

Ballot secrecy: Security definition, sufficient conditions, and analysis of Helios

Ben Smyth

Interdisciplinary Centre for Security, Reliability and
Trust, University of Luxembourg, Luxembourg

March 22, 2019

Abstract

We propose a definition of ballot secrecy as an indistinguishability game in the computational model of cryptography. Our definition improves upon earlier definitions to ensure ballot secrecy is preserved in the presence of an adversary that controls ballot collection. We also propose a definition of ballot independence as an adaptation of an indistinguishability game for asymmetric encryption. We prove relations between our definitions. In particular, we prove ballot independence is sufficient for ballot secrecy in voting systems with zero-knowledge tallying proofs. Moreover, we prove that building voting systems from non-malleable asymmetric encryption schemes suffices for ballot secrecy, thereby eliminating the expense of ballot-secrecy proofs for a class of encryption-based voting systems. We demonstrate applicability of our results by analysing the Helios voting system and its mixnet variant. Our analysis reveals that Helios does not satisfy ballot secrecy in the presence of an adversary that controls ballot collection. The vulnerability cannot be detected by earlier definitions of ballot secrecy, because they do not consider such adversaries. We adopt non-malleable ballots as a fix and prove that the fixed system satisfies ballot secrecy.

1 Introduction

An election is a decision-making procedure to choose representatives [LG84, Saa95, Gum05, AH10]. Choices should be made by voters with equal influence, and this must be ensured by voting systems, as prescribed by the United Nations [UN48], the Organisation for Security & Cooperation in Europe [OSC90], and the Organization of American States [OAS69]. Historically, “Americans [voted] with their voices – *viva voce* – or with their hands or with their feet. Yea or nay. Raise your hand. All in favor of Jones, stand on this side of the

town common; if you support Smith, line up over there” [Lep08]. Thus, ensuring that only voters voted and did so with equal influence was straightforward. Indeed, the election outcome could be determined by anyone present, simply by considering at most one vote per voter and disregarding non-voters. Yet, voting systems must also ensure choices are made freely, as prescribed by the aforementioned organisations [UN48, OSC90, OAS69]. Mill eloquently argues that choices cannot be expressed freely in public: “The unfortunate voter is in the power of some opulent man; the opulent man informs him how he must vote. Conscience, virtue, moral obligation, religion, all cry to him, that he ought to consult his own judgement, and faithfully follow its dictates. The consequences of pleasing, or offending the opulent man, stare him in the face...the moral obligation is disregarded, a faithless, a prostitute, a pernicious vote is given” [Mil30].

The need for free-choice started a movement towards voting as a private act, i.e., “when numerous social constraints in which citizens are routinely and universally enmeshed – community of religious allegiances, the patronage of big men, employers or notables, parties, ‘political machines’ – are kept at bay,” and “this idea has become the current *doxa* of democracy-builders worldwide” [BBP07]. The most widely used embodiment of this idea is the Australian system, which demands that votes be marked on uniform ballots in polling booths and deposited into ballot boxes. Uniformity is intended to enable free-choice during distribution, collection and tallying of ballots, and the isolation of polling booths is intended to facilitate free-choice whilst marking.¹ Moreover, the Australian system can assure that only voters vote and do so with equal influence. Indeed, observers can check that ballots are only distributed to voters and at most one ballot is deposited by each voter. Furthermore, observers can check that spoiled ballots are discarded and that votes expressed in the remaining ballots correspond to the election outcome. Albeit, assurance is limited by an observer’s ability to monitor [Bjo04, Kel12, Nor15] and the ability to transfer that assurance is limited to the observer’s “good word or sworn testimony” [NA03].

Beyond the paper-based Australian system, electronic voting systems are emerging. Unfortunately, these electronic systems are routinely broken in ways that violate free-choice, e.g., [KSRW04, GH07, Bow07, WWH⁺10, WWIH12, SFD⁺14], or permit undue influence, e.g., [KSRW04, UK07, Bow07, Ger09, JS12]. Breaks can be avoided by proving that systems satisfy formal notions of voters voting freely and of detecting undue influence. Universal verifiability formalises the latter notion, and we propose a definition of ballot secrecy that formalises the former. Our definition is presented in the computational, game-based model of cryptography, whereby a benign challenger, a malicious adversary and a voting system engage in a series of interactions which task the adversary to break security.

Ballot secrecy formalises a notion of free-choice,² assuming voters’ ballots

¹Earlier systems merely required ballots to be marked in polling booths and deposited into ballot boxes, which permitted non-uniform ballots, including ballots of different colours and sizes, that could be easily identified as party tickets [Bre06].

are constructed and tallied in the prescribed manner.

- **Ballot secrecy.** A voter’s vote is not revealed to anyone.

We capture ballot secrecy as a game that proceeds as follows. First, the adversary picks a pair of votes v_0 and v_1 . Secondly, the challenger constructs a ballot for vote v_β (in the manner prescribed by the voting system), where β is a bit chosen uniformly at random. That ballot is given to the adversary. The adversary and challenger repeat the process to construct further ballots, using the same bit β . Thirdly, the adversary constructs a set of ballots, which may include ballots constructed by the adversary and ballots constructed by the challenger. Thus, the game captures a setting where the adversary casts ballots on behalf of some voters and controls the votes cast by the remaining voters. Fourthly, the challenger tallies the set of ballots (in the manner prescribed by the voting system) to determine the election outcome, which is given to the adversary. Finally, the adversary is tasked with determining if $\beta = 0$ or $\beta = 1$. To avoid trivial distinctions, we require that the aforementioned votes (controlled by the adversary) remain constant regardless of whether $\beta = 0$ or $\beta = 1$. If the adversary wins the game, then a voter’s vote can be revealed, otherwise, it cannot, i.e., the voting system provides ballot secrecy. Our game improves upon games by Bernhard *et al.* [BCP⁺11, BPW12b, SB13, SB14, BCG⁺15b] to ensure ballot secrecy is preserved in the presence of an adversary that controls ballot collection (i.e., the bulletin board and the communication channel), whereas games by Bernhard *et al.* do not.

Beyond ballot secrecy, voting systems should satisfy properties including universal verifiability, which requires systems to produce evidence that can be checked to determine whether votes expressed in ballots correspond to the election outcome, thereby enabling the detection of undue influence. Smyth, Frink & Clarkson [SFC17] capture universal verifiability as a game that tasks the adversary to falsify evidence that causes checks to succeed when the outcome does not correspond to the votes expressed in collected ballots, or that cause checks to fail when the outcome does correspond to the votes expressed. Thus, winning the game signifies the existence of a scenario in which a spurious outcome will be accepted or a legitimate outcome rejected. By comparison, when no winning adversary exists, anyone can determine whether the election outcome is correct. Universal verifiability and ballot secrecy are orthogonal properties of voting systems that can be studied, formulated and analysed independently. We shall largely avoid discussion of verifiability, except to introduce general concepts, to highlight features needed solely for verifiability, and to simplify proofs.

We introduce two voting systems to demonstrate how ballot secrecy and universal verifiability can be achieved. The first (**Nonce**) instructs each voter to cast a ballot comprising of their vote paired with a nonce (which is collected and stored on a bulletin board) and instructs the tallier to publish the election

²*Ballot secrecy* and *privacy* occasionally appear as synonyms in the literature. We favour ballot secrecy to avoid confusion with other privacy notions, such as receipt-freeness and coercion resistance, which we will briefly discuss in Section 7.

outcome corresponding to votes (stored on that board). The second (*Enc2Vote*) instructs voters to cast asymmetric encryptions of their votes and instructs the tallier to decrypt the encrypted votes and publish the outcome corresponding to those votes. Universal verifiability is ensured by the former system, because anyone can recompute the election outcome to check that it corresponds to votes expressed in collected ballots. But, ballot secrecy is not, because voters' votes are revealed. By comparison, secrecy is ensured by the latter system, because asymmetric encryption can ensure that votes cannot be recovered from ballots and the tallying procedure ensures that individual votes are not revealed. But, universal verifiability is not ensured. Indeed, spurious election outcomes need not correspond to the encrypted votes. Thus, *Enc2Vote* ensures secrecy not verifiability, and *Nonce* achieves the reverse. More advanced voting systems must simultaneously satisfy both secrecy and verifiability, and we will consider the Helios voting system.

Helios is an open-source, web-based electronic voting system [AMPQ09], which has been used in binding elections. In particular, the International Association of Cryptologic Research (IACR) has used Helios annually since 2010 to elect board members [BVQ10, HBH10],³ the Association for Computing Machinery (ACM) used Helios for their 2014 general election [Sta14], the Catholic University of Louvain used Helios to elect their university president in 2009 [AMPQ09], and Princeton University has used Helios since 2009 to elect student governments.⁴ Helios is intended to satisfy universal verifiability whilst maintaining ballot secrecy. For ballot secrecy, each voter is instructed to encrypt their vote using an asymmetric homomorphic encryption scheme. Encrypted votes are homomorphically combined and the homomorphic combination is decrypted to reveal the outcome. Alternatively, a mixnet is applied to the encrypted votes and the mixed encrypted votes are decrypted to reveal the outcome [Adi08, BGP11]. We continue to refer to the former voting system as Helios and, henceforth, refer to the latter variant as *Helios Mixnet*. For universal verifiability, the encryption step is accompanied by a non-interactive zero-knowledge proof demonstrating correct computation. This ensures homomorphic combinations of encrypted votes and mixed encrypted votes can be decrypted, hence, the outcome can be recovered. Helios additionally requires proof that ciphertexts encrypt votes. This prevents an adversarial voter crafting a ciphertext that could be combined with others to derive an election outcome in the voter's favour. (E.g., votes might be switched between candidates.) The decryption step is similarly accompanied by a non-interactive zero-knowledge proof to prevent spurious outcomes.

Contribution and structure Contributions are summarised in the bullet points below and described in detail by the surrounding text:

- A definition of ballot secrecy that overcomes limitations of previous works

³<https://www.iacr.org/elections/>, accessed 21 Sep 2017.

⁴<http://heliosvoting.wordpress.com/2009/10/13/helios-deployed-at-princeton/> and <https://princeton.heliosvoting.org/>, accessed 21 Sep 2017.

by considering an adversary that controls ballot collection.

Section 3 briefly explains the pitfalls of existing ballot secrecy definitions, introduces our game-based definition of ballot secrecy, adapts formalisations of non-malleability and indistinguishability for asymmetric encryption to derive two equivalent game-based definitions of ballot independence, and proves relations between definitions. In particular, ballot independence is shown to be sufficient for ballot secrecy in a class of voting systems with zero-knowledge tallying proofs, and it is shown to be necessary, but not sufficient, in general.

- A Helios case study that discovers an attack, which cannot be detected by previous work, and a proof that secrecy is satisfied after applying a fix.

Section 4 shows that our definition of ballot secrecy can be used to identify a known vulnerability in Helios, discovers that its patched successor does not defend against that vulnerability in the presence of an adversary that controls ballot collection, explains why earlier definitions of ballot secrecy by Bernhard *et al.* cannot detect that vulnerability, identifies a new exploit that enables an adversary to determine if a voter did not vote for the adversary's preferred candidate, discusses non-malleable ballots as a fix, and uses our sufficient condition to prove that secrecy is satisfied when the fix is applied.

- A proof that voting systems built from non-malleable encryption satisfy ballot secrecy if tallying is additive, which trivialises secrecy proofs.

Section 5 proves that ballot independence cannot be harmed by tallying, if all ballots are tallied correctly; shows that universally-verifiable voting systems tally ballots correctly; proves `Enc2Vote` satisfies ballot independence, assuming the underlying asymmetric encryption scheme is non-malleable; and combines those results to show that proofs of ballot secrecy are trivial for a class of universally-verifiable, encryption-based voting systems.

- A Helios Mixnet case study that demonstrates the triviality of ballot secrecy proofs for systems built from non-malleable encryption.

Section 6 presents an analysis of Helios Mixnet and demonstrates that our results do indeed make proofs of ballot secrecy trivial, by showing that the combination of universal verifiability and non-malleable encryption suffice for ballot secrecy in Helios Mixnet.

- A proof that interactive ballot construction is necessary for stronger notions of privacy such as receipt-freeness and coercion resistance.

Section 7 proves that receipt-freeness cannot be satisfied by a class of voting systems with non-interactive ballot construction from only public inputs and a vote; presents a generic construction for voting systems with interactive ballot construction from systems without interaction; proves that the construction

produces systems satisfying receipt-freeness, if the underlying voting system satisfies ballot secrecy; and argues that coercion resistance cannot be satisfied without private inputs during ballot construction. The remaining sections present syntax (§2), related work (§8), and a brief conclusion (§9); Sidebar 1 introduces game-based security definitions and recalls notation; and the appendices define cryptographic primitives and relevant security definitions (Appendix A) and present further supplementary material. (Readers familiar with games might like to skip Sidebar 1, and some readers might like to study the related work before our definition of ballot secrecy.)

2 Election scheme syntax

We recall syntax (Definition 1) for voting systems that consist of the following three steps. First, a tallier generates a key pair. Secondly, each voter constructs and casts a ballot for their vote. These ballots are collected and recorded on a bulletin board. Finally, the tallier tallies the collected ballots and announces the outcome as a frequency distribution of votes. The chosen representative is derived from this distribution, e.g., as the candidate with the most votes.⁵

Definition 1 (Election scheme [SFC17]). *An election scheme is a tuple of probabilistic polynomial-time algorithms (Setup, Vote, Tally) such that:*⁶

Setup, denoted $(pk, sk, mb, mc) \leftarrow \text{Setup}(\kappa)$, is run by the tallier. The algorithm takes a security parameter κ as input and outputs a key pair pk, sk , a maximum number of ballots mb , and a maximum number of candidates mc .

Vote, denoted $b \leftarrow \text{Vote}(pk, v, nc, \kappa)$, is run by voters. The algorithm takes as input a public key pk , a voter’s vote v , some number of candidates nc , and a security parameter κ . Vote v should be selected from a sequence $1, \dots, nc$ of candidates. The algorithm outputs a ballot b or error symbol \perp .

Tally, denoted $(\mathbf{v}, pf) \leftarrow \text{Tally}(sk, \mathbf{bb}, nc, \kappa)$, is run by the tallier. The algorithm takes as input a private key sk , a bulletin board \mathbf{bb} , some number of candidates nc , and a security parameter κ , where \mathbf{bb} is a set. The algorithm outputs an election outcome \mathbf{v} and a non-interactive tallying proof pf , where \mathbf{v} is a vector of length nc and each index v of that vector should indicate the number of votes for candidate v .

⁵Smyth, Frink & Clarkson use the syntax to model first-past-the-post voting systems [SFC17] and Smyth shows ranked-choice voting systems can be modelled too [Smy17]. Both works consider a single tallier and we discuss distributing the tallier’s role in Section 8.

⁶The syntax bounds the number of ballots mb , respectively candidates mc , to broaden the correctness definition’s scope (indeed, Helios requires mb and mc to be less than or equal to the size of the underlying encryption scheme’s message space); represents votes as integers, rather than alphanumeric strings, for brevity; and employs sets, rather than multisets or lists, to preclude the construction of schemes vulnerable to attacks that arise due to duplicate ballots [BS15, §2.1 & §4.3] (systems vulnerable to such attacks cannot be modelled using the syntax).

Sidebar 1 Preliminaries: Games and notation

A game formulates a series of interactions between a benign challenger, a malicious adversary, and a cryptographic scheme. The adversary wins by completing a task that captures an execution of the scheme in which security is broken, i.e., winning captures what should be unachievable. Tasks can generally be expressed as *indistinguishability* or *reachability* requirements. For example, universal verifiability can be expressed as the inability to reach a state that causes a voting system’s checks to succeed for invalid election outcomes, or fail for valid outcomes. Moreover, ballot secrecy can be expressed as the inability to distinguish between an instance of a voting system in which voters cast some votes, from another instance in which the voters cast a permutation of those votes.

Formally, games are probabilistic algorithms that output booleans. We let $A(x_1, \dots, x_n; r)$ denote the output of probabilistic algorithm A on inputs x_1, \dots, x_n and coins r , and we let $A(x_1, \dots, x_n)$ denote $A(x_1, \dots, x_n; r)$, where coins r are chosen uniformly at random from the coin space of algorithm A . Moreover, we let $x \leftarrow T$ denote assignment of T to x , and $x \leftarrow_R S$ denote assignment to x of an element chosen uniformly at random from set S . Using our notation, we can formulate game $\text{Exp}(H, S, \mathcal{A})$ – which tasks an adversary \mathcal{A} to distinguish between a function H and a simulator S – as follows: $m \leftarrow \mathcal{A}(); \beta \leftarrow_R \{0, 1\};$ **if** $\beta = 0$ **then** $x \leftarrow H(m);$ **else** $x \leftarrow S(m);$ $g \leftarrow \mathcal{A}(x);$ **return** $g = \beta$. Adversaries are *stateful*, i.e., information persists across invocations of an adversary in a game. In particular, adversaries can access earlier assignments. For instance, the adversary’s second instantiation in game Exp has access to any assignments made during its first instantiation. An adversary *wins* a game by causing it to output true (\top) and the adversary’s *success* in a game $\text{Exp}(\cdot)$, denoted $\text{Succ}(\text{Exp}(\cdot))$, is the probability that the adversary wins, that is, $\text{Succ}(\text{Exp}(\cdot)) = \Pr[\text{Exp}(\cdot) = \top]$. We focus on computational security, rather than information-theoretic security, and tolerate breaks by adversaries in non-polynomial time and breaks with negligible success, since such breaks are infeasible in practice.

Game Exp captures a single interaction between the challenger and the adversary. We can extend games with oracles to capture arbitrarily many interactions. For instance, we can formulate a strengthening of Exp as follows: $\beta \leftarrow_R \{0, 1\}; g \leftarrow \mathcal{A}^\mathcal{O}(x);$ **return** $g = \beta$, where $\mathcal{A}^\mathcal{O}$ denotes \mathcal{A} ’s access to oracle \mathcal{O} and $\mathcal{O}(m)$ computes **if** $\beta = 0$ **then** $x \leftarrow H(m);$ **else** $x \leftarrow S(m);$ **return** x . Oracles may access game parameters such as bit β .

Beyond the above notation, we let $x[i]$ denote component i of vector x and let $|x|$ denote the length of vector x . Moreover, we write $(x_1, \dots, x_{|T|}) \leftarrow T$ for $x \leftarrow T; x_1 \leftarrow x[1]; \dots; x_{|T|} \leftarrow x[|T|]$, when T is a vector, and $x, x' \leftarrow_R S$ for $x \leftarrow_R S; x' \leftarrow_R S$.

Election schemes must satisfy correctness, that is, there exists a negligible function negl , such that for all security parameters κ , integers n_b and n_c , and votes $v_1, \dots, v_{n_b} \in \{1, \dots, n_c\}$, it holds that, given a zero-filled vector \mathbf{v} of length n_c , we have:

```

Pr[(pk, sk, mb, mc) ← Setup(κ);
  for 1 ≤ i ≤ nb do
  | bi ← Vote(pk, vi, nc, κ);
  | v[vi] ← v[vi] + 1;
(v', pf) ← Tally(sk, {b1, ..., bnb}, nc, κ) :
nb ≤ mb ∧ nc ≤ mc ⇒ v = v' > 1 - negl(κ).

```

The syntax provides a language to model voting systems and the correctness condition ensures that such systems function, i.e., election outcomes correspond to votes expressed in ballots, when ballots are constructed and tallied in the prescribed manner. We will use our syntax to express a privacy property of election schemes, moreover, we will model and analyse voting systems, including Helios and Helios Mixnet.

Election schemes may also include an algorithm `Verify`, which is used to audit an election. We omit that algorithm from Definition 1, because we focus on ballot secrecy, rather than verifiability.

Our syntax captures voting systems including Helios and Helios Mixnet (§4 & §6). It also captures Helios-C [CGGI14] and the system by Juels, Catalano & Jakobsson [JCJ10] (once extended to include voter credentials [SFC17]), so Belenios and Civitas should be included as well. These schemes are rather dominant, so the class is interesting and rather general. Notable voting systems that are excluded include: systems relying on features implemented with paper (e.g., [Nef04, CRS05, CCC+08]); systems which require partially private ballots (e.g., [MN06]); and schemes relying on partially honest bulletin boards (e.g., [CHI+]).

3 Privacy

Some scenarios inevitably reveal voters' votes: Unanimous outcomes reveal how everyone voted and, more generally, outcomes can be coupled with partial knowledge of voters' votes to deduce voters' votes. For example, suppose Alice, Bob and Mallory participate in a referendum and the outcome has frequency two for 'yes' and one for 'no.' Mallory and Alice can deduce Bob's vote by pooling knowledge of their own votes. Similarly, Mallory and Bob can deduce Alice's vote. Furthermore, Mallory can deduce that Alice and Bob both voted yes, if she voted no. For simplicity, our informal definition of ballot secrecy (§1) deliberately omitted side-conditions which exclude these inevitable revelations and which are necessary for satisfiability.⁷ We now refine that definition as follows:

A voter's vote is not revealed to anyone, except when the vote can be deduced from the election outcome and any partial knowledge of voters' votes.

