ is a modified version of El Gamal [65] invented by Juels et al. [87, §4] that can be seen as a simplified version of Cramer–Shoup [57], Σ_1 is the proof of key construction by Gennaro et al. [69], Σ_4 is the conjunction [54] of two Schnorr proofs [119], Σ_5 is the PET by MacKenzie et al. [102], and \mathcal{H} is a random oracle. Juels et al. leave Σ_2, Σ_3 and Σ_6 unspecified.

Juels et al. [87] do not mandate particular cryptographic primitives, so Definition 31 might be seen more as an instantiation of their scheme than an exact recollection of it. We assume that the primitives in Definition 31 satisfy the properties required by generalized JCJ. We leave formally proving this assumption as future work. Under this assumption, Lemma 26 demonstrates that JCJ is an election scheme.

Proof of Proposition 9. Let $\text{JCJ}(\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5, \Sigma_6, \mathcal{H}) = (\text{Setup}, \text{Register}, \text{Vote}, \text{Tally}, \text{Verify}), \text{FS}(\Sigma_1, \mathcal{H}) = (\text{ProveKey}, \text{VerKey}), \text{FS}(\Sigma_2, \mathcal{H}) = (\text{ProveCiph}, \text{VerCiph}), \text{FS}(\Sigma_3, \mathcal{H}) = (\text{ProveBind}, \text{VerBind}), \text{FS}(\Sigma_4, \mathcal{H}) = (\text{ProveDec}, \text{VerDec}), \text{FS}(\Sigma_5, \mathcal{H}) = (\text{ProvePET}, \text{VerPET}),$ and $\text{FS}(\Sigma_6, \mathcal{H}) = (\text{ProveMix}, \text{VerMix})$. Moreover, let $\beta_1 = 1$ and $\beta_2 = 2$. We construct an adversary \mathcal{A} (Figure 3) against the universal verifiability experiment.

Let k be a security parameter such that Γ has a message space over \mathbb{Z}_m^* for some integer m such that $1, 2 \in \mathbb{Z}_m^*$. Suppose an execution of Exp-UV-Int computes

```
(PKT) ← A(k);
for 1 ≤ i ≤ nV do (pki, ski) ← Register(PKT, k);
L ← {pk1, ..., pknV};
M ← {(pk1, sk1), ..., (pknV, sknV)};
(BB, nC, X, P) ← A(M);
Y ← correct-tally(PKT, BB, M, nC, k);
```

By definition of function *correct-tally*, we have $\mathbf{Y} = (1, 0)$. Thus, $\mathbf{X} \neq \mathbf{Y}$. Let us prove that $\text{Verify}(PK_T, BB, L, n_C, \mathbf{X}, \mathbf{P}, k) = 1$.

By definition of \mathcal{A} , we have PK_T parses as (pk_T, \mathbf{m}, ρ) , where ρ is constructed by the adversary using algorithm *ProveKey*. It follows by completeness of $(\text{ProveKey}, \text{VerKey})$ that $\text{VerKey}((k, pk_T, \mathbf{m}), \rho, k) = 1$. By definition of \mathcal{A} , we also have $n_C = 2$, and, since 1 and 2 are elements of Γ 's message space, we have $n_C \leq |\mathbf{m}|$. Moreover, \mathbf{X} parses as

Fig. 3 Adversary against JCJ

Given a security parameter k as input, adversary \mathcal{A} computes $(PK_T, SK_T, m_B, m_C) \leftarrow \text{Setup}(k); n_V \leftarrow 1$ and outputs (PK_T, n_V) . Moreover, given a set of credentials M , adversary \mathcal{A} parses M as set $\{(pd_1, d_1)\}$, PK_T as a vector (pk_T, \mathbf{m}, ρ) , and SK_T as a vector (pk_T, sk_T) , computes

```
1 %number of candidates
2 nC ← 2;
3 %authorized ballot for choice 1
4 b1 ← Vote(d1, PKT, nC, β1, k);
5 %unauthorized ballot for choice 2
6 (pd2, d2) ← Register(PKT, k);
7 b2 ← Vote(d2, PKT, nC, β2, k);
8 %bulletin board
9 BB ← {b1, b2};
```

selects permutation π on $\{1, 2\}$ such that $b_{\pi(1)} < b_{\pi(2)}$, initializes vectors $\mathbf{C}_1, \mathbf{C}_2, \mathbf{r}_1$ and \mathbf{r}_2 of length 2, initializes vectors \mathbf{C}_3 and \mathbf{r}_3 of length 1, fills $\mathbf{r}_1, \mathbf{r}_2$ and \mathbf{r}_3 with coins, selects permutations χ and χ' on $\{1, 2\}$ such that χ is the identity function and χ' is not, be coins, computes

```
10 %proof of duplicate elimination
11 σ1 ← ProvePET((pkT, bπ(1)[2], bπ(2)[2], 0), skT, k);
12 σ2 ← ProvePET((pkT, bπ(2)[2], bπ(1)[2], 0), skT, k);
13 Pdupl ← ((0, σ1), (0, σ2));
14 %mix ciphertexts in ballots with
15 %distinct permutations
16 C1[1] ← bχ(π(1))[1] ⊗ Enc(pkT, c; r1[1]);
17 C1[2] ← bχ(π(2))[1] ⊗ Enc(pkT, c; r1[2]);
18 Pmix,1 ← ProveMix((pkT, (bπ(1)[1], bπ(2)[1]), C1), (r1, χ), k);
19 C2[1] ← bχ'(π(1))[2] ⊗ Enc(pkT, c; r2[1]);
20 C2[2] ← bχ'(π(2))[2] ⊗ Enc(pkT, c; r2[2]);
21 Pmix,2 ← ProveMix((pkT, (bπ(1)[2], bπ(2)[2]), C2), (r2, χ'), k);
22 %mix public credentials
23 C3[1] ← pd1 ⊗ Enc(pkT, c; r2[1]);
24 Pmix,3 ← ProveMix((pkT, (pd1), C3), (r3, χ), k);
25 %proof of ineligible ballots
26 τ1 ← ProvePET((pkT, C2[1], C3[1], π(1) - 1), skT, k);
27 τ2 ← ProvePET((pkT, C2[2], C3[1], π(2) - 1), skT, k);
28 Pinelig ← ((π(1) - 1, τ1), (π(2) - 1, τ2));
29 %tally
30 X ← (0, 1);
31 %proof of decryption
32 σ ← ProveDec((pkT, C1[π(2)], β2), skT, k);
33 Pdec ← ((β2, σ));
34 %proof of tallying
35 P ← (Pdupl, C1, Pmix,1, C2, Pmix,2, C3, Pmix,3, Pinelig, Pdec);
```

and outputs $(BB, n_C, \mathbf{X}, \mathbf{P})$.

a vector of length n_C and \mathbf{P} parses as a vector $(\mathbf{P}_{\text{dupl}}, \mathbf{C}_1, P_{\text{mix},1}, \mathbf{C}_2, P_{\text{mix},2}, \mathbf{C}_3, P_{\text{mix},3}, \mathbf{P}_{\text{inelig}}, \mathbf{P}_{\text{dec}})$. Thus, the

initial checks performed by algorithm Verify succeed and we proceed by proving that checks performed in Steps 1–5 of Verify also succeed.

By definition of \mathcal{A} , we have $BB = \{b_1, b_1\}$, where b_1 , respectively b_2 , is computed using algorithm Vote on inputs including private credential d_1 and choice β_1 , respectively d_2 and β_2 , where d_2 is the private credential constructed by adversary \mathcal{A} . Therefore, by definition of Vote, for all $i \in \{1, 2\}$ we have:

$$\begin{aligned} b_i[1] &= \text{Enc}(pk_T, \beta_i; r_{i,1}), \\ b_i[2] &= \text{Enc}(pk_T, d_i; r_{i,2}), \end{aligned}$$

$b_i[3]$ is an output of $\text{ProveCiph}((pk_T, b_i[1], \{1, 2\}), (\beta_i, r_{i,1}), k)$, and $b_i[4]$ is an output of $\text{ProveBind}((pk_T, b_i[1], b_i[2]), (\beta_i, r_{i,1}, d_i, r_{i,2}), k)$, where $r_{i,1}$ and $r_{i,2}$ are coins chosen uniformly at random by Vote.

Suppose a subset of BB is computed as per Step 1 of algorithm Tally. By completeness of $(\text{ProveCiph}, \text{VerCiph})$ and $(\text{ProveBind}, \text{VerBind})$, that subset is $\{b_{\pi(1)}, b_{\pi(2)}\}$, where permutation π is selected by adversary \mathcal{A} . Thus, the check holds in Step 1 of Verify.

We have \mathbf{P}_{dupl} is a vector of length 2 such that $\mathbf{P}_{\text{dupl}}[1]$ parses as a vector $(0, \sigma_1)$, where σ_1 is an output of $\text{ProvePET}((pk_T, b_{\pi(1)}[2], b_{\pi(2)}[2], 0), sk_T, k)$. By correctness of Γ , we have $\text{Dec}(sk_T, b_{\pi(1)}[2]) = d_{\pi(1)}$ and $\text{Dec}(sk_T, b_{\pi(2)}[2]) = d_{\pi(2)}$. And, since d_1 and d_2 were selected uniformly at random from \mathfrak{m} , we have $d_1 \neq d_2$, with probability greater than negligible, because $n_C \leq |\mathfrak{m}|$. Hence, $\text{Dec}(sk_T, b_{\pi(1)}[2]) \neq \text{Dec}(sk_T, b_{\pi(2)}[2])$, with probability greater than negligible. Moreover, by completeness of $(\text{ProvePET}, \text{VerPET})$, we have $\text{VerPET}((pk_T, b_{\pi(1)}[2], b_{\pi(2)}[2], 0), \sigma_1, k) = 1$, with probability greater than negligible. Similarly, $\mathbf{P}_{\text{dupl}}[1]$ parses as a vector $(0, \sigma_2)$ and $\text{VerPET}((pk_T, b_{\pi(2)}[2], b_{\pi(1)}[2], 0), \sigma_2, k) = 1$, with probability greater than negligible. Thus, checks hold in Step 2 of Verify, with probability greater than negligible.

Suppose \mathbf{BB} is computed as per Step 2 of the tallying algorithm. Hence, $\mathbf{BB} = (b_{\pi(1)}, b_{\pi(2)})$. By completeness of $(\text{ProveMix}, \text{VerMix})$, we have $\text{VerMix}((pk_T, (b_{\pi(1)}[1], b_{\pi(2)}[1]), \mathbf{C}_1), P_{\text{mix},1}, k) = 1$, $\text{VerMix}((pk_T, (b_{\pi(1)}[2], b_{\pi(2)}[2]), \mathbf{C}_2), P_{\text{mix},2}, k) = 1$, and $\text{VerMix}((pk_T, (pd_1), \mathbf{C}_3), P_{\text{mix},3}, k) = 1$. Thus, checks hold in Step 3 of Verify.

We have for all $i \in \{1, 2\}$ that $\mathbf{C}_2[i] = b_{\chi'(\pi(i))}[2] \otimes \text{Enc}(pk_T, \epsilon; \mathbf{r}_2[i])$. And, since Γ is homomorphic and ϵ is an identity element, we have $\mathbf{C}_2[i] = \text{Enc}(pk_T, d_{\chi'(\pi(i))}; r_{\pi(i),1} \oplus \mathbf{r}_2[i])$, hence, $\text{Dec}(sk_T, \mathbf{C}_2[i]) = d_{\chi'(\pi(i))}$. Similarly, we have $\mathbf{C}_3[1] = pd_1 \otimes \text{Enc}(pk_T, \epsilon; \mathbf{r}_2[1])$, where pd_1 is a ciphertext on $d_1 \in \mathfrak{m}$ constructed by algorithm Register. Hence, $\text{Dec}(sk_T, \mathbf{C}_3[1]) = d_1$. It follows that $\text{Dec}(sk_T, \mathbf{C}_2[1]) \neq \text{Dec}(sk_T, \mathbf{C}_3[1]) \wedge \text{Dec}(sk_T, \mathbf{C}_2[2]) = \text{Dec}(sk_T, \mathbf{C}_3[1])$ iff π is an identity function. We have $\mathbf{P}_{\text{inelig}} = ((\pi(1) - 1, \tau_1), (\pi(2) - 1, \tau_2))$, where τ_1 and τ_2 are constructed by the adversary. It follows by completeness of $(\text{ProvePET}, \text{VerPET})$ that $\text{VerPET}((pk_T, \mathbf{C}_2[1], \mathbf{C}_3[1], \pi(1) - 1), \tau_1, k) = 1$ and $\text{VerPET}((pk_T, \mathbf{C}_2[2], \mathbf{C}_3[1], \pi(2) - 1), \tau_2, k) = 1$. Thus, checks hold in Step 4 of Verify.

Suppose \mathbf{C}'_1 is computed as per Step 4 of the tallying algorithm. Hence, $\mathbf{C}'_1 = (\mathbf{C}_1[\pi(2)])$. We have \mathbf{P}_{dec} parses as a vector $((\beta_2, \sigma))$, where σ is constructed by the adversary using algorithm ProveDec on inputs including $\mathbf{C}_1[\pi(2)]$ and β_2 . Moreover, since π is a permutation on $\{1, 2\}$ and χ is an identity function, we have $\chi(\pi(\pi(2))) = 2$, therefore, $\mathbf{C}_1[\pi(2)] = b_2[1] \otimes \text{Enc}(pk_T, \epsilon; \mathbf{r}_1[2])$. And, since Γ is homomorphic and ϵ is an identity element, we have $\mathbf{C}_1[\pi(2)] = \text{Enc}(pk_T, \beta_2; r_{2,1} \oplus \mathbf{r}_1[2])$, hence, $\text{Dec}(sk_T, \mathbf{C}_1[\pi(2)]) = \beta_2$. Therefore, by completeness of $(\text{ProveDec}, \text{VerDec})$, we have $\text{VerDec}((pk_T, \mathbf{C}_1[\pi(2)], \beta_2), \sigma, k) = 1$. Furthermore, since $\mathbf{X} = (0, 1)$, we have for all $1 \leq \beta \leq n_C$ that $\exists^{\mathbf{X}|\beta} \beta = \beta_2$. Thus, checks hold in Step 5 of Verify.

We have shown that checks performed in Steps 1–5 of algorithm Verify all succeed, thus, $\text{Verify}(PK_T, BB, L, n_C, \mathbf{X}, \mathbf{P}, k) = 1$, concluding our proof. \square

APPENDIX K

PROOF: JCJ'16 IS VERIFIABLE

We formalize a variant of the generic construction for JCJ-like election schemes that uses a mixnet capable of proving that the relation between encrypted choices and encrypted credentials is maintained.

Definition 32. Let $(\text{Gen}, \text{Enc}, \text{Dec})$ be a homomorphic asymmetric encryption scheme and Σ be a sigma protocol for a binary relation R . Suppose $(pk, sk, \mathfrak{m}) = \text{Gen}(k; r)$, for some security parameter k and coins r . We say Σ is a mixnet on pairs if $((pk, \mathbf{c}_1, \mathbf{c}'_1, \mathbf{c}_2, \mathbf{c}'_2), (\mathbf{r}_1, \mathbf{r}_2, \chi)) \in R \Leftrightarrow \bigwedge_{1 \leq i \leq |\mathbf{c}_1|, j \in \{1, 2\}} \mathbf{c}'_j[i] = \mathbf{c}_j[\chi(i)] \otimes \text{Enc}(pk, \epsilon; \mathbf{r}_j[i]) \wedge |\mathbf{c}_1| = |\mathbf{c}'_1| = |\mathbf{c}_2| = |\mathbf{c}'_2| = |\mathbf{r}_1| = |\mathbf{r}_2|$, where $\mathbf{c}_1, \mathbf{c}'_1, \mathbf{c}_2$ and \mathbf{c}'_2 are vectors of ciphertexts encrypted under pk , \mathbf{r}_1 and \mathbf{r}_2 are vectors of coins, χ is a permutation on $\{1, \dots, |\mathbf{c}_1|\}$, and ϵ is an identity element of the encryption scheme's message space with respect to \odot .

Definition 33. Suppose $\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5, \Sigma_6$, and \mathcal{H} satisfy the preconditions of Definition 30. Further suppose Σ_7 is a mixnet on pairs. Let $\Gamma = (\text{Gen}, \text{Enc}, \text{Dec})$, $\text{FS}(\Sigma_6, \mathcal{H}) = (\text{ProveMix}, \text{VerMix})$, and $\text{FS}(\Sigma_7, \mathcal{H}) = (\text{ProveMixPair}, \text{VerMixPair})$. Moreover, let ϵ be an identity element of Γ 's message space with respect to \odot . We define $\widehat{\text{JCJ}}(\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5, \Sigma_6, \Sigma_7, \mathcal{H})$ as $\text{JCJ}(\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5, \Sigma_6, \mathcal{H}) = (\text{Setup}, \text{Register}, \text{Vote}, \text{Tally}, \text{Verify})$ after the following modifications. First, Tally computes $P_{\text{mix},1}$ as $P_{\text{mix},1} \leftarrow \text{ProveMixPair}((pk_T, (b'_1[1], \dots, b'_{|\mathbf{BB}|}[1]), \mathbf{C}_1), (b'_1[2], \dots, b'_{|\mathbf{BB}|}[2]), \mathbf{C}_2), (\mathbf{r}_1, \mathbf{r}_2, \chi), k)$, and $P_{\text{mix},2}$ as $P_{\text{mix},2} \leftarrow \perp$. Secondly, Verify replaces checks using VerMix with the following check $\text{VerMixPair}((pk_T, (b'_1[1], \dots, b'_{|\mathbf{BB}|}[1]), \mathbf{C}_1, (b'_1[2], \dots, b'_{|\mathbf{BB}|}[2]), \mathbf{C}_2), P_{\text{mix},1}, k) = 1 \wedge \text{VerMix}((pk_T, (pd_1, \dots, pd_{|L|}), \mathbf{C}_3), P_{\text{mix},3}, k) = 1$.

Lemmata 26 can be adapted to show that $\widehat{\text{JCJ}}$ is a construction for election schemes.

Election schemes constructed from $\widehat{\text{JCJ}}$ satisfy individual (§K-A), universal (§K-B) and eligibility (§K-C) verifiability, hence, such schemes satisfy election verifiability with internal

authentication (§K-D), assuming that the cryptographic primitives satisfy certain properties that we identify.