⁷Voting systems that announce chosen representatives (e.g., [BY86, HK02, HK04, DK05]), rather than frequency distributions of votes, could offer stronger notions of privacy.

This refinement ensures the aforementioned examples are not violations of ballot secrecy. By comparison, if Mallory votes yes and she can deduce the vote of Alice, without knowledge of Bob’s vote, then ballot secrecy is violated.

We could formulate ballot secrecy as the following game: First, the adversary picks a pair of votes v_0 and v_1 . Secondly, the challenger constructs a ballot b_1 for vote v_β and a second ballot b_2 for $v_{1-\beta}$, where β is a bit chosen uniformly at random. Those ballots are given to the adversary. Thirdly, the adversary constructs ballots b_3, \dots, b_n . Fourthly, the challenger tallies all the ballots (i.e., b_1, \dots, b_n) to determine the election outcome, which the adversary is given. Finally, the adversary is tasked with determining bit β . This game challenges the adversary to determine if the first ballot is for v_0 and the second is for v_1 , or vice-versa. Intuitively, a losing adversary cannot distinguish ballots; seemingly suggesting that Alice voting ‘yes’ is indistinguishable from Bob voting ‘no.’

The first release of Helios is not secure with respect to the aforementioned game, due to a vulnerability identified by Cortier & Smyth [CS13, CS11]. Indeed, an adversary can observe a ballot constructed by the challenger, compute a meaningfully related ballot (from a malleable Helios ballot), and exploit the relation to win the game. This vulnerability can be attributed to tallying meaningfully related ballots; omitting such ballots from tallying, i.e., *ballot weeding*, is postulated as a defence [CS11, SC11, Smy12, CS13, SB13, BCG⁺15b, BCG⁺15a]. Variants of Helios with ballot weeding seem secure with respect to this game. Unfortunately, ballot weeding mechanisms can be subverted by intercepting ballots or by re-ordering ballots. For instance, Smyth, Frink & Clarkson show how re-ordering ballots can subvert weeding mechanisms in a manner that violates universal verifiability [SFC17], and we will see that ballot secrecy can be violated too (§4.2). Given that current definitions cannot detect such vulnerabilities (§8), we should conclude that they are unsuitable. Indeed, the challenger tallying *all* ballots introduces an implicit trust assumption: ballots are *recorded-as-cast*, i.e., cast ballots are preserved with integrity through the ballot collection process.⁸ Thus, vulnerabilities that manipulate the ballot collection process cannot be detected, including vulnerabilities that can be exploited to distinguish Alice voting ‘yes’ from Bob voting ‘no.’ To overcome this shortcoming, we formulate a new definition of ballot secrecy in which the adversary controls the ballot collection process, i.e., the bulletin board and the communication channel.

3.1 Ballot secrecy

We formalise ballot secrecy (Definition 2) as a game that tasks the adversary to: select two lists of votes; construct a bulletin board from ballots for votes in one of those lists, which list is decided by a coin flip; and (non-trivially) determine the result of the coin flip from the resulting election outcome and tallying proof. That is, the game tasks the adversary to distinguish between an instance of the voting system for one list of votes, from another instance with the other list of votes, when the votes cast from each list are permutations of each

⁸The recorded-as-cast notion was introduced by Adida & Neff [AN06, §2].

other (hence, the distinction is non-trivial). The game proceeds as follows: The challenger generates a key pair (Line 1), the adversary chooses some number of candidates (Line 2), and the challenger flips a coin (Line 3) and initialises a set to record lists of votes (Line 4). The adversary computes a bulletin board from ballots for votes in one of two possible lists (Line 5), where the lists are chosen by the adversary, the choice between lists is determined by the coin flip, and the ballots (for votes in one of the lists) are constructed by an left-right oracle (further ballots may be constructed by the adversary).⁹ The challenger tallies the bulletin board to derive the election outcome and tallying proof (Line 6), which are given to the adversary and the adversary is tasked with determining the result of the coin flip (Line 7 & 8).

Definition 2 (Ballot-Secrecy). *Let $\Gamma = (\text{Setup}, \text{Vote}, \text{Tally})$ be an election scheme, \mathcal{A} be an adversary, κ be a security parameter, and Ballot-Secrecy be the following game.*

Ballot-Secrecy($\Gamma, \mathcal{A}, \kappa$) =

- 1 $(pk, sk, mb, mc) \leftarrow \text{Setup}(\kappa)$;
- 2 $nc \leftarrow \mathcal{A}(pk, \kappa)$;
- 3 $\beta \leftarrow_R \{0, 1\}$;
- 4 $L \leftarrow \emptyset$;
- 5 $\mathbf{bb} \leftarrow \mathcal{A}^{\mathcal{O}}()$;
- 6 $(\mathbf{v}, pf) \leftarrow \text{Tally}(sk, \mathbf{bb}, nc, \kappa)$;
- 7 $g \leftarrow \mathcal{A}(\mathbf{v}, pf)$;
- 8 **return** $g = \beta \wedge \text{balanced}(\mathbf{bb}, nc, L) \wedge 1 \leq nc \leq mc \wedge |\mathbf{bb}| \leq mb$;

Predicate balanced and oracle \mathcal{O} are defined as follows:

- $\text{balanced}(\mathbf{bb}, nc, L)$ holds if for all votes $v \in \{1, \dots, nc\}$ we have $|\{b \mid b \in \mathbf{bb} \wedge \exists v_1 . (b, v, v_1) \in L\}| = |\{b \mid b \in \mathbf{bb} \wedge \exists v_0 . (b, v_0, v) \in L\}|$; and
- $\mathcal{O}(v_0, v_1)$ computes $b \leftarrow \text{Vote}(pk, v_\beta, nc, \kappa)$; $L \leftarrow L \cup \{(b, v_0, v_1)\}$ and outputs b , where $v_0, v_1 \in \{1, \dots, nc\}$.

We say Γ satisfies Ballot-Secrecy, if for all probabilistic polynomial-time adversaries \mathcal{A} , there exists a negligible function negl , such that for all security parameters κ , we have $\text{Succ}(\text{Ballot-Secrecy}(\Gamma, \mathcal{A}, \kappa)) \leq \frac{1}{2} + \text{negl}(\kappa)$.

An election scheme satisfies Ballot-Secrecy when algorithm `Vote` outputs ballots that do not reveal votes and algorithm `Tally` outputs election outcomes and proofs that do not reveal the relation between votes expressed in collected ballots and the outcome.

Game Ballot-Secrecy tasks the adversary to compute a bulletin board, from ballots constructed by a left-right oracle for votes in one of two possible lists, and determine which list was used from the election outcome and proof generated from tallying that board. The choice between lists is determined by the result

⁹Bellare *et al.* introduced left-right oracles in the context of symmetric encryption [BDJR97] and Bellare & Rogaway provide a tutorial on their use [BR05].

β of a coin flip, and the left-right oracle outputs a ballot for vote v_β on input of a pair of votes v_0, v_1 . Hence, the left-right oracle constructs ballots for votes in one of two possible lists, where the lists are chosen by the adversary, and the bulletin board may contain those ballots in addition to ballots constructed by the adversary.

Election schemes reveal the number of votes for each candidate (i.e., the election outcome). Hence, to avoid trivial distinctions in game **Ballot-Secrecy**, we require that runs of the game are *balanced*: “left” and “right” inputs to the left-right oracle are equivalent, when the corresponding outputs appear on the bulletin board. For example, suppose the inputs to the left-right oracle are $(v_{1,0}, v_{1,1}), \dots, (v_{n,0}, v_{n,1})$ and the corresponding outputs are b_1, \dots, b_n , further suppose the bulletin board is $\{b_1, \dots, b_\ell\}$ such that $\ell \leq n$. That game is balanced if the “left” inputs $v_{1,0}, \dots, v_{\ell,0}$ are a permutation of the “right” inputs $v_{1,1}, \dots, v_{\ell,1}$. The balanced condition prevents trivial distinctions.¹⁰ For instance, an adversary that computes a bulletin board containing only the ballot output by a left-right oracle query with input $(1, 2)$ cannot win the game, because it is unbalanced. Albeit, that adversary could trivially determine whether $\beta = 0$ or $\beta = 1$, given the tally of that board.

Intuitively, if the adversary wins game **Ballot-Secrecy**, then there exists a strategy to distinguish ballots. Indeed, such an adversary can distinguish between an instance of the voting system in which voters cast some votes, from another instance in which voters cast a permutation of those votes, thus, voters’ votes are revealed. Otherwise, the adversary is unable to distinguish between a voter casting a ballot for vote v_0 and another voter casting a ballot for vote v_1 , hence, voters’ votes cannot be revealed.

Proving ballot secrecy is time consuming. Indeed, a proof for the simple **Enc2Vote** scheme consumes around four and a half pages [QS18b, Appendix C.6]. To reduce the expense of such proofs, we introduce ballot independence (§3.2) and prove that it suffices for ballot secrecy (§3.3). Using this sufficient condition, the four and a half page proof can be reduced to around one page [Smy18a, §4.3].

3.2 Ballot independence

Ballot independence [Gen95, CS13, CGMA85] is seemingly related to ballot secrecy.

- *Ballot independence.* Observing another voter’s interaction with the voting system does not allow a voter to cast a meaningfully related vote.

Our informal definition essentially states that an adversary is unable to construct a ballot meaningfully related to a non-adversarial ballot, i.e., ballots are non-malleable. Hence, we can formalise ballot independence as a straightforward adaptation of the non-malleability definition for asymmetric encryption by

¹⁰A weaker balanced condition might be sufficient for alternative formalisations of election schemes. For instance, voting systems which only announce the winning candidate could be analysed using a balanced condition asserting that the winning candidate was input on both the “left” and “right.”

Bellare & Sahai [BS99].¹¹ Such a formalisation captures an intuitive notion of ballot independence, but it is complex and proofs of non-malleability are relatively difficult. Bellare & Sahai observe similar complexities and show that their definition is equivalent to a simpler, indistinguishability notion. In a similar direction, we derive a definition of ballot independence, called indistinguishability under chosen vote attack (IND-CVA), as a straightforward adaptation of their indistinguishability notion for asymmetric encryption.

Definition 3 (IND-CVA). *Let $\Gamma = (\text{Setup}, \text{Vote}, \text{Tally})$ be an election scheme, \mathcal{A} be an adversary, κ be the security parameter, and IND-CVA be the following game.*

IND-CVA($\Gamma, \mathcal{A}, \kappa$) =

```

    ( $pk, sk, mb, mc$ )  $\leftarrow$  Setup( $\kappa$ );
    ( $v_0, v_1, nc$ )  $\leftarrow$   $\mathcal{A}(pk, \kappa)$ ;
     $\beta \leftarrow_R \{0, 1\}$ ;
     $b \leftarrow$  Vote( $pk, v_\beta, nc, \kappa$ );
     $bb \leftarrow$   $\mathcal{A}(b)$ ;
    ( $v, pf$ )  $\leftarrow$  Tally( $sk, bb, nc, \kappa$ );
     $g \leftarrow$   $\mathcal{A}(v)$ ;
    return  $g = \beta \wedge b \notin bb \wedge 1 \leq v_0, v_1 \leq nc \leq mc \wedge |bb| \leq mb$ ;

```

We say Γ satisfies indistinguishability under chosen vote attack (IND-CVA), if for all probabilistic polynomial-time adversaries \mathcal{A} , there exists a negligible function negl , such that for all security parameters κ , we have $\text{Succ}(\text{IND-CVA}(\Gamma, \mathcal{A}, \kappa)) \leq \frac{1}{2} + \text{negl}(\kappa)$.

Game IND-CVA is satisfied if the adversary cannot determine whether the challenge ballot b is for one of two possible votes v_0 and v_1 . In addition to the challenge ballot, the adversary is given the election outcome derived by tallying a bulletin board computed by the adversary. To avoid trivial distinctions, the adversary's bulletin board should not contain the challenge ballot. Intuitively, the adversary wins if there exists a strategy to construct related ballots, since this strategy enables the adversary to construct a ballot b' , related to the challenge ballot b , and determine if b is for v_0 or v_1 from the outcome derived by tallying a bulletin board containing b' .

Comparing IND-CVA and IND-PA0. The main distinction between indistinguishability for asymmetric encryption (IND-PA0) and indistinguishability for election schemes (IND-CVA) is as follows: IND-PA0 performs a parallel decryption, whereas IND-CVA performs a single tally. Hence, indistinguishability for encryption reveals plaintexts corresponding to ciphertexts, whereas indistinguishability for elections reveals the number of votes for each candidate.

¹¹Non-malleability was first formalised by Dolev, Dwork & Naor in the context of asymmetric encryption [DDN91, DDN00]; the definition by Bellare & Sahai builds upon their results and results by Bellare *et al.* [BDPR98].

We present an alternative definition of ballot independence (Appendix B), based upon the definition of non-malleability for asymmetric encryption by Bellare & Sahai, and prove that the definition is equivalent to IND-CVA.

3.3 Secrecy and independence coincide

The main distinctions between our ballot secrecy (Ballot-Secrecy) and ballot independence (IND-CVA) games are as follows.

1. The challenger produces one challenge ballot for the adversary in game IND-CVA, whereas the left-right oracle produces arbitrarily many challenge ballots for the adversary in game Ballot-Secrecy.
2. The adversary in game Ballot-Secrecy has access to a tallying proof, but the adversary in game IND-CVA does not.
3. The winning condition in game Ballot-Secrecy requires the bulletin board to be balanced, whereas the bulletin board must not contain the challenge ballot in game IND-CVA.

The second point distinguishes our games and shows Ballot-Secrecy is at least as strong as IND-CVA. Hence, non-malleable ballots are necessary in election schemes satisfying ballot secrecy.

Theorem 1 (Ballot-Secrecy \Rightarrow IND-CVA). *Given an election scheme Γ satisfying Ballot-Secrecy, we have Γ satisfies IND-CVA.*

A proof of Theorem 1 and all further proofs, except where otherwise stated, appear in Appendix C.

Tallying proofs may reveal voters' votes. For example, a variant of Enc2Vote might define tallying proofs that map ballots to votes. Since proofs are available to the adversary in game Ballot-Secrecy, but not in game IND-CVA, it follows that Ballot-Secrecy is strictly stronger than IND-CVA.

Proposition 2 (IND-CVA $\not\Rightarrow$ Ballot-Secrecy). *There exists an election scheme satisfying IND-CVA, but not Ballot-Secrecy.*

Proposition 2 follows from our informal reasoning and we omit a formal proof.

Game Ballot-Secrecy is generally (strictly) stronger than game IND-CVA. Nonetheless, we show that our games coincide for election schemes without tallying proofs (Definition 4), assuming tallying is additive (Definition 5), that is, the sum of election outcomes derived by tallying distinct subsets of a bulletin board is equivalent to the election outcome derived by tallying the union of those subsets.

Definition 4. *An election scheme $\Gamma = (\text{Setup}, \text{Vote}, \text{Tally})$ is without tallying proofs, if there exists a constant symbol ϵ such that for all private keys sk , sets \mathbf{bb} , integers nc , security parameters κ , and computations $(\mathbf{v}, pf) \leftarrow \text{Tally}(sk, \mathbf{bb}, nc, \kappa)$, we have $pf = \epsilon$.*

Definition 5 (Additivity). *Let $\Gamma = (\text{Setup}, \text{Vote}, \text{Tally})$ be an election scheme, \mathcal{A} be an adversary, κ be a security parameter, and Additivity be the following game.*

Additivity($\Gamma, \mathcal{A}, \kappa$) =
 $(pk, sk, mb, mc) \leftarrow \text{Setup}(\kappa);$
 $(nc, \mathbf{bb}, \mathbf{bb}') \leftarrow \mathcal{A}(pk, \kappa);$
 $(\mathbf{v}, pf) \leftarrow \text{Tally}(sk, \mathbf{bb}, nc, \kappa);$
 $(\mathbf{v}_0, pf_0) \leftarrow \text{Tally}(sk, \mathbf{bb} \setminus \mathbf{bb}', nc, \kappa);$
 $(\mathbf{v}_1, pf_1) \leftarrow \text{Tally}(sk, \mathbf{bb} \cap \mathbf{bb}', nc, \kappa);$
return $\mathbf{v} = \mathbf{v}_0 + \mathbf{v}_1 \wedge nc \leq mc \wedge |\mathbf{bb}| \leq mb;$

We say Γ satisfies Additivity, if for all probabilistic polynomial-time adversaries \mathcal{A} , there exists a negligible function negl , such that for all security parameters κ , we have $\text{Succ}(\text{Additivity}(\Gamma, \mathcal{A}, \kappa)) > 1 - \text{negl}(\kappa)$.

Proposition 3 (Ballot-Secrecy = IND-CVA, without proofs). *Let Γ be an election scheme without tallying proofs. Suppose Γ satisfies Additivity. We have Γ satisfies Ballot-Secrecy iff Γ satisfies IND-CVA.*

Our equivalence result generalises to election schemes with zero-knowledge tallying proofs, i.e., schemes that compute proofs using non-interactive zero-knowledge proof systems.

Definition 6 (Zero-knowledge tallying proofs). *An election scheme $\Gamma = (\text{Setup}, \text{Vote}, \text{Tally})$ has zero-knowledge tallying proofs, if there exists a non-interactive zero-knowledge proof system $(\text{Prove}, \text{Verify})$, such that for all security parameters κ , integers nc , bulletin boards \mathbf{bb} , outputs (pk, sk, mb, mc) of $\text{Setup}(\kappa)$, and outputs (\mathbf{v}, pf) of $\text{Tally}(sk, \mathbf{bb}, nc, \kappa)$, we have pf is an output of algorithm Prove parameterised with statement $(pk, \mathbf{bb}, nc, \mathbf{v})$ and coins chosen uniformly at random.*

Theorem 4 (Ballot-Secrecy = IND-CVA, with ZK proofs). *Let Γ be an election scheme with zero-knowledge tallying proofs. Suppose Γ satisfies Additivity. We have Γ satisfies Ballot-Secrecy iff Γ satisfies IND-CVA.*

Additivity is implied by universal verifiability (Lemmata 8 & 31). Thus, a special case of Theorem 4 requires the election scheme to satisfy universal verifiability, which is useful to simplify its application. Indeed, we exploit this result in the following section to prove Ballot-Secrecy.

4 Case study I: Helios

Helios can be informally modelled as the following election scheme (further details appear in Sidebar 2):

Setup generates a key pair for an asymmetric additively-homomorphic encryption scheme, proves correct key generation in zero-knowledge, and outputs the key pair and proof.

Vote enciphers the vote’s bitstring encoding to a tuple of ciphertexts, proves in zero-knowledge that each ciphertext is correctly constructed and that the vote is selected from the sequence of candidates, and outputs the ciphertexts coupled with the proofs.

Tally selects ballots from the bulletin board for which proofs hold, homomorphically combines the ciphertexts in those ballots, decrypts the homomorphic combination to reveal the election outcome, and announces the outcome, along with a zero-knowledge proof of correct decryption.

Helios was first released in 2009 as *Helios 2.0*,¹² the current release is *Helios 3.1.4*,¹³ and a new release is planned.¹⁴ Henceforth, we’ll refer to the planned release as *Helios’12*.

4.1 Helios 2.0 & Helios 3.1.4

Cortier & Smyth show that Helios 2.0 does not satisfy ballot secrecy (§3) and neither does Helios 3.1.4.¹⁵ Thus, we would not expect our definition of ballot secrecy to hold. Indeed, we adopt formal descriptions of Helios 2.0 and Helios 3.1.4 by Smyth, Frink & Clarkson [SFC17] (Appendix D) and use those descriptions to prove that **Ballot-Secrecy** is not satisfied.

Theorem 5. *Neither Helios 2.0 nor Helios 3.1.4 satisfy Ballot-Secrecy.*

Cortier & Smyth attribute the vulnerability to tallying meaningfully related ballots. Indeed, Helios uses malleable ballots: Given a ballot $c_1, \dots, c_{nc-1}, \sigma_1, \dots, \sigma_{nc}$, we have $c_{\chi(1)}, \dots, c_{\chi(nc-1)}, \sigma_{\chi(1)}, \dots, \sigma_{\chi(nc-1)}, \sigma_{nc}$ is a ballot for all permutations χ on $\{1, \dots, nc - 1\}$. Thus, ballots are malleable, which is incompatible with ballot secrecy (§3.3).

Proof sketch. Suppose an adversary queries the left-right oracle with (distinct) inputs $v_0, v_1 \in \{1, \dots, nc - 1\}$ to derive a ballot for v_β , where integer $nc \geq 3$ is chosen by the adversary and bit β is chosen by the challenger. Further suppose the adversary picks a permutation χ on $\{1, \dots, nc - 1\}$, abuses malleability to derive a related ballot b for $\chi(v_\beta)$, and outputs bulletin board $\{b\}$. The board is balanced, because it does not contain the ballot output by the oracle. Suppose the adversary performs the following computation on input of election outcome \mathbf{v} : if $\mathbf{v}[\chi(v_0)] = 1$, then output 0, otherwise, output 1. Since b is a ballot for $\chi(v_\beta)$, it follows by correctness that $\mathbf{v}[\chi(v_0)] = 1$ iff $\beta = 0$, and $\mathbf{v}[\chi(v_1)] = 1$ iff $\beta = 1$, hence, the adversary wins the game. \square

¹²<https://github.com/benadida/helios/releases/tag/2.0>, released 25 Jul 2009, accessed 21 Sep 2017.

¹³<https://github.com/benadida/helios-server/releases/tag/v3.1.4>, released 10 Mar 2011, last patched 27 Oct 2017, accessed 17 Jan 2018.

¹⁴<https://web.archive.org/web/20171026064140/http://documentation.heliosvoting.org/verification-specs/helios-v4>, published c. 2012, accessed 30 Apr 2018.

¹⁵Helios 3.1.4 mitigates against a universal-verifiability vulnerability in Helios 2.0, by checking that tallied ballots are constructed from suitability cryptographic parameters [CE16, §4.1]. The vulnerabilities described herein use well-formed parameters, hence, the additional checks do not preclude vulnerabilities and we refer the reader to the original description for details.

Sidebar 2 Helios: Ballot construction and tallying

Algorithm `Vote` inputs a vote v selected from candidates $1, \dots, nc$ and computes ciphertexts c_1, \dots, c_{nc-1} such that if $v \neq nc$, then ciphertext c_v contains plaintext 1 and the remaining ciphertexts contain plaintext 0, otherwise, all ciphertexts contain plaintext 0. The algorithm also computes zero-knowledge proofs $\sigma_1, \dots, \sigma_{nc}$ demonstrating correct computation. Proof σ_j demonstrates that ciphertext c_j contains 0 or 1, where $1 \leq j \leq nc - 1$, and proof σ_{nc} demonstrates that the homomorphic combination of ciphertexts $c_1 \otimes \dots \otimes c_{nc-1}$ contains 0 or 1. The algorithm outputs the ciphertexts and proofs.

Algorithm `Tally` inputs a bulletin board \mathbf{bb} ; selects all the ballots $b_1, \dots, b_k \in \mathbf{bb}$ for which proofs hold, i.e., ballots $b_i = \text{Enc}(pk, m_{i,1}), \dots, \text{Enc}(pk, m_{i,nc-1}), \sigma_{i,1}, \dots, \sigma_{i,nc}$ such that proofs $\sigma_{i,1}, \dots, \sigma_{i,nc}$ hold, where $1 \leq i \leq k$; forms a matrix of the encapsulated ciphertexts, i.e.,

$$\begin{array}{c} \text{Enc}(pk, m_{1,1}), \dots, \text{Enc}(pk, m_{1,nc-1}) \\ \vdots \\ \text{Enc}(pk, m_{k,1}), \dots, \text{Enc}(pk, m_{k,nc-1}); \end{array}$$

homomorphically combines the ciphertexts in each column to derive the encrypted outcome, i.e.,

$$\text{Enc}(pk, \sum_{i=1}^k m_{i,1}), \dots, \text{Enc}(pk, \sum_{i=1}^k m_{i,nc-1});$$

decrypts the homomorphic combinations to reveal the frequency of votes $1, \dots, nc - 1$, i.e.,

$$\sum_{i=1}^k m_{i,1}, \dots, \sum_{i=1}^k m_{i,nc-1};$$

computes the frequency of vote nc by subtracting the frequency of any other vote from the number of ballots for which proofs hold, i.e., $k - \sum_{j=1}^{nc-1} \sum_{i=1}^k m_{i,j}$; and announces the outcome as those frequencies, along with a zero-knowledge proof demonstrating correctness of decryption.

For simplicity, our proof sketch considers an adversary that omits ballots from the bulletin board. Voters might detect such an adversary, because Helios satisfies individual verifiability, hence, voters can discover if their ballot is omitted. Our proof sketch can be extended to avoid such detection: Let b_1 be the ballot output by the left-right oracle in the proof sketch and suppose b_2 is the ballot output by a (second) left-right oracle query with inputs v_1 and v_0 . Further suppose the adversary outputs (the balanced) bulletin board $\{b, b_1, b_2\}$ and performs the following computation on input of election outcome \mathbf{v} : if \mathbf{v} corresponds to votes $v_0, v_1, \chi(v_0)$, then output 0, otherwise, output 1, where χ is the permutation chosen by the adversary. Hence, the adversary wins the game.

Beyond ballot secrecy, Bernhard, Pereira & Warinschi show that Helios 3.1.4

does not satisfy universal verifiability [BPW12a].¹⁶ They attribute vulnerabilities to application of the Fiat-Shamir transformation without inclusion of statements in hashes (i.e., weak Fiat-Shamir), and propose including statements in hashes (i.e., applying the Fiat-Shamir transformation) as a defence.

4.2 Helios’12

Helios’12 is intended to mitigate against the aforementioned vulnerabilities in Helios 3.1.4. In particular, the specification incorporates the Fiat-Shamir transformation (rather than weak Fiat-Shamir), and there are plans to incorporate ballot weeding (i.e., to omit meaningfully related ballots from tallying).¹⁷

Bernhard, Pereira & Warinschi [BPW12a], Bernhard [Ber14, §6.11] and Bernhard *et al.* [BCG⁺15a, §D.3] show that Helios’12 satisfies various notions of ballot secrecy. These notions all assume ballots are recorded-as-cast. Unfortunately, ballot secrecy is not satisfied without this assumption, because Helios 2.0, Helios 3.1.4 and Helios’12 all use malleable ballots in elections with more than two candidates.^{18,19}

Remark 6. *Helios’12 does not satisfy* Ballot-Secrecy.

Proof sketch. Neither ballot weeding nor the Fiat-Shamir transformation eliminate the vulnerability we identified in Helios 3.1.4,²⁰ because related ballots need not be tallied, as shown in the proof sketch of Theorem 5. Hence, we conclude by that theorem. \square

Remark 6 shows that Helios’12 does not defend against a known Helios 2.0 vulnerability in the presence of an adversary that controls ballot collection. We also derive a new exploit (as the following example demonstrates) by extrapolating from the proof sketch of Theorem 5 and an attack by Cortier & Smyth that asserts the following: given a ballot b for vote v , we can abuse malleability to derive a ballot b' for vote v' [CS13, §3.2.2]. Suppose Alice, Bob and Charlie are voters, and Mallory is an adversary that wants to convince herself that Alice did not vote for a candidate v . Further suppose Alice casts a ballot b_1 for vote v_1 , Bob casts a ballot b_2 , and Charlie casts a ballot b_3 . Moreover, suppose that either Bob or Charlie vote for v . (Thereby avoiding scenarios without any votes for candidate v , i.e., scenarios which inevitably permit Mallory to convince herself that Alice did not vote for candidate v .) Let us assume that votes for v'

¹⁶Beyond secrecy and verifiability, eligibility is not satisfied [SP13, SP15, MS19].

¹⁷Cf. <https://github.com/benadida/helios-server/issues/8> and <https://github.com/benadida/helios-server/issues/35>, accessed 21 Sep 2017.

¹⁸Proofs by Bernhard, Pereira & Warinschi and Bernhard *et al.* are limited to two candidate elections, for which Helios’12 uses non-malleable ballots.

¹⁹We state Remark 6 informally, because we have not formalised a description of Helios’12. Such a description can be derived as a straightforward variant of Helios 3.1.4 that uses ballot weeding and applies the Fiat-Shamir transformation (rather than the weak Fiat-Shamir transformation). But, these details provide little value, so we do not pursue them.

²⁰The proof sketch of Theorem 5 violates the recorded-as-cast assumption, since the ballot output by the left-right oracle does not appear on the bulletin board.

are not expected. Mallory proceeds as follows: she intercepts ballot b_1 , abuses malleability to derive a ballot b such that $v = v_1$ implies b is a vote for v' , and casts ballot b . It follows that the tallier will compute the election outcome from bulletin board $\{b, b_2, b_3\}$. (Omitting meaningfully related ballots before tallying does not eliminate the vulnerability, because none of the tallied ballots are related.) If the outcome does not contain any votes for v' , then Mallory is convinced that Alice did not vote for v . Notions of ballot secrecy used by Bernhard, Pereira & Warinschi [BPW12a], Bernhard [Ber14, §6.11] and Bernhard *et al.* [BCG⁺15a, §D.3] cannot detect this new exploit nor the known Helios 2.0 vulnerability (in the presence of an adversary that controls ballot collection), because interception is not possible when ballots are recorded-as-cast.²¹

The exploit is reliant on a particular candidate not receiving any votes. This is trivial to capture in the context of game **Ballot-Secrecy**, because the bulletin board is computed by an adversary that casts ballots on behalf of some voters and controls the votes cast by the remaining voters. Beyond the game, candidates will presumably vote for themselves. Thus, for first-past-the-post elections, the exploit's success is probably limited to elections in which voters vote in constituencies and each polling station announces its own outcome (cf. Cortier & Smyth [CS13, §3.3]).

We have seen that an attack against Helios'12 is feasible. Consequently, we cannot prove that Helios'12 satisfies **Ballot-Secrecy**. Whether attacks are likely is subjective, and instead of offering an opinion, let us consider a tale of when an attack might be used in the real-world: Alice meets Charlie, who Alice's family believe is a supporter of a undesirable party. Alice is torn; her love for Charlie is hurting her family. She devises a plan to ensure Charlie is welcomed with open arms: She uses the attack to prove that Charlie did not support the undesirable candidate in the town's election. The concerns of Alice's family are alleviated, Charlie is welcomed, and Alice is finally content. Those that believe in the existence of a real-world Alice will surely accept the attack as practical, because observation of a real-world attack surely suffices to demonstrate practicality.

Attacks that omit ballots from the bulletin board can be detected by voters that cast the omitted ballots, because Helios satisfies individual verifiability. Nonetheless, it is well-known that many voters do not perform necessary checks [BBH⁺17, §2.1.6], hence, the attack is particular effective against voters that do not perform checks. Moreover, even if an attack is detected, recovery is problematic, because there is no convincing evidence that any malpractice has taken place, hence, victims have little recourse.

Ballot weeding. Ballot weeding mechanisms have been proposed, e.g., [CS11, SC11, Smy12, CS13, SB13, BW14, BCG⁺15b, BCG⁺15a], but the specification for Helios'12 does not yet define a particular mechanism. One candidate mechanism omits any ballot containing a previously observed hash from the tallying

²¹This observation suggests that recorded-as-cast is unsatisfiable: An adversary that can intercept ballots can always prevent the collection of ballots. Nevertheless, the definition of recorded-as-cast is informal, thus ambiguity should be expected and some interpretation of the definition should be satisfiable.

procedure. Another – already in use by the IACR – omits any ballot containing a previously observed hash from the bulletin board.²² (More precisely, the mechanism stores the hashes used by non-interactive zero-knowledge proofs in a hashtable and any ballot containing a previously stored hash is omitted from the bulletin board.) These mechanisms can be subverted by excluding ballots (Remark 6). Moreover, similarly to our extended proof sketch of Theorem 5 (§4.1), we can extend our proof sketch of Remark 6 to avoid voter detection, because the former mechanism includes all ballots on the bulletin board and (silently) omits ballots during tallying, and the latter can be disregarded by an adversary that controls ballot collection (hence, the bulletin board).

4.3 Helios’16

We have seen that non-malleable ballots are necessary for ballot secrecy (§3.3), hence, future Helios releases should adopt non-malleable ballots. Smyth, Frink & Clarkson make progress in this direction by proposing Helios’16 [SFC17], a variant of Helios 3.1.4 that uses the Fiat-Shamir transformation and is intended, but not proven, to use non-malleable ballots. We recall their formal description in Appendix D, and using that formalisation we prove that Helios’16 satisfies secrecy.

Theorem 7. *Helios’16 satisfies Ballot-Secrecy.*

Proof sketch. We prove that Helios’16 has zero-knowledge tallying proofs and, since universal verifiability is satisfied [SFC17], we have **Additivity** too (Lemmata 8 & 31). Hence, by Theorem 4, it suffices to show that Helios’16 satisfies **IND-CVA**, which we prove by reduction to the security of the underlying encryption scheme, using an extractor that exists for the proof system used to construct ballots. \square

A formal proof of Theorem 7 appears in Appendix D.1, assuming the random oracle model [BR93]. This proof, coupled with the proof of verifiability by Smyth, Frink & Clarkson [SFC17], provides strong motivation for future Helios releases being based upon Helios’16, since it is the only variant of Helios which is proven to satisfy both ballot secrecy and verifiability (security results are summarised in Table 1).²³

5 Simplifying ballot-secrecy proofs

We have seen that our definitions of ballot secrecy and ballot independence coincide when tallying is additive and tallying proofs are zero-knowledge (Theorem 4). Building upon this result and Proposition 9, we show that tallying

²²David Bernhard, email communication, c. 2014 and 19 Sep 2017.

²³Earlier versions of Helios have been shown to satisfy definitions of ballot secrecy by Bernhard *et al.*, but not notions of verifiability (the analysis by Küsters *et al.* [KTV12b] does not detect vulnerabilities identified by Bernhard *et al.* [BPW12a] and Chang-Fong & Essex [CE16], possibly because their analysis “does not formalize all the cryptographic primitives used by Helios” [SFC17, §9]).

	Helios 2.0	Helios 3.1.4	Helios'12	Helios'16
Secrecy	✗	✗	✗	✓
Verifiability	✗	✗	✗	✓

Cortier & Smyth identify a secrecy vulnerability in Helios 2.0 and Helios 3.1.4 [CS13], and we have seen that the vulnerability is exploitable in Helios'12 when the adversary controls ballot collection (§4.2). Moreover, we have proved that Helios'16 satisfies ballot secrecy (Theorem 7). Bernhard, Pereira & Warinschi identify universal-verifiability vulnerabilities in Helios 2.0 and Helios 3.1.4 [BPW12a], Chang-Fong & Essex identify vulnerabilities in Helios 2.0 [CE16], and Smyth, Frink, & Clarkson identify a vulnerability in Helios'12 [SFC17]. Moreover, Smyth, Frink, & Clarkson prove that Helios'16 satisfies individual and universal verifiability.

Table 1: Summary of Helios security results

cannot harm secrecy when ballots are tallied correctly and tallying proofs are zero-knowledge. That is, (Setup, Vote, Tally) satisfies Ballot-Secrecy if and only if (Setup, Vote, Tally') does, assuming algorithms Tally and Tally' both tally ballots correctly and tallying proofs are zero-knowledge.²⁴

Smyth, Frink & Clarkson capture the notion of tallying ballots correctly using function *correct-outcome* [SFC17]. That function uses a *counting quantifier*: A predicate $(\exists^{=\ell} x : P(x))$ that holds exactly when there are ℓ distinct values for x such that $P(x)$ is satisfied [Sch05].²⁵ Using the counting quantifier, function *correct-outcome* is defined such that $\text{correct-outcome}(pk, nc, \mathbf{bb}, \kappa)[v] = \ell$ iff $\exists^{=\ell} b \in \mathbf{bb} \setminus \{\perp\} : \exists r : b = \text{Vote}(pk, v, nc, \kappa; r)$, where $\text{correct-outcome}(pk, nc, \mathbf{bb}, \kappa)$ is a vector of length nc and $1 \leq v \leq nc$. Hence, component v of vector $\text{correct-outcome}(pk, nc, \mathbf{bb}, \kappa)$ equals ℓ iff there exist ℓ ballots for vote v on the bulletin board. The function requires ballots be interpreted for only one candidate, which can be ensured by injectivity.

Definition 7 (HK-Injectivity). *An election scheme (Setup, Vote, Tally) satisfies honest-key injectivity (HK-Injectivity), if for all probabilistic polynomial-time adversaries \mathcal{A} , security parameters κ and computations $(pk, sk, mb, mc) \leftarrow \text{Setup}(\kappa); (nc, v, v') \leftarrow \mathcal{A}(pk, \kappa); b \leftarrow \text{Vote}(pk, nc, v, \kappa); b' \leftarrow \text{Vote}(pk, nc, v', \kappa)$ such that $v \neq v' \wedge b \neq \perp \wedge b' \neq \perp$, we have $b \neq b'$.*

Equipped with notions of injectivity and of tallying ballots correctly, we formalise a soundness condition asserting that an election scheme tallies ballots correctly (Definition 8), which allows us to formally state that tallying cannot harm ballot independence (Proposition 9).

Definition 8 (Tally-Soundness). *Let $\Gamma = (\text{Setup}, \text{Vote}, \text{Tally})$ be an election scheme, \mathcal{A} be an adversary, κ be a security parameter, and Tally-Soundness be the following game.*

²⁴A more general result also holds: (Setup, Vote, Tally) satisfies ballot secrecy iff (Setup, Vote, Tally') satisfies ballot secrecy, assuming algorithms Tally and Tally' are indistinguishable, in particular, they tally ballots in the same way. However, election schemes that tally ballots incorrectly are not useful, so we forgo generality for practicality.

²⁵Variable x is bound by the quantifier and integer ℓ is free.

Tally-Soundness($\Gamma, \mathcal{A}, \kappa$) =

$(pk, sk, mb, mc) \leftarrow \text{Setup}(\kappa);$
 $(nc, \mathbf{bb}) \leftarrow \mathcal{A}(pk, \kappa);$
 $(\mathbf{v}, pf) \leftarrow \text{Tally}(sk, \mathbf{bb}, nc, \kappa);$
return $\mathbf{v} = \text{correct-outcome}(pk, nc, \mathbf{bb}, \kappa) \wedge nc \leq mc \wedge |\mathbf{bb}| \leq mb;$

We say Γ satisfies tally soundness (Tally-Soundness), if Γ satisfies HK-Injectivity and for all probabilistic polynomial-time adversaries \mathcal{A} , there exists a negligible function negl , such that for all security parameters κ , we have $\text{Succ}(\text{Tally-Soundness}(\Gamma, \mathcal{A}, \kappa)) > 1 - \text{negl}(\kappa)$.

Lemma 8. Tally-Soundness *implies* Additivity.

Proposition 9. Let $\Gamma = (\text{Setup}, \text{Vote}, \text{Tally})$ and $\Gamma' = (\text{Setup}, \text{Vote}, \text{Tally}')$ be election schemes. Suppose Γ and Γ' satisfy Tally-Soundness. We have Γ satisfies IND-CVA iff Γ' satisfies IND-CVA.

Proof. Tally soundness assures us that algorithms Tally and Tally' produce indistinguishable election outcomes, thus they are interchangeable in the context of game IND-CVA. \square

It follows from Proposition 9 that tally soundness suffices for ballot independence of scheme (Setup, Vote, Tally), if there exists an algorithm Tally' such that (Setup, Vote, Tally') is an election scheme satisfying tally soundness and ballot independence. We demonstrate the existence of such an algorithm with respect to election scheme Enc2Vote,²⁶ thereby showcasing the applicability of Proposition 9 for a class of encryption-based election schemes.

Definition 9 (Enc2Vote). Given an asymmetric encryption scheme $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$, we define $\text{Enc2Vote}(\Pi) = (\text{Setup}, \text{Vote}, \text{Tally})$ such that:

- $\text{Setup}(\kappa)$ computes $(pk, sk, \mathbf{m}) \leftarrow \text{Gen}(\kappa); pk' \leftarrow (pk, \mathbf{m}); sk' \leftarrow (pk, sk)$, derives mc as the largest integer such that $\{0, \dots, mc\} \subseteq \{0\} \cup \mathbf{m}$ and for all $m_0, m_1 \in \{1, \dots, mc\}$ we have $|m_0| = |m_1|$, and outputs $(pk', sk', p(\kappa), mc)$, where p is a polynomial function.
- $\text{Vote}(pk', v, nc, \kappa)$ parses pk' as pair (pk, \mathbf{m}) , outputting \perp if parsing fails or $v \notin \{1, \dots, nc\} \vee \{1, \dots, nc\} \not\subseteq \mathbf{m}$, computes $b \leftarrow \text{Enc}(pk, v)$, and outputs b .
- $\text{Tally}(sk', \mathbf{bb}, nc, \kappa)$ initialises \mathbf{v} as a zero-filled vector of length nc , parses sk' as pair (pk, sk) , outputting (\mathbf{v}, \perp) if parsing fails, computes **for** $b \in \mathbf{bb}$ **do** $v \leftarrow \text{Dec}(sk, b)$; **if** $1 \leq v \leq nc$ **then** $\mathbf{v}[v] \leftarrow \mathbf{v}[v] + 1$, and outputs (\mathbf{v}, ϵ) , where ϵ is a constant symbol.