A. Individual verifiability

Proposition 27. *Suppose $\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5, \Sigma_6, \Sigma_7$ and \mathcal{H} satisfy the preconditions of Definition 33. Further suppose that Γ is collision-free for its message space. We have $\widehat{\text{JCJ}}(\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5, \Sigma_6, \Sigma_7, \mathcal{H})$ satisfies individual verifiability.*

Proof. Let $\widehat{\text{JCJ}}(\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5, \Sigma_6, \Sigma_7, \mathcal{H}) = (\text{Setup}, \text{Register}, \text{Vote}, \text{Tally}, \text{Verify})$, $\Gamma = (\text{Gen}, \text{Enc}, \text{Dec})$, and $\text{FS}(\Sigma_1, \mathcal{H}) = (\text{ProveKey}, \text{VerKey})$. Suppose k is a security parameter, $PK_{\mathcal{T}}$ is a public key, n_C is an integer, and β and β' are choices. Further suppose (pk, sk) and (pk', sk') are outputs of $\text{Register}(PK_{\mathcal{T}}, k)$, b is an output of $\text{Vote}(sk, PK_{\mathcal{T}}, n_C, \beta, k)$, and b' is an output of $\text{Vote}(sk', PK_{\mathcal{T}}, n_C, \beta', k)$, such that $b \neq \perp$ and $b' \neq \perp$. By definition of Vote , we have $PK_{\mathcal{T}}$ is a vector (pk_T, m, ρ) and $\text{VerKey}((k, pk_T, m), \rho, k) = 1$. Moreover, $b[2]$ is an output of $\text{Enc}(pk_T, sk)$ and $b'[2]$ is an output of $\text{Enc}(pk_T, sk')$, where $sk, sk' \in m$. Furthermore, the ciphertexts are constructed using coins chosen uniformly at random—i.e., the coins used by $b[2]$ and $b'[2]$ will be distinct with overwhelming probability. Since Γ is collision-free for m , we have $b[2] \neq b'[2]$ and $b \neq b'$ with overwhelming probability, concluding our proof. \square

B. Universal verifiability.

Lemma 28. *Suppose $\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5, \Sigma_6, \Sigma_7$, and \mathcal{H} satisfy the preconditions of Definition 30. Further suppose Γ is collision-free for its message space. We have $\text{JCJ}(\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5, \Sigma_6, \Sigma_7, \mathcal{H})$ satisfies Injectivity.*

The proof of Lemma 28 is similar to the proof of Lemma 19.

Proof sketch. Generalized JCJ ballots contain encrypted choices, hence, collision-freeness of the encryption scheme ensures that distinct choices are not mapped to the same ballot. \square

Proposition 29. *Suppose $\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5, \Sigma_6, \Sigma_7$, and \mathcal{H} satisfy the preconditions of Definition 33. Further suppose that Γ is perfectly correct, perfectly homomorphic, and collision-free for its message space, the sigma protocols satisfy special soundness and special honest verifier zero-knowledge, and \mathcal{H} is a random oracle. We have $\widehat{\text{JCJ}}(\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5, \Sigma_6, \Sigma_7, \mathcal{H})$ satisfies universal verifiability.*

Proof. Let $\widehat{\text{JCJ}}(\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5, \Sigma_6, \Sigma_7, \mathcal{H}) = (\text{Setup}, \text{Register}, \text{Vote}, \text{Tally}, \text{Verify})$, $\text{FS}(\Sigma_1, \mathcal{H}) = (\text{ProveKey}, \text{VerKey})$, $\text{FS}(\Sigma_2, \mathcal{H}) = (\text{ProveCiph}, \text{VerCiph})$, $\text{FS}(\Sigma_3, \mathcal{H}) = (\text{ProveBind}, \text{VerBind})$, $\text{FS}(\Sigma_4, \mathcal{H}) = (\text{ProveDec}, \text{VerDec})$, $\text{FS}(\Sigma_5, \mathcal{H}) = (\text{ProvePET}, \text{VerPET})$, $\text{FS}(\Sigma_6, \mathcal{H}) = (\text{ProveMix}, \text{VerMix})$, and $\text{FS}(\Sigma_7, \mathcal{H}) = (\text{ProveMixPair}, \text{VerMixPair})$.

Suppose an execution of $\text{Exp-UV-Int}(\Pi, \mathcal{A}, k)$ computes

```
(PKT, nV) ← A(k);
for 1 ≤ i ≤ nV do (pdi, di) ← Register(PKT, k);
L ← {pd1, ..., pdnV};
M ← {(pd1, d1), ..., (pdnV, dnV)};
(BB, nC, X, P) ← A(M);
Y ← correct-tally(PKT, BB, M, nC, k);
```

such that $\text{Verify}(PK_{\mathcal{T}}, BB, L, n_C, X, P, k) = 1$. By definition of algorithm Verify , we have $PK_{\mathcal{T}}$ parses as a vector (pk_T, m, ρ) , X parses as a vector of length n_C , and P parses as a vector $(P_{\text{dupl}}, C_1, P_{\text{mix},1}, C_2, P_{\text{mix},2}, C_3, P_{\text{mix},3}, P_{\text{inelig}}, P_{\text{dec}})$. Moreover, $\text{VerKey}((k, pk_T, m), \rho, k) = 1$ and $n_C \leq |m|$. By simulation sound extractability, we are assured that pk_T is an output of Gen with overwhelming probability—i.e., there exists r and $SK_{\mathcal{T}}$ such that $(pk_T, SK_{\mathcal{T}}, m) = \text{Gen}(k; r)$. By definition of Register , we have for all $1 \leq i \leq n_V$ that d_i is chosen uniformly at random from m and there exists coins s_i such that $pd_i = \text{Enc}(pk_T, d_i; s_i)$.

Let $\{b_1, \dots, b_\ell\}$ be the largest subset of BB such that for all $1 \leq i \leq \ell$ we have b_i is a vector of length 4 and $\text{VerCiph}((pk_T, b_i[1]\{1, \dots, n_C\}), b_i[3], k) = 1 \wedge \text{VerBind}((pk_T, b_i[1], b_i[2]), b_i[4], k) = 1$. We have for all choices $\beta \in \{1, \dots, n_C\}$, private credentials d , coins r , and ballots $b = \text{Vote}(d, PK_{\mathcal{T}}, n_C, \beta, k; r)$ that $b \notin BB \setminus \{b_1, \dots, b_\ell\}$ with overwhelming probability, since such an occurrence would imply a contradiction: $\{b_1, \dots, b_\ell\}$ is not the largest subset of BB satisfying the conditions of the Tally algorithm. It follows that:

$$\begin{aligned} \text{correct-tally}(PK_{\mathcal{T}}, M, BB, n_C, k) \\ = \text{correct-tally}(PK_{\mathcal{T}}, M, \{b_1, \dots, b_\ell\}, n_C, k) \end{aligned} \quad (5)$$

A proof of (5) follows from the definition of function correct-tally .

By Step 1 of algorithm Verify , if $\{b_1, \dots, b_\ell\} = \emptyset$, then X is a zero-filled vector. And, by definition of function correct-tally and (5), Y is a vector of length n_C such that $\bigwedge_{j=1}^{n_C} Y[j] = 0$. Thus, $X = Y$, concluding our proof. Otherwise ($\{b_1, \dots, b_\ell\} \neq \emptyset$), we proceed as follows.

By simulation sound extractability, we have, with overwhelming probability, that for all $1 \leq i \leq \ell$ there exists choice $\beta_i \in \{1, \dots, n_C\}$, message $d'_i \in m$, and coins $r_{i,1}$ and $r_{i,2}$, such that

$$\begin{aligned} b_i[1] &= \text{Enc}(pk_T, \beta_i; r_{i,1}), \\ b_i[2] &= \text{Enc}(pk_T, d'_i; r_{i,2}), \end{aligned}$$

$b_i[3]$ is an output of $\text{ProveCiph}((pk_T, b_i[1], \{1, \dots, n_C\}), (\beta_i, r_{i,1}), k)$, and $b_i[4]$ is an output of $\text{ProveBind}((pk_T, b_i[1], b_i[2]), (\beta_i, r_{i,1}, d'_i, r_{i,2}), k)$. Moreover, by inspection of Vote , we have

$$\forall i \in \{1, \dots, \ell\}, \exists r : b_i = \text{Vote}(d'_i, PK_{\mathcal{T}}, n_C, \beta_i, k; r) \quad (6)$$

Thus, $\{b_1, \dots, b_\ell\}$ is a set of ballots, and we will now consider which ballots are authorized.

By Step 2 of algorithm `Verify`, we have \mathbf{P}_{dupl} is a vector of length ℓ and for all $1 \leq i \leq \ell$ either: i) $\mathbf{P}_{\text{dupl}}[i]$ parses as a vector (j, σ) , $\text{VerPET}((pk_T, b_i[2], b_j[2], 1), \sigma, k) = 1$, and $j \in \{1, \dots, i-1, i+1, \dots, \ell\}$, therefore, by simulation sound extractability, we have $\text{Dec}(sk_T, b_i[2]) = \text{Dec}(sk_T, b_j[2])$, or ii) $\mathbf{P}_{\text{dupl}}[i]$ parses as a vector $(0, \sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_\ell)$ and for all $j \in \{1, \dots, i-1, i+1, \dots, \ell\}$ we have $\text{VerPET}((pk_T, b_i[2], b_j[2], 0), \sigma_j, k) = 1$ and, by simulation sound extractability, we have $\text{Dec}(sk_T, b_i[2]) \neq \text{Dec}(sk_T, b_j[2])$. Although, key pair pk_T and sk_T may not have been constructed with coins chosen uniformly at random, and similarly ciphertexts $b_1[2], \dots, b_\ell[2]$ may not have been constructed with coins chosen uniformly at random, we nevertheless have for all $1 \leq i \leq \ell$ that if $\mathbf{P}_{\text{dupl}}[i]$ parses as a vector (j, σ) such that $j \in \{1, \dots, i-1, i+1, \dots, \ell\}$, then $d'_i = d'_j$, otherwise, $d'_i \neq d'_j$ for all $j \in \{1, \dots, i-1, i+1, \dots, \ell\}$, with overwhelming probability, because Γ is perfectly correct. Let \mathbf{BB} be computed as per Step 2 of the tallying algorithm. Suppose $\mathbf{BB} = (b'_1, \dots, b'_{|\mathbf{BB}|})$. Hence, there trivially exists an injective function $\lambda : \{1, \dots, |\mathbf{BB}|\} \rightarrow \{1, \dots, \ell\}$ such that for all $1 \leq i \leq |\mathbf{BB}|$ we have $b'_i = b_{\lambda(i)}$, moreover, for all $j \in \{1, \dots, i-1, i+1, \dots, |\mathbf{BB}|\}$ we have $d'_{\lambda(i)} \neq d'_{\lambda(j)}$. It follows that

$$\begin{aligned} \forall i \in \lambda(\{1, \dots, |\mathbf{BB}|\}) : \\ \neg \exists j, \beta, r : b_j = \text{Vote}(d_i, PK_{\mathcal{T}}, n_C, \beta, k; r) \\ \wedge j \in \{1, \dots, i-1, i+1, \dots, \ell\} \quad (7) \end{aligned}$$

Moreover,

$$\begin{aligned} \forall i \in \{1, \dots, \ell\} \setminus \lambda(\{1, \dots, |\mathbf{BB}|\}) : \\ \exists j, \beta, r : b_j = \text{Vote}(d_i, PK_{\mathcal{T}}, n_C, \beta, k; r) \\ \wedge j \in \{1, \dots, i-1, i+1, \dots, \ell\} \quad (8) \end{aligned}$$

Thus, $\{b_i \mid i \in \lambda(\{1, \dots, |\mathbf{BB}|\})\}$ is the largest subset of ballots from $\{b_1, \dots, b_\ell\}$ such that each ballot was constructed using a distinct private credential.

By Step 3 of algorithm `Verify`, we have $\text{VerMixPair}((pk_T, (b'_1[1], \dots, b'_{|\mathbf{BB}|}[1]), \mathbf{C}_1, (b'_1[2], \dots, b'_{|\mathbf{BB}|}[2]), \mathbf{C}_2), P_{\text{mix},1}, k) = 1 \wedge \text{VerMix}((pk_T, (pd_{\pi(1)}, \dots, pd_{\pi(|L|)}), \mathbf{C}_3), P_{\text{mix},3}, k) = 1$, where π is a permutation on $\{1, \dots, |L|\}$ such that $pd_{\pi(1)} < \dots < pd_{\pi(|L|)}$. And, by simulation sound extractability, there exists vectors $\mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_3$, a permutation χ on $\{1, \dots, |\mathbf{BB}|\}$, and a permutation χ' on $\{1, \dots, n_V\}$, such that for all $1 \leq i \leq |\mathbf{BB}|$ we have $\mathbf{C}_1[i] = b'_{\chi(i)}[1] \otimes \text{Enc}(pk_T, \mathbf{e}; \mathbf{r}_1[i])$ and $\mathbf{C}_2[i] = b'_{\chi(i)}[2] \otimes \text{Enc}(pk_T, \mathbf{e}; \mathbf{r}_2[i])$, and for all $1 \leq i \leq n_V$ we have $\mathbf{C}_3[i] = pd_{\chi'(\pi(i))} \otimes \text{Enc}(pk_T, \mathbf{e}; \mathbf{r}_3[i])$. Although, key pair pk_T may not have been constructed with coins chosen uniformly at random, we nevertheless have for all $1 \leq i \leq |\mathbf{BB}|$ that

$$\begin{aligned} \mathbf{C}_1[i] &= \text{Enc}(pk_T, \beta_{\lambda(\chi(i))}; r_{\lambda(\chi(i)),1} \oplus \mathbf{r}_1[i]) \\ \mathbf{C}_2[i] &= \text{Enc}(pk_T, d'_{\lambda(\chi(i))}; r_{\lambda(\chi(i)),2} \oplus \mathbf{r}_2[i]) \end{aligned}$$

and for all $1 \leq i \leq n_V$ that

$$\mathbf{C}_3[i] = \text{Enc}(pk_T, d_{\chi'(\pi(i))}; s_{\chi'(\pi(i))} \oplus \mathbf{r}_3[i])$$

because Γ is perfectly homomorphic, and \mathbf{e} is an identity element.

By Step 4 of algorithm `Verify`, we have $\mathbf{P}_{\text{inelig}}$ is a vector of length $|\mathbf{C}_2|$ and for all $1 \leq i \leq |\mathbf{C}_2|$ either: i) $\mathbf{P}_{\text{inelig}}[i]$ parses as a vector (j, σ) , $\text{VerPET}((pk_T, \mathbf{C}_2[i], \mathbf{C}_3[j], 1), \sigma, k) = 1$, and $j \in \{1, \dots, |\mathbf{C}_3|\}$, therefore, by simulation sound extractability, we have $\text{Dec}(sk_T, \mathbf{C}_2[i]) = \text{Dec}(sk_T, \mathbf{C}_3[j])$, or ii) $\mathbf{P}_{\text{inelig}}[i]$ parses as a vector $(0, \sigma_1, \dots, \sigma_{|\mathbf{C}_3|})$ and for all $1 \leq j \leq |\mathbf{C}_3|$ we have $\text{VerPET}((pk_T, \mathbf{C}_2[i], \mathbf{C}_3[j], 0), \sigma_j, k) = 1$, therefore, by simulation sound extractability, we have $\text{Dec}(sk_T, \mathbf{C}_2[i]) \neq \text{Dec}(sk_T, \mathbf{C}_3[j])$. Although, key pair pk_T and sk_T may not have been constructed with coins chosen uniformly at random, and similarly ciphertexts $\mathbf{C}_2[1], \dots, \mathbf{C}_2[|\mathbf{BB}|], \mathbf{C}_3[1], \dots, \mathbf{C}_3[n_V]$ may not have been constructed with coins chosen uniformly at random, we nevertheless have for all $1 \leq i \leq |\mathbf{C}_2|$ that if $\mathbf{P}_{\text{inelig}}[i]$ parses as a vector (j, σ) such that $j \in \{1, \dots, |\mathbf{C}_3|\}$, then $d'_{\lambda(\chi(i))} = d_{\chi'(\pi(j))}$, otherwise, $d'_{\lambda(\chi(i))} \notin \{d_1, \dots, d_{n_V}\}$, with overwhelming probability, because Γ is perfectly correct. Let \mathbf{C}'_1 be computed as per Step 4 of algorithm `Tally`. Hence, there trivially exists an injective function $\lambda' : \{1, \dots, |\mathbf{C}'_1|\} \rightarrow \{1, \dots, |\mathbf{C}_1|\}$ such that for all $1 \leq i \leq |\mathbf{C}'_1|$ we have $\mathbf{C}'_1[i] = \mathbf{C}_1[\lambda'(i)]$, moreover, $d'_{\lambda(\chi(\lambda'(i)))} \in \{d_1, \dots, d_{n_V}\}$. It follows that

$$\forall i \in \lambda(\chi(\lambda'(\{1, \dots, |\mathbf{C}'_1|\}))) : d_i \in \{d_1, \dots, d_{n_V}\} \quad (9)$$

Moreover,

$$\begin{aligned} \forall i \in \{1, \dots, \ell\} \setminus \lambda(\chi(\lambda'(\{1, \dots, |\mathbf{C}'_1|\}))) : \\ d_i \notin \{d_1, \dots, d_{n_V}\} \quad (10) \end{aligned}$$

Thus, $\{b_i \mid i \in \lambda(\chi(\lambda'(\{1, \dots, |\mathbf{C}'_1|\})))\}$ is the largest subset of ballots from $\{b_1, \dots, b_\ell\}$ such that each ballot was constructed using a distinct private credential from M .

By (6) – (10), the set of authorized ballots in $\{b_1, \dots, b_\ell\}$ is

$$BB^* = \left\{ b_i \mid i \in \lambda(\chi(\lambda'(\{1, \dots, |\mathbf{C}'_1|\}))) \right\}$$

therefore, since $\perp \notin \{b_1, \dots, b_\ell\}$, we have

$$\begin{aligned} &\text{authorized}(PK_{\mathcal{T}}, \{b_1, \dots, b_\ell\} \setminus \{\perp\}, M, n_C, k) \\ &= \text{authorized}(PK_{\mathcal{T}}, \{b_1, \dots, b_\ell\}, M, n_C, k) \\ &= \text{authorized}(PK_{\mathcal{T}}, BB^*, M, n_C, k) \\ &= BB^* \end{aligned}$$

Hence, by (5) and definition of *correct-tally*, and since $\mathbf{Y} = \text{correct-tally}(PK_{\mathcal{T}}, BB, M, n_C, k)$, it follows for all $\beta \in \{1, \dots, n_C\}$ that $\exists \mathbf{Y}^{[\beta]} b \in BB^* : \exists sk, r : b = \text{Vote}(sk, PK_{\mathcal{T}}, n_C, \beta, k; r)$, therefore, $\exists \mathbf{Y}^{[\beta]} i \in \lambda(\chi(\lambda'(\{1, \dots, |\mathbf{C}'_1|\}))) : \beta = \beta_i$ and, equivalently,

$$\exists \mathbf{Y}^{[\beta]} i \in \{1, \dots, |\mathbf{C}'_1|\} : \beta = \beta_{\lambda(\chi(\lambda'(i)))} \quad (11)$$

Thus, $\beta_{\lambda(\chi(\lambda'(1)))}, \dots, \beta_{\lambda(\chi(\lambda'(|\mathbf{C}'_1|)))}$ are the choices used to construct authorized recorded ballots.