²⁶Our presentation of Enc2Vote extends the presentation by Quaglia & Smyth [QS18b, Definition 7] to make the plaintext space explicit. We also embed the public key inside the private key. (Quaglia & Smyth's formalisation of Enc2Vote builds upon constructions by Bernhard *et al.* [SB14, SB13, BPW12b, BCP⁺11].)

To ensure $\text{Enc2Vote}(\Pi)$ is an election scheme, we require asymmetric encryption scheme Π to produce distinct ciphertexts with overwhelming probability [Smy18a, Lemma 3]. Hence, we must restrict the class of asymmetric encryption schemes used to instantiate Enc2Vote . We could consider a broad class of schemes that produce distinct ciphertexts with overwhelming probability, but we favour the narrower class of schemes satisfying IND-PA0, since we require IND-PA0 for ballot secrecy.

Lemma 10. *Given an asymmetric encryption scheme Π satisfying IND-CPA, we have $\text{Enc2Vote}(\Pi)$ is an election scheme. Moreover, if Π has perfect correctness, then $\text{Enc2Vote}(\Pi)$ satisfies HK-Injectivity.*

A proof of Lemma 10 follows from Smyth’s correctness proof [Smy18a, Lemma 4] and Quaglia & Smyth’s proof of a slightly stronger notion of HK-Injectivity [QS18b, Lemma 2].

Intuitively, given a non-malleable asymmetric encryption scheme Π , election scheme $\text{Enc2Vote}(\Pi)$ derives ballot secrecy from Π until tallying and tallying maintains ballot secrecy by returning only the number of votes for each candidate. Smyth presents a formal proof of ballot secrecy [Smy18a, Proposition 5], hence, ballot independence is satisfied too (Theorem 1).

Corollary 11. *Given an asymmetric encryption scheme Π satisfying IND-PA0, we have $\text{Enc2Vote}(\Pi)$ satisfies IND-CVA.*

The reverse implication of Corollary 11 does not hold. Indeed, we have the following (stronger) result.

Proposition 12. *There exists an asymmetric encryption scheme Π such that $\text{Enc2Vote}(\Pi)$ satisfies Ballot-Secrecy, but Π does not satisfy IND-PA0.*

To capitalise on Proposition 9, we demonstrate that Enc2Vote produces election schemes satisfying tallying soundness (Lemma 13), assuming “ill-formed” ciphertexts are distinguishable from “well-formed” ciphertexts, and combine our results to derive Theorem 14.

Definition 10. *Given an asymmetric encryption scheme $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$, we say Π satisfies well-definedness, if for all probabilistic polynomial-time adversaries \mathcal{A} , there exists a negligible function negl , such that for all security parameters κ , we have $\Pr[(pk, sk, \mathbf{m}) \leftarrow \text{Gen}(\kappa); c \leftarrow \mathcal{A}(pk, \mathbf{m}, \kappa) : \text{Dec}(sk, c) \neq \perp \Rightarrow \exists m, r. m \in \mathbf{m} \wedge c = \text{Enc}(pk, m; r) \wedge c \neq \perp] > 1 - \text{negl}(\kappa)$.*

Lemma 13. *Given a perfectly-correct asymmetric encryption scheme Π satisfying well-definedness and IND-CPA, we have $\text{Enc2Vote}(\Pi)$ satisfies Tally-Soundness.*

Theorem 14. *Let Π be an asymmetric encryption scheme, $\text{Enc2Vote}(\Pi) = (\text{Setup}, \text{Vote}, \text{Tally})$, and $\Gamma = (\text{Setup}, \text{Vote}, \text{Tally}')$ for some algorithm Tally' such that Γ is an election scheme with zero-knowledge tallying proofs. Suppose Π is perfectly correct and satisfies IND-PA0 and well-definedness. Further suppose Γ satisfies Tally-Soundness. We have Γ satisfies Ballot-Secrecy.*

Proof. Election scheme $\text{Enc2Vote}(\Pi)$ satisfies Tally-Soundness (Lemma 13) and IND-CVA (Corollary 11). Thus, Γ satisfies IND-CVA (Proposition 9) and Ballot-Secrecy (Theorem 4 & Lemma 8). \square

We show that tally soundness is implied by universal verifiability in Appendix E. Thus, a special case of Theorem 14 requires universal verifiability rather than tally soundness. It follows that ballot secrecy is satisfied by verifiable election schemes that produce ballots by encrypting votes with asymmetric encryption schemes satisfying well-definedness and IND-PA0. Thereby making proofs of ballot secrecy trivial for a class of encryption-based election schemes. Indeed, we exploit this result in the following section to show that the combination of non-malleable encryption and universal verifiability suffice for ballot secrecy.

6 Case study II: Helios Mixnet

Helios Mixnet can be informally modelled as the following election scheme:

Setup generates a key pair for an asymmetric homomorphic encryption scheme, proves correct key generation in zero-knowledge, and outputs the key pair and proof.

Vote enciphers the vote to a ciphertext, proves correct ciphertext construction in zero-knowledge, and outputs the ciphertext coupled with the proof.

Tally selects ballots from the bulletin board for which proofs hold, mixes the ciphertexts in those ballots, decrypts the ciphertexts output by the mix to reveal the election outcome (i.e., the frequency distribution of votes), and announces that outcome, along with zero-knowledge proofs demonstrating correct decryption.

Neither Adida [Adi08] nor Bulens, Giry & Pereira [BGP11] have released an implementation of Helios Mixnet.²⁷ Tsoukalas *et al.* [TPLT13] released *Zeus* as a fork of Helios spliced with mixnet code to derive an implementation,²⁸ and Yingtong Li released *helios-server-mixnet* as an extension of Zeus with threshold asymmetric encryption and some other minor changes.^{29,30}

We can derive Helios Mixnet from $\text{Enc2Vote}(\Pi)$ by replacing its tallying algorithm, and by using an asymmetric encryption scheme $\Pi = (\text{Gen}, \text{Enc},$

²⁷The planned implementation of Helios Mixnet (<https://web.archive.org/web/20160912182802/https://documentation.heliosvoting.org/verification-specs/mixnet-support>, published c. 2010, accessed 30 Apr 2018, & <https://web.archive.org/web/20110119223848/http://documentation.heliosvoting.org/verification-specs/helios-v3-1>, published Dec 2010, accessed 30 Apr 2018) has not been released.

²⁸<https://github.com/grnet/zeus>, accessed 15 Sep 2017.

²⁹<https://github.com/RunasSudo/helios-server-mixnet>, accessed 15 Sep 2017.

³⁰The problem of malleable Helios ballots (§4) was discussed with the developers of Zeus and helios-server-mixnet, and they explained that their systems use non-malleable ballots (Email communication, Oct & Dec 2017).

Dec), where algorithm **Gen** proves correct key generation, and algorithm **Enc** verifies such proofs, enciphers plaintexts to ciphertexts using a second encryption scheme, proves correct ciphertext construction, and outputs the ciphertext coupled with the proof. However, a blight arises when **Enc2Vote** is instantiated with encryption schemes that prove correct key generation. To avoid this blight, we extend **Enc2Vote** with such proofs and show that results in Section 5 still hold (Appendix F). This leads us to treat Helios Mixnet as an election scheme built from asymmetric encryption schemes $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ and $\Pi_0 = (\text{Gen}, \text{Enc}', \text{Dec}')$ such that:

- **Setup**(κ) selects coins s uniformly at random, computes $(pk, sk, \mathbf{m}) \leftarrow \text{Gen}(\kappa; s)$ and a proof ρ of correct key generation using sk and s as the witness, derives mc as the largest integer such that $\{0, \dots, mc\} \subseteq \{0\} \cup \mathbf{m}$, computes $pk' \leftarrow (pk, \mathbf{m}, \rho)$; $sk' \leftarrow (pk, sk)$, and outputs $(pk', sk', p(\kappa), mc)$, where p is a polynomial function.
- **Vote**(pk, v, nc, κ) parses pk' as a vector (pk, \mathbf{m}, ρ) , outputting \perp if parsing fails, ρ does not verify, $v \notin \{1, \dots, nc\}$, or $\{0, \dots, nc\} \not\subseteq \mathbf{m}$, computes $b \leftarrow \text{Enc}(pk, v)$, and outputs b .

where

- **Enc**(pk, v) selects coins r uniformly at random, computes $c \leftarrow \text{Enc}'(pk, v; r)$ and a proof σ of correct ciphertext construction using v and r as the witness, and outputs (c, σ) .
- **Dec**(sk, b) parses b as a pair (c, σ) , outputting \perp if parsing fails or σ does not verify, computes $v \leftarrow \text{Dec}'(sk, c)$, and outputs v .

It follows that our results (§5) can be applied to trivialise a ballot-secrecy proof: It is known that Π is a non-malleable encryption scheme [BPW12a, Theorem 2], assuming the proof system used by algorithm **Enc** satisfies simulation sound extractability and Π_0 satisfies IND-CPA. Moreover, we have Π satisfies well-definedness, by the former assumption. Furthermore, Smyth has shown that universal verifiability is satisfied [Smy18b], hence, Tally-Soundness is satisfied too. Thus, Ballot-Secrecy is satisfied. Thereby providing evidence that our results do indeed make ballot-secrecy proofs trivial.

To formally state our ballot secrecy result, we adopt a set **HeliosM'17** of election schemes that includes Helios Mixnet and prove ballot secrecy for every scheme in that set.

Theorem 15. *Each election scheme in set **HeliosM'17** satisfies Ballot-Secrecy.*

A proof of Theorem 15 along with a definition of **HeliosM'17** appear in Appendix G.

7 Stronger privacy notions

Ballot secrecy formalises a notion of free-choice assuming ballots are constructed and tallied in the prescribed manner. Moreover, our definition of ballot secrecy assumes the adversary’s capabilities are limited to casting ballots on behalf of some voters and controlling the votes cast by any remaining voters. We have seen that Helios’16 satisfies our definition, but ballot secrecy does not ensure free-choice when an adversary is able to communicate with voters nor when voters deviate from the prescribed voting procedure to follow instructions provided by an adversary. Indeed, the coins used for encryption serve as proof of how a voter voted in Helios and the voter may communicate those coins to the adversary. Stronger notions of free-choice, such as receipt-freeness [DKR06, MN06, KZZ15, CCFG16] and coercion resistance [JCJ05, GGR09, UM10, KTV12a], are needed in the presence of such adversaries.

Receipt-freeness formalises a notion of free-choice in the presence of an adversary that can communicate with voters.

- Receipt-freeness. A voter cannot collaborate with a conspirator to produce information which can be used to prove how they voted.

Free-choice may be compromised in receipt-free voting systems if voters deviate from the prescribed voting procedure.³¹ Coercion-resistance formalises a stronger notion of free-choice assuming that not only can voters deviate, but the adversary can instruct voters how to deviate.

- Coercion resistance. A voter can deviate from a coercer’s instructions, to cast their vote, without detection.

The distinction between receipt-freeness and coercion resistance is subtle: “receipt-freeness deals with a coercer who is only concerned with deducing information about how someone voted from receipts and public information, but who does not give detailed instructions on how to cast the vote. Coercion resistance, on the other hand, includes dealing with a coercer who gives details not just on which candidate to vote for but also on how to cast the vote” [HS12, §1.1]. Both receipt-freeness and coercion resistance retain the assumption that voters’ ballots are tallied in the prescribed manner, and receipt-freeness additionally assumes voters’ ballots are constructed in the prescribed manner. Moreover, both require side conditions to exclude inevitable revelations (§3).

7.1 Receipt-freeness

We cast the definition of receipt-freeness by Delaune, Kremer & Ryan [DKR06, DKR09] from the symbolic model to the computational model of cryptography,

³¹For receipt-freeness to be an effective notion of free-choice it might be necessary for the voting system to prevent voters deviating from the prescribed voting procedure. This might be achieved by physically securing devices that can be used to cast ballots.

in the context of our election scheme syntax. Moreover, we extend their work to consider arbitrarily many voters, rather than just the two considered by the original definition. The resulting formalisation (Definition 11) is a pair of games.

The first game (**Receipt-Freeness-A**) is an extension of **Ballot-Secrecy** that tasks the adversary to: select a list of their preferred votes and a list of voter preferred votes (although it is somewhat unnatural for the adversary to specify voter preferred votes, this is useful to quantify over all possible voter preferences); construct a bulletin board from either (i) ballots for their preferred votes and coins used to construct those ballots, or (ii) ballots for voter preferred votes and simulated coins; and non-trivially determine whether their preferred or voter preferred votes were used, from the resulting election outcome and tallying proof. The game proceeds as per game **Ballot-Secrecy**, except a new oracle is used (Line 5). That oracle constructs ballots for either the adversary's preferred votes (using algorithm **Vote**) or for voter preferred votes (using algorithm **V**). Those ballots are output along with either the coins used by algorithm **Vote** or coins simulated by algorithm **V**. Intuitively, algorithm **V** provides a strategy for voters to cast voter preferred votes whilst convincing the adversary that its preferred votes were cast, hence, information cannot be produced to prove how voters voted.

The second game (**Receipt-Freeness-B**) tasks the adversary to compute inputs (including a public key) to algorithms **Vote** and **V** that cause the algorithms to output ballots that can be distinguished, thereby over-approximating the requirement that algorithm **V** must produce a ballot for voter preferred votes. (Indeed, the adversary can use algorithm **Tally** to determine whether a ballot is for the expected vote.) The game proceeds as follows: The adversary computes inputs to algorithms **Vote** and **V** (Line 1); the challenger flips a coin (Line 2) and computes a ballot using one of the algorithms (Lines 3–6), where the choice between algorithms is determined by the coin flip; and the adversary is tasked with determining the result of the coin flip from the ballot (Lines 7 & 8).

Definition 11 (Receipt-Freeness). *Let $\Gamma = (\text{Setup}, \text{Vote}, \text{Tally})$ be an election scheme, \mathcal{V} be an algorithm, \mathcal{A} be an adversary, κ be a security parameter, and Receipt-Freeness-A be the following game.*

Receipt-Freeness-A($\Gamma, \mathcal{V}, \mathcal{A}, \kappa$) =

- 1 $(pk, sk, mb, mc) \leftarrow \text{Setup}(\kappa)$;
- 2 $nc \leftarrow \mathcal{A}(pk, \kappa)$;
- 3 $\beta \leftarrow_R \{0, 1\}$;
- 4 $L \leftarrow \emptyset$;
- 5 $\mathbf{bb} \leftarrow \mathcal{A}^\mathcal{O}()$;
- 6 $(\mathbf{v}, pf) \leftarrow \text{Tally}(sk, \mathbf{bb}, nc, \kappa)$;
- 7 $g \leftarrow \mathcal{A}(\mathbf{v}, pf)$;
- 8 **return** $g = \beta \wedge \text{balanced}(\mathbf{bb}, nc, L) \wedge 1 \leq nc \leq mc \wedge |\mathbf{bb}| \leq mb$;

Oracle \mathcal{O} is defined as follows:

- $\mathcal{O}(v_0, v_1)$ chooses coins r uniformly at random from the coin space of algorithm **Vote**, computes **if** $\beta = 0$ **then** $b \leftarrow \text{Vote}(pk, v_0, nc, \kappa; r)$ **else**

$(b, r) \leftarrow V(pk, v_1, v_0, nc, \kappa)$ and $L \leftarrow L \cup \{(b, v_0, v_1)\}$, and outputs (b, r) , where $v_0, v_1 \in \{1, \dots, nc\}$.

Moreover, let \mathcal{B} be an adversary and Receipt-Freeness-B be the following game.

Receipt-Freeness-B($\Gamma, \mathcal{V}, \mathcal{B}, \kappa$) =

```

1  $(pk, v_0, v_1, nc) \leftarrow \mathcal{B}(\kappa)$ ;
2  $\beta \leftarrow_R \{0, 1\}$ ;
3 if  $\beta = 0$  then
4   |  $b \leftarrow \text{Vote}(pk, v_0, nc, \kappa)$ 
5 else
6   |  $(b, r) \leftarrow V(pk, v_0, v_1, nc, \kappa)$ 
7  $g \leftarrow \mathcal{B}(b)$ ;
8 return  $g = \beta \wedge 1 \leq v_0, v_1 \leq nc$ ;

```

We say Γ satisfies Receipt-Freeness, if there exists a probabilistic polynomial-time algorithm \mathcal{V} such that for all probabilistic polynomial-time adversaries \mathcal{A} and \mathcal{B} , there exists a negligible function negl and for all security parameters κ , we have $\text{Succ}(\text{Receipt-Freeness-A}(\Gamma, \mathcal{V}, \mathcal{A}, \kappa)) \leq \frac{1}{2} + \text{negl}(\kappa)$ and $\text{Succ}(\text{Receipt-Freeness-B}(\Gamma, \mathcal{V}, \mathcal{B}, \kappa)) \leq \frac{1}{2} + \text{negl}(\kappa)$.

An election scheme satisfies receipt-freeness when ballot secrecy is preserved even when coins used to construct ballots are revealed.

Similarly to ballot secrecy (§3), receipt-freeness tolerates inevitable revelations, e.g., unanimous election outcomes. It follows that an election scheme for one candidate satisfies our definition of receipt-freeness, because that scheme will always produce unanimous election outcomes.

Proposition 16. *Given an election scheme Γ for one candidate (i.e., for all security parameters the maximum number of candidates is one), we have Γ satisfies Ballot-Secrecy and Receipt-Freeness.*

Beyond the uninteresting case (Proposition 16), we prove that Receipt-Freeness cannot be satisfied by election schemes for more than one candidate (Theorem 17), assuming the election scheme is perfectly correct or satisfies tally soundness.

Theorem 17. *Let Γ be an election scheme for more than one candidate (i.e., there exists a security parameter such that the maximum number of candidates is greater than one). Suppose Γ is perfectly correct or Γ satisfies Tally-Soundness. We have Γ does not satisfy Receipt-Freeness.*

A special case of our theorem holds for universal verifiability, rather than tally soundness, because universal verifiability is strictly stronger (Appendix E).

It follows from Theorem 17 that our syntax cannot be used to construct (interesting) election schemes satisfying stronger notions of privacy such as Receipt-Freeness. Nonetheless, algorithm `Vote` could be distributed between the voter and some other (possibly untrusted) parties to enable stronger privacy notions. Indeed, a variant of Helios that delegates ballot construction to a trusted

party would satisfy receipt-freeness, since voters cannot access coins used by the trusted party to construct ballots, hence, cannot communicate such coins to the adversary. More generally, an election scheme satisfying ballot secrecy can delegate ballot construction to achieve receipt-freeness.³²

Definition 12. *An election scheme with an interactive voting algorithm is a tuple of probabilistic polynomial-time algorithms $(\text{Setup}, \text{VoteClient}, \text{VoteServer}, \text{Tally})$ such that $(\text{Setup}, \text{VoteClient}^{\text{VoteServer}}, \text{Tally})$ is an election scheme.*

Definition 13. *An election scheme with an interactive voting algorithm $(\text{Setup}, \text{VoteClient}, \text{VoteServer}, \text{Tally})$ satisfies Receipt-Freeness, if election scheme $(\text{Setup}, \text{VoteClient}^{\text{VoteServer}}, \text{Tally})$ satisfies Receipt-Freeness.*

Receipt-freeness for election schemes with interactive voting algorithms must hold when coins used by algorithm VoteClient are revealed (coins used by algorithm VoteServer are not revealed),³³ whereas receipt-freeness for (regular) election schemes must hold when coins used by algorithm Vote are revealed.

Definition 14. *Given an election scheme $\Gamma = (\text{Setup}, \text{Vote}, \text{Tally})$, we define $\text{BS2RF}(\Gamma) = (\text{Setup}, \text{VoteClient}, \text{VoteServer}, \text{Tally})$ such that*

- $\text{VoteClient}(pk, v, nc, \kappa)$ computes $b \leftarrow \text{VoteServer}(pk, v, nc, \kappa)$ and outputs b .
- $\text{VoteServer}(pk, v, nc, \kappa)$ computes $b \leftarrow \text{Vote}(pk, v, nc, \kappa)$ and outputs b .

Proposition 18. *Given an election scheme Γ satisfying Ballot-Secrecy, we have $\text{BS2RF}(\Gamma)$ satisfies Receipt-Freeness.*

It follows from Proposition 18 that ballot construction can be delegated to a trusted party to achieve receipt-freeness.

We have seen that election schemes for more than one candidate cannot satisfy Receipt-Freeness (Theorem 17), assuming the election scheme is perfectly correct or satisfies tally soundness. Intuitively, it follows that coercion resistance cannot be satisfied by such election schemes either, because coercion resistance strengthens receipt-freeness. Nevertheless, we have seen that Receipt-Freeness can be satisfied by election schemes with interactive voting algorithms (Proposition 18) and we now consider whether coercion resistance can hold too.

7.2 Coercion resistance

Coercion resistance requires a mechanism to deviate from a coercer's instructions. Intuitively, no such mechanism exists for election schemes with interactive

³²Similarly to oracle notation, we write $\text{VoteClient}^{\text{VoteServer}}$ to denote algorithm VoteClient 's access to algorithm VoteServer .

³³Formally, given an election scheme $(\text{Setup}, \text{VoteClient}^{\text{VoteServer}}, \text{Tally})$, the oracle in game Receipt-Freeness chooses coins r uniformly at random from the coin space of algorithm VoteClient and computes the if-then-else branch as **if** $\beta = 0$ **then** $b \leftarrow \text{VoteClient}^{\text{VoteServer}}(pk, v_0, nc, \kappa; r)$ **else** $(b, r) \leftarrow V(pk, v_1, v_0, nc, \kappa)$.

voting algorithms, because there exists instructions for which deviations are impossible. Indeed, the coercer can instruct the voter to run algorithms `VoteClient` and `VoteServer` themselves (i.e., without the aid of another party) and furnish the coercer with the coins used, thereby enabling the coercer to check whether voters followed their instructions. It follows that our syntax cannot be used to construct (interesting) schemes satisfying coercion resistance.

Algorithm `VoteServer` can be extended with private inputs which cannot be simulated by voters, thereby enabling voters to deviate from the coercer’s instructions. Alternatively, the variant of our syntax with voter credentials [SFC17, Definition 6] can be used; that variant extends algorithm `Vote` to input a private credential (essentially resulting in interactive ballot construction). The latter has been used by Smyth, Frink & Clarkson [SFC17, §6] to model the coercion-resistant voting system by Juels, Catalano & Jakobsson [JCJ05, JCJ10], thus, the syntax with voter credentials is compatible with stronger privacy notions.

8 Related work

Discussion of ballot secrecy originates from Chaum [Cha81] and the earliest definitions of ballot secrecy are due to Benaloh *et al.* [BY86, BT94, Ben96].³⁴ More recently, Bernhard *et al.* propose a series of ballot secrecy definitions: They consider election schemes without tallying proofs [BCP⁺11, BPW12b] and, subsequently, schemes with tallying proofs [BPW12a, SB13, SB14, BCG⁺15b]. The definition by Bernhard, Pereira & Warinschi computes tallying proofs using algorithm `Tally` or a simulator [BPW12a], but that definition is too weak and some strengthening is required [BCG⁺15b, §3.4 & §4]. (Cortier *et al.* [CGGI13a, CGGI13b] propose a variant of the ballot secrecy definition by Bernhard, Pereira & Warinschi, which is also too weak [BCG⁺15b].) By comparison, the definition by Smyth & Bernhard computes tallying proofs using only algorithm `Tally` [SB13], but their definition is too strong [BCG⁺15b, §3.5] and a weakening is required [SB14]. We prove that our ballot secrecy definition is strictly stronger than the weakened definition by Smyth & Bernhard (Appendix H). The relative merits of definitions by Smyth & Bernhard [SB14, Definition 5] and by Bernhard *et al.* [BCG⁺15b, Definition 7] are unknown, in particular, it is unknown whether one definition is stronger than the other.

Discussion of ballot independence originates from Gennaro [Gen95] and the apparent relationship between ballot secrecy and ballot independence has been considered. In particular, Benaloh shows that a simplified version of his voting system allows the administrator’s private key to be recovered by an adversary who casts a ballot as a function of other voters’ ballots [Ben96, §2.9]. More generally, Sako & Kilian [SK95, §2.4], Michels & Horster [MH96, §3], Wikström [Wik06, Wik08, Wik16] and Cortier & Smyth [CS13, CS11] discuss how malleable ballots can be abused to compromise ballot secrecy. The first definition of ballot independence seems to be due to Smyth & Bernhard [SB13, SB14].

³⁴Quaglia & Smyth present a tutorial-style introduction to modelling ballot secrecy [QS18c], and Bernhard *et al.* survey ballot secrecy definitions [BCG⁺15b, BCG⁺15a].

Moreover, Smyth & Bernhard formally prove relations between their definitions of ballot secrecy and ballot independence. Independence has also been studied beyond elections, e.g., [CGMA85], and the possibility of compromising security in the absence of independence has been considered, e.g., [CR87, PP89, Pfi94, DDN91, DDN00, Gen00].

All of the ballot secrecy definitions by Bernhard *et al.* [BCP⁺11, BPW12b, BPW12a, SB13, SB14, BCG⁺15b] and the ballot independence definition by Smyth & Bernhard [SB13, SB14] focus on detecting vulnerabilities exploitable by adversaries that control some voters. Vulnerabilities that require control of the bulletin board or the communication channel are not detected, i.e., the bulletin board is implicitly assumed to operate in accordance with the election scheme’s rules and the communication channel is implicitly assumed to be secure. This introduces a trust assumption. Under this assumption, Smyth & Bernhard prove that non-malleable ballots are not necessary for ballot secrecy [SB13, §4.3], and Bernhard, Pereira & Warinschi [BPW12a], Bernhard [Ber14] and Bernhard *et al.* [BCG⁺15a, BCG⁺15b] prove that Helios’12 satisfies various notions of ballot secrecy. By comparison, we prove that non-malleable ballots are necessary for ballot secrecy without this trust assumption. Hence, Helios’12 does not satisfy our definition of ballot secrecy. Thus, our definition of ballot secrecy improves upon definitions by Bernhard *et al.* by detecting more vulnerabilities.

Confidence in our ballot secrecy definition might be improved by proving equivalence with a simulation-based definition of ballot secrecy. However, it is unclear how to formulate a suitable simulation-based definition. Bernhard *et al.* propose an ideal functionality that “collects all votes from the voters, then computes and announces the [election outcome]” [BCG⁺15b, §1],³⁵ but a voting system satisfying ballot secrecy need not be equivalent, because ballot secrecy does not guarantee correct computation of the election outcome. Equivalence can perhaps be shown between their ideal functionality and voting systems satisfying ballot secrecy and some soundness condition (e.g., Tally-Soundness). Albeit, voting systems that bound the number of ballots or candidates, e.g., Helios, may not be equivalent, because soundness conditions (such as Tally-Soundness) need only provide guarantees when operating within the aforementioned bounds. Thus, bounds need to be taken into consideration. Moreover, no voting system is equivalent to the ideal functionality by Bernhard *et al.* in the presence of an adversary that controls the bulletin board, because such an adversary can discard ballots, which creates a trivial distinction.³⁶ Nonetheless, equivalence can perhaps be shown for verifiable voting systems in cases where the number of ballots and candidates are bounded, and all voters successfully verify the presence of their ballot, but such a result has no practical relevance,

³⁵In the context of voting systems that announce the chosen representative (rather than the frequency distribution of votes), a stronger ideal functionality might announce the chosen representative.

³⁶The real functionality by Bernhard *et al.* does not capture adversaries that control ballot collection. Thus, the relation they prove between their game-based and simulation-based definitions of ballot secrecy does not preclude vulnerabilities exploitable by such adversaries. Indeed, proving such relations does not guarantee the absence of vulnerabilities.

because it is well-known that many voters do not perform checks necessary for verifiability [BBH⁺17, §2.1.6]. Alternatively, the ideal functionality could be weakened such that a discarded ballot in the real functionality corresponds to a discarded vote in the ideal functionality. We leave further exploration of simulation-based definitions of ballot secrecy as a possible extension for future work.

Bulens, Giry & Pereira pose the investigation of voting systems which allow casting of meaningfully related ballots as a means for voters to delegate candidate selection, whilst preserving ballot secrecy, as an interesting research question [BGP11, §3.2]. Desmedt & Chaidos claim to provide such a system [DC12]. Smyth & Bernhard critique their work and argue that the security results do not support their claims [SB13, §5.1]. We have shown that non-malleable ballots are necessary to satisfy **Ballot-Secrecy** (§3.3), providing a negative result for the question posed by Bulens, Giry & Pereira.

Some of the ideas presented in this article previously appeared in my technical report [Smy14] and an extended version of that technical report by Bernhard & Smyth [BS15]. In particular, the limitations of ballot secrecy definitions by Bernhard *et al.* were identified and Definition 2 is based upon my earlier definition [Smy14]. The main distinction between Definition 2 and the earlier definition [Smy14, Definition 3] is syntax: This article adopts syntax for election schemes from Smyth, Frink & Clarkson [SFC17], whereas the earlier definition adopts syntax by Smyth & Bernhard [SB14, SB13]. The change in syntax is motivated by the superiority of syntax by Smyth, Frink & Clarkson. Moreover, we can capitalise upon results by Smyth, Frink & Clarkson [SFC17] and Quaglia & Smyth [QS18b].

Following the initial release of these results [Smy15, Smy16], Cortier *et al.* [CSD⁺17] presented a machine-checked proof that variants of Helios satisfy the notion of ballot secrecy by Bernhard *et al.* [BCG⁺15b]. As discussed above, that notion is too weak: It cannot detect vulnerabilities that require control of ballot collection. Thus, our proof is more appropriate. Nonetheless, their proof builds upon similar ideas. In particular, their proof is dependent upon non-malleable ballots and zero-knowledge tallying proofs.

Quaglia & Smyth [QS18a] extend game **Ballot-Secrecy** to syntax with voter credentials [SFC17, Definition 6]. Moreover, they define a transformation to that syntax from our syntax (Definition 1) and prove that their transformation preserves secrecy and verifiability. Furthermore, they apply their transformation to Helios. The resulting scheme is similar Helios-C [CGGI14], albeit Quaglia & Smyth observe that Helios-C does not satisfy ballot secrecy when the adversary controls ballot collection, whereas the scheme derived by applying their transformation does.

Beyond the computational model of security, Delaune, Kremer & Ryan formulate a definition of ballot secrecy in the applied pi calculus [DKR06, DKR09] and Smyth *et al.* show that this definition is amenable to automated reasoning [DRS08, KSR10, Smy11, BS16, BS18]. An alternative definition is proposed by Cremers & Hirschi, along with sufficient conditions which are also amenable to automated reasoning [CH17]. Albeit, the scope of automated reasoning is

limited by analysis tools (e.g., ProVerif [BSCS16]), because the function symbols and equational theory used to model cryptographic primitives might not be suitable for automated analysis (cf. [DKRS11, PB12, ABR12]).

Ballot secrecy does not provide assurances when deviations from the prescribed tallying procedure are possible. Indeed, ballots can be tallied individually to reveal votes. Hence, the tallier must be trusted. Alternatively, we can design election schemes that distribute the tallier’s role amongst several talliers and ensure free-choice assuming at least one tallier tallies ballots in the prescribed manner. Extending our results in this direction is an opportunity for future work. Ultimately, we would prefer not to trust talliers. Unfortunately, this is only known to be possible for decentralised voting systems, e.g., [Sch99, KY02, Gro04, HRZ10, KSRH12, KRA18], which are designed such that ballots cannot be tallied individually, but are unsuitable for large-scale elections.

McCarthy, Smyth & Quaglia [MSQ14] have shown that auction schemes can be constructed from election schemes, and Quaglia & Smyth [QS18b] provide a generic construction for auction schemes from election schemes. Moreover, Quaglia & Smyth adapt our definition of ballot secrecy to a definition of bid secrecy, and prove that auction schemes produced by their construction satisfy bid secrecy. (Similarly, they adapt the definition of verifiability by Smyth, Frink & Clarkson [SFC17] to a definition of verifiability for auctions, and prove that their construction produces schemes satisfying verifiability.) Thus, this research has applications beyond voting.

9 Conclusion

This work was initiated by a desire to eliminate the trust assumptions placed upon the bulletin board and the communication channel in definitions of ballot secrecy by Bernhard *et al.* and the definition of ballot independence by Smyth & Bernhard. This necessitated the introduction of new security definitions. The definition of ballot secrecy was largely constructed from intuition, with inspiration from indistinguishability games for asymmetric encryption and existing definitions of ballot secrecy. Moreover, the definition was guided by the desire to strengthen existing definitions of ballot secrecy. The definition of ballot independence was inspired by the realisation that independence requires non-malleable ballots, which enabled definitions to be constructed as straightforward adaptations of non-malleability and indistinguishability definitions for asymmetric encryption. The former adaptation being a more natural formulation of ballot independence and the latter being simpler.

Relationships between security definitions aid our understanding and offer insights that facilitate the construction of secure schemes. This prompted the study of relations between our definitions of ballot secrecy and ballot independence, resulting in a proof that non-malleable ballots are necessary for ballot secrecy. We also proved non-malleable ballots are sufficient for ballot secrecy in election schemes with zero-knowledge tallying proofs. Moreover, we established

a separation result demonstrating that our ballot secrecy definition is strictly stronger than our ballot independence definition.

In light of the revelation that non-malleable ballots are necessary for ballot secrecy, and in the knowledge that ballots are malleable in Helios, it was discovered that Helios'12 does not defend against the vulnerability identified by Cortier & Smyth. We also revealed a new exploit against Helios'12 that allows an adversary to determine if a voter did not vote for the adversary's preferred candidate. This naturally led to the consideration of whether definitions of ballot secrecy by Bernhard *et al.* could detect these vulnerabilities and to the conclusion that they could not, because the vulnerabilities require the adversary to control ballot collection, which is prohibited by those definitions.

We have considered vulnerabilities that are only exploitable by an adversary that controls ballot collection. Hence, the vulnerability in Helios'12 can be vacuously eliminated by trusting the tallier to collect ballots. However, Smyth, Frink & Clarkson have shown that Helios'12 does not satisfy universal verifiability, which must hold without trusting the tallier. Thus, even if we are willing to accept additional trust assumptions to ensure ballot secrecy, we cannot accept such trust assumptions to ensure universal verifiability, because they defy the purpose of verifiability. An alternative solution is necessary and non-malleable ballots are proposed. Moreover, we prove that Helios'16 uses non-malleable ballots and a proof that Helios'16 satisfies ballot secrecy follows from our results. This proof is particularly significant due to the use of Helios in binding elections, and we encourage developers to base future releases on this variant, since it is the only variant of Helios which is proven to satisfy both ballot secrecy and verifiability.

Proving ballot secrecy is expensive: It requires a significant devotion of time by experts. Indeed, Cortier *et al.* devoted one person-year to their machine-checked proof. Thus, sufficient conditions for ballot secrecy are highly desirable. We have established that non-malleable ballots are sufficient for ballot secrecy in election schemes with zero-knowledge tallying proofs and this simplified our ballot-secrecy proof for Helios'16. We have also established that building election schemes from non-malleable asymmetric encryption schemes suffices for ballot secrecy if tallying is additive (a condition implied by verifiability), and this trivialised our ballot-secrecy proof for Helios Mixnet. Thereby demonstrating the possibility of simple, inexpensive proofs.

This article delivers a definition of ballot secrecy that has been useful in detecting subtle vulnerabilities in voting systems, and has led to the development of election schemes that are proven secure. Thereby demonstrating the necessity of appropriate security definitions and accompanying analysis to ensure security of voting systems, especially those used in binding elections. I hope this article will simplify future proofs of ballot secrecy and, ultimately, aid democracy-builders in deploying their systems securely.

Acknowledgements

I am grateful to David Bernhard and to Elizabeth Quaglia for extensive discussions that helped improve this article and, more generally, my knowledge of cryptography. In addition, I am grateful to Constantin Cătălin Drăgan for explaining subtleties of his work, to Maxime Meyer for his careful proofreading, and to the anonymous reviewers who provided constructive criticism. This work received financial support from the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013) / ERC project *CRYSP* (259639) and from the Luxembourg National Research Fund (FNR) under the FNR-INTER-VoteVerif project. Work was partly performed at INRIA.

A Cryptographic primitives

A.1 Asymmetric encryption

Definition 15 (Asymmetric encryption scheme [KL07]). *An asymmetric encryption scheme is a tuple of probabilistic polynomial-time algorithms $(\text{Gen}, \text{Enc}, \text{Dec})$, such that:*³⁷

- **Gen**, denoted $(pk, sk, \mathbf{m}) \leftarrow \text{Gen}(\kappa)$, inputs a security parameter κ and outputs a key pair (pk, sk) and message space \mathbf{m} .
- **Enc**, denoted $c \leftarrow \text{Enc}(pk, m)$, inputs a public key pk and message $m \in \mathbf{m}$, and outputs a ciphertext c .
- **Dec**, denoted $m \leftarrow \text{Dec}(sk, c)$, inputs a private key sk and ciphertext c , and outputs a message m or an error symbol. We assume **Dec** is deterministic.

Moreover, the scheme must be correct: there exists a negligible function negl , such that for all security parameters κ and messages m , we have $\Pr[(pk, sk, \mathbf{m}) \leftarrow \text{Gen}(\kappa); c \leftarrow \text{Enc}(pk, m) : m \in \mathbf{m} \Rightarrow \text{Dec}(sk, c) = m] > 1 - \text{negl}(\kappa)$. A scheme has perfect correctness if the probability is 1.

Definition 16 (Homomorphic encryption [SFC17]). *An asymmetric encryption scheme $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ is homomorphic, with respect to ternary operators \odot , \oplus , and \otimes ,³⁸ if there exists a negligible function negl , such that for all security parameters κ , we have the following.³⁹ First, for all messages m_1 and m_2 we have $\Pr[(pk, sk, \mathbf{m}) \leftarrow \text{Gen}(\kappa); c_1 \leftarrow \text{Enc}(pk, m_1); c_2 \leftarrow \text{Enc}(pk, m_2) : m_1, m_2 \in$*

³⁷Our definition differs from Katz and Lindell’s original definition [KL07, Definition 10.1] in that we formally state the plaintext space.

³⁸For brevity, we write Π is a homomorphic asymmetric encryption scheme as opposed to the more verbose Π is a homomorphic asymmetric encryption scheme, with respect to ternary operators \odot , \oplus , and \otimes .

³⁹We 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.

$\mathbf{m} \Rightarrow \text{Dec}(sk, c_1 \otimes_{pk} c_2) = \text{Dec}(sk, c_1) \odot_{pk} \text{Dec}(sk, c_2)] > 1 - \text{negl}(\kappa)$. Secondly, for all messages m_1 and m_2 , and all coins r_1 and r_2 , we have $\Pr[(pk, sk, \mathbf{m}) \leftarrow \text{Gen}(\kappa) : m_1, m_2 \in \mathbf{m} \Rightarrow \text{Enc}(pk, m_1; r_1) \otimes_{pk} \text{Enc}(pk, m_2; r_2) = \text{Enc}(pk, m_1 \odot_{pk} m_2; r_1 \oplus_{pk} r_2)] > 1 - \text{negl}(\kappa)$. We say Π is additively homomorphic, if for all security parameters κ , key pairs pk, sk , and message spaces \mathbf{m} , such that there exists coins r and $(pk, sk, \mathbf{m}) = \text{Gen}(\kappa; r)$, we have \odot_{pk} is the addition operator in group (\mathbf{m}, \odot_{pk}) .

Definition 17 (IND-CPA [BDPR98]). Let $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ be an asymmetric encryption scheme, \mathcal{A} be an adversary, κ be the security parameter, and IND-CPA be the following game.⁴⁰

IND-CPA(Π, \mathcal{A}, κ) =
 $(pk, sk, \mathbf{m}) \leftarrow \text{Gen}(\kappa);$
 $(m_0, m_1) \leftarrow \mathcal{A}(pk, \mathbf{m}, \kappa);$
 $\beta \leftarrow_R \{0, 1\};$
 $c \leftarrow \text{Enc}(pk, m_\beta);$
 $g \leftarrow \mathcal{A}(c);$
return $g = \beta;$

In the above game, we require $m_0, m_1 \in \mathbf{m}$ and $|m_0| = |m_1|$. We say Π satisfies IND-CPA, if for all probabilistic polynomial-time adversaries \mathcal{A} , there exists a negligible function negl , such that for all security parameters κ , we have $\text{Succ}(\text{IND-CPA}(\Pi, \mathcal{A}, \kappa)) \leq \frac{1}{2} + \text{negl}(\kappa)$.

Definition 18 (IND-PA0 [BS99]). Let $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ be an asymmetric encryption scheme, \mathcal{A} be an adversary, κ be the security parameter, and IND-PA0 be the following game.

IND-PA0(Π, \mathcal{A}, κ) =
 $(pk, sk, \mathbf{m}) \leftarrow \text{Gen}(\kappa);$
 $(m_0, m_1) \leftarrow \mathcal{A}(pk, \mathbf{m}, \kappa);$
 $\beta \leftarrow_R \{0, 1\};$
 $c \leftarrow \text{Enc}(pk, m_\beta);$
 $\mathbf{c} \leftarrow \mathcal{A}(c);$
 $\mathbf{m} \leftarrow (\text{Dec}(sk, \mathbf{c}[1]), \dots, \text{Dec}(sk, \mathbf{c}[|\mathbf{c}|]));$
 $g \leftarrow \mathcal{A}(\mathbf{m});$
return $g = \beta \wedge \bigwedge_{1 \leq i \leq |\mathbf{c}|} c \neq \mathbf{c}[i];$

In the above game, we require $m_0, m_1 \in \mathbf{m}$ and $|m_0| = |m_1|$. We say Π satisfies IND-PA0, if for all probabilistic polynomial-time adversaries \mathcal{A} , there exists a negligible function negl , such that for all security parameters κ , we have $\text{Succ}(\text{IND-PA0}(\Pi, \mathcal{A}, \kappa)) \leq \frac{1}{2} + \text{negl}(\kappa)$.