By Step 5 of algorithm Verify, we have \mathbf{P}_{dec} is a vector $((\beta'_1, \sigma_1), \dots, (\beta'_{|\mathbf{C}'_1|}, \sigma_{|\mathbf{C}'_1|}))$ such that for all $1 \leq i \leq |\mathbf{C}'_1|$ we have $\text{VerDec}((pk_T, \mathbf{C}'_1[i], \beta'_i), \sigma_i, k) = 1$ and for all $1 \leq \beta \leq n_C$ we have $\exists^{|\mathbf{X}^{[\beta]}|} j \in \{1, \dots, |\mathbf{C}'_1|\} : \beta = \beta'_j$. And, by simulation sound extractability, we have $\beta'_j = \beta_{\lambda(\chi(\lambda'(j)))}$. Thus, we have $\mathbf{X} = \mathbf{Y}$ by (11), concluding our proof. \square

Proposition 30. *Suppose $\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5, \Sigma_6, \Sigma_7$ and \mathcal{H} satisfy the preconditions of Definition 30. Further suppose Γ is perfectly correct and Σ_2 and Σ_5 satisfy special soundness and special honest verifier zero-knowledge. We have $\text{JCJ}(\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5, \Sigma_6, \Sigma_7, \mathcal{H})$ satisfies Completeness.*

Proof. Let $\text{JCJ}(\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5, \Sigma_6, \Sigma_7, \mathcal{H}) = (\text{Setup}, \text{Register}, \text{Vote}, \text{Tally}, \text{Verify}), \text{FS}(\Sigma_1, \mathcal{H}) = (\text{ProveKey}, \text{VerKey}), \text{FS}(\Sigma_2, \mathcal{H}) = (\text{ProveCiph}, \text{VerCiph}), \text{FS}(\Sigma_4, \mathcal{H}) = (\text{ProveDec}, \text{VerDec}), \text{FS}(\Sigma_5, \mathcal{H}) = (\text{ProvePET}, \text{VerPET}), \text{FS}(\Sigma_6, \mathcal{H}) = (\text{ProveMix}, \text{VerMix}),$ and $\text{FS}(\Sigma_7, \mathcal{H}) = (\text{ProveMixPair}, \text{VerMixPair})$.

Suppose k is a security parameter and \mathcal{A} is a PPT adversary. Further suppose (PK_T, SK_T, m_B, m_C) is an output of $\text{Setup}(k)$, n_V is an output of $\mathcal{A}(PK_T, k)$, $(pd_1, d_1), \dots, (pd_{n_V}, d_{n_V})$ are outputs of $\text{Register}(PK_T, k)$, $L = \{pd_1, \dots, pd_{n_V}\}$, $M = \{(pk_1, sk_1), \dots, (pk_{n_V}, sk_{n_V})\}$, (BB, n_C) is an output of $\mathcal{A}(M)$, and (\mathbf{X}, \mathbf{P}) is an output of $\text{Tally}(SK_T, BB, L, n_C, k)$. If $|BB| \not\leq m_B \vee n_C \not\leq m_C$, then we conclude immediately, otherwise $(|BB| \leq m_B \wedge n_C \leq m_C)$, we proceed as follows.

By definition of Setup, $PK_T = (pk, m, \rho)$, $SK_T = (pk, sk)$, and $m_C = |m|$, where $(pk, sk, m) = \text{Gen}(k; r)$ and ρ is an output of $\text{ProveKey}((k, pk, m), (sk, r), k)$ for some coins r chosen uniformly at random by algorithm Setup. By definition of algorithm Tally, \mathbf{X} is a vector of length n_C and \mathbf{P} is a vector $(\mathbf{P}_{\text{dupl}}, \mathbf{C}_1, P_{\text{mix},1}, \mathbf{C}_2, P_{\text{mix},2}, \mathbf{C}_3, P_{\text{mix},3}, \mathbf{P}_{\text{inelig}}, \mathbf{P}_{\text{dec}})$. It follows that algorithm Verify can parse PK_T, \mathbf{X} and \mathbf{P} successfully. Moreover, by completeness of $(\text{ProveKey}, \text{VerKey})$, we have $\text{VerKey}((k, pk, m), \rho, k) = 1$, with overwhelming probability.

Suppose subset $\{b_1, \dots, b_\ell\}$ is computed as per Step 1 of algorithm Tally. Hence, $\{b_1, \dots, b_\ell\}$ is the largest subset of BB such that $b_1 < \dots < b_\ell$ and for all $1 \leq i \leq \ell$ we have b_i is a vector of length 4, $\text{VerCiph}((pk_T, b_i[1], \{1, \dots, n_C\}), b_i[3], k) = 1$, and $\text{VerBind}((pk_T, b_i[1], b_i[2]), b_i[4], k) = 1$. (Condition $b_1 < \dots < b_\ell$ ensures that algorithms Tally and Verify compute b_1, \dots, b_ℓ in the same order, which is necessary to ensure that proofs constructed by Tally in relation to a particular ballot, are checked by Verify in relation to that ballot.) We have $\{b_1, \dots, b_\ell\} = \emptyset$ implies \mathbf{X} is a zero-filled vector, because \mathbf{X} is initialized as a zero-filled vector. Thus, the check holds in Step 1 of Verify.

Since Σ_2 satisfies special soundness and special honest verifier zero-knowledge, we have by simulation sound extractability that for all $1 \leq i \leq \ell$ there exists messages

$\beta_i, d'_i \in \mathfrak{m}$ and coins $r_{i,1}$ and $r_{i,2}$, such that

$$\begin{aligned} b_i[1] &= \text{Enc}(pk_T, \beta_i; r_{i,1}) \\ b_i[2] &= \text{Enc}(pk_T, d'_i; r_{i,2}) \end{aligned}$$

with overwhelming probability.

Suppose \mathbf{P}_{dupl} is computed as per Step 2 of algorithm Tally. Hence, \mathbf{P}_{dupl} is a vector of length ℓ such that for all $1 \leq i \leq \ell$ we have either: i) $\mathbf{P}_{\text{dupl}}[i]$ is a vector (j, σ) , $\text{VerPET}((pk_T, b_i[2], b_j[2], 1), \sigma, k) = 1$, and $j \in \{1, \dots, i-1, i+1, \dots, \ell\}$ or ii) for all $j \in \{1, \dots, i-1, i+1, \dots, \ell\}$ we have $\text{VerPET}((pk_T, b_i[2], b_j[2], 1), \sigma, k) \neq 1$ for some output σ of $\text{ProvePET}((pk_T, b_i[2], b_j[2], 1), sk_T, k)$, and $\mathbf{P}_{\text{dupl}}[i]$ is a vector $(0, \sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_\ell)$ such that σ_j is an output of $\text{ProvePET}((pk_T, b_i[2], b_j[2], 0), sk_T, k)$ for all $j \in \{1, \dots, i-1, i+1, \dots, \ell\}$. In the former case, relevant checks trivially hold in Step 2 of Verify. Let us show that relevant checks hold in the latter case too. Although ciphertexts $b_1[2], \dots, b_\ell[2]$ may not have been constructed with coins chosen uniformly at random, we nevertheless have for all $1 \leq i \leq \ell$ that $\text{Dec}(sk, b_i[2]) \neq \perp$, because Γ is perfectly correct. Suppose $\mathbf{P}_{\text{dupl}}[i] = (0, \sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_\ell)$ in the latter case. Since Σ_5 satisfies special soundness and special honest verifier zero-knowledge, we have by simulation sound extractability that $\text{Dec}(sk, b_i[2]) \neq \text{Dec}(sk, b_j[2])$ for all integers $j \in \{1, \dots, i-1, i+1, \dots, \ell\}$, with overwhelming probability. Therefore, by completeness of $(\text{ProvePET}, \text{VerPET})$, we have $\text{VerPET}((pk_T, b_i[2], b_j[2], 0), \sigma_j, k) = 1$ for all $j \in \{1, \dots, i-1, i+1, \dots, \ell\}$, with overwhelming probability. Thus, the relevant checks hold in Step 2 of Verify, with overwhelming probability.

Suppose \mathbf{BB} is computed as per Step 2 of algorithm Tally. Moreover, suppose $\mathbf{BB} = (b'_1, \dots, b'_{|\mathbf{BB}|})$. Further suppose vectors \mathbf{C}_1 and \mathbf{C}_2 are computed as per Step 3 of algorithm Tally. Hence, for all $1 \leq i \leq |\mathbf{BB}|$ we have

$$\begin{aligned} \mathbf{C}_1[i] &= b'_{\chi(i)}[1] \otimes \text{Enc}(pk_T, \mathbf{e}; \mathbf{r}_1[i]) \text{ and} \\ \mathbf{C}_2[i] &= b'_{\chi(i)}[2] \otimes \text{Enc}(pk_T, \mathbf{e}; \mathbf{r}_2[i]), \end{aligned}$$

where χ is a permutation on $\{1, \dots, |\mathbf{BB}|\}$, and \mathbf{r}_1 and \mathbf{r}_2 are vectors of coins. Let $\mathbf{BB}_1 = (b'_1[1], \dots, b'_{|\mathbf{BB}|}[1])$ and $\mathbf{BB}_2 = (b'_1[2], \dots, b'_{|\mathbf{BB}|}[2])$. Suppose $P_{\text{mix},1}$ is computed as per Step 3 of algorithm Tally. Hence, $P_{\text{mix},1}$ is an output of $\text{ProveMixPair}((pk_T, \mathbf{BB}_1, \mathbf{C}_1, \mathbf{BB}_2, \mathbf{C}_2), (\mathbf{r}_1, \mathbf{r}_2, \chi), k)$. By the completeness of $(\text{ProveMixPair}, \text{VerMixPair})$, we have $\text{VerMixPair}((pk_T, \mathbf{BB}_1, \mathbf{C}_1, \mathbf{BB}_2, \mathbf{C}_2), P_{\text{mix},1}, k) = 1$, with overwhelming probability. Similarly, suppose $L = \{pd_1, \dots, pd_{|L|}\}$ such that $pd_1 < \dots < pd_{|L|}$. Moreover, suppose vector \mathbf{C}_3 is computed as per Step 3 of algorithm Tally. Hence, for all $1 \leq i \leq |L|$ we have

$$\mathbf{C}_3[i] = pd_{\chi'(i)} \otimes \text{Enc}(pk_T, \mathbf{e}; \mathbf{r}_3[i]),$$

where χ' is a permutation on $\{1, \dots, |L|\}$ and \mathbf{r}_3 is a vector of coins chosen uniformly at random by algorithm Tally. Suppose $P_{\text{mix},3}$ is also computed as per Step 3 of algorithm Tally. Hence, $P_{\text{mix},3}$ is an output of $\text{ProveMix}((pk_T, (pd_1, \dots, pd_{|L|}), \mathbf{C}_3), (\mathbf{r}_3, \chi'), k)$.

By the completeness of (ProveMix, VerMix), we have $\text{VerMix}((pk_T, (pd_1, \dots, pd_{|L|}), \mathbf{C}_3), P_{mix,3}, k) = 1$, with overwhelming probability. It follows that checks hold in Step 3 of Verify, with overwhelming probability.

Suppose $\mathbf{P}_{\text{inelig}}$ is computed as per Step 4 of algorithm Tally. Hence, $\mathbf{P}_{\text{inelig}}$ is a vector of length $|\mathbf{C}_2|$ such that for all $1 \leq i \leq |\mathbf{C}_2|$ we have either: i) $\mathbf{P}_{\text{inelig}}[i]$ is a vector (j, σ) , $\text{VerPET}((pk_T, \mathbf{C}_2[i], \mathbf{C}_3[j], 1), \sigma, k) = 1$, and $j \in \{1, \dots, |\mathbf{C}_3|\}$, or ii) for all $j \in \{1, \dots, |\mathbf{C}_3|\}$ we have $\text{VerPET}((pk_T, \mathbf{C}_2[i], \mathbf{C}_3[j], 1), \sigma, k) \neq 1$ for some output σ of $\text{ProvePET}((pk_T, \mathbf{C}_2[i], \mathbf{C}_3[j], 1), sk_T, k)$, and $\mathbf{P}_{\text{inelig}}[i]$ is a vector $(0, \sigma_1, \dots, \sigma_{|\mathbf{C}_3|})$ such that for all $j \in \{1, \dots, |\mathbf{C}_3|\}$ we have σ_j is an output of $\text{ProvePET}((pk_T, \mathbf{C}_2[i], \mathbf{C}_3[j], 0), sk_T, k)$. In the former case, relevant checks trivially hold in Step 4 of Verify. Let us show that relevant checks hold in the latter case too. We have for all $1 \leq i \leq |L|$ that $pd_{\chi'(i)}$ is a ciphertext on $d_{\chi'(i)} \in \mathbf{m}$ constructed using some coins r_i chosen uniformly at random by algorithm Register. Thus, for all $1 \leq i \leq |L|$ we have $\mathbf{C}_3[i] = \text{Enc}(pk_T, d_{\chi'(i)}; r_i \oplus \mathbf{r}_3[i])$, therefore, $\text{Dec}(sk, \mathbf{C}_3[j]) \neq \perp$, with overwhelming probability, because Γ is homomorphic and ϵ is an identity element. Moreover, we have for all $1 \leq i \leq |\mathbf{BB}|$ that $\mathbf{C}_2[i] = \text{Enc}(pk_T, d'_{\lambda(\chi(i))}; r_{\lambda(\chi(i)),2} \oplus \mathbf{r}_2[i])$, with overwhelming probability, because Γ is homomorphic and ϵ is an identity element. And, since Γ is perfectly correct, we have $\text{Dec}(sk, \mathbf{C}_2[i]) \neq \perp$ for all $1 \leq i \leq |\mathbf{BB}|$. (The homomorphic property of Γ is insufficient to infer $\text{Dec}(sk, \mathbf{C}_2[i]) \neq \perp$, because ciphertext $b'_{\chi(i)}[2]$ may not have been constructed using coins chosen uniformly at random.) Suppose $\mathbf{P}_{\text{inelig}}[i] = (0, \sigma_1, \dots, \sigma_{|\mathbf{C}_3|})$ in the latter case. Since Σ_5 satisfies special soundness and special honest verifier zero-knowledge, we have by simulation sound extractability that $\text{Dec}(sk, \mathbf{C}_2[i]) \neq \text{Dec}(sk, \mathbf{C}_3[j])$ for all $1 \leq j \leq |L|$, with overwhelming probability. Therefore, by completeness of (ProvePET, VerPET), we have $\text{VerPET}((pk_T, \mathbf{C}_2[i], \mathbf{C}_3[j], 0), \sigma_j, k) = 1$ for all $1 \leq j \leq |L|$, with overwhelming probability. Thus, the relevant checks hold in Step 4 of Verify, with overwhelming probability.

Suppose \mathbf{C}'_1 is computed as per Step 4 of algorithm Tally. And \mathbf{P}_{dec} is computed as per Step 5 of algorithm Tally. Hence, \mathbf{P}_{dec} is a vector $((\beta_1, \sigma_1), \dots, (\beta_{|\mathbf{C}'_1|}, \sigma_{|\mathbf{C}'_1|}))$ such that for all $1 \leq i \leq |\mathbf{C}'_1|$ we have $\beta_i = \text{Dec}(sk_T, \mathbf{C}'_1[i])$ and σ_i is an output of $\text{ProveDec}(pk_T, \mathbf{C}'_1[i], \beta_i, sk_T, k)$, therefore, by completeness of (ProveDec, VerDec), we have $\text{VerDec}((pk_T, \mathbf{C}'_1[i], \beta_i), \sigma_i, k) = 1$, with overwhelming probability. Moreover, since \mathbf{X} is derived by initializing \mathbf{X} as a zero-filled vector of length n_C and computing **for** $1 \leq i \leq |\mathbf{C}'_1|$ **do** $\mathbf{X}[\beta] \leftarrow \mathbf{X}[\beta] + 1$, we have for all $1 \leq \beta \leq n_C$ that $\exists = \mathbf{X}[\beta] j \in \{1, \dots, |\mathbf{C}'_1|\} : \beta = \beta_j$. It follows that checks hold in Step 5 of Verify, with overwhelming probability.

Since all the above checks succeed, Verify outputs 1, with overwhelming probability, concluding our proof. \square

C. Eligibility Verifiability

We derive an asymmetric encryption scheme from generalized JCJ (Definition 34) which satisfies IND-PA0 (Proposi-

tion 31), and prove that eligibility verifiability follows (Proposition 32).

Definition 34. Suppose $\Gamma = (\text{Gen}, \text{Enc}, \text{Dec})$ is a multiplicatively homomorphic asymmetric encryption scheme, Σ_3 proves conjunctive plaintext knowledge, and \mathcal{H} is a random oracle. Let $\text{FS}(\Sigma_3, \mathcal{H}) = (\text{ProveBind}, \text{VerBind})$. We define $\Gamma\text{-JCJ}(\Gamma, \Sigma_3, \mathcal{H}) = (\text{Gen}', \text{Enc}', \text{Dec}')$ as follows:

- $\text{Gen}'(k)$ selects coins r uniformly at random, computes $(pk, sk, \mathbf{m}) \leftarrow \text{Gen}(k; r)$; $\mathbf{m}' \leftarrow \{(m_1, m_2) \mid m_1, m_2 \in \mathbf{m}\}$, and outputs (pk, sk, \mathbf{m}') .
- $\text{Enc}'(pk, \mathbf{m})$ parses \mathbf{m} as a vector of length 2, outputting \perp if parsing fails; selects coins r_1 and r_2 uniformly at random; computes $c_1 \leftarrow \text{Enc}(pk, \mathbf{m}[1]; r_1)$; $c_2 \leftarrow \text{Enc}(pk, \mathbf{m}[2]; r_2)$; $\tau \leftarrow \text{ProveBind}((pk, c_1, c_2), (\mathbf{m}[1], r_1, \mathbf{m}[2], r_2), k)$; and outputs (c_1, c_2, τ) .
- $\text{Dec}'(sk, c)$ parses c as (c_1, c_2, τ) , outputting \perp if parsing fails or $\text{VerBind}((pk, c_1, c_2), \tau, k) \neq 1$; computes $m_1 \leftarrow \text{Dec}(sk, c_1)$; $m_2 \leftarrow \text{Dec}(sk, c_2)$; and outputs (m_1, m_2) .

Proposition 31. Let Γ be a multiplicatively homomorphic asymmetric encryption scheme, Σ_3 be a sigma protocol that proves conjunctive plaintext knowledge, and \mathcal{H} be a random oracle. Suppose Γ satisfies IND-CPA and Σ_3 satisfies special soundness and special honest verifier zero-knowledge. We have $\Gamma\text{-JCJ}(\Gamma, \Sigma_3, \mathcal{H})$ satisfies IND-PA0.

A proof of Proposition 31 is similar to the proof of [24, Theorem 5.1], so we omit formalizing a proof.

Proposition 32. Suppose $\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5, \Sigma_6, \Sigma_7$ and \mathcal{H} satisfy the preconditions of Definition 33. Further suppose that satisfies IND-CPA, and Σ_1 and Σ_3 satisfy special soundness and special honest verifier zero-knowledge. We have $\text{JCJ}(\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5, \Sigma_6, \Sigma_7, \mathcal{H})$ satisfies Exp-EV-Int-Weak.

Proof. Let $\Gamma = (\text{Gen}, \text{Enc}, \text{Dec})$, $\Gamma\text{-JCJ}(\Gamma, \Sigma_3, \mathcal{H}) = (\text{Gen}', \text{Enc}', \text{Dec}')$, $\text{FS}(\Sigma_2, \mathcal{H}) = (\text{ProveCiph}, \text{VerCiph})$, and $\text{FS}(\Sigma_3, \mathcal{H}) = (\text{ProveBind}, \text{VerBind})$. By Theorem 14, there exists a simulator for proof system (ProveBind, VerBind). Let SimProveBind be such a simulator. Similarly, let SimProveKey be a simulator for $\text{FS}(\Sigma_1, \mathcal{H})$. Moreover, let ϵ be an identity element of Γ 's message space with respect to \odot .

By Lemma 25, it suffices to show that Exp-EV-Int-Weak' is satisfied. We proceed by contradiction. Suppose Exp-EV-Int-Weak' is not satisfied, hence, there exists a PPT adversary \mathcal{A} that wins Exp-EV-Int-Weak' with non-negligible probability. We construct an adversary \mathcal{B} against $\Gamma\text{-JCJ}(\Gamma, \Sigma_3, \mathcal{H})$.

- $\mathcal{B}(pk, \mathbf{m}, k)$ computes $d_0 \leftarrow_R \mathbf{m}$; $d_1 \leftarrow_R \mathbf{m}$ and outputs $((\epsilon, d_0), (\epsilon, d_1))$.
- $\mathcal{B}(c)$ parses c as a vector of length 3, computes $\rho \leftarrow \text{SimProveKey}((k, pk, \mathbf{m}), k)$; $PK_{\mathcal{T}} \leftarrow (pk, \mathbf{m}, \rho)$; $(n_C, \beta, b) \leftarrow \mathcal{A}(PK_{\mathcal{T}}, c[2], k)$, and outputs $((b[1], b[2], b[4]))$, responding to \mathcal{A} 's oracle calls $R'(\beta, n_C)$ as follows, namely, if $\beta \notin \{1, \dots, n_C\} \vee \{1, \dots, n_C\} \not\subseteq \mathbf{m}$,

then return \perp , otherwise, select coins r_1 uniformly at random, compute

$$\begin{aligned} c_1 &\leftarrow \text{Enc}(pk, \beta; r_1); \\ c_2 &\leftarrow \text{Enc}(pk, \epsilon) \otimes \mathbf{c}[2]; \\ \sigma &\leftarrow \text{ProveCiph}((pk, c_1, \{1, \dots, n_C\}), (\beta, r_1), k); \\ \tau &\leftarrow \text{SimProveBind}((pk, c_1, \mathbf{c}[2]), k); \\ b &\leftarrow (c_1, c_2, \sigma, \tau); \end{aligned}$$

and return b .

- $\mathcal{B}(\mathbf{m})$ parses $\mathbf{m}[1]$ as a vector (β, d) and if $d = d_0$, then outputs 0, otherwise, outputs 1.

We prove \mathcal{B} wins IND-PA0 with non-negligible probability.

Suppose (pk, sk, \mathbf{m}) is an output of $\text{Gen}(k)$, $(\mathbf{m}_0, \mathbf{m}_1)$ is an output of $\mathcal{B}(pk, \mathbf{m}, k)$, and \mathbf{c} is an output of $\text{Enc}'(pk, \mathbf{m}_\alpha)$, for some bit α chosen uniformly at random. By definition of \mathcal{B} and Enc' , we have \mathbf{c} is a vector such that $\mathbf{c}[2]$ is an output of $\text{Enc}'(pk, \mathbf{m}_\alpha[2])$, where $\mathbf{m}_\alpha[2] \in \mathbf{m}$ was chosen uniformly at random by \mathcal{B} . Further suppose we run $\mathcal{B}(\mathbf{c})$. Hence, we compute $\rho \leftarrow \text{SimProveKey}((k, pk, \mathbf{m}), k); PK_{\mathcal{T}} \leftarrow (pk, \mathbf{m}, \rho); (n_C, \beta, b) \leftarrow \mathcal{A}(PK_{\mathcal{T}}, \mathbf{c}[2], k)$. It is straightforward to see that \mathcal{B} simulates \mathcal{A} 's challenger to \mathcal{A} , because proofs output by SimProveKey are indistinguishable from proofs produced by proof system $\text{FS}(\Sigma_1, \mathcal{H})$ and $\mathbf{c}[2]$ corresponds to a public credential. Let us assume that \mathcal{B} simulates \mathcal{A} 's oracle to \mathcal{A} too. Hence, since \mathcal{A} is a winning adversary, we have b is an output of $\text{Vote}(\mathbf{m}_\alpha[2], PK_{\mathcal{T}}, n_C, \beta, k)$ such that $b \neq \perp$ and b was not simulated by \mathcal{B} in response to an oracle call by \mathcal{A} . By definition of Vote , we have $b[1] = \text{Enc}'(pk, \beta; r_1)$, $b[2] = \text{Enc}'(pk, \mathbf{m}_\alpha[2]; r_2)$, and $b[4]$ is an output of $\text{ProveBind}((pk, \mathbf{c}[1], \mathbf{c}[2]), (\beta, r_1, \mathbf{m}_\alpha[2], r_2), k)$, for some coins r_1 and r_2 . Moreover, we have $\beta \in \{1, \dots, n_C\}$ and $\{1, \dots, n_C\} \not\subseteq \mathbf{m}$, hence, $\beta \in \mathbf{m}$. Suppose the run of $\mathcal{B}(\mathbf{c})$ concludes by outputting $((b[1], b[2], b[4]))$. By completeness of $(\text{ProveBind}, \text{VerBind})$, we have $\text{VerBind}((pk, b[1], b[2]), b[4], k) = 1$. Hence, $\text{Dec}(sk, (b[1], b[2], b[4])) = (\beta, \mathbf{m}_\alpha[2])$. Further suppose g is an output of $\mathcal{A}((\beta, \mathbf{m}_\alpha[2]))$. Thus, by definition of \mathcal{A} , we have $g = \alpha$. Moreover, we have $\mathbf{c} \neq (b[1], b[2], b[4])$, because $\mathbf{c}[1]$ is not revealed to \mathcal{A} , hence, \mathcal{A} cannot construct $\mathbf{c}[1]$, due to the precondition that Γ satisfies IND-CPA. It remains to prove that \mathcal{B} simulates \mathcal{A} 's oracle to \mathcal{A} .

An oracle call $R'(\beta, n_C)$ outputs \perp if $\beta \notin \{1, \dots, n_C\} \vee \{1, \dots, n_C\} \not\subseteq \mathbf{m}$, and it is trivial to see that \mathcal{B} simulates \mathcal{A} 's oracle to \mathcal{A} in this case. Otherwise, $R'(\beta, n_C)$ computes $b \leftarrow \text{Vote}(\mathbf{m}_\alpha[2], PK_{\mathcal{T}}, n_C, \beta, k)$ and outputs b . By definition of Vote , we have b is a vector of length 4. It is trivial to see that \mathcal{B} simulates the computation of $b[1]$ and $b[3]$. Moreover, if \mathcal{B} simulates the computation of $b[2]$, then \mathcal{B} simulates the computation of $b[4]$ too, because proofs output by SimProveBind are indistinguishable from proofs output by ProveBind . Thus, it remains to prove that \mathcal{B} simulates the computation of $b[2]$. It suffices to show that selecting coins r_2 uniformly at random and computing $c_2 \leftarrow \text{Enc}'(pk, \mathbf{m}_\alpha[2]; r_2)$ is indistinguishable from computing $c_2 \leftarrow \text{Enc}'(pk, \epsilon) \otimes \mathbf{c}[2]$, i.e., $c_2 \leftarrow \text{Enc}'(pk, \epsilon; r'_2) \otimes \text{Enc}'(pk, \mathbf{m}_\alpha[2]; r)$, where coins r'_2 and r are selected uniformly at random. Since Γ is a homomorphic asymmetric encryption scheme and ϵ is an identity element, we have

$c_2 \leftarrow \text{Enc}'(pk, \epsilon; r'_2) \otimes \text{Enc}'(pk, \mathbf{m}_\alpha[2]; r)$ is indistinguishable from $c_2 \leftarrow \text{Enc}'(pk, \mathbf{m}_\alpha[2]; r'_2 \oplus r)$, which is indistinguishable from $c_2 \leftarrow \text{Enc}'(pk, \mathbf{m}_\alpha[2]; r_2)$, because coins $r'_2 \oplus r$ are indistinguishable from coins (e.g., r_2) selected uniformly at random, thereby concluding our proof. \square

D. Proof: Theorem 10

By Propositions 27, 29, 30, & 32 and Lemma 28, election schemes constructed from $\widehat{\text{JCJ}}$ satisfy election verifiability with internal authentication:

Corollary 33. *Suppose $\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5, \Sigma_6, \Sigma_7$ and \mathcal{H} satisfy the preconditions of Definition 33. Further suppose that Γ is perfectly correct, perfectly homomorphic, and collision-free for its message space. Moreover, suppose Γ satisfies IND-CPA. Furthermore, suppose the sigma protocols satisfy special soundness and special honest verifier zero-knowledge, and \mathcal{H} is a random oracle. We have $\widehat{\text{JCJ}}(\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5, \Sigma_6, \Sigma_7, \mathcal{H})$ satisfies election verifiability with internal authentication.*

Proof of Theorem 10. Let $\text{JCJ}'16$ be the set of election schemes derived from $\widehat{\text{JCJ}}(\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5, \Sigma_6, \Sigma_7, \mathcal{H})$, where primitives $\Gamma, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5, \Sigma_6, \Sigma_7$ and \mathcal{H} satisfy the conditions identified in Corollary 33. Hence, Theorem 10 is an immediate consequence of Corollary 33. \square

APPENDIX L

JUELS ET AL. DEFINITIONS

Juels et al. [87, §2] define an election scheme as a tuple of (Register, Vote, Tally, Verify) PPT algorithms:

- **Register**, denoted $(pk, sk) \leftarrow \text{Register}(SK_{\mathcal{R}}, i, k_1)$, is executed by the registrars. Register takes as input the private key $SK_{\mathcal{R}}$ of the registrars, a voter's identity i , and security parameter k_1 . It outputs a credential pair (pk, sk) .
- **Vote**, denoted $b \leftarrow \text{Vote}(sk, PK_{\mathcal{T}}, n_C, \beta, k_2)$, is executed by voters. Vote takes as input a voter's private credential sk , the public key $PK_{\mathcal{T}}$ of the tallier, the number of candidates n_C , the voter's choice β , and security parameter k_2 . It outputs a ballot b .
- **Tally**, denoted $(\mathbf{X}, P) \leftarrow \text{Tally}(SK_{\mathcal{T}}, BB, n_C, \{pk_i\}_{i=1}^{n_V}, k_3)$, is executed by the tallier. Tally takes as input the private key $SK_{\mathcal{T}}$ of the tallier, the bulletin board BB , the number of candidates n_C , the set containing voters' public credentials, and security parameter k_3 . It outputs the tally \mathbf{X} and a proof P that the tally is correct.
- **Verify**, denoted $v \leftarrow \text{Verify}(PK_{\mathcal{R}}, PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, P)$, can be executed by anyone to audit the election. Verify takes as input the public key $PK_{\mathcal{R}}$ of the registrars, the public key $PK_{\mathcal{T}}$ of the tallier, the bulletin board BB , the number of candidates n_C , and a candidate proof P of correct tallying. It outputs a bit v , which is 1 if the tally successfully verifies and 0 on failure.

The above definition fixes an apparent oversight in JCJ's presentation: we supply the registrars' public key as input to

the verification algorithm, because that key would be required by Verify to check the signature on the electoral roll.

Juels et al. [87, §3] formalize *correctness* and *verifiability* to capture their notion of election verifiability. We rename those to *JCJ-correctness* and *JCJ-verifiability* to avoid ambiguity. For readability, the definitions we give below contain subtle differences from the original presentation. For example, we sometimes use for loops instead of pattern matching.

JCJ-correctness asserts that an adversary cannot modify or eliminate votes of honest voters, and stipulates that at most one ballot is tallied per voter. Intuitively, the security definition challenges the adversary to ensure that verification succeeds and the tally⁶² does not include some honest votes or contains too many votes. The definition of JCJ-correctness fixes apparent errors in the original presentation: the adversary is given the credentials for corrupt voters and distinct security parameters are supplied to the Register and Vote algorithms. An implicit assumption is also omitted: $\{\beta_i\}_{i \in \mathcal{V} \setminus \mathcal{V}'}$ is a multiset of valid votes, that is, for all $\beta \in \{\beta_i\}_{i \in \mathcal{V} \setminus \mathcal{V}'}$ we have $1 \leq \beta \leq n_C$. Without this assumption the security definition cannot be satisfied by many election schemes, including the election scheme by Juels et al.

Definition 35 (JCJ-correctness). *An election scheme $\Pi = (\text{Register}, \text{Vote}, \text{Tally}, \text{Verify})$ satisfies JCJ-correctness if for all PPT adversary \mathcal{A} , there exists a negligible function μ , such that for all positive integers n_C and n_V , and security parameters k_1 , k_2 , and k_3 , we have $\text{Succ}(\text{Exp-JCJ-Cor}(\Pi, \mathcal{A}, n_C, n_V, k_1, k_2, k_3)) \leq \mu(k_1, k_2, k_3)$, where Exp-JCJ-Cor is defined as follows:*⁶³

```

Exp-JCJ-Cor( $\Pi, \mathcal{A}, n_C, n_V, k_1, k_2, k_3$ ) =
1  $\mathcal{V} \leftarrow \{1, \dots, n_V\}$ ;
2 for  $i \in \mathcal{V}$  do  $(pk_i, sk_i) \leftarrow \text{Register}(SK_{\mathcal{R}}, i, k_1)$ ;
3  $\mathcal{V}' \leftarrow \mathcal{A}(\{pk_i\}_{i=1}^{n_V})$ ;
4 for  $i \in \mathcal{V} \setminus \mathcal{V}'$  do  $\beta_i \leftarrow \mathcal{A}()$ ;
5  $BB \leftarrow \{\text{Vote}(sk_i, PK_{\mathcal{T}}, n_C, \beta_i, k_2)\}_{i \in \mathcal{V} \setminus \mathcal{V}'}$ ;
6  $(\mathbf{X}, P) \leftarrow \text{Tally}(SK_{\mathcal{T}}, BB, n_C, \{pk_i\}_{i=1}^{n_V}, k_3)$ ;
7  $BB \leftarrow BB \cup \mathcal{A}(BB, \{(pk_i, sk_i)\}_{i \in \mathcal{V} \cap \mathcal{V}'})$ ;
8  $(\mathbf{X}', P') \leftarrow \text{Tally}(SK_{\mathcal{T}}, BB, n_C, \{pk_i\}_{i=1}^{n_V}, k_3)$ ;
9 if  $\text{Verify}(PK_{\mathcal{R}}, PK_{\mathcal{T}}, BB, n_C, \mathbf{X}', P') = 1$ 
   $\wedge (\{\beta_i\}_{i \in \mathcal{V} \setminus \mathcal{V}'} \not\subseteq \langle \mathbf{X}' \rangle \vee |\langle \mathbf{X}' \rangle| - |\langle \mathbf{X} \rangle| > |\mathcal{V}'|)$  then
10 | return 1
11 else
12 | return 0

```

The JCJ-correctness definition implicitly assumes that the tally and associated proof are honestly computed using the Tally algorithm. By comparison, the definition of JCJ-verifiability (Definition 36) does not use this assumption, hence, JCJ-verifiability is intended to assert that voters and auditors can check whether votes have been recorded and tallied correctly. Intuitively, the adversary is assumed to control the tallier and voters, and the security definition challenges the adversary to concoct an election (that is, the adversary generates a bulletin board BB , a tally \mathbf{X} , and a proof of

tallying P) such that verification succeeds and tally \mathbf{X} differs tally \mathbf{X}' derived from honestly tallying the bulletin board BB . It follows that there is at most one verifiable tally that can be derived.

Definition 36 (JCJ-verifiability). *An election scheme $\Pi = (\text{Register}, \text{Vote}, \text{Tally}, \text{Verify})$ satisfies JCJ-verifiability if for all PPT adversary \mathcal{A} , there exists a negligible function μ , such that for all positive integers n_C and n_V , and security parameters k_1 and k_3 , we have $\text{Succ}(\text{Exp-JCJ-Ver}(\Pi, \mathcal{A}, n_C, n_V, k_1, k_2, k_3)) \leq \mu(k_1, k_2, k_3)$, where Exp-JCJ-Ver is defined as follows:*

```

Exp-JCJ-Ver( $\Pi, \mathcal{A}, n_C, n_V, k_1, k_2, k_3$ ) =
1 for  $1 \leq i \leq n_V$  do  $(pk_i, sk_i) \leftarrow \text{Register}(SK_{\mathcal{R}}, i, k_1)$ ;
2  $(BB, \mathbf{X}, P) \leftarrow \mathcal{A}(SK_{\mathcal{T}}, \{(pk_i, sk_i)\}_{i=1}^{n_V})$ ;
3  $(\mathbf{X}', P') \leftarrow \text{Tally}(SK_{\mathcal{T}}, BB, n_C, \{pk_i\}_{i=1}^{n_V}, k_3)$ ;
4 if  $\text{Verify}(PK_{\mathcal{R}}, PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, P) = 1 \wedge \mathbf{X} \neq \mathbf{X}'$ 
  then
5 | return 1
6 else
7 | return 0

```

APPENDIX M

PROOFS: JUELS ET AL. ADMIT ATTACKS

This appendix contains proofs demonstrating that the definition of election verifiability by Juels et al. [87] admits collusion and biasing attacks (§VIII). We have reported these findings to the original authors.^{64,65}

A. Proof: Proposition 11

Suppose $\Pi = (\text{Register}, \text{Vote}, \text{Tally}, \text{Verify})$ is an election scheme satisfying JCJ-correctness and JCJ-verifiability. Further suppose $\text{Stuff}(\Pi, \beta, \kappa) = (\text{Register}, \text{Vote}, \text{Tally}_S, \text{Verify}_S)$, for some integers $\beta, \kappa \in \mathbb{N}$. We prove that $\text{Stuff}(\Pi, \beta, \kappa)$ satisfies JCJ-correctness and JCJ-verifiability.