⁴⁰Our definition of an asymmetric encryption scheme explicitly defines the plaintext space, whereas Bellare *et al.* [BDPR98] leave the plaintext space implicit; this change is reflected in our definition of IND-CPA. Moreover, we provide the adversary with the message space and security parameter. We adapt IND-PA0 similarly.

A.2 Proof systems

Definition 19 (Non-interactive proof system [SFC17]). *A non-interactive proof system for a relation R is a tuple of algorithms (Prove, Verify), such that:*

- **Prove**, denoted $\sigma \leftarrow \text{Prove}(s, w, \kappa)$, is executed by a prover to prove $(s, w) \in R$.
- **Verify**, denoted $v \leftarrow \text{Verify}(s, \sigma, \kappa)$, 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 negl , such that for all statement and witnesses $(s, w) \in R$ and security parameters κ , we have $\Pr[\sigma \leftarrow \text{Prove}(s, w, \kappa) : \text{Verify}(s, \sigma, \kappa) = 1] > 1 - \text{negl}(\kappa)$. A system has perfect completeness if the probability is 1.

Definition 20 (Fiat-Shamir transformation [FS87]). *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 non-interactive proof system (Prove, Verify), defined as follows:*

$\text{Prove}(s, w, \kappa) =$
 $(\text{comm}, t) \leftarrow \text{Comm}(s, w, \kappa);$
 $\text{chal} \leftarrow \mathcal{H}(\text{comm}, s);$
 $\text{resp} \leftarrow \text{Resp}(\text{chal}, t, \kappa);$
return $(\text{comm}, \text{resp});$

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

A string m can be included in the hashes computed by algorithms Prove and Verify. That is, the hashes are computed in both algorithms as $\text{chal} \leftarrow \mathcal{H}(\text{comm}, s, m)$. We write $\text{Prove}(s, w, m, \kappa)$ and $\text{Verify}(s, (\text{comm}, \text{resp}), m, \kappa)$ for invocations of Prove and Verify which include string m .

Definition 21 (Zero-knowledge [QS18b]). *Let $\Delta = (\text{Prove}, \text{Verify})$ be a non-interactive proof system for a relation R , derived by application of the Fiat-Shamir transformation [FS87] to a random oracle \mathcal{H} and a sigma protocol. Moreover, let \mathcal{S} be an algorithm, \mathcal{A} be an adversary, κ be a security parameter, and ZK be the following game.*

$\text{ZK}(\Delta, \mathcal{A}, \mathcal{H}, \mathcal{S}, \kappa) =$
 $\beta \leftarrow_R \{0, 1\};$
 $g \leftarrow \mathcal{A}^{\mathcal{H}, \mathcal{P}}(\kappa);$
return $g = \beta;$

Oracle \mathcal{P} is defined on inputs $(s, w) \in R$ as follows:

- $\mathcal{P}(s, w)$ computes **if** $\beta = 0$ **then** $\sigma \leftarrow \text{Prove}(s, w, \kappa)$ **else** $\sigma \leftarrow \mathcal{S}(s, \kappa)$ **and** outputs σ .

And algorithm \mathcal{S} can patch random oracle \mathcal{H} .⁴¹ We say Δ satisfies zero-knowledge, if there exists a probabilistic polynomial-time algorithm \mathcal{S} , such that for all probabilistic polynomial-time algorithm adversaries \mathcal{A} , there exists a negligible function negl , and for all security parameters κ , we have $\text{Succ}(\text{ZK}(\Delta, \mathcal{A}, \mathcal{H}, \mathcal{S}, \kappa)) \leq \frac{1}{2} + \text{negl}(\kappa)$. An algorithm \mathcal{S} for which zero-knowledge holds is called a simulator for (Prove, Verify).

Definition 22 (Simulation sound extractability [SFC17, BPW12a, Gro06]). Suppose Σ is a sigma protocol for relation R , \mathcal{H} is a random oracle, and (Prove, 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 (Prove, Verify) and \mathcal{H} can be patched by \mathcal{S} . Proof system (Prove, Verify) satisfies simulation sound extractability if there exists a probabilistic polynomial-time algorithm \mathcal{K} , such that for all probabilistic polynomial-time adversaries \mathcal{A} and coins r , there exists a negligible function negl , such that for all security parameters κ , we have:⁴²

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

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

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

Theorem 19 (from [BPW12a]). Let Σ be a sigma protocol for relation R , and let \mathcal{H} be a random oracle. Suppose Σ satisfies special soundness and special honest verifier zero-knowledge. Non-interactive proof system $\text{FS}(\Sigma, \mathcal{H})$ satisfies zero-knowledge and simulation sound extractability.

The Fiat-Shamir transformation may include a string in the hashes computed by functions Prove and Verify. Simulators can be generalised to include such a string too. We write $\mathcal{S}(s, m, \kappa)$ for invocations of simulator \mathcal{S} which include string m . And remark that Theorem 19 can be extended to this generalisation.

⁴¹Random oracles can be *programmed* or *patched*. We will not need the details of how patching works, so we omit them here; see Bernhard et al. [BPW12a] for a formalisation.

⁴²We extend set membership notation to vectors: we write $x \in \mathbf{x}$ if x is an element of the set $\{\mathbf{x}[i] : 1 \leq i \leq |\mathbf{x}|\}$.

B Ballot independence: Non-malleability game

We formalise an alternative definition of ballot independence as a non-malleability game, called comparison based non-malleability under chosen vote attack (CNM-CVA), as a straightforward adaptation of the non-malleability definition for asymmetric encryption by Bellare & Sahai [BS99].

Definition 23 (CNM-CVA). *Let $\Gamma = (\text{Setup}, \text{Vote}, \text{Tally})$ be an election scheme, \mathcal{A} be an adversary, κ be a security parameter, and cnm-cva and $\text{cnm-cva-}\$$ be the following games.*

$$\begin{array}{ll}
 \text{cnm-cva}(\Gamma, \mathcal{A}, \kappa) = & \text{cnm-cva-}\$(\Gamma, \mathcal{A}, \kappa) = \\
 (pk, sk, mb, mc) \leftarrow \text{Setup}(\kappa); & (pk, sk, mb, mc) \leftarrow \text{Setup}(\kappa); \\
 (V, nc) \leftarrow \mathcal{A}(pk, \kappa); & (V, nc) \leftarrow \mathcal{A}(pk, \kappa); \\
 v \leftarrow_R V; & v, v' \leftarrow_R V; \\
 b \leftarrow \text{Vote}(pk, v, nc, \kappa); & b \leftarrow \text{Vote}(pk, v', nc, \kappa); \\
 (R, \mathbf{bb}) \leftarrow \mathcal{A}(b); & (R, \mathbf{bb}) \leftarrow \mathcal{A}(b); \\
 (\mathbf{v}, pf) \leftarrow \text{Tally}(sk, \mathbf{bb}, nc, \kappa); & (\mathbf{v}, pf) \leftarrow \text{Tally}(sk, \mathbf{bb}, nc, \kappa); \\
 \mathbf{return} R(v, \mathbf{v}) \wedge b \notin \mathbf{bb} & \mathbf{return} R(v, \mathbf{v}) \wedge b \notin \mathbf{bb} \\
 \wedge V \subseteq \{1, \dots, nc\} & \wedge V \subseteq \{1, \dots, nc\} \\
 \wedge 1 \leq nc \leq mc \wedge |\mathbf{bb}| \leq mb; & \wedge 1 \leq nc \leq mc \wedge |\mathbf{bb}| \leq mb;
 \end{array}$$

In the above games, we require that relation R is computable in polynomial time. We say Γ satisfies comparison based non-malleability under chosen vote attack (CNM-CVA), if for all probabilistic polynomial-time adversaries \mathcal{A} , there exists a negligible function negl , such that for all security parameters κ , we have $\text{Succ}(\text{cnm-cva}(\Gamma, \mathcal{A}, \kappa)) - \text{Succ}(\text{cnm-cva-}\$(\Gamma, \mathcal{A}, \kappa)) \leq \text{negl}(\kappa)$.

CNM-CVA is satisfied if no adversary can distinguish between games cnm-cva and $\text{cnm-cva-}\$$. That is, for all adversaries, the adversary wins cnm-cva iff the adversary loses $\text{cnm-cva-}\$$, with negligible probability. The first three steps of games cnm-cva and $\text{cnm-cva-}\$$ are identical, thus, these steps cannot be distinguished. Then, game $\text{cnm-cva-}\$$ performs an additional step: the challenger samples a second vote v' from vote space V . Thereafter, game cnm-cva , respectively game $\text{cnm-cva-}\$$, proceeds as follows: the challenger constructs a challenge ballot b for v , respectively v' ; the adversary is given ballot b and computes a relation R and bulletin board \mathbf{bb} ; and the challenger tallies \mathbf{bb} to derive election outcome \mathbf{v} and evaluates whether $R(v, \mathbf{v})$ holds. Hence, CNM-CVA is satisfied if there is no advantage of the relation computed by an adversary given a challenge ballot for v , over the relation computed by the adversary given a challenge ballot for v' . That is, for all adversaries, we have with negligible probability that the relation evaluated by the challenger in cnm-cva holds iff the relation evaluated in $\text{cnm-cva-}\$$ does not hold. It follows that no adversary can meaningfully relate ballots. On the other hand, if CNM-CVA is not satisfied, then there exists a strategy to construct related ballots.

CNM-CVA avoids crediting the adversary for trivial and unavoidable relations which hold if the challenge ballot appears on the bulletin board. For example,

suppose the adversary is given a challenge ballot for v in cnm-cva , respectively v' in $\text{cnm-cva-}\$$. This adversary could output a bulletin board containing only the challenge ballot and a relation R such that $R(v, \mathbf{v})$ holds if $\mathbf{v}[v] = 1$, hence, the relation evaluated in cnm-cva holds, whereas the relation evaluated in $\text{cnm-cva-}\$$ does not hold, but the adversary loses in both games because the challenge ballot appears on the bulletin board. By contrast, if the adversary can derive a ballot meaningfully related to the challenge ballot, then the adversary can win the game. For instance, Cortier & Smyth [CS13, CS11] identify a class of vulnerabilities against voting systems, which can be exploited as follows: an adversary observes a voter's ballot, casts a meaningfully related ballot, and abuses the relation to recover the voter's vote from the election outcome.

Comparing CNM-CVA and CNM-CPA. Unsurprisingly, the distinction between non-malleability for asymmetric encryption (CNM-CPA) and non-malleability for election schemes (CNM-CVA) is similar to the distinction between indistinguishability for asymmetric encryption and indistinguishability for election schemes (§3.2), namely, CNM-CPA performs a parallel decryption, whereas CNM-CVA performs a single tally.

Our ballot independence games are adaptations of standard security definitions for asymmetric encryption: CNM-CVA is based on non-malleability for encryption and IND-CVA is based on indistinguishability for encryption. Bellare & Sahai [BS99] have shown that non-malleability is equivalent to indistinguishability for encryption and their proof can be adapted to show that CNM-CVA and IND-CVA are equivalent.

Theorem 20 (CNM-CVA = IND-CVA). *Given an election scheme Γ , we have Γ satisfies CNM-CVA iff Γ satisfies IND-CVA.*

Proof. For the *if* implication, suppose Γ does not satisfy CNM-CVA, hence, there exists a probabilistic polynomial-time adversary \mathcal{A} , such that for all negligible functions negl , there exists a security parameter κ and $\text{Succ}(\text{cnm-cva}(\Gamma, \mathcal{A}, \kappa)) - \text{Succ}(\text{cnm-cva-}\$(\Gamma, \mathcal{A}, \kappa)) > \text{negl}(\kappa)$. We construct an adversary \mathcal{B} against game IND-CVA from \mathcal{A} .

- $\mathcal{B}(pk, \kappa)$ computes $(V, nc) \leftarrow \mathcal{A}(pk, \kappa)$; $v, v' \leftarrow_R V$ and outputs (v, v', nc) .
- $\mathcal{B}(b)$ computes $(R, \mathbf{bb}) \leftarrow \mathcal{A}(b)$ and outputs \mathbf{bb} .
- $\mathcal{B}(\mathbf{v})$ outputs 0 if $R(v, \mathbf{v})$ holds and 1 otherwise.

If the challenger selects $\beta = 0$ in game IND-CVA, then adversary \mathcal{B} simulates \mathcal{A} 's challenger to \mathcal{A} in cnm-cva and \mathcal{B} 's success (which requires $R(v, \mathbf{v})$ to hold) is $\text{Succ}(\text{cnm-cva}(\Gamma, \mathcal{A}, \kappa))$. Otherwise ($\beta = 1$), adversary \mathcal{B} simulates \mathcal{A} 's challenger to \mathcal{A} in $\text{cnm-cva-}\$$ and, since \mathcal{B} will evaluate $R(v, \mathbf{v})$, \mathcal{B} 's success (which requires $R(v, \mathbf{v})$ not to hold) is $1 - \text{Succ}(\text{cnm-cva-}\$(\Gamma, \mathcal{A}, \kappa))$. It follows that $\text{Succ}(\text{IND-CVA}(\Gamma, \mathcal{B}, \kappa)) = \frac{1}{2} \cdot (\text{Succ}(\text{cnm-cva}(\Gamma, \mathcal{A}, \kappa)) + 1 - \text{Succ}(\text{cnm-cva-}\$(\Gamma, \mathcal{A}, \kappa)))$,

therefore, $2 \cdot \text{Succ}(\text{IND-CVA}(\Gamma, \mathcal{B}, \kappa)) - 1 = \text{Succ}(\text{cnm-cva}(\Gamma, \mathcal{A}, \kappa)) - \text{Succ}(\text{cnm-cva-}\$(\Gamma, \mathcal{A}, \kappa))$. Thus, $\text{Succ}(\text{IND-CVA}(\Gamma, \mathcal{B}, \kappa)) > \frac{1}{2} + \frac{1}{2} \cdot \text{negl}(\kappa)$, concluding our proof of the *if* implication.

For the *only if* implication, suppose Γ does not satisfy IND-CVA, hence, there exists a probabilistic polynomial-time adversary \mathcal{A} , such that for all negligible functions negl , there exists a security parameter κ and $\text{Succ}(\text{IND-CVA}(\Gamma, \mathcal{A}, \kappa)) > \frac{1}{2} + \text{negl}(\kappa)$. We construct an adversary \mathcal{B} against CNM-CVA from \mathcal{A} .

- $\mathcal{B}(pk, \kappa)$ computes $(v_0, v_1, nc) \leftarrow \mathcal{A}(pk, \kappa)$ and outputs $(\{v_0, v_1\}, nc)$.
- $\mathcal{B}(b)$ computes $\mathbf{bb} \leftarrow \mathcal{A}(b)$, picks coins r uniformly at random, derives a relation R such that $R(v, \mathbf{v})$ holds if there exists a bit g such that $v = v_g \wedge g = \mathcal{A}(\mathbf{v}; r)$ and fails otherwise, and outputs (R, \mathbf{bb}) .

Adversary \mathcal{B} simulates \mathcal{A} 's challenger to \mathcal{A} in game IND-CVA. Indeed, the challenge ballot is equivalently computed. As is the election outcome. The computation $\mathcal{A}(\mathbf{v}; r)$ is not black-box, but this does not matter: it is still invoked exactly once in the game. Let us consider adversary \mathcal{B} 's success against cnm-cva and $\text{cnm-cva-}\$$.

- Game cnm-cva samples a single vote v from V . By inspection of cnm-cva and IND-CVA, we have $\text{Succ}(\text{cnm-cva}(\Gamma, \mathcal{B}, \kappa)) = \text{Succ}(\text{IND-CVA}(\Gamma, \mathcal{A}, \kappa))$, hence, $\text{Succ}(\text{cnm-cva}(\Gamma, \mathcal{B}, \kappa)) - \frac{1}{2} > \text{negl}(\kappa)$.
- Game $\text{cnm-cva-}\$$ samples votes v and v' from V . Vote v is independent of \mathcal{A} 's perspective, indeed, an equivalent formulation of $\text{cnm-cva-}\$$ could sample v after \mathcal{A} has terminated and immediately before evaluating the adversary's relation. By inspection of $\text{cnm-cva-}\$$ and IND-CVA, we have $\text{Succ}(\text{cnm-cva-}\$(\Gamma, \mathcal{B}, \kappa)) = \frac{1}{2} \cdot \text{Succ}(\text{IND-CVA}(\Gamma, \mathcal{A}, \kappa)) + \frac{1}{2} \cdot (1 - \text{Succ}(\text{IND-CVA}(\Gamma, \mathcal{A}, \kappa))) = \frac{1}{2}$.

It follows that $\text{Succ}(\text{cnm-cva}(\Gamma, \mathcal{B}, \kappa)) - \text{Succ}(\text{cnm-cva-}\$(\Gamma, \mathcal{B}, \kappa)) > \text{negl}(\kappa)$. \square

C Proofs

C.1 Proof of Theorem 1

Suppose Γ does not satisfy ballot independence, hence, there exists a probabilistic polynomial-time adversary \mathcal{A} , such that for all negligible functions negl , there exists a security parameter κ and $\text{Succ}(\text{IND-CVA}) > \frac{1}{2} + \text{negl}(\kappa)$. We construct an adversary \mathcal{B} against Ballot-Secrecy from \mathcal{A} .

- $\mathcal{B}(pk, \kappa)$ computes $(v_0, v_1, nc) \leftarrow \mathcal{A}(pk, \kappa)$ and outputs nc .
- $\mathcal{B}()$ computes $b \leftarrow \mathcal{O}(v_0, v_1)$; $\mathbf{bb} \leftarrow \mathcal{A}(b)$ and outputs \mathbf{bb} .
- $\mathcal{B}(\mathbf{v}, pf)$ computes $g \leftarrow \mathcal{A}(\mathbf{v})$ and outputs g .

Adversary \mathcal{B} simulates \mathcal{A} 's challenger to \mathcal{A} . Indeed, the challenge ballot and election outcome are equivalently computed. Moreover, the challenge ballot does not appear on the bulletin board, hence, the bulletin board is balanced. It follows that $\text{Succ}(\text{IND-CVA}(\Gamma, \mathcal{A}, \kappa)) = \text{Succ}(\text{Ballot-Secrecy}(\Gamma, \mathcal{B}, \kappa))$, hence, $\text{Succ}(\text{Ballot-Secrecy}(\Gamma, \mathcal{B}, \kappa)) > \frac{1}{2} + \text{negl}(\kappa)$, concluding our proof. \square

C.2 Proof of Proposition 3

In essence, the proof follows from Theorem 4. Albeit, formally, a few extra steps are required. In particular, the definition of an election scheme with zero-knowledge proofs demands that tallying proofs must be computed by a zero-knowledge non-interactive proof system, but an election scheme without tallying proofs need not compute proofs with such a system. Thus, we must introduce an election scheme with zero-knowledge proofs and prove that it is equivalent to the election scheme without proofs. This is trivial, so we do not pursue the details. \square

C.3 Proof of Theorem 4

Game **Ballot-Secrecy** computes the election outcome from ballots constructed by the oracle and ballots constructed by the adversary. Intuitively, such an outcome can be equivalently computed as follows:

$$\begin{aligned} (\mathbf{v}, pf) &\leftarrow \text{Tally}(sk, \mathbf{bb} \setminus \{b \mid (b, v_0, v_1) \in L\}, nc, \kappa); \\ (\mathbf{v}', pf') &\leftarrow \text{Tally}(sk, \mathbf{bb} \cap \{b \mid (b, v_0, v_1) \in L\}, nc, \kappa); \\ \mathbf{v} &\leftarrow \mathbf{v} + \mathbf{v}'; \end{aligned}$$

Yet, a poorly designed tallying algorithm might not ensure equivalence. In particular, ballots constructed by the adversary can cause the algorithm to behave unexpectedly. (Such algorithms are nonetheless compatible with our correctness requirement, because correctness does not consider an adversary.) Nevertheless, the equivalence holds when **Additivity** is satisfied. Moreover, the above computation is equivalent to the following:

$$\begin{aligned} (\mathbf{v}, pf) &\leftarrow \text{Tally}(sk, \mathbf{bb} \setminus \{b \mid (b, v_0, v_1) \in L\}, nc, \kappa); \\ \mathbf{for} \ b \in \mathbf{bb} \wedge (b, v_0, v_1) \in L \ \mathbf{do} & \\ \quad \left[\begin{array}{l} (\mathbf{v}', pf') \leftarrow \text{Tally}(sk, \{b\}, nc, \kappa); \\ \mathbf{v} \leftarrow \mathbf{v} + \mathbf{v}'; \end{array} \right. & \end{aligned}$$

Furthermore, by correctness of the election scheme, the above for-loop can be equivalently computed as follows:

$$\begin{aligned} \mathbf{for} \ b \in \mathbf{bb} \wedge (b, v_0, v_1) \in L \ \mathbf{do} & \\ \quad \left[\mathbf{v}[v_\beta] \leftarrow \mathbf{v}[v_\beta] + 1; \right. & \end{aligned}$$

Indeed, for each $b \in \mathbf{bb} \wedge (b, v_0, v_1) \in L$, we have b is an output of $\text{Vote}(pk, v_\beta, nc, \kappa)$, hence, $\text{Tally}(sk, \{b\}, nc, \kappa)$ outputs (\mathbf{v}, pf) such that \mathbf{v} is a zero-filled vector, except for index v_β which contains one, and this suffices to ensure equivalence. In addition, for any adversary that wins game **Ballot-Secrecy**, we are assured that $\text{balanced}(\mathbf{bb}, nc, L)$ holds, hence, the above for-loop can be computed as

for $b \in \mathbf{bb} \wedge (b, v_0, v_1) \in L$ **do**
 $\quad \lfloor \mathbf{v}[v_0] \leftarrow \mathbf{v}[v_0] + 1;$
 or
for $b \in \mathbf{bb} \wedge (b, v_0, v_1) \in L$ **do**
 $\quad \lfloor \mathbf{v}[v_1] \leftarrow \mathbf{v}[v_1] + 1;$

without weakening the game. Thus, perhaps surprisingly, tallying ballots constructed by the oracle does not provide the adversary with an advantage (in determining whether $\beta = 0$ or $\beta = 1$) and we can omit such ballots from tallying in game **Ballot-Secrecy**. Thus, game **Ballot-Secrecy** is equivalent to game **BS**, which modifies the tallying procedure as described. Moreover, that game simulates tallying proofs.