We show that $\text{Stuff}(\Pi, \beta, \kappa)$ satisfies JCJ-correctness by contradiction. Suppose $\text{Succ}(\text{Exp-JCJ-Cor}(\text{Stuff}(\Pi, \beta, \kappa), \mathcal{A}, n_C, n_V, k_1, k_2, k_3))$ is non-negligible for some k_1, k_2, k_3, n_C, n_V , and \mathcal{A} . Hence, there exists an execution of the experiment

$$\text{Exp-JCJ-Cor}(\text{Stuff}(\Pi, \beta, \kappa), \mathcal{A}, n_C, n_V, k_1, k_2, k_3)$$

that satisfies

$$\text{Verify}_S(PK_{\mathcal{R}}, PK_{\mathcal{T}}, BB, n_C, \mathbf{X}', P') = 1$$

$$\wedge (\{\beta_i\}_{i \in \mathcal{V} \setminus \mathcal{V}'} \not\subseteq \langle \mathbf{X}' \rangle \vee |\langle \mathbf{X}' \rangle| - |\langle \mathbf{X} \rangle| > |\mathcal{V}'|)$$

with non-negligible probability, where $\{\beta_i\}_{i \in \mathcal{V} \setminus \mathcal{V}'}$ is the set of honest votes, (\mathbf{X}, P) is the tally of honest votes, (\mathbf{X}', P')

62. Juels et al. translate tallies \mathbf{X} into a multisets $\langle \mathbf{X} \rangle$ representing the tally as follows: $\langle \mathbf{X} \rangle = \bigcup_{1 \leq j \leq |\mathbf{X}|} \underbrace{\{j, \dots, j\}}_{\mathbf{X}[j] \text{ times}}$.

63. We write $\mu(k_1, k_2, k_3)$ for the smallest value in $\{\mu(k_1), \mu(k_2), \mu(k_3)\}$ (cf. [87, pp45]).

64. Dario Catalano, personal communication, Paris, France, 10 October 2013.

65. Markus Jakobsson, personal communication, New Orleans, USA, 27 June 2013.

is the tally of all votes, \mathcal{V} is a set of corrupt voter identities, and BB is the bulletin board. Further suppose BB_0 is the bulletin board BB before adding stuffed ballots. By definition of Tally_S , there exist computations

$$(\mathbf{Y}, Q) \leftarrow \text{Tally}(SK_{\mathcal{T}}, BB_0, n_C, \{pk_i\}_{i=1}^{n_V}, k_3)$$

and

$$(\mathbf{Y}', Q') \leftarrow \text{Tally}(SK_{\mathcal{T}}, BB, n_C, \{pk_i\}_{i=1}^{n_V}, k_3)$$

such that $\mathbf{X} = \text{Add}(\mathbf{Y}, \beta, \kappa)$, $\mathbf{X}' = \text{Add}(\mathbf{Y}', \beta, \kappa)$, and $P' = Q'$. Since $\kappa \in \mathbb{N}$, we have $\langle \mathbf{Y}' \rangle \subseteq \langle \mathbf{X}' \rangle$. Moreover, $|\langle \mathbf{X} \rangle| = |\langle \mathbf{Y} \rangle| + \kappa$ and $|\langle \mathbf{X}' \rangle| = |\langle \mathbf{Y}' \rangle| + \kappa$, hence,

$$|\langle \mathbf{Y}' \rangle| - |\langle \mathbf{Y} \rangle| = |\langle \mathbf{X}' \rangle| - |\langle \mathbf{X} \rangle|.$$

By definition of Verify_S and since $\mathbf{Y}' = \text{Sub}(\mathbf{X}', \beta, \kappa)$, there exists a computation

$$v \leftarrow \text{Verify}_0(PK_{\mathcal{R}}, PK_{\mathcal{T}}, BB, n_C, \mathbf{Y}', Q')$$

such that $v = 1$. It follows that

$$\begin{aligned} \text{Verify}(PK_{\mathcal{R}}, PK_{\mathcal{T}}, BB, n_C, \mathbf{Y}', Q') &= 1 \\ \wedge (\{\beta_i\}_{i \in \mathcal{V} \setminus \mathcal{V}'} \notin \langle \mathbf{Y}' \rangle \vee |\langle \mathbf{Y}' \rangle| - |\langle \mathbf{Y} \rangle| > |\mathcal{V}'|) \end{aligned}$$

with non-negligible probability and, furthermore, we have $\text{Succ}(\text{Exp-JCJ-Cor}(\Pi, \mathcal{A}, n_C, n_V, k_1, k_2, k_3))$ is non-negligible, thereby deriving a contradiction.

We show that $\text{Stuff}(\Pi, \beta, \kappa)$ satisfies JCJ-verifiability by contradiction. Suppose $\text{Succ}(\text{Exp-JCJ-Ver}(\text{Stuff}(\Pi, \beta, \kappa), \mathcal{A}, n_C, n_V, k_1, k_2, k_3))$ is non-negligible for some k_1, k_3, n_C, n_V , and \mathcal{A} . Hence, there exists an execution of the experiment $\text{Exp-JCJ-Ver}(\text{Stuff}(\Pi, \beta, \kappa), \mathcal{A}, n_C, n_V, k_1, k_2, k_3)$ which satisfies

$$\text{Verify}(PK_{\mathcal{R}}, PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, P) = 1 \wedge \mathbf{X} \neq \mathbf{X}'$$

with non-negligible probability, where (BB, \mathbf{X}, P) is an election concocted by the adversary and (\mathbf{X}', P') is produced by tallying BB . By definition of Tally_S , there exists a computation

$$(\mathbf{Y}', Q') \leftarrow \text{Tally}(SK_{\mathcal{T}}, BB, n_C, \{pk_i\}_{i=1}^{n_V}, k_3)$$

such that $\mathbf{X}' = \text{Add}(\mathbf{Y}', \beta, \kappa)$ and $P' = Q'$. By definition of Verify_S , there exists a computation

$$v \leftarrow \text{Verify}(PK_{\mathcal{R}}, PK_{\mathcal{T}}, BB, n_C, \text{Sub}(\mathbf{X}, \beta, \kappa), P)$$

such that $v = 1$. Let the adversary \mathcal{B} be defined as follows: given input K and S , the adversary \mathcal{B} computes

$$(BB, \mathbf{X}, P) \leftarrow \mathcal{A}(K, S)$$

and outputs $(BB, \text{Sub}(\mathbf{X}, \beta, \kappa), P)$. We have an execution of the experiment $\text{Exp-JCJ-Ver}(\text{Stuff}(\Pi, \beta, \kappa), \mathcal{B}, n_C, n_V, k_1, k_2, k_3)$ that concocts the election $(BB, \text{Sub}(\mathbf{X}, \beta, \kappa), P)$ and tallying BB produces (\mathbf{Y}', Q') such that

$$\text{Verify}(PK_{\mathcal{R}}, PK_{\mathcal{T}}, BB, n_C, \text{Sub}(\mathbf{X}, \beta, \kappa), P) = 1$$

with non-negligible probability. Moreover, since $\mathbf{X} \neq \mathbf{X}'$ and $\mathbf{Y}' = \text{Sub}(\mathbf{X}', \beta, \kappa)$, we have $\text{Sub}(\mathbf{X}, \beta, \kappa) \neq \mathbf{Y}'$ with

non-negligible probability. It follows immediately that $\text{Succ}(\text{Exp-JCJ-Cor}(\Pi, \mathcal{B}, n_C, n_V, k_1, k_2, k_3))$ is non-negligible, thus deriving a contradiction and concluding our proof. \square

B. Proof: Proposition 12

We define key leakage before proving Proposition 12.

Definition 37 (Key leakage). *An election scheme $\Pi = (\text{Register}, \text{Vote}, \text{Tally}, \text{Verify})$ does not leak the tallier's private key if for all positive integers n_C and n_V , security parameters k_1 and k_3 , and PPT adversary \mathcal{A} , we have $\text{Succ}(\text{Exp-leak}(\Pi, \mathcal{A}, k_1, k_3, n_C, n_V))$ is negligible, where $\text{Exp-leak}(\cdot)$ is defined as follows:*

$$\text{Exp-leak}(\Pi, \mathcal{A}, k_1, k_3, n_C, n_V) =$$

```

1 for  $1 \leq i \leq n_V$  do  $(pk_i, sk_i) \leftarrow \text{Register}(SK_{\mathcal{R}}, i, k_1)$ ;
2  $BB \leftarrow \mathcal{A}(sk_1, \dots, sk_{n_V})$ ;
3  $(\mathbf{X}, P) \leftarrow \text{Tally}(SK_{\mathcal{T}}, BB, n_C, \{pk_i\}_{i=1}^{n_V}, k_3)$ ;
4 if  $\mathcal{A}(\mathbf{X}, P) = SK_{\mathcal{T}}$  then
5 | return 1
6 else
7 | return 0

```

Proof of Proposition 12. Suppose $\Pi = (\text{Register}, \text{Vote}, \text{Tally}, \text{Verify})$ is an election scheme satisfying JCJ-correctness and JCJ-verifiability that does not leak the tallier's private key, and $\text{Backdoor}(\Pi) = (\text{Register}, \text{Vote}, \text{Tally}_B, \text{Verify}_B)$. Since Π does not leak the tallier's private key, the experiment $\text{Exp-JCJ-Cor}(\Pi, \mathcal{A}, n_C, n_V, k_1, k_2, k_3)$ produces a bulletin board BB such that $SK_{\mathcal{T}} \in BB$ with negligible probability, for all k_1, k_2, k_3, n_C, n_V and \mathcal{A} . It follows that an execution of the experiment $\text{Exp-JCJ-Cor}(\text{Backdoor}(\Pi), \mathcal{A}, n_C, n_V, k_1, k_2, k_3)$ is an execution of the experiment $\text{Exp-JCJ-Cor}(\Pi, \mathcal{A}, n_C, n_V, k_1, k_2, k_3)$ with non-negligible probability and, hence, $\text{Backdoor}(\Pi)$ satisfies JCJ-correctness.

We show that $\text{Backdoor}(\Pi)$ satisfies JCJ-verifiability by contradiction. Suppose $\text{Succ}(\text{Exp-JCJ-Ver}(\text{Backdoor}(\Pi), \mathcal{A}, n_C, n_V, k_1, k_2, k_3))$ is non-negligible for some k_1, k_3, n_C, n_V , and \mathcal{A} . Hence, there exists an execution of the experiment $\text{Exp-JCJ-Ver}(\text{Backdoor}(\Pi), \mathcal{A}, n_C, n_V, k_1, k_2, k_3)$ which satisfies

$$\text{Verify}(PK_{\mathcal{R}}, PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, P) = 1 \wedge \mathbf{X} \neq \mathbf{X}'$$

with non-negligible probability, where (BB, \mathbf{X}, P) is an election concocted by the adversary and (\mathbf{X}', P') is produced by tallying BB . If there is one and only one \mathbf{Y} such that $(SK_{\mathcal{T}}, \mathbf{Y}) \in BB$, then $\mathbf{X}' = \mathbf{Y}$ by definition of Tally and $\mathbf{X} = \mathbf{Y}$ by definition of Verify , otherwise, the execution of the experiment $\text{Exp-JCJ-Cor}(\text{Backdoor}(\Pi), \mathcal{A}, n_C, n_V, k_1, k_2, k_3)$ is an execution of the experiment $\text{Exp-JCJ-Cor}(\Pi, \mathcal{A}, n_C, n_V, k_1, k_2, k_3)$ and, hence,

$$\begin{aligned} \text{Succ}(\text{Exp-JCJ-Ver}(\text{Backdoor}(\Pi), \mathcal{A}, n_C, n_V, k_1, k_2, k_3)) \\ = \text{Succ}(\text{Exp-JCJ-Ver}(\Pi, \mathcal{A}, n_C, n_V, k_1, k_2, k_3)). \end{aligned}$$

In both cases we derive a contradiction, thereby concluding our proof. \square

C. Proof sketch: Proposition 13

Suppose $\Pi = (\text{Register}, \text{Vote}, \text{Tally}, \text{Verify})$ is an election scheme satisfying JCJ-correctness and JCJ-verifiability. Further suppose $\text{Bias}(\Pi, Z) = (\text{Register}, \text{Vote}, \text{Tally}, \text{Verify}_R)$, for some set of vectors Z . By definition of Verify_R , we have

$$\text{Verify}_R(PK_{\mathcal{R}}, PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, P) = 1$$

implies the existence of a computation

$$v \leftarrow \text{Verify}(PK_{\mathcal{R}}, PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, P)$$

such that $v = 1$ with non-negligible probability, for all $PK_{\mathcal{T}}, BB, n_C, \mathbf{X}$, and P . It follows that

$$\begin{aligned} & \text{Succ}(\text{Exp-JCJ-Cor}(\text{Bias}(\Pi), \mathcal{A}, n_C, n_V, k_1, k_2, k_3)) \\ & \leq \text{Succ}(\text{Exp-JCJ-Cor}(\Pi, \mathcal{A}, n_C, n_V, k_1, k_2, k_3)) \end{aligned}$$

and

$$\begin{aligned} & \text{Succ}(\text{Exp-JCJ-Ver}(\text{Bias}(\Pi), \mathcal{A}, n_C, n_V, k_1, k_2, k_3)) \\ & \leq \text{Succ}(\text{Exp-JCJ-Ver}(\Pi, \mathcal{A}, n_C, n_V, k_1, k_2, k_3)) \end{aligned}$$

for all k_1, k_2, k_3, n_C, n_V , and \mathcal{A} . Hence, $\text{Bias}(\Pi, Z)$ satisfies JCJ-correctness and JCJ-verifiability. \square

APPENDIX N GLOBAL VERIFIABILITY

Küsters et al. [97] propose a definition, called Protocols, to describe any kind of protocol, not just electronic voting protocols. Their definition is independent of any particular computational model, assuming the model provides a notion of processes. These processes must be able to perform internal computation and communicate with each other, and must define a family of probability distributions over runs, indexed by a security parameter.⁶⁶

A. Protocols

We consider the following simplified definition of Protocols.

Definition 38 (Protocol). *A Protocol is a tuple of sets of processes Π_1, \dots, Π_n and processes $\hat{\pi}_1, \dots, \hat{\pi}_n$. Protocols must satisfy the following conditions: each process in $\hat{\pi}_1, \dots, \hat{\pi}_n$ has a special output channel which no process can input on, and $\Pi_i = \{\hat{\pi}_1\}$ or $\Pi_i = \Pi(\hat{\pi}_i)$ for all $1 \leq i \leq n$.*⁶⁷

Processes $\hat{\pi}_1, \dots, \hat{\pi}_n$ capture protocol participants. And sets of processes Π_1, \dots, Π_n capture adversarial behavior. In particular, if $\Pi_i = \{\hat{\pi}_i\}$, then an adversary following the protocol is captured. Otherwise, an adversary controlling the channels in $\hat{\pi}_i$ is captured.^{68,69}

An *instance of Protocol* $(\Pi_1, \dots, \Pi_n, \hat{\pi}_1, \dots, \hat{\pi}_n)$ is the composition of processes π_1, \dots, π_n , where $\pi_i \in \Pi_i$. Process π_i is *honest* in such an instance, if $\hat{\pi}_i = \pi_i$. Each instance of a Protocol defines a set of *runs*. We say an instance of a Protocol produces a run, if the run belongs to that set. A process is honest in a run produced by an instance of a Protocol, if the process is honest in the instance.

a) *Comparison with the original definition:* Definition 38 modifies the original definition [97, §2] as follows. First, we omit agents, since they are only used to refer to a process's owner. Secondly, we omit the finite set of channels used by agents and we omit functions to compute the channels of a particular agent, because these sets can be derived from processes. Thirdly, we restrict Protocols to some processes $\hat{\pi}_1, \dots, \hat{\pi}_n$, whereas the original definition considers sets of processes $\hat{\Pi}_1, \dots, \hat{\Pi}_n$. Fourthly, we require a stronger assumption on the sets of processes: we require $\Pi_i = \{\hat{\pi}_1\}$ or $\Pi_i = \Pi(\hat{\pi}_i)$, whereas the original definition requires $\Pi_i \subseteq \Pi(\hat{\pi}_i)$. Fifthly, we forbid the sets of processes from using special channels. (This restriction does not appear in the original definition, but it is necessary to ensure that global verifiability is satisfiable by interesting protocols.⁷⁰) Finally, we permit channels to be shared between processes. (Thus, we drop the implicit assumption that communication is authenticated. And we permit broadcast channels.⁷¹)

B. Global verifiability

A *goal* of a Protocol is a subset of the sets of runs produced by instances of the Protocol. Processes can accept runs by outputting on their special channels. Global verifiability is intended to ensure that processes only accept runs when the goal has been achieved in those runs. We consider the following simplified definition of global verifiability.⁷²

66. Neither syntax nor semantics are defined for processes, leading to informality that prohibits rigorous analysis [88, §1.4], [92].

67. Let $\text{In}(\pi)$ denote the input channels of process π and $\text{Out}(\pi)$ denote the output channels of process π , excluding π 's special output channel. Moreover, let $\Pi(I, O)$ denote the set of all processes with input channels in I and output channels in O . We abbreviate $\Pi(\text{In}(\pi), \text{Out}(\pi))$ as $\Pi(\pi)$.

68. The adversary model captured by Protocols is unrealistic. In particular, communication between a process in Π_i and a process in Π_j is prohibited, if process $\hat{\pi}_i$ cannot input (respectively output) on a channel that process $\hat{\pi}_j$ can output (respectively input) on. Consequently, the definition of global verifiability cannot detect some attacks. For instance, given a Protocol $P = (\dots, \hat{\pi})$, let $\text{Accept}(P) = (\dots, \hat{\pi}')$ such that process $\hat{\pi}'$ awaits input on a channel that is not used by any other process in P , if an input is received, then the process outputs on $\hat{\pi}'$'s special channel, otherwise, the process executes $\hat{\pi}$. Hence, $\text{Accept}(P)$ accepts all runs that input on the public channel introduced by Accept . Thus, $\text{Accept}(P)$ should not satisfy any definition of verifiability. Yet, the adversary model prohibits input on the channel introduced by Accept , therefore, Protocols P and $\text{Accept}(P)$ are identical from the adversary's perspective. It follows that: given a Protocol $P = (\dots, \hat{\pi})$ and goal γ of P , such that γ is globally verifiable by $\hat{\pi}$, it holds that γ is globally verifiable by $\text{Accept}(P)$. This problem can be overcome by assuming a single, shared broadcast channel between all processes.

69. Analysts must take care to avoid unnecessarily restricting the adversary. For instance, a Protocol could define sets of processes $\{\hat{\pi}_1\}, \dots, \{\hat{\pi}_n\}$, which restricts the adversary to following the protocol.

70. A goal is *not* globally verifiable, if the Protocol produces a run that does not achieve the goal, but is nevertheless accepted. Given that acceptance is captured by outputting on special channels and the original definition permits the adversary to output on such channels, global verifiability is unsatisfiable for interesting protocols. Insisting that $\text{Out}(\pi)$ excludes π 's special output channel suffices to overcome this problem.

71. Broadcast channels are necessary to ensure a realistic adversary model.

72. We omit the definition's notion of Completeness for brevity.

Definition 39 (Global verifiability). *Given a Protocol P , goal γ of P , and process $\hat{\pi}$ of P , we say γ is globally verifiable by $\hat{\pi}$, if for all instances Λ of P parameterized by k , there exists a negligible function μ such that for all security parameters k and (efficient) runs r of Λ that include an output on $\hat{\pi}$'s special channel, we have $r \notin \gamma$, with probability less than or equal to $\mu(k)$.*

Our simplified definition refines the original definition by incorporating our simplified syntax and considering a tighter security bound. Moreover, we require that runs are efficient. (This is necessary to ensure that global verifiability is satisfiable by interesting protocols.)

C. Protocols generalize election schemes

We propose a translation from election schemes to Protocols.⁷³

Definition 40. *Suppose $\Pi = (\text{Setup}, \text{Vote}, \text{Tally}, \text{Verify})$ is an election scheme. Let processes Auditor, Board, Tallier, Voter that depend upon security parameter k , integer n_C , and well-formed choice β , be defined as follows.*

- Tallier. *Computes $(PK_{\mathcal{T}}, SK_{\mathcal{T}}, m_B, m_C) \leftarrow \text{Setup}(k)$ and outputs public key $PK_{\mathcal{T}}$ to all other processes. (We omit explicitly inputting the public key in other processes for brevity.) Inputs a bulletin board BB from process Board, computes $(\mathbf{X}, P) \leftarrow \text{Tally}(PK_{\mathcal{T}}, SK_{\mathcal{T}}, BB, n_C, k)$, and outputs (\mathbf{X}, P) to process Auditor.*
- Voter. *Computes $b \leftarrow \text{Vote}(PK_{\mathcal{T}}, n_C, \beta, k)$ and outputs b to process Board.*
- Board. *Initializes a bulletin board as the empty set—i.e., computes $BB \leftarrow \emptyset$. Inputs ballots from Voter processes and adds ballots to the bulletin board—i.e., computes $BB \leftarrow BB \cup \{b\}$ for every ballot b input. And, concurrently, outputs bulletin boards to processes Auditor, Tallier, and Voter.*
- Auditor. *Inputs tally and proof (\mathbf{X}, P) from process Tallier and a bulletin board BB from process Board, computes $v \leftarrow \text{Verify}(PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, P, k)$, and outputs on its special channel if $v = 1$.*

Let $\text{Voter}_1, \dots, \text{Voter}_{n_V}$ be n_V instances of process Voter. We assume a single, shared broadcast channel.⁷⁴ We define function Alg2Prot such that $\text{Alg2Prot}(\Pi)$ outputs sets of processes $\{\text{Auditor}\}, \Pi(\text{Board}), \Pi(\text{Tallier}), \Pi(\text{Voter}_1), \dots, \Pi(\text{Voter}_{n_V})$ and processes Auditor, Board, Tallier, $\text{Voter}_1, \dots, \text{Voter}_{n_V}$.

We can use function Alg2Prot to derive Protocols representing Helios and Nonce.

APPENDIX O

GOAL γ_{GV} BY KÜSTERS ET AL.

We consider a simplified case of a goal proposed by Küsters et al. [101, §5.2].

Definition 41. *Suppose r is a run of some instance of a Protocol. Let n_h be the number of honest voters in r and*

$\beta_1, \dots, \beta_{n_h}$ be the choices of honest voters in r . Let n_d be the number of dishonest voters in r . We say that we are satisfied with r , if a tally is published in r and that tally contains $n_d + n_h$ choices including $\beta_1, \dots, \beta_{n_h}$.

Given a Protocol, we define γ_{GV} as the following set of runs: for all instances Λ of the Protocol and for each run r produced by Λ , we include r in γ_{GV} , if we are satisfied with r .

Our simplified definition is a special case of the original: set γ_{GV} contains runs in which no choices of honest voters may be excluded from the tally.⁷⁵ Hence, goal γ_{GV} is a slightly more formal presentation of goal γ_l for $l = 0$ (§IX).

A. Unsatisfiable by Helios'16 and Nonce

In Section VII, we claimed that Helios'16 and Nonce do not satisfy global verifiability using goal γ_{GV} . We now prove the validity of our claims.

Proposition 34. *Suppose Π is Helios'16 and $\text{Alg2Prot}(\Pi) = (\{\text{Auditor}\}, \Pi(\text{Board}), \Pi(\text{Tallier}), \Pi(\text{Voter}), \dots)$, we have γ_{GV} is not globally verifiable by Auditor.*

We prove Proposition 34 by demonstrating that a voter may not participate.

Proof. Let $\overline{\text{Voter}}$ be the process that inputs a public key from process Tallier and terminates.⁷⁶ We have $\text{Board} \in \Pi(\text{Board})$, $\text{Tallier} \in \Pi(\text{Tallier})$, and $\overline{\text{Voter}} \in \Pi(\text{Voter})$, hence, Auditor, Board, Tallier, $\overline{\text{Voter}}$ is an instance of $\text{Alg2Prot}(\Pi)$. Let us consider the following run r of this instance:

- 1) Tallier computes $(PK_{\mathcal{T}}, SK_{\mathcal{T}}, m_B, m_C) \leftarrow \text{Setup}(k)$.
- 2) Tallier sends public key $PK_{\mathcal{T}}$ to all other processes.
- 3) Board initializes bulletin board $BB \leftarrow \emptyset$.
- 4) Board sends BB to Auditor and Tallier.
- 5) Tallier computes $(\mathbf{X}, P) \leftarrow \text{Tally}(PK_{\mathcal{T}}, SK_{\mathcal{T}}, BB, n_C, k)$.
- 6) Tallier sends (\mathbf{X}, P) to Auditor.
- 7) Auditor computes $v \leftarrow \text{Verify}(PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, P, k)$.

By definition of Tally, we have \mathbf{X} is a zero-filled vector of length n_C . Moreover, \mathbf{P} is a vector such that $|\mathbf{P}| = 1 \wedge n_C = 1$ or $|\mathbf{P}| = n_C - 1 \wedge n_C > 1$. It follows that $v = 1$, by definition of Verify, hence,

- 8) Auditor outputs on its special channel.

Since this is a run with one (dishonest) voter, goal γ_{GV} teaches us to expect tally \mathbf{X} to contain one choice. Given that \mathbf{X} does not contain any choices, we have γ_{GV} is unsatisfied in r , hence, $r \notin \gamma_{GV}$, but, nonetheless, Auditor outputs on its special channel. Thereby concluding our proof. \square

^{73.} We should not expect a translation from Protocols to election schemes, because Protocols are more general.

^{74.} Thus, we avoid constructing Protocols that consider unrealistic adversaries.

^{75.} We remark that Küsters et al. only define when we are satisfied with run r and do not define γ_{GV} as a set. Nonetheless, we believe our definition captures their intent.

^{76.} Inputting prevents a deadlock that can occur when communication is symmetric.

Proposition 35. *Suppose $\text{Alg2Prot}(\text{Nonce}) = (\{\text{Auditor}\}, \Pi(\text{Board}), \Pi(\text{Tallier}), \Pi(\text{Voter}), \dots)$, we have γ_{GV} is not globally verifiable by Auditor.*

We prove Proposition 35 by demonstrating that a malicious bulletin board can inject ballots.

Proof. Let $\overline{\text{Board}}$ be the process that inputs a public key from process Tallier, computes $b \leftarrow \text{Vote}(PK_{\mathcal{T}}, n_C, \beta, k)$ for some well-formed choice β , initializes bulletin board $BB \leftarrow \{b\}$, and outputs BB to processes Auditor and Tallier. We have $\overline{\text{Board}} \in \Pi(\text{Board})$, hence, Auditor, $\overline{\text{Board}}$, Tallier, Voter is an instance of $\text{Alg2Prot}(\text{Nonce})$, when $n_V = 1$. Let us consider the following run r of this instance:

- 1) Tallier computes $(PK_{\mathcal{T}}, SK_{\mathcal{T}}, m_B, m_C) \leftarrow \text{Setup}(k)$.
- 2) Tallier sends public key $PK_{\mathcal{T}}$ to all other processes.
- 3) $\overline{\text{Board}}$ computes $b \leftarrow \text{Vote}(PK_{\mathcal{T}}, n_C, \beta, k)$ for some well-formed choice β and initializes bulletin board $BB \leftarrow \{b\}$.
- 4) $\overline{\text{Board}}$ sends BB to Auditor and Tallier.
- 5) Tallier computes $(\mathbf{X}, P) \leftarrow \text{Tally}(PK_{\mathcal{T}}, SK_{\mathcal{T}}, BB, n_C, k)$.
- 6) Tallier sends (\mathbf{X}, P) to Auditor.
- 7) Auditor computes $v \leftarrow \text{Verify}(PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, P, k)$.

By definition of Tally and Verify, we have $v = 1$. It follows that:

- 8) Auditor outputs on its special channel.

Since this is a run with one honest voter, goal γ_{GV} teaches us to expect tally \mathbf{X} to contain the honest voter's vote β' . It follows that γ_{GV} is unsatisfied in r if $\beta \neq \beta'$, which occurs with non-negligible probability for $n_C > 1$, thereby concluding our proof. \square

B. Admitting attacks

In Section VIII-C, we claimed that global verifiability instantiated with goal γ_{GV} admits revelation attacks. We now prove the validity of our claim.

Definition 42 (Coin leakage). *An election scheme $\Pi = (\text{Setup}, \text{Vote}, \text{Tally}, \text{Verify})$ does not leak coins used to construct ballots, if for all PPT adversaries \mathcal{A} , there exists a negligible function μ , such that for all security parameters k , we have $\Pr[(PK_{\mathcal{T}}, n_C, \beta) \leftarrow \mathcal{A}(k); b \leftarrow \text{Vote}(PK_{\mathcal{T}}, n_C, \beta, k); r \leftarrow \mathcal{A}(b) : 1 \leq \beta \leq n_C \wedge b = \text{Vote}(PK_{\mathcal{T}}, n_C, \beta, k; r)] \leq \mu(k)$.*

Intuitively, the definition captures the idea that ballots do not leak coins. Of course, coins may be leaked indirectly. For example, by the software, hardware, or voter that constructed the ballot.

Proposition 36. *Suppose $\Pi = (\dots, \text{Verify})$ is an election scheme that does not leak coins used to construct ballots, $\text{Alg2Prot}(\Pi) = (\{A\}, \dots)$, and $\text{Alg2Prot}(\text{Replace}(\Pi)) = (\{\text{Auditor}\}, \Pi(\text{Board}), \Pi(\text{Tallier}), \Pi(\text{Voter}_1), \dots, \Pi(\text{Voter}_{n_V}), \dots)$. If γ_{GV} is globally verifiable by \mathcal{A} , then γ_{GV} is globally verifiable by Auditor.*

Proof. Suppose γ_{GV} is not globally verifiable by Auditor, hence, there exists an instance of protocol $\text{Alg2Prot}(\text{Replace}(\Pi))$ such that for all negligible functions μ there exists a security parameter k and run $r \notin \gamma_{GV}$ of the instance that includes an output on the special channel belonging to process Auditor, with probability greater than $\mu(k)$. By definition of goal γ_{GV} , either no tally is published in run r or the run publishes a tally that does not contain $n_d + n_h$ choices including $\beta_1, \dots, \beta_{n_h}$, where n_h is the number of honest voters, $\beta_1, \dots, \beta_{n_h}$ are the choices of honest voters, and n_d is the number of dishonest voters. By definition of function Alg2Prot , process Auditor only outputs on its special channel if a tally is published, hence, we consider only the latter case. Moreover, the process only outputs on its special channel if $\text{Verify}_R(PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, P, k) = 1$, where Verify_R is the algorithm introduced by function Replace , $PK_{\mathcal{T}}$, \mathbf{X} and P are outputs of a process in $\Pi(\text{Tallier})$, BB is an output of a process in $\Pi(\text{Board})$, and n_C is the number of candidates.

It follows from the definition of function Replace that protocols $\text{Alg2Prot}(\Pi)$ and $\text{Alg2Prot}(\text{Replace}(\Pi))$ are equivalent (which would permit an immediate conclusion), unless $BB = \{b_1, \dots, b_\ell, (\alpha_1, \alpha'_1, r_1), \dots, (\alpha_k, \alpha'_k, r_k)\}$ such that $\bigwedge_{1 \leq i \leq k} b_i = \text{Vote}(PK_{\mathcal{T}}, n_C, \alpha_i, k; r_i) \wedge 1 \leq \alpha_i, \alpha'_i \leq n_C$. Moreover, we have $\text{Verify}(PK_{\mathcal{T}}, BB, n_C, \mathbf{X}^*, P, k) = 1$, where tally \mathbf{X}^* is derived from \mathbf{X} by replacing choices $\alpha'_1, \dots, \alpha'_k$ with $\alpha_1, \dots, \alpha_k$. Since γ_{GV} is globally verifiable by \mathcal{A} , tally \mathbf{X}^* contains $n_d + n_h$ choices including the choices of honest voters, namely, $\beta_1, \dots, \beta_{n_h}$. By comparison, \mathbf{X} does not. Since \mathbf{X} and \mathbf{X}^* contain the same number of choices, there exists an honest choice $\beta \in \{\beta_1, \dots, \beta_{n_h}\}$ which is replaced by a distinct choice $\alpha \in \{\alpha'_1, \dots, \alpha'_k\}$. Suppose that honest choice belongs to a process $\text{Voter} \in \bigcup_{1 \leq i \leq n_V} \Pi(\text{Voter}_i)$. By definition of process Voter , the choice is encapsulated inside a ballot $b = \text{Vote}(PK_{\mathcal{T}}, n_C, \beta, k; r)$, where coins r are chosen uniformly at random. Since Π does not leak coins used to construct ballots, it follows that coins r cannot appear on bulletin board BB , thereby deriving a contradiction and concluding our proof. \square

Proposition 37. *Suppose Π is an election scheme that does not leak coins used to construct ballots, $\text{Alg2Prot}(\Pi) = (\{A\}, \dots)$, and $\text{Alg2Prot}(\text{Drop}(\Pi)) = (\{B\}, \dots)$. If γ_{GV} is globally verifiable by \mathcal{A} , then γ_{GV} is globally verifiable by \mathcal{B} .*

A proof of Proposition 37 can be constructed on the basis that algorithm Vote does not leak coins. That idea has already been demonstrated in our proof of Proposition 36, so we omit a formal proof.

APPENDIX P

GOAL δ_{GV} AND RELATION WITH ELECTION VERIFIABILITY

The definition below expresses the goal introduced in Section VII: the goal that is satisfied in a run if ballots b_1, \dots, b_n for choices β_1, \dots, β_n appear in the run, such that b_1, \dots, b_n are included on the bulletin board and no further ballots are included, and the run produces a tally for choices β_1, \dots, β_n .

Definition 43. Suppose r is a run of some instance of a Protocol. Further suppose BB is the bulletin board in r and b_1, \dots, b_n are ballots for well-formed choices β_1, \dots, β_n in r , such that $b_1, \dots, b_n \in BB \setminus \{\perp\}$ and no further ballots appear in BB . We say that we are satisfied with r , if a tally is published in r and that tally is for choices β_1, \dots, β_n .

Given a Protocol, we define δ_{GV} as the following set of runs: for all instances Λ of the Protocol and for each run r produced by Λ , we include r in δ_{GV} , if we are satisfied with r .

We show election verifiability implies global verifiability using goal δ_{GV} (Proposition 38). It follows that Helios'16, respectively Nonce, satisfy global verifiability using goal δ_{GV} by Theorem 5, respectively Proposition 1. We also show that global verifiability implies universal verifiability (Proposition 39), but not individual verifiability, with that goal.

Proposition 38. Suppose Π is an election scheme and $\text{Alg2Prot}(\Pi) = (\{\text{Auditor}\}, \dots)$. If Π satisfies Ver-Ext, then δ_{GV} is globally verifiable by Auditor.

Proof. Suppose δ_{GV} is not globally verifiable by Auditor. Hence, there exists an instance Λ of $\text{Alg2Prot}(\Pi)$ parameterized by k , such that for all negligible functions μ , there exists a security parameter k and an (efficient) run $r \notin \gamma$ of Λ that includes an output on Auditor's special channel, with probability greater than $\mu(k)$. The instance Λ of $\text{Alg2Prot}(\Pi)$ is a composition of processes that includes Auditor. Since process Auditor outputs on its special channel, run r constructs public key $PK_{\mathcal{T}}$, bulletin board BB , tally \mathbf{X} , and proof P , and inputs these values to process Auditor such that $\text{Verify}(PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, P, k) = 1$, where n_C is an integer. Let b_1, \dots, b_n be ballots for well-formed choices β_1, \dots, β_n in r such that $b_1, \dots, b_n \in BB$ and no further ballots appear in BB . (We implicitly assume that ballots are outputs of algorithm Vote.) We proceed by distinguishing two cases.

- Case I: Π satisfies individual verifiability. In this case, ballots b_1, \dots, b_n are pairwise distinct. It follows that $\text{correct-tally}(PK_{\mathcal{T}}, BB, M, n_C, k)$ is a tally for choices β_1, \dots, β_n . Yet, since δ_{GV} is not globally verifiable by Auditor, tally \mathbf{X} is not for choices β_1, \dots, β_n . Thus, the adversary that constructs $PK_{\mathcal{T}}$, BB , \mathbf{X} and P in the same way as they are constructed in r , and outputs $(PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, P)$, wins the universal verifiability game.
- Case II: Π satisfies universal verifiability. In this case, $\mathbf{X} = \text{correct-tally}(PK_{\mathcal{T}}, BB, M, n_C, k)$. Yet, since δ_{GV} is not globally verifiable by Auditor, tally \mathbf{X} is not for choices β_1, \dots, β_n . Hence, there exists distinct integers i and j such that $b_i = b_j \wedge b_i \neq \perp \wedge b_j \neq \perp$. Thus, the adversary that constructs $PK_{\mathcal{T}}$ in the same way as it is constructed in r and outputs $(PK_{\mathcal{T}}, n_C, \beta_i, \beta_j)$, wins the individual verifiability game. \square

Proposition 39. Suppose Π is an election scheme and $\text{Alg2Prot}(\Pi) = (\{\text{Auditor}\}, \Pi(\text{Board}), \Pi(\text{Tallier}), \Pi(\text{Voter}_1),$

$\dots, \Pi(\text{Voter}_{n_V}), \dots)$. If δ_{GV} is globally verifiable by Auditor, then Π satisfies Exp-UV-Ext.

Proof. Suppose Π does not satisfy universal verifiability. Hence, there exists a PPT adversary \mathcal{A} , such that for all negligible functions μ , there exists a security parameter k and $\text{Succ}(\text{Exp-UV-Ext}(\Pi, \mathcal{A}, k)) > \mu(k)$.