Definition 24. *Let $\Gamma = (\text{Setup}, \text{Vote}, \text{Tally})$ be an election scheme with zero-knowledge tallying proofs, \mathcal{A} be an adversary, and κ be a security parameter. Moreover, let \mathcal{S} be the simulator for the non-interactive zero-knowledge proof system used by algorithm **Tally** to compute tallying proofs. We define game **BS** as follows.*

$\text{BS}(\Gamma, \mathcal{A}, \mathcal{S}, \kappa) =$
 $(pk, sk, mb, mc) \leftarrow \text{Setup}(\kappa);$
 $nc \leftarrow \mathcal{A}(pk, \kappa);$
 $\beta \leftarrow_R \{0, 1\};$
 $L \leftarrow \emptyset;$
 $\mathbf{bb} \leftarrow \mathcal{A}^{\mathcal{O}}();$
 $(\mathbf{v}, pf) \leftarrow \text{Tally}(sk, \mathbf{bb} \setminus \{b \mid (b, v_0, v_1) \in L\}, nc, \kappa);$
for $b \in \mathbf{bb} \wedge (b, v_0, v_1) \in L$ **do**
 $\quad \lfloor \mathbf{v}[v_0] \leftarrow \mathbf{v}[v_0] + 1;$
 $pf \leftarrow \mathcal{S}((pk, nc, \mathbf{bb}, \mathbf{v}), \kappa);$
 $g \leftarrow \mathcal{A}(\mathbf{v}, pf);$
return $g = \beta \wedge \text{balanced}(\mathbf{bb}, nc, L) \wedge 1 \leq nc \leq mc \wedge |\mathbf{bb}| \leq mb;$

*Predicate **balanced** and oracle \mathcal{O} are defined as per game **Ballot-Secrecy**.*

Lemma 21. *Let $\Gamma = (\text{Setup}, \text{Vote}, \text{Tally})$ be an election scheme with zero-knowledge tally proofs, \mathcal{S} be the simulator for that proof system, \mathcal{A} be an adversary, and κ be a security parameter. If Γ satisfies **Additivity**, then $\text{Succ}(\text{BS}(\Gamma, \mathcal{A}, \mathcal{S}, \kappa)) = \text{Succ}(\text{Ballot-Secrecy}(\Gamma, \mathcal{A}, \kappa))$.*

Proof. The challengers in games **BS** and **Ballot-Secrecy** both compute public keys using the same algorithm and provide those keys, along with the security parameter, as input to the first adversary call, thus, these inputs and corresponding outputs are equivalent. Moreover, the left-right oracle is the same in both **BS** and **Ballot-Secrecy**, hence, the bulletin board output by the second adversary call is equivalent in both games. Furthermore, predicate *balanced* is satisfied in **BS** iff it is satisfied in **Ballot-Secrecy**, hence, if predicate *balanced* is not satisfied, then $\text{Succ}(\text{BS}(\Gamma, \mathcal{A}, \mathcal{S}, \kappa)) = \text{Succ}(\text{Ballot-Secrecy}(\Gamma, \mathcal{A}, \kappa))$, concluding our proof. Otherwise, it suffices to show that the outcome and tallying

proof are equivalently computed in **BS** and **Ballot-Secrecy**, since this ensures the inputs to the third adversary call are equivalent, thus the corresponding outputs are equivalent too, which suffices to conclude.

In **Ballot-Secrecy**, the outcome is computed by tallying the bulletin board. By comparison, in **BS**, the outcome is computed by tallying the ballots on the bulletin board that were constructed by the adversary (i.e., ballots in $\mathbf{bb} \setminus \{b \mid (b, v_0, v_1) \in L\}$, where \mathbf{bb} is the bulletin board and L is the set constructed by the oracle) and by simulating the tallying of any remaining ballots (i.e., ballots constructed by the oracle, namely, ballots in $\mathbf{bb} \cap \{b \mid (b, v_0, v_1) \in L\}$). Suppose (pk, sk, mb, mc) is an output of $\text{Setup}(\kappa)$ and nc is an integer. Since Γ satisfies **Additivity**, computing \mathbf{v} as

$$(\mathbf{v}, pf) \leftarrow \text{Tally}(sk, \mathbf{bb}, nc, \kappa);$$

is equivalent to computing \mathbf{v} as

$$\begin{aligned} (\mathbf{v}, pf) &\leftarrow \text{Tally}(sk, \mathbf{bb} \setminus \{b \mid (b, v_0, v_1) \in L\}, nc, \kappa); \\ (\mathbf{v}', pf') &\leftarrow \text{Tally}(sk, \mathbf{bb} \cap \{b \mid (b, v_0, v_1) \in L\}, nc, \kappa); \\ \mathbf{v} &\leftarrow \mathbf{v} + \mathbf{v}'; \end{aligned}$$

and as

$$\begin{aligned} (\mathbf{v}, pf) &\leftarrow \text{Tally}(sk, \mathbf{bb} \setminus \{b \mid (b, v_0, v_1) \in L\}, nc, \kappa); \\ (\mathbf{v}', pf') &\leftarrow \text{Tally}(sk, \emptyset, nc, \kappa); \\ \mathbf{for } b \in \mathbf{bb} \wedge (b, v_0, v_1) \in L \mathbf{ do} & \\ \quad \left[\begin{array}{l} (\mathbf{v}'', pf'') \leftarrow \text{Tally}(sk, \{b\}, nc, \kappa); \\ \mathbf{v}' \leftarrow \mathbf{v}' + \mathbf{v}''; \end{array} \right. & \\ \mathbf{v} &\leftarrow \mathbf{v} + \mathbf{v}'; \end{aligned}$$

Moreover, by correctness of Γ , we have $\text{Tally}(sk, \emptyset, nc, \kappa)$ outputs (\mathbf{v}', pf') such that \mathbf{v}' is a zero-filled vector. Hence, the above computation is equivalent to

$$\begin{aligned} (\mathbf{v}, pf) &\leftarrow \text{Tally}(sk, \mathbf{bb} \setminus \{b \mid (b, v_0, v_1) \in L\}, nc, \kappa); \\ \mathbf{for } b \in \mathbf{bb} \wedge (b, v_0, v_1) \in L \mathbf{ do} & \\ \quad \left[\begin{array}{l} (\mathbf{v}', pf') \leftarrow \text{Tally}(sk, \{b\}, nc, \kappa); \\ \mathbf{v} \leftarrow \mathbf{v} + \mathbf{v}'; \end{array} \right. & \end{aligned}$$

Thus, to prove the outcome is computed equivalently in **Ballot-Secrecy** and **BS**, it suffices to prove that the simulations are valid, i.e., computing the above for-loop is equivalent to computing $\mathbf{for } b \in \mathbf{bb} \wedge (b, v_0, v_1) \in L \mathbf{ do } \mathbf{v}[v_0] \leftarrow \mathbf{v}[v_0] + 1$.

In **BS**, we have for all $(b, v_0, v_1) \in L$ that b is an output of $\text{Vote}(pk, v_\beta, nc, \kappa)$ such that $v_0, v_1 \in \{1, \dots, nc\}$, where β is the bit chosen by the challenger. Moreover, by correctness of Γ , we have $\text{Tally}(sk, \{b\}, nc, \kappa)$ outputs (\mathbf{v}', pf') such that \mathbf{v}' is a zero-filled vector, except for index v_β , which contains one, hence, computing $\mathbf{v} \leftarrow \mathbf{v} + \mathbf{v}'$ inside the for-loop is equivalent to computing $\mathbf{v}[v_\beta] \leftarrow \mathbf{v}[v_\beta] + 1$ inside the loop. Furthermore, since predicate *balanced* holds in **BS**, we have for all $v \in \{1, \dots, nc\}$ that $|\{b \mid b \in \mathbf{bb} \wedge \exists v_1 . (b, v, v_1) \in L\}| = |\{b \mid b \in \mathbf{bb} \wedge \exists v_0 . (b, v_0, v) \in L\}|$. Hence, in **BS**, computing

$$\mathbf{for } b \in \mathbf{bb} \wedge (b, v_0, v_1) \in L \mathbf{ do } \mathbf{v}[v_\beta] \leftarrow \mathbf{v}[v_\beta] + 1;$$

is equivalent to computing

for $b \in \mathfrak{bb} \wedge (b, v_0, v_1) \in L$ **do** $\mathfrak{v}[v_0] \leftarrow \mathfrak{v}[v_0] + 1$;

Thus, the simulation is valid in BS.

In **Ballot-Secrecy**, the tallying proof is computed by tallying the bulletin board. By comparison, in **BS**, the tallying proof is computed by simulator \mathcal{S} . Since Γ has zero-knowledge tallying proofs, there exists a non-interactive proof system (Prove, Verify) such that for all (\mathfrak{v}, pf) output by $\text{Tally}(sk, \mathfrak{bb}, nc, \kappa)$, we have pf is an output of algorithm **Prove** parameterised with statement $(pk, \mathfrak{bb}, nc, \mathfrak{v})$ and coins chosen uniformly at random. Moreover, since \mathcal{S} is a simulator for (Prove, Verify), proofs output by algorithm **Prove** are indistinguishable from outputs of simulator \mathcal{S} . Thus, tallying proofs are equivalently computed in **Ballot-Secrecy** and **BS**, thereby concluding our proof. \square

Let **BS-0**, respectively **BS-1**, be the game derived from **BS** by replacing $\beta \leftarrow_R \{0, 1\}$ with $\beta \leftarrow 0$, respectively $\beta \leftarrow 1$. These games are trivially related to **BS**, namely, $\text{Succ}(\text{BS}(\Gamma, \mathcal{A}, \mathcal{S}, \kappa)) = \frac{1}{2} \cdot \text{Succ}(\text{BS-0}(\Gamma, \mathcal{A}, \mathcal{S}, \kappa)) + \frac{1}{2} \cdot \text{Succ}(\text{BS-1}(\Gamma, \mathcal{A}, \mathcal{S}, \kappa))$. Moreover, let **BS-1:0** be the game derived from **BS-1** by replacing $g = \beta$ with $g = 0$ and let G_j be the game derived from **BS-1:0** by removing $\beta \leftarrow 1$ and redefining oracle \mathcal{O} such that $\mathcal{O}(v_0, v_1)$ computes, on inputs $v_0, v_1 \in \{1, \dots, nc\}$, the following:

```

if  $|L| < j$  then
  |  $b \leftarrow \text{Vote}(pk, v_1, nc, \kappa)$ ;
else
  |  $b \leftarrow \text{Vote}(pk, v_0, nc, \kappa)$ ;
 $L \leftarrow L \cup \{(b, v_0, v_1)\}$ ;
return  $b$ ;

```

Games G_0, G_1, \dots are distinguished from games **BS-0** and **BS-1:0** by their left-right oracles. In particular, the first j left-right oracle queries in G_j construct ballots for the oracle's "right" input and any remaining queries construct ballots for the oracle's "left" input, whereas the left-right oracle in **BS-0**, respectively **BS-1:0**, always constructs ballots for the oracle's "left," respectively "right," input. We relate game **BS-1:0** to **BS-1**, games **BS-0** and **BS-1:0** to the hybrid games G_0, G_1, \dots , and we prove Theorem 4 using these relations.

Lemma 22. *Let Γ be an election scheme with zero-knowledge tally proofs, \mathcal{S} be the simulator for that proof system, \mathcal{A} be an adversary, and κ be a security parameter. If adversary \mathcal{A} wins game **Ballot-Secrecy** against election scheme Γ , then $\text{Succ}(\text{BS-1}(\Gamma, \mathcal{A}, \mathcal{S}, \kappa)) = 1 - \text{Succ}(\text{BS-1:0}(\Gamma, \mathcal{A}, \mathcal{S}, \kappa))$.*

Lemma 23. *Let Γ be an election scheme with zero-knowledge tally proofs, \mathcal{S} be the simulator for that proof system, \mathcal{A} be an adversary, and κ be a security parameter. If Γ satisfies **Additivity**, then $\text{Succ}(\text{BS-0}(\Gamma, \mathcal{A}, \mathcal{S}, \kappa)) = \text{Succ}(G_0(\Gamma, \mathcal{A}, \mathcal{S}, \kappa))$ and $\text{Succ}(\text{BS-1:0}(\Gamma, \mathcal{A}, \mathcal{S}, \kappa)) = \text{Succ}(G_q(\Gamma, \mathcal{A}, \mathcal{S}, \kappa))$, where q is an upper-bound on \mathcal{A} 's left-right oracle queries.*

Proof. Games **BS-0** and G_0 , respectively **BS-1:0** and G_q , are identical up to the oracle, except **BS-0** computes $\beta \leftarrow 0$ and checks $g = \beta$, which is equivalent to

G_0 checking $g = 0$, respectively BS-1:0 computes $\beta \leftarrow 1$ and checks $g = 0$, which is equivalent to G_q checking $g = 0$. Thus, it suffices to show that oracle outputs in BS-0 are equivalent to oracle outputs in G_0 , and similarly for BS-1:0 and G_q . Left-right oracles queries $\mathcal{O}(v_0, v_1)$ in games BS-0 and G_0 output ballots for vote v_0 , hence, outputs are equivalent in both games. Oracle outputs in BS-1:0 and G_q are similarly equivalent, in particular, left-right oracles queries $\mathcal{O}(v_0, v_1)$ in both games output ballots for vote v_1 , because q is an upper-bound on the left-right oracle queries, therefore, $|L| < q$ in G_q , concluding our proof. \square

Proof of Theorem 4. By Theorem 1, it suffices to prove that ballot independence implies ballot secrecy. Suppose Γ does not satisfy ballot secrecy, hence, there exists a probabilistic polynomial-time adversary \mathcal{A} , such that for all negligible functions negl , there exists a security parameter κ and

$$\frac{1}{2} + \text{negl}(\kappa) < \text{Succ}(\text{Ballot-Secrecy}(\Gamma, \mathcal{A}, \kappa))$$

Let $\Gamma = (\text{Setup}, \text{Vote}, \text{Tally})$. Since Γ has zero-knowledge tallying proofs, tallying proofs output by Tally are computed by a non-interactive zero-knowledge proof system. Let algorithm \mathcal{S} be the simulator for that proof system. By Lemma 21, we have

$$= \text{Succ}(\text{BS}(\Gamma, \mathcal{A}, \mathcal{S}, \kappa))$$

By definition of BS-0 and BS-1, we have

$$= \frac{1}{2} \cdot (\text{Succ}(\text{BS-0}(\Gamma, \mathcal{A}, \mathcal{S}, \kappa)) + \text{Succ}(\text{BS-1}(\Gamma, \mathcal{A}, \mathcal{S}, \kappa)))$$

And, by Lemma 22, we have

$$\begin{aligned} &= \frac{1}{2} \cdot (\text{Succ}(\text{BS-0}(\Gamma, \mathcal{A}, \mathcal{S}, \kappa)) + 1 - \text{Succ}(\text{BS-1:0}(\Gamma, \mathcal{A}, \mathcal{S}, \kappa))) \\ &= \frac{1}{2} + \frac{1}{2} \cdot (\text{Succ}(\text{BS-0}(\Gamma, \mathcal{A}, \mathcal{S}, \kappa)) - \text{Succ}(\text{BS-1:0}(\Gamma, \mathcal{A}, \mathcal{S}, \kappa))) \end{aligned}$$

Let q be an upper-bound on \mathcal{A} 's left-right oracle queries. Hence, by Lemma 23, we have

$$= \frac{1}{2} + \frac{1}{2} \cdot (\text{Succ}(G_0(\Gamma, \mathcal{A}, \mathcal{S}, \kappa)) - \text{Succ}(G_q(\Gamma, \mathcal{A}, \mathcal{S}, \kappa)))$$

which can be rewritten as the telescoping series

$$= \frac{1}{2} + \frac{1}{2} \cdot \sum_{1 \leq j \leq q} \text{Succ}(G_{j-1}(\Gamma, \mathcal{A}, \mathcal{S}, \kappa)) - \text{Succ}(G_j(\Gamma, \mathcal{A}, \mathcal{S}, \kappa))$$

Let $j \in \{1, \dots, q\}$ be such that $\text{Succ}(\mathbf{G}_{j-1}(\Gamma, \mathcal{A}, \mathcal{S}, \kappa)) - \text{Succ}(\mathbf{G}_j(\Gamma, \mathcal{A}, \mathcal{S}, \kappa))$ is the largest term in that series. Hence,

$$\leq \frac{1}{2} + \frac{1}{2} \cdot q \cdot (\text{Succ}(\mathbf{G}_{j-1}(\Gamma, \mathcal{A}, \mathcal{S}, \kappa)) - \text{Succ}(\mathbf{G}_j(\Gamma, \mathcal{A}, \mathcal{S}, \kappa)))$$

Thus,

$$\frac{1}{2} + \frac{1}{q} \cdot \text{negl}(\kappa) \leq \frac{1}{2} + \frac{1}{2} \cdot (\text{Succ}(\mathbf{G}_{j-1}(\Gamma, \mathcal{A}, \mathcal{S}, \kappa)) - \text{Succ}(\mathbf{G}_j(\Gamma, \mathcal{A}, \mathcal{S}, \kappa)))$$

From \mathcal{A} , we construct the following adversary \mathcal{B} against IND-CVA:

- $\mathcal{B}(pk, \kappa)$ computes $nc \leftarrow \mathcal{A}(pk, \kappa); L \leftarrow \emptyset$ and runs $\mathcal{A}^{\mathcal{O}}()$, handling \mathcal{A} 's oracle queries $\mathcal{O}(v_0, v_1)$ as follows: if $|L| < j$, then compute $b \leftarrow \text{Vote}(pk, v_1, nc, \kappa); L \leftarrow L \cup \{b, v_0, v_1\}$ and return b to \mathcal{A} , otherwise, assign $v_0^c \leftarrow v_0; v_1^c \leftarrow v_1$, and output (v_0, v_1, nc) .
- $\mathcal{B}(b)$ assigns $L \leftarrow L \cup \{(b, v_0^c, v_1^c)\}$; returns b to \mathcal{A} and handles any further oracle queries $\mathcal{O}(v_0, v_1)$ as follows, namely, compute $b \leftarrow \text{Vote}(pk, v_0, nc, \kappa); L \leftarrow L \cup \{(b, v_0, v_1)\}$ and return b to \mathcal{A} ; assigns \mathcal{A} 's output to \mathbf{bb} ; and outputs $\mathbf{bb} \setminus \{b \mid (b, v_0, v_1) \in L\}$.

- $\mathcal{B}(\mathbf{v})$ computes

for $b \in \mathbf{bb} \wedge (b, v_0, v_1) \in L$ **do**
 $\lfloor \mathbf{v}[v_0] \leftarrow \mathbf{v}[v_0] + 1;$
 $pf \leftarrow \mathcal{S}((pk, \mathbf{bb}, nc, \mathbf{v}), \kappa);$
 $g \leftarrow \mathcal{A}(\mathbf{v}, pf);$

and outputs g .