Let $\overline{\text{Tallier}}$ be the process that computes $(PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, P) \leftarrow \mathcal{A}(k)$, outputs $PK_{\mathcal{T}}$ to all other processes and outputs BB and (\mathbf{X}, P) to process Auditor (given that the channel is shared, process Auditor will input BB as if it were output by Board). We have $\overline{\text{Tallier}} \in \Pi(\text{Tallier})$, hence, Auditor, Board, $\overline{\text{Tallier}}$, $\text{Voter}_1, \dots, \text{Voter}_{n_V}$ is an instance of $\text{Alg2Prot}(\Pi)$. Let us consider the following run r of this instance. Suppose $\overline{\text{Tallier}}$ computes $(PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, P) \leftarrow \mathcal{A}(k)$, sends $PK_{\mathcal{T}}$ to all other processes, and sends BB and (\mathbf{X}, P) to process Auditor. Further suppose Auditor computes $v \leftarrow \text{Verify}(PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, P, k)$. Since \mathcal{A} is a winning adversary, we have $\text{Verify}(PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, P, k) = 1$ and $\mathbf{X} \neq \text{correct-tally}(PK_{\mathcal{T}}, BB, n_C, k)$. Hence, Auditor outputs on its special channel. Suppose \mathcal{A} produces ballots $b_1, \dots, b_n \in BB$ for well-formed choices β_1, \dots, β_n . Without loss of generality, we can assume b_1, \dots, b_n are pairwise distinct, since there exists an equivalent adversary that can simulate \mathcal{A} 's output if this assumption does not hold. It follows that $\text{correct-tally}(PK_{\mathcal{T}}, BB, n_C, k)$ is the tally of choices β_1, \dots, β_n . Yet, tally \mathbf{X} is not for the choices in $\text{correct-tally}(PK_{\mathcal{T}}, BB, n_C, k)$. Thereby concluding our proof. \square

To show global verifiability using goal δ_{GV} does not imply individual verifiability, we adapt our toy scheme from nonces (§II-C) such that nonces are chosen from a space parametrised by the public key, rather than the security parameter, and verification fails when the public key is not equal to the security parameter. It follows that the two schemes are equivalent when the public key is the security parameter. Yet, there exists the possibility to cause collisions using maliciously generated public keys.

Definition 44. Election scheme Nonce' is defined as follows:

- $\text{Setup}(k)$ outputs $(k, k, p_1(k), p_2(k))$, where p_1 and p_2 may be any polynomial functions.
- $\text{Vote}(k, n_C, \beta, k')$ selects a nonce r uniformly at random from \mathbb{Z}_{2^k} and outputs (r, β) .
- $\text{Tally}(k, BB, n_C, k')$ computes a vector \mathbf{X} of length n_C , such that \mathbf{X} is a tally of the votes on BB for which the nonce is in \mathbb{Z}_{2^k} , and outputs (\mathbf{X}, \perp) .
- $\text{Verify}(k, BB, n_C, \mathbf{X}, P, k')$ outputs 1 if $(\mathbf{X}, P) = \text{Tally}(k, BB, n_C, k') \wedge k = k'$ and 0 otherwise.

Collisions resulting from maliciously generated public keys cannot be detected by global verifiability using goal δ_{GV} , because global verifiability only requires the properties of goal δ_{GV} to hold on runs in which auditing succeeds. Thus, Nonce' satisfies global verifiability using goal δ_{GV} , but not individual verifiability. We prove a more general result: for any goal

such that Nonce satisfies global verifiability, we have Nonce' satisfies global verifiability too.

Proposition 40. *Suppose $\text{Alg2Prot}(\text{Nonce}) = (\{\text{Auditor}\}, \dots)$, $\text{Alg2Prot}(\text{Nonce}') = (\{\text{Auditor}'\}, \dots)$, and γ is a goal. If γ is globally verifiable by Auditor, then γ is globally verifiable by Auditor'.*

Proof sketch. By definition of global verifiability, we need only consider runs produced by instances of $\text{Alg2Prot}(\text{Nonce}')$ that include an output on the special channel of process Auditor'. In these runs, the public key input by Auditor' is equal to the security parameter. And $\text{Alg2Prot}(\text{Nonce}')$ is equivalent to $\text{Alg2Prot}(\text{Nonce})$ on such runs, concluding our proof. \square

Corollary 41. *Given $\text{Alg2Prot}(\text{Nonce}') = (\{\text{Auditor}'\}, \dots)$, we have δ_{GV} is globally verifiable by Auditor'.*

Proof sketch. Given that Nonce satisfies Ver-Ext (Proposition 1), we have δ_{GV} is globally verifiable by Auditor (Proposition 38). Hence, δ_{GV} is globally verifiable by Auditor' too (Proposition 40). \square

Proposition 42. *Nonce' does not satisfy Exp-IV-Ext.*

Proof sketch. An adversary that outputs $(1, 1, 1, 1)$ can cause a collision with non-negligible probability. \square

REFERENCES

- [1] Ben Adida. *Advances in Cryptographic Voting Systems*. PhD thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, 2006.
- [2] Ben Adida. Helios: Web-based Open-Audit Voting. In *USENIX Security'08: 17th USENIX Security Symposium*, pages 335–348. USENIX Association, 2008.
- [3] Ben Adida. Helios deployed at Princeton. <http://heliosvoting.wordpress.com/2009/10/13/helios-deployed-at-princeton/> (accessed 7 May 2014), 2009.
- [4] Ben Adida. Helios v4 Verification Specs. Helios documentation, <http://documentation.heliosvoting.org/verification-specs/helios-v4> (accessed 2 May 2014), 2014. A snapshot of the specification on 8 Oct 2013 is available from <https://web.archive.org/web/20131018033747/http://documentation.heliosvoting.org/verification-specs/helios-v4>.
- [5] Ben Adida, Olivier de Marneffe, Olivier Pereira, and Jean-Jacques Quisquater. Electing a University President Using Open-Audit Voting: Analysis of Real-World Use of Helios. In *EVT/WOTE'09: Electronic Voting Technology Workshop/Workshop on Trustworthy Elections*. USENIX Association, 2009.
- [6] Ben Adida and C. Andrew Neff. Ballot casting assurance. In *EVT'06: Electronic Voting Technology Workshop*, Berkeley, CA, USA, 2006. USENIX Association.
- [7] Jee Hea An, Yevgeniy Dodis, and Tal Rabin. On the security of joint signature and encryption. In *EUROCRYPT'02: 21st International Conference on the Theory and Applications of Cryptographic Techniques*, volume 2332 of *LNCS*, pages 83–107. Springer, 2002.
- [8] Myrto Arapinis, Sergiu Bursuc, and Mark Ryan. Reduction of Equational Theories for Verification of Trace Equivalence: Re-encryption, Associativity and Commutativity. In *POST'12: First Conference on Principles of Security and Trust*, volume 7215 of *LNCS*, pages 169–188. Springer, 2012.
- [9] Michael Backes, Cătălin Hrițcu, and Matteo Maffei. Automated Verification of Remote Electronic Voting Protocols in the Applied Pi-calculus. In *CSF'08: 21st Computer Security Foundations Symposium*, pages 195–209. IEEE Computer Society, 2008.
- [10] Mihir Bellare, Anand Desai, David Pointcheval, and Phillip Rogaway. Relations Among Notions of Security for Public-Key Encryption Schemes. In *CRYPTO'98: 18th International Cryptology Conference*, volume 1462 of *LNCS*, pages 26–45. Springer, 1998.
- [11] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *CCS '93: 1st ACM Conference on Computer and Communications Security*, pages 62–73, New York, NY, USA, 1993. ACM.
- [12] Mihir Bellare and Amit Sahai. Non-malleable Encryption: Equivalence between Two Notions, and an Indistinguishability-Based Characterization. In *CRYPTO'99: 19th International Cryptology Conference*, volume 1666 of *LNCS*, pages 519–536. Springer, 1999.
- [13] Mihir Bellare and Amit Sahai. Non-malleable encryption: Equivalence between two notions, and an indistinguishability-based characterization. Cryptology ePrint Archive, Report 2006/228, 2006.
- [14] Josh Benaloh. *Verifiable Secret-Ballot Elections*. PhD thesis, Department of Computer Science, Yale University, 1996.
- [15] Josh Benaloh. Simple Verifiable Elections. In *EVT'06: Electronic Voting Technology Workshop*. USENIX Association, 2006.
- [16] Josh Benaloh. Ballot Casting Assurance via Voter-Initiated Poll Station Auditing. In *EVT'07: Electronic Voting Technology Workshop*. USENIX Association, 2007.
- [17] Josh Benaloh and Dwight Tuinstra. Receipt-free secret-ballot elections (extended abstract). In *STOC '94: Twenty-sixth Annual ACM Symposium on Theory of Computing*, pages 544–553, New York, NY, USA, 1994. ACM Press.
- [18] Josh Benaloh, Serge Vaudenay, and Jean-Jacques Quisquater. Final Report of IACR Electronic Voting Committee. International Association for Cryptologic Research. http://www.iacr.org/elections/eVoting/finalReportHelios_2010-09-27.html (accessed 7 May 2014), Sept 2010.
- [19] Josh Benaloh and Moti Yung. Distributing the Power of a Government to Enhance the Privacy of Voters. In *PODC'86: 5th Principles of Distributed Computing Symposium*, pages 52–62. ACM Press, 1986.
- [20] David Bernhard. *Zero-Knowledge Proofs in Theory and Practice*. PhD thesis, Department of Computer Science, University of Bristol, 2014.
- [21] David Bernhard, Véronique Cortier, David Galindo, Olivier Pereira, and Bogdan Warinschi. A comprehensive analysis of game-based ballot privacy definitions. Cryptology ePrint Archive, Report 2015/255 (version 20150319:100626), 2015.
- [22] David Bernhard, Véronique Cortier, David Galindo, Olivier Pereira, and Bogdan Warinschi. SoK: A comprehensive analysis of game-based ballot privacy definitions. In *S&P'15: 36th Security and Privacy Symposium*, pages 499–516. IEEE Computer Society, 2015.
- [23] David Bernhard, Olivier Pereira, and Bogdan Warinschi. How Not to Prove Yourself: Pitfalls of the Fiat-Shamir Heuristic and Applications to Helios. In *ASIACRYPT'12: 18th International Conference on the Theory and Application of Cryptology and Information Security*, volume 7658 of *LNCS*, pages 626–643. Springer, 2012.
- [24] David Bernhard, Olivier Pereira, and Bogdan Warinschi. On Necessary and Sufficient Conditions for Private Ballot Submission. Cryptology ePrint Archive, Report 2012/236 (version 20120430:154117b), 2012.
- [25] David Bernhard and Bogdan Warinschi. Cryptographic Voting — A Gentle Introduction. In *Foundations of Security Analysis and Design VII*, volume 8604 of *LNCS*, pages 167–211. Springer, 2014.
- [26] Bruno Blanchet, Ben Smyth, and Vincent Cheval. *ProVerif 1.91: Automatic Cryptographic Protocol Verifier, User Manual and Tutorial*, 2015. Originally appeared as Bruno Blanchet and Ben Smyth (2011) ProVerif 1.85: Automatic Cryptographic Protocol Verifier, User Manual and Tutorial.

- [27] Dan Boneh, Emily Shen, and Brent Waters. Strongly unforgeable signatures based on computational diffie-hellman. In *PKC'06: 9th International Workshop on Practice and Theory in Public Key Cryptography*, pages 229–240. Springer, 2006.
- [28] Jonathan Bootle, Andrea Cerulli, Pyrros Chaidos, and Jens Groth. Efficient Zero-Knowledge Proof Systems. In *Foundations of Security Analysis and Design VIII*, volume 9808 of *LNCS*, pages 1–31. Springer, 2016.
- [29] Debra Bowen. Secretary of State Debra Bowen Moves to Strengthen Voter Confidence in Election Security Following Top-to-Bottom Review of Voting Systems. California Secretary of State, press release DB07:042 <http://www.sos.ca.gov/voting-systems/oversight/ttbr/db07-042-ttbr-system-decisions-release.pdf> (accessed 7 May 2014), August 2007.
- [30] Bundesverfassungsgericht (Germany's Federal Constitutional Court). *Use of voting computers in 2005 Bundestag election unconstitutional*, March 2009. Press release 19/2009 <https://www.bundesverfassungsgericht.de/SharedDocs/Pressemitteilungen/EN/2009/bvg09-019.html> (accessed 17 Jan 2018).
- [31] Ran Canetti. Universally composable security: a new paradigm for cryptographic protocols. In *FOCS'01: 42nd IEEE Symposium on Foundations of Computer Science*, pages 136–145, October 2001.
- [32] Nicholas Chang-Fong and Aleksander Essex. The Cloudier Side of Cryptographic End-to-end Verifiable Voting: A Security Analysis of Helios. In *ACSAC'16: 32nd Annual Conference on Computer Security Applications*, pages 324–335. ACM Press, 2016.
- [33] Melissa Chase, Markulf Kohlweiss, Anna Lysyanskaya, and Sarah Meiklejohn. Malleable signatures: New definitions and delegatable anonymous credentials. In *CSF'14: 27th Computer Security Foundations Symposium*. IEEE Computer Society, 2014. To appear.
- [34] David Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2):84–88, 1981.
- [35] David Chaum. Secret-ballot receipts: True voter-verifiable elections. *IEEE Security and Privacy*, 2(1):38–47, 2004.
- [36] David Chaum, Richard Carback, Jeremy Clark, Aleksander Essex, Stefan Popoveniuc, Ronald L. Rivest, Peter Y. A. Ryan, Emily Shen, and Alan T. Sherman. Scantegrity II: End-to-end verifiability for optical scan election systems using invisible ink confirmation codes. In *EVT'08: Electronic Voting Technology Workshop*, pages 14:1–14:13. USENIX Association, 2008.
- [37] David Chaum, Jan-Hendrik Evertse, Jeroen van de Graaf, and René Peralta. Demonstrating Possession of a Discrete Logarithm Without Revealing It. In *CRYPTO'86: 6th International Cryptology Conference*, volume 263 of *LNCS*, pages 200–212. Springer, 1987.
- [38] David Chaum and Torben P. Pedersen. Wallet Databases with Observers. In *CRYPTO'92: 12th International Cryptology Conference*, volume 740 of *LNCS*, pages 89–105. Springer, 1993.
- [39] David Chaum, Peter Y. A. Ryan, and Steve Schneider. A Practical Voter-Verifiable Election Scheme. In *ESORICS'05: 10th European Symposium On Research In Computer Security*, volume 3679 of *LNCS*, pages 118–139. Springer, 2005.
- [40] Benoît Chevallier-Mames, Pierre-Alain Fouque, David Pointcheval, Julien Stern, and Jacques Traoré. On Some Incompatible Properties of Voting Schemes. In *WOTE'06: Workshop on Trustworthy Elections*, 2006.
- [41] Benoît Chevallier-Mames, Pierre-Alain Fouque, David Pointcheval, Julien Stern, and Jacques Traoré. On Some Incompatible Properties of Voting Schemes. In David Chaum, Markus Jakobsson, Ronald L. Rivest, and Peter Y. A. Ryan, editors, *Towards Trustworthy Elections: New Directions in Electronic Voting*, volume 6000 of *LNCS*, pages 191–199. Springer, 2010.
- [42] Michael Clarkson, Brian Hay, Meador Inge, abhi shelat, David Wagner, and Alec Yasinsac. Software review and security analysis of Scytl remote voting software. Report commissioned by Florida Division of Elections. Available from <http://election.dos.state.fl.us/voting-systems/pdf/FinalReportSept19.pdf>. Filed September 19, 2008.
- [43] Michael R. Clarkson, Stephen Chong, and Andrew C. Myers. Civitas: Toward a Secure Voting System. In *S&P'08: 29th Security and Privacy Symposium*, pages 354–368. IEEE Computer Society, 2008.
- [44] Josh Daniel Cohen and Michael J. Fischer. A Robust and Verifiable Cryptographically Secure Election Scheme. In *FOCS'85: 26th IEEE Symposium on Foundations of Computer Science*, pages 372–382. IEEE Computer Society, 1985.
- [45] Véronique Cortier, Fabienne Eigner, Steve Kremer, Matteo Maffei, and Cyrille Wiedling. Type-Based Verification of Electronic Voting Protocols. In *POST'15: 4th Conference on Principles of Security and Trust*, LNCS, pages 303–323. Springer, 2015.
- [46] Véronique Cortier, David Galindo, Stéphane Glondu, and Malika Izabachène. A generic construction for voting correctness at minimum cost - Application to Helios. *Cryptology ePrint Archive*, Report 2013/177 (version 20130521:145727), 2013.
- [47] Véronique Cortier, David Galindo, Stéphane Glondu, and Malika Izabachène. Election Verifiability for Helios under Weaker Trust Assumptions. In *ESORICS'14: 19th European Symposium on Research in Computer Security*, volume 8713 of *LNCS*, pages 327–344. Springer, 2014.
- [48] Véronique Cortier, David Galindo, Stéphane Glondu, and Malika Izabachène. Election Verifiability for Helios under Weaker Trust Assumptions. Technical Report RR-8555, INRIA, 2014.
- [49] Veronique Cortier, David Galindo, Ralf Küsters, Johannes Müller, and Tomasz Truderung. SoK: Verifiability Notions for E-Voting Protocols. In *S&P'16: 37th IEEE Symposium on Security and Privacy*, pages 779–798. IEEE Computer Society, 2016.
- [50] Veronique Cortier, David Galindo, Ralf Küsters, Johannes Müller, and Tomasz Truderung. Verifiability Notions for E-Voting Protocols. *Cryptology ePrint Archive*, Report 2016/287 (version 20160317:161048), 2016.
- [51] Véronique Cortier and Joseph Lallemand. Voting: You Can't Have Privacy without Individual Verifiability. Technical Report hal-01858034, INRIA, 2018.
- [52] Véronique Cortier and Ben Smyth. Attacking and fixing Helios: An analysis of ballot secrecy. In *CSF'11: 24th Computer Security Foundations Symposium*, pages 297–311. IEEE Computer Society, 2011.
- [53] Véronique Cortier and Ben Smyth. Attacking and fixing Helios: An analysis of ballot secrecy. *Journal of Computer Security*, 21(1):89–148, 2013.
- [54] Ronald Cramer, Ivan Damgård, and Berry Schoenmakers. Proofs of Partial Knowledge and Simplified Design of Witness Hiding Protocols. In *CRYPTO'94: 14th International Cryptology Conference*, volume 839 of *LNCS*, pages 174–187. Springer, 1994.
- [55] Ronald Cramer, Matthew K. Franklin, Berry Schoenmakers, and Moti Yung. Multi-Authority Secret-Ballot Elections with Linear Work. In *EUROCRYPT'96: 15th International Conference on the Theory and Applications of Cryptographic Techniques*, volume 1070 of *LNCS*, pages 72–83. Springer, 1996.
- [56] Ronald Cramer, Rosario Gennaro, and Berry Schoenmakers. A Secure and Optimally Efficient Multi-Authority Election Scheme. In *EUROCRYPT'97: 16th International Conference on the Theory and Applications of Cryptographic Techniques*, volume 1233 of *LNCS*, pages 103–118. Springer, 1997.
- [57] Ronald Cramer and Victor Shoup. A Practical Public Key Cryptosystem Provably Secure Against Adaptive Chosen Ciphertext Attack. In *CRYPTO'98: 18th International Cryptology Conference*, volume 1462 of *LNCS*, pages 13–25. Springer, 1998.
- [58] Ivan Damgård. On Σ -protocols, 2010. Available from <http://www.daimi.au.dk/~ivan/Sigma.pdf>.
- [59] Ivan Damgård, Mads Jurik, and Jesper Buus Nielsen. A Generalization of Paillier's Public-Key System with Applications to Electronic Voting. *International Journal of Information Security*, 9(6):371–385, 2010.

- [60] Stéphanie Delaune, Steve Kremer, and Mark D. Ryan. Verifying privacy-type properties of electronic voting protocols. *Journal of Computer Security*, 17(4):435–487, July 2009.
- [61] Stéphanie Delaune, Steve Kremer, Mark D. Ryan, and Graham Steel. Formal analysis of protocols based on TPM state registers. In *CSF'11: 24th Computer Security Foundations Symposium*, pages 66–80. IEEE Computer Society, 2011.
- [62] Jannik Dreier, Rosario Giustolisi, Ali Kassem, Pascal Lafourcade, and Gabriele Lenzini. On the Verifiability of (Electronic) Exams. Technical Report TR-2014-2, Verimag, 2014.
- [63] Jannik Dreier, Hugo Jonker, and Pascal Lafourcade. Defining verifiability in e-auction protocols. In *ASIACCS'13: 8th ACM Symposium on Information, Computer and Communications Security*, pages 547–552. ACM Press, 2013.
- [64] Jannik Dreier, Hugo Jonker, and Pascal Lafourcade. A framework for analyzing verifiability in traditional and electronic exams. In *ISPEC'15: 11th International Conference on Information Security Practice and Experience*, volume 9065 of *LNCS*, pages 514–529. Springer, 2015.
- [65] Taher ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE Transactions on Information Theory*, 31(4):469–472, 1985.
- [66] Amos Fiat and Adi Shamir. How To Prove Yourself: Practical Solutions to Identification and Signature Problems. In *CRYPTO'86: 6th International Cryptology Conference*, volume 263 of *LNCS*, pages 186–194. Springer, 1987.
- [67] Atsushi Fujioka, Tatsuaki Okamoto, and Kazuo Ohta. A Practical Secret Voting Scheme for Large Scale Elections. In *AUSCRYPT'92: Workshop on the Theory and Application of Cryptographic Techniques*, volume 718 of *LNCS*, pages 244–251. Springer, 1992.
- [68] Jun Furukawa and Kazuo Sako. An efficient scheme for proving a shuffle. In *CRYPTO'01: 21st International Cryptology Conference*, volume 2139 of *LNCS*. Springer, 2001.
- [69] Rosario Gennaro, Stanisław Jarecki, Hugo Krawczyk, and Tal Rabin. Secure distributed key generation for discrete-log based cryptosystems. In *EUROCRYPT'99: 18th International Conference on the Theory and Applications of Cryptographic Techniques*, pages 295–310. Springer, 1999.
- [70] Oded Goldreich. *Foundations of Cryptography: Basic Tools*. Cambridge University Press, Cambridge, UK, 2001.
- [71] Shafi Goldwasser and Silvio Micali. Probabilistic Encryption & How to Play Mental Poker Keeping Secret All Partial Information. In *STOC'82: 14th Annual ACM Symposium on Theory of Computing*, pages 365–377. ACM Press, 1982.
- [72] Shafi Goldwasser and Silvio Micali. Probabilistic Encryption. *Journal of Computer and System Sciences*, 28(2):270–299, 1984.
- [73] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18(1):186–208, February 1989.
- [74] Jens Groth. Evaluating security of voting schemes in the universal composability framework. In *Applied Cryptography and Network Security*, volume 3089 of *LNCS*, pages 46–60. Springer, 2004.
- [75] Jens Groth. Simulation-Sound NIZK Proofs for a Practical Language and Constant Size Group Signatures. In *ASIACRYPT'02: 12th International Conference on the Theory and Application of Cryptology and Information Security*, volume 4284 of *LNCS*, pages 444–459. Springer, 2006.
- [76] Stuart Haber, Josh Benaloh, and Shai Halevi. The Helios e-Voting Demo for the IACR. International Association for Cryptologic Research. <http://www.iacr.org/elections/eVoting/heliosDemo.pdf> (accessed 7 May 2014), May 2010.
- [77] Carmit Hazay and Yehuda Lindell. Sigma protocols and efficient zero-knowledge. In *Efficient Secure Two-Party Protocols*, Information Security and Cryptography, chapter 6, pages 147–175. Springer, 2010.
- [78] Martin Hirt. Receipt-Free K -out-of- L Voting Based on ElGamal Encryption. In David Chaum, Markus Jakobsson, Ronald L. Rivest, and Peter Y. A. Ryan, editors, *Towards Trustworthy Elections: New Directions in Electronic Voting*, volume 6000 of *LNCS*, pages 64–82. Springer, 2010.
- [79] Martin Hirt and Kazuo Sako. Efficient Receipt-Free Voting Based on Homomorphic Encryption. In *EUROCRYPT'06: 25th International Conference on the Theory and Applications of Cryptographic Techniques*, volume 1807 of *LNCS*, pages 539–556. Springer, 2006.
- [80] Benjamin Hosp and Poorvi L. Vora. An information-theoretic model of voting systems. In *WOTE'06: Workshop on Trustworthy Elections*, 2006.
- [81] Benjamin Hosp and Poorvi L. Vora. An information-theoretic model of voting systems. *Mathematical and Computer Modelling*, 48(9-10):1628–1645, 2008.
- [82] Markus Jakobsson and Ari Juels. Mix and Match: Secure Function Evaluation via Ciphertexts. In *ASIACRYPT'00: 6th International Conference on the Theory and Application of Cryptology and Information Security*, volume 1976 of *LNCS*, pages 162–177. Springer, 2000.
- [83] Markus Jakobsson, Ari Juels, and Ronald L. Rivest. Making Mix Nets Robust for Electronic Voting by Randomized Partial Checking. In *11th USENIX Security Symposium*, pages 339–353, 2002.
- [84] Douglas W. Jones and Barbara Simons. *Broken Ballots: Will Your Vote Count?*, volume 204 of *CSLI Lecture Notes*. Center for the Study of Language and Information, Stanford University, 2012.
- [85] Ari Juels, Dario Catalano, and Markus Jakobsson. Coercion-Resistant Electronic Elections. Cryptology ePrint Archive, Report 2002/165, 2002.
- [86] Ari Juels, Dario Catalano, and Markus Jakobsson. Coercion-Resistant Electronic Elections. In *WPES'05: 4th Workshop on Privacy in the Electronic Society*, pages 61–70. ACM Press, 2005. See also <http://www.colombia.rsa.com/rsalabs/staff/bios/ajuels/publications/Coercion/Coercion.pdf> (accessed 7 May 2014).
- [87] Ari Juels, Dario Catalano, and Markus Jakobsson. Coercion-Resistant Electronic Elections. In David Chaum, Markus Jakobsson, Ronald L. Rivest, and Peter Y. A. Ryan, editors, *Towards Trustworthy Elections: New Directions in Electronic Voting*, volume 6000 of *LNCS*, pages 37–63. Springer, 2010.
- [88] Jonathan Katz and Yehuda Lindell. *Introduction to Modern Cryptography*. Chapman & Hall/CRC, 2007.
- [89] Shahram Khazaei and Douglas Wikström. Randomized partial checking revisited. In *CT-RSA'13: The Cryptographers' Track at the RSA Conference*, volume 7779 of *LNCS*, pages 115–128. Springer, 2013.
- [90] Aggelos Kiayias. Electronic voting. In *Handbook of Financial Cryptography and Security*, chapter 3. Chapman and Hall/CRC, 2010.
- [91] Aggelos Kiayias, Thomas Zacharias, and Bingsheng Zhang. End-to-end verifiable elections in the standard model. In *EUROCRYPT'15: 34th International Conference on the Theory and Applications of Cryptographic Techniques*, volume 9057 of *LNCS*, pages 468–498. Springer, 2015.
- [92] Neal Koblitz and Alfred Menezes. Another look at security definitions. <http://www.cacr.math.uwaterloo.ca/~ajmenezes/anotherlook/definitions.shtml>, 2011.
- [93] Steve Kremer and Mark D. Ryan. Analysis of an Electronic Voting Protocol in the Applied Pi Calculus. In *ESOP'05: 14th European Symposium on Programming*, volume 3444 of *LNCS*, pages 186–200. Springer, 2005.
- [94] Steve Kremer, Mark D. Ryan, and Ben Smyth. Election verifiability in electronic voting protocols. In *ESORICS'10: 15th European Symposium on Research in Computer Security*, volume 6345 of *LNCS*, pages 389–404. Springer, 2010.
- [95] Ralf Küsters and Johannes Müller. Cryptographic Security Analysis of E-voting Systems: Achievements, Misconceptions, and Limitations. In *E-Vote-ID'17: 2nd International Joint Conference on Electronic Voting*, pages 21–41. Springer, 2017.

- [96] Ralf Küsters and Tomasz Truderung. An Epistemic Approach to Coercion-Resistance for Electronic Voting Protocols. In *S&P'09: 30th IEEE Symposium on Security and Privacy*, pages 251–266. IEEE Computer Society, 2009.
- [97] Ralf Küsters, Tomasz Truderung, and Andreas Vogt. Accountability: Definition and relationship to verifiability. In *CCS'10: 17th ACM Conference on Computer and Communications Security*, pages 526–535. ACM Press, 2010.
- [98] Ralf Küsters, Tomasz Truderung, and Andreas Vogt. Verifiability, Privacy, and Coercion-Resistance: New Insights from a Case Study. In *S&P'11: 32nd IEEE Symposium on Security and Privacy*, pages 538–553. IEEE Computer Society, 2011. Full version available at <http://infsec.uni-trier.de/publications/paper/KuestersTruderungVogt-TR-2011.pdf>.
- [99] Ralf Küsters, Tomasz Truderung, and Andreas Vogt. A Game-Based Definition of Coercion-Resistance and its Applications. *Journal of Computer Security*, 20(6):709–764, 2012.
- [100] Ralf Küsters, Tomasz Truderung, and Andreas Vogt. Clash Attacks on the Verifiability of E-Voting Systems. In *S&P'12: 33rd IEEE Symposium on Security and Privacy*, pages 395–409. IEEE Computer Society, 2012.
- [101] Ralf Küsters, Tomasz Truderung, and Andreas Vogt. Accountability: Definition and relationship to verifiability. Cryptology ePrint Archive, Report 2010/236 (version 20150202:163211), 2015.
- [102] Philip MacKenzie, Thomas Shrimpton, and Markus Jakobsson. Threshold password-authenticated key exchange. In *CRYPTO'02: 22nd International Cryptology Conference*, volume 2442 of *LNCS*, pages 385–400. Springer, 2002.
- [103] Adam McCarthy, Ben Smyth, and Elizabeth A. Quaglia. Hawk and Aucitas: e-auction schemes from the Helios and Civitas e-voting schemes. In *FC'14: 18th International Conference on Financial Cryptography and Data Security*, volume 8437 of *LNCS*, pages 51–63. Springer, 2014.
- [104] Maxime Meyer and Ben Smyth. Exploiting re-voting in the helios election system. *Information Processing Letters*, (143):14–19, 2019.
- [105] Tal Moran and Moni Naor. Receipt-Free Universally-Verifiable Voting with Everlasting Privacy. In *CRYPTO'06: 26th International Cryptology Conference*, volume 4117 of *LNCS*, pages 373–392. Springer, 2006.
- [106] C. Andrew Neff. A verifiable secret shuffle and its application to e-voting. In *CCS'01: 8th ACM Conference on Computer and Communications Security*, pages 116–125. ACM Press, 2001.
- [107] C. Andrew Neff. Practical high certainty intent verification for encrypted votes. Unpublished manuscript, 2004.
- [108] NIST. Secure Hash Standard (SHS). FIPS PUB 180-4, Information Technology Laboratory, National Institute of Standards and Technology, March 2012.
- [109] Helios Princeton Elections. <https://princeton.heliosvoting.org/> (accessed 7 May 2014), 2012.
- [110] Miriam Paiola and Bruno Blanchet. Verification of Security Protocols with Lists: From Length One to Unbounded Length. In *POST'12: First Conference on Principles of Security and Trust*, volume 7215 of *LNCS*, pages 69–88. Springer, 2012.
- [111] Participants of the Dagstuhl Conference on Frontiers of E-Voting. *Dagstuhl Accord*, 2007. <http://www.dagstuhlaccord.org/> (accessed 7 May 2014).
- [112] Olivier Pereira. Internet Voting with Helios. In *Real-World Electronic Voting: Design, Analysis and Deployment*, chapter 11. CRC Press, 2016.
- [113] R. A. Peters. A secure bulletin board. Master's thesis, Technische Universiteit Eindhoven, June 2005.
- [114] Elizabeth A. Quaglia and Ben Smyth. A short introduction to secrecy and verifiability for elections. arXiv, Report 1702.03168, 2017.
- [115] Elizabeth A. Quaglia and Ben Smyth. Authentication with weaker trust assumptions for voting systems. In *AFRICACRYPT'18: 10th International Conference on Cryptology in Africa*, volume 10831 of *LNCS*, pages 322–343. Springer, 2018.
- [116] Elizabeth A Quaglia and Ben Smyth. Secret, verifiable auctions from elections. *Theoretical Computer Science*, 730:44–92, 2018.
- [117] Kazue Sako and Joe Kilian. Secure Voting Using Partially Compatible Homomorphisms. In *CRYPTO'94: 14th International Cryptology Conference*, volume 839 of *LNCS*, pages 411–424. Springer, 1994.
- [118] Daniel Sandler and Dan S. Wallach. Casting votes in the Auditorium. In *EVT'07: Electronic Voting Technology Workshop*, Berkeley, CA, USA, 2007. USENIX Association.
- [119] Claus-Peter Schnorr. Efficient Identification and Signatures for Smart Cards. In *CRYPTO'89: 9th International Cryptology Conference*, volume 435 of *LNCS*, pages 239–252. Springer, 1990.
- [120] Berry Schoenmaker. Lecture notes: Cryptographic protocols, 2016. Version 1.21 <http://www.win.tue.nl/~berry/2DMI00/LectureNotes.pdf>. (accessed 19 December 2016).
- [121] Berry Schoenmakers. Voting Schemes. In Mikhail J. Atallah and Marina Blanton, editors, *Algorithms and Theory of Computation Handbook, Second Edition, Volume 2: Special Topics and Techniques*, chapter 15. CRC Press, 2009.
- [122] Nicole Schweikardt. Arithmetic, first-order logic, and counting quantifiers. *ACM Transactions on Computational Logic*, 6(3):634–671, July 2005.
- [123] Ben Smyth. *Formal verification of cryptographic protocols with automated reasoning*. PhD thesis, School of Computer Science, University of Birmingham, 2011.
- [124] Ben Smyth. Replay attacks that violate ballot secrecy in Helios. Cryptology ePrint Archive, Report 2012/185, 2012.
- [125] Ben Smyth. First-past-the-post suffices for ranked voting. <https://bensmyth.com/publications/2017-FPTP-suffices-for-ranked-voting/>, 2017.
- [126] Ben Smyth. Ballot secrecy: Security definition, sufficient conditions, and analysis of Helios. Cryptology ePrint Archive, Report 2015/942, 2018.
- [127] Ben Smyth. A foundation for secret, verifiable elections. Cryptology ePrint Archive, Report 2018/225, 2018.
- [128] Ben Smyth. Verifiability of helios mixnet. In *Voting'18: 3rd Workshop on Advances in Secure Electronic Voting*, LNCS. Springer, 2018. To appear.
- [129] Ben Smyth. Verifiability of helios mixnet. Cryptology ePrint Archive, Report 2018/017, 2018.
- [130] Ben Smyth, Myrto Arapinis, and Mark D Ryan. Translating between equational theories for automated reasoning. *FCS'13: Workshop on Foundations of Computer Security*, 2013.
- [131] Ben Smyth and David Bernhard. Ballot secrecy and ballot independence coincide. In *ESORICS'13: 18th European Symposium on Research in Computer Security*, volume 8134 of *LNCS*, pages 463–480. Springer, 2013.
- [132] Ben Smyth and David Bernhard. Ballot secrecy and ballot independence: definitions and relations. Cryptology ePrint Archive, Report 2013/235, 2014.
- [133] Ben Smyth and David Bernhard. Ballot secrecy with malicious bulletin boards. Technical Report 2014/822, Cryptology ePrint Archive, 2014.
- [134] Ben Smyth and Véronique Cortier. A note on replay attacks that violate privacy in electronic voting schemes. Technical Report RR-7643, INRIA, June 2011.
- [135] Ben Smyth and Yoshikazu Hanatani. Non-malleable encryption with proofs of plaintext knowledge and applications to voting. Cryptology ePrint Archive, Report 2019/026, 2018.
- [136] Ben Smyth, Yoshikazu Hanatani, and Hirofumi Muratani. NM-CPA secure encryption with proofs of plaintext knowledge. In *IWSEC'15: 10th International Workshop on Security*, volume 9241 of *LNCS*. Springer, 2015.

- [137] Ben Smyth and Alfredo Pironti. Truncating TLS Connections to Violate Beliefs in Web Applications. In *WOOT'13: 7th USENIX Workshop on Offensive Technologies*. USENIX Association, 2013. (First appeared at Black Hat USA 2013.).
- [138] Ben Smyth and Alfredo Pironti. Truncating TLS Connections to Violate Beliefs in Web Applications. Technical Report hal-01102013, INRIA, 2015.
- [139] Ben Smyth, Mark D. Ryan, Steve Kremer, and Mounira Kourjeh. Towards automatic analysis of election verifiability properties. In *ARSPA-WITS'10: Joint Workshop on Automated Reasoning for Security Protocol Analysis and Issues in the Theory of Security*, volume 6186 of *LNCS*, pages 165–182. Springer, 2010.
- [140] CACM Staff. ACM's 2014 General Election: Please Take This Opportunity to Vote. *Communications of the ACM*, 57(5):9–17, May 2014.
- [141] Georgios Tsoukalas, Kostas Papadimitriou, Panos Louridas, and Panayiotis Tsanakas. From Helios to Zeus. *Journal of Election Technology and Systems*, 1(1), 2013.
- [142] UK Electoral Commission. *Key issues and conclusions: May 2007 electoral pilot schemes*, May 2007. http://www.electoralcommission.org.uk/__data/assets/electoral_commission_pdf_file/0015/13218/Keyfindingsandrecommendationssummarypaper_27191-20111__E__N__S__W__.pdf (accessed 7 May 2014).
- [143] Filip Zagórski, Richard T. Carback, David Chaum, Jeremy Clark, Aleksander Essex, and Poorvi L. Vora. Remotegrity: Design and Use of an End-to-End Verifiable Remote Voting System. In *ACNS'13: 11th International Conference on Applied Cryptography and Network Security*, volume 7954 of *LNCS*, pages 441–457. Springer, 2013.