We prove that \mathcal{B} wins IND-CVA with success of at least $\frac{1}{2} + \frac{1}{2} \cdot (\text{Succ}(\mathbf{G}_{j-1}(\Gamma, \mathcal{A}, \mathcal{S}, \kappa)) - \text{Succ}(\mathbf{G}_j(\Gamma, \mathcal{A}, \mathcal{S}, \kappa)))$.

Suppose (pk, sk, mb, mc) is an output of $\text{Setup}(\kappa)$. Further suppose we run $\mathcal{B}(pk, \kappa)$. It is straightforward to see that \mathcal{B} simulates the challenger and oracle in both \mathbf{G}_{j-1} and \mathbf{G}_j to \mathcal{A} . In particular, \mathcal{B} simulates query $\mathcal{O}(v_0, v_1)$ by computing $b \leftarrow \text{Vote}(pk, v_1, nc, \kappa)$ for the first $j - 1$ queries. Since \mathbf{G}_{j-1} and \mathbf{G}_j are equivalent to adversaries that make fewer than j left-right oracle queries, adversary \mathcal{A} must make at least j queries to ensure $\text{Succ}(\mathbf{G}_{j-1}(\Gamma, \mathcal{A}, \mathcal{S}, \kappa)) - \text{Succ}(\mathbf{G}_j(\Gamma, \mathcal{A}, \mathcal{S}, \kappa))$ is non-negligible. Hence, $\mathcal{B}(pk, \kappa)$ terminates with non-negligible probability. Suppose adversary \mathcal{B} terminates by outputting (v_0, v_1, nc) , where v_0, v_1 correspond to the inputs of the j th oracle query by \mathcal{A} . Further suppose b is an output of $\text{Vote}(pk, v_\beta, nc, \kappa)$, where β is a bit. If $\beta = 0$, then $\mathcal{B}(b)$ simulates the oracle in \mathbf{G}_{j-1} to \mathcal{A} , otherwise, $\mathcal{B}(b)$ simulates the oracle in \mathbf{G}_j to \mathcal{A} . In particular, $\mathcal{B}(b)$ responds to the j th oracle query with ballot b for v_β , thus, simulating the challenger in \mathbf{G}_{j-1} when $\beta = 0$, respectively \mathbf{G}_j when $\beta = 1$. And $\mathcal{B}(b)$ responds to any further oracle queries $\mathcal{O}(v_0, v_1)$ with ballots for v_0 . Suppose \mathbf{bb} is an output of \mathcal{A} , thus $\mathcal{B}(b)$ outputs $\mathbf{bb} \setminus \{b \mid (b, v_0, v_1) \in L\}$.

Further suppose (\mathbf{v}, pf) is an output of $\text{Tally}(sk, \mathbf{bb} \setminus \{b \mid (b, v_0, v_1) \in L\}, nc, \kappa)$ and g is an output of $\mathcal{B}(\mathbf{v})$. It is trivial to see that $\mathcal{B}(\mathbf{v})$ simulates \mathcal{A} 's challenger. Thus, either

1. $\beta = 0$ and \mathcal{B} simulates G_{j-1} to \mathcal{A} , thus, $g = \beta$ with at least the probability that \mathcal{A} wins G_{j-1} ; or
2. $\beta = 1$ and \mathcal{B} simulates G_j to \mathcal{A} , thus, $g \neq 0$ with at least the probability that \mathcal{B} loses G_j and, since \mathcal{A} wins game Ballot-Secrecy, we have g is a bit, hence, $g = \beta$.

It follows that the success of adversary \mathcal{B} is at least $\frac{1}{2} \cdot \text{Succ}(G_{j-1}(\Gamma, \mathcal{A}, \mathcal{S}, \kappa)) + \frac{1}{2} \cdot (1 - \text{Succ}(G_j(\Gamma, \mathcal{A}, \mathcal{S}, \kappa)))$, thus we conclude our proof. \square

C.4 Proof of Lemma 8

Suppose (pk, sk, mb, mc) is an output of $\text{Setup}(\kappa)$ and $(nc, \mathbf{bb}, \mathbf{bb}')$ is an output of $\mathcal{A}(pk, \kappa)$ such that $nc \leq mc$ and $|\mathbf{bb}| \leq mb$. Further suppose (\mathbf{v}, pf) is an output $\text{Tally}(sk, \mathbf{bb}, nc, \kappa)$, (\mathbf{v}_0, pf_0) is an output of $\text{Tally}(sk, \mathbf{bb} \setminus \mathbf{bb}', nc, \kappa)$, and (\mathbf{v}_1, pf_1) is an output of $\text{Tally}(sk, \mathbf{bb} \cap \mathbf{bb}', nc, \kappa)$. Let $\mathbf{v}^+ = \mathbf{v}_0 + \mathbf{v}_1$. Moreover, let $\mathbf{v}^* = \text{correct-outcome}(pk, nc, \mathbf{bb}, \kappa)$, $\mathbf{v}_0^* = \text{correct-outcome}(pk, nc, \mathbf{bb} \setminus \mathbf{bb}', \kappa)$, and $\mathbf{v}_1^* = \text{correct-outcome}(pk, nc, \mathbf{bb} \cap \mathbf{bb}', \kappa)$. By definition of function *correct-outcome*, we have $\mathbf{v}^* = \mathbf{v}_0^* + \mathbf{v}_1^*$. Moreover, by Tally-Soundness, we have $\mathbf{v} = \mathbf{v}^*$, $\mathbf{v}_0 = \mathbf{v}_0^*$, and $\mathbf{v}_1 = \mathbf{v}_1^*$, with overwhelming probability. It follows by transitivity that $\mathbf{v} = \mathbf{v}^+$, with overwhelming probability.

C.5 Proof of Proposition 12

We present a construction (Definition 25) for encryption schemes (Lemma 24) which are clearly not secure (Lemma 25). Nevertheless, the construction produces encryption schemes that are sufficient for ballot secrecy (Lemma 26). The proof of Proposition 12 follows from Lemmata 24–26.

Definition 25. *Given an asymmetric encryption scheme $\Pi = (\text{Gen}_\Pi, \text{Enc}_\Pi, \text{Dec}_\Pi)$ and a constant symbol ω , let $\text{Leak}(\Pi, \omega) = (\text{Gen}_\Pi, \text{Enc}_\Pi, \text{Dec})$ such that $\text{Dec}(sk, c)$ proceeds as follows: if $c = \omega$, then output sk , otherwise, compute $m \leftarrow \text{Dec}_\Pi(sk, c)$ and output m .*

Lemma 24. *Given an asymmetric encryption scheme Π and a constant symbol ω such that Π 's ciphertext space does not contain ω , we have $\text{Leak}(\Pi, \omega)$ is an asymmetric encryption scheme.*

Proof sketch. The proof follows immediately from correctness of the underlying encryption scheme, because constant symbol ω does not appear in the scheme's ciphertext space. \square

Lemma 25. *Given an asymmetric encryption scheme Π and a constant symbol ω such that Π 's ciphertext space does not contain ω and Π 's message space is larger than one for some security parameter, we have $\text{Leak}(\Pi, \omega)$ does not satisfy IND-PA0.*

Proof sketch. The proof is trivial: an adversary can output two distinct messages and a vector containing constant symbol ω during the first two adversary calls, learn the private key from the parallel decryption, and use the key to recover the plaintext from the challenge ciphertext, which allows the adversary to win the game. \square

Lemma 26. *Let $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ be an asymmetric encryption scheme and ω be a constant symbol. Suppose Π 's ciphertext space does not contain ω and Π 's message space is smaller than the private key. Further suppose $\text{Enc2Vote}(\Pi)$ satisfies Ballot-Secrecy. We have $\text{Enc2Vote}(\text{Leak}(\Pi, \omega))$ satisfies Ballot-Secrecy.*

Proof. Let $\text{Enc2Vote}(\Pi) = (\text{Setup}, \text{Vote}, \text{Tally})$ and let $\text{Enc2Vote}(\text{Leak}(\Pi, \omega)) = (\text{Setup}', \text{Vote}', \text{Tally}')$. By definition of $\text{Enc2Vote}(\Pi)$ and Leak , we have $\text{Setup} = \text{Setup}'$ and $\text{Vote} = \text{Vote}'$. Suppose \mathbf{m} is Π 's message space. By definition of Leak , we have \mathbf{m} is $\text{Leak}(\Pi, \omega)$'s message space too. Moreover, since $|\mathbf{m}|$ is smaller than the private key, we have for all security parameters κ , bulletin boards \mathbf{bb} , and number of candidates nc , that $nc \leq |\mathbf{m}|$ implies

$$\Pr[(pk, sk, \mathbf{m}) \leftarrow \text{Gen}(\kappa); (\mathbf{v}, pf) \leftarrow \text{Tally}(sk, \mathbf{bb}, nc, \kappa); \\ (\mathbf{v}', pf') \leftarrow \text{Tally}'(sk, \mathbf{bb}, nc, \kappa) : \mathbf{v} = \mathbf{v}' \wedge pf = pf'] = 1,$$

because $\text{Enc2Vote}(\Pi)$ ensures that \mathbf{v}' is not influenced by decrypting ω (witness that decrypting ω outputs sk such that $sk > |\mathbf{m}| \geq nc$) and pf is a constant symbol. It follows for all adversaries \mathcal{A} and security parameters κ that games $\text{Ballot-Secrecy}(\text{Enc2Vote}(\Pi), \mathcal{A}, \kappa)$ and $\text{Ballot-Secrecy}(\text{Enc2Vote}(\text{Leak}(\Pi, \omega)), \mathcal{A}, \kappa)$ are equivalent, hence, we have $\text{Succ}(\text{Ballot-Secrecy}(\text{Enc2Vote}(\Pi), \mathcal{A}, \kappa)) = \text{Succ}(\text{Ballot-Secrecy}(\text{Enc2Vote}(\text{Leak}(\Pi, \omega)), \mathcal{A}, \kappa))$. Moreover, since $\text{Enc2Vote}(\Pi)$ satisfies Ballot-Secrecy, it follows that $\text{Enc2Vote}(\text{Leak}(\Pi, \omega))$ satisfies Ballot-Secrecy too. \square

Proof of Proposition 12. Let Π be an asymmetric encryption scheme and ω be a constant symbol. Suppose Π 's ciphertext space does not contain ω . Further suppose Π 's message space is larger than one for some security parameter, but smaller than the private key. We have $\text{Enc2Vote}(\text{Leak}(\Pi, \omega))$ is an asymmetric encryption scheme (Lemma 24) such that $\text{Enc2Vote}(\text{Leak}(\Pi, \omega))$ satisfies Ballot-Secrecy (Lemma 26), but $\text{Leak}(\Pi, \omega)$ does not satisfy IND-PA0 (Lemma 25), concluding our proof. \square

C.6 Proof of Lemma 13

Let $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ and $\text{Enc2Vote}(\Pi) = (\text{Setup}, \text{Vote}, \text{Tally})$. Election scheme $\text{Enc2Vote}(\Pi)$ satisfies HK-Injectivity (Lemma 10). Suppose $\text{Enc2Vote}(\Pi)$ does not satisfy Tally-Soundness, hence, there exists a probabilistic polynomial-time adversary \mathcal{A} , such that for all negligible functions negl , there exists a security parameter κ and $\text{Succ}(\text{Tally-Soundness}(\text{Enc2Vote}(\Pi), \mathcal{A}, \kappa)) \leq 1 - \text{negl}(\kappa)$. Further suppose (pk', sk, mb, mc) is an output of $\text{Setup}(\kappa)$, (nc, \mathbf{bb}) is an output of

$\mathcal{A}(pk, \kappa)$, and (\mathbf{v}, pf) is an output of $\text{Tally}(sk, \mathbf{bb}, nc, \kappa)$. By definition of algorithm **Setup**, we have pk' is a pair (pk, \mathbf{m}) such that (pk, sk, \mathbf{m}) is an output of $\text{Gen}(\kappa)$, and mc is the largest integer such that $\{0, \dots, mc\} \subseteq \{0\} \cup \mathbf{m}$. Moreover, since \mathcal{A} is a winning adversary, we have $nc \leq mc$. By definition of algorithm **Tally**, we have \mathbf{v} is initialised as a zero-filled vector of length nc and updated by computing **for** $b \in \mathbf{bb}$ **do** $v \leftarrow \text{Dec}(sk, b)$; **if** $1 \leq v \leq nc$ **then** $\mathbf{v}[v] \leftarrow \mathbf{v}[v] + 1$. Since Π satisfies well-definedness and error symbol \perp is not an integer, that computation is equivalent to

```

for  $b \in \mathbf{bb} \wedge \exists m, r . m \in \mathbf{m} \wedge b = \text{Enc}(pk, m; r) \wedge b \neq \perp$  do
   $v \leftarrow \text{Dec}(sk, b)$ ;
  if  $1 \leq v \leq nc$  then
     $\mathbf{v}[v] \leftarrow \mathbf{v}[v] + 1$ ;

```

with overwhelming probability. Although each ciphertext $\text{Enc}(pk, m; r) \in \mathbf{bb}$ may not have been computed using coins r chosen uniformly at random, we nonetheless have $\text{Dec}(sk, \text{Enc}(pk, m; r)) = m$, because Π is perfectly correct. Hence, the above computation is equivalent to

```

for  $b \in \mathbf{bb} \wedge \exists v, r . v \in \mathbf{m} \wedge b = \text{Enc}(pk, v; r) \wedge b \neq \perp$  do
  if  $1 \leq v \leq nc$  then
     $\mathbf{v}[v] \leftarrow \mathbf{v}[v] + 1$ ;

```

Thus, for all $v \in \{1, \dots, nc\}$, we have $\mathbf{v}[v] = \ell$ if and only if $\exists^\ell b \in \mathbf{bb} \setminus \{\perp\} : \exists r : b = \text{Enc}(pk, v; r)$, with overwhelming probability. It follows by definition of **Vote** that for all $v \in \{1, \dots, nc\}$ we have

$$\mathbf{v}[v] = \ell \text{ iff } \exists^\ell b \in \mathbf{bb} \setminus \{\perp\} : \exists r : b = \text{Vote}(pk, v, nc, \kappa; r)$$

with overwhelming probability. Thereby contradicting our assumption that \mathcal{A} is a winning adversary, since $\mathbf{v} = \text{correct-outcome}(pk, nc, \mathbf{bb}, \kappa)$, with overwhelming probability, which concludes our proof. \square

C.7 Proof of Proposition 16

Let $\Gamma = (\text{Setup}, \text{Vote}, \text{Tally})$. Suppose (pk, sk, mb, mc) is an output of $\text{Setup}(\kappa)$ and nc is an output of $\mathcal{A}(pk, \kappa)$ such that $1 \leq nc \leq mc$, for some security parameter κ and some adversary \mathcal{A} against game **Ballot-Secrecy** or game **Receipt-Freeness-A**. Since $mc = 1$ by hypothesis, we have $nc = 1$ too. For game **Receipt-Freeness-A**, let algorithm \mathbf{V} be such that $\mathbf{V}(pk, v_0, v_1, nc, \kappa)$ chooses coins r uniformly at random from the coin space of algorithm **Vote**, computes $b \leftarrow \text{Vote}(pk, v_0, nc, \kappa)$, and outputs (b, r) . Suppose β is a bit chosen uniformly at random and \mathbf{bb} is an output of $\mathcal{A}()$. By inspection of the oracle definition in each game, we have for every oracle call $\mathcal{O}(v_0, v_1)$ that $v_0 = v_1$, moreover, the ballot output by the oracle is independent of bit β . It follows that the adversary loses each game, hence Γ satisfies **Ballot-Secrecy**. To show that **Receipt-Freeness** is satisfied too, we must consider game **Receipt-Freeness-B**.

Suppose (pk, v_0, v_1, nc) is an output of $\mathcal{B}(\kappa)$ such that $1 \leq v_0, v_1 \leq nc$, for some adversary \mathcal{B} against game **Receipt-Freeness-B**. Let β be a bit chosen

uniformly at random. If $\beta = 0$, then suppose b is an output of $\text{Vote}(pk, v_0, nc, \kappa)$, otherwise ($\beta = 1$), suppose (b, r) is an output of $\mathbf{V}(pk, v_0, v_1, nc, \kappa)$, hence, b is an output of $\text{Vote}(pk, v_0, nc, \kappa)$ by definition of \mathbf{V} . Thus, the ballot computed by the challenger is independent of bit β . It follows that the adversary loses game Receipt-Freeness-B, hence, Γ satisfies Receipt-Freeness, thereby concluding our proof. \square

C.8 Proof of Theorem 17

Let $\Gamma = (\text{Setup}, \text{Vote}, \text{Tally})$. Suppose to the contrary that Γ satisfies Receipt-Freeness, hence, there exists a probabilistic polynomial-time algorithm \mathbf{V} such that for all probabilistic polynomial-time adversaries \mathcal{A} and \mathcal{B} , there exists a negligible function negl and for all security parameters κ , we have $\text{Succ}(\text{Receipt-Freeness-A}(\Gamma, \mathbf{V}, \mathcal{A}, \kappa)) + \text{Succ}(\text{Receipt-Freeness-B}(\Gamma, \mathbf{V}, \mathcal{B}, \kappa)) \leq \frac{1}{2} + \text{negl}(\kappa)$. Further suppose κ is such that the maximum number of candidates output by algorithm $\text{Setup}(\kappa)$ is greater than one. Moreover, suppose \mathcal{A} is the following adversary.

- $\mathcal{A}(pk, \kappa)$ computes $nc \leftarrow 2$ and outputs nc .
- $\mathcal{A}()$ computes $v_0 \leftarrow 1; v_1 \leftarrow 2; (b, r) \leftarrow \mathcal{O}(v_0, v_1)$ and outputs \emptyset .
- $\mathcal{A}(\mathbf{v}, pf)$ outputs 0 if $v = \text{Vote}(pk, v_0, nc, \kappa; r)$ and 1 otherwise.

Since Γ satisfies Receipt-Freeness, it follows that $\mathbf{V}(pk, v_1, v_0, nc, \kappa)$ outputs (b, r) such that $b = \text{Vote}(pk, v_0, nc, \kappa; r)$, with overwhelming probability. Suppose \mathcal{B} is the following adversary.

- $\mathcal{B}(\kappa)$ computes $(pk, sk, mb, mc) \leftarrow \text{Setup}(\kappa); v_0 \leftarrow 1; v_1 \leftarrow 2; nc \leftarrow 2$ and outputs (pk, v_0, v_1, nc) .
- $\mathcal{B}(b)$ computes $(\mathbf{v}, pf) \leftarrow \text{Tally}(sk, \{b\}, nc, \kappa)$ and outputs 0 if $\mathbf{v} = (1, 0)$ and 1 otherwise.

We prove that \mathcal{B} wins Receipt-Freeness-B($\Gamma, \mathbf{V}, \mathcal{B}, \kappa$).

Suppose (pk, v_0, v_1, nc) is an output of $\mathcal{B}(\kappa)$ and β is a bit chosen uniformly at random. If $\beta = 0$, then further suppose b is an output of $\text{Vote}(pk, v_0, nc, \kappa)$, and g is an output of $\mathcal{B}(b)$. Outputs (\mathbf{v}, pf) of $\text{Tally}(sk, \{b\}, nc, \kappa)$ are such that $\mathbf{v} = (1, 0)$ by correctness, which ensures $g = \beta$ by definition of \mathcal{B} . Otherwise ($\beta = 1$), suppose (b, r) is an output of $\mathbf{V}(pk, v_0, v_1, nc, \kappa)$, hence, $b = \text{Vote}(pk, v_1, nc, \kappa; r)$, and g is an output of $\mathcal{B}(b)$. Outputs (\mathbf{v}, pf) of $\text{Tally}(sk, \{b\}, nc, \kappa)$ are such that $\mathbf{v} = (0, 1)$ by perfect correctness and $\mathbf{v} = \text{correct-outcome}(pk, nc, \{b\}, \kappa) = (0, 1)$ by Tally-Soundness, which ensures $g = \beta$ by definition of \mathcal{B} . (Correctness, rather than perfect correctness, is insufficient, because algorithm \mathbf{V} may not have chosen coins r uniformly at random.) Thus, $\text{Succ}(\text{Receipt-Freeness-A}(\Gamma, \mathbf{V}, \mathcal{A}, \kappa)) + \text{Succ}(\text{Receipt-Freeness-B}(\Gamma, \mathbf{V}, \mathcal{B}, \kappa)) \leq \frac{1}{2} + \text{negl}(\kappa)$, concluding our proof. \square

C.9 Proof sketch of Proposition 18

Let $\Gamma = (\text{Setup}, \text{Vote}, \text{Tally})$ and $\text{BS2RF}(\Gamma) = (\text{Setup}, \text{VoteClient}, \text{VoteServer}, \text{Tally})$. Moreover, let $V(pk, v, v', nc, \kappa)$ compute $b \leftarrow \text{VoteServer}(pk, v, nc, \kappa)$ and output (b, ϵ) , where ϵ is selected from the empty coin space. Since algorithms V and $\text{VoteClient}(pk, v, nc, \kappa)$ compute ballots in the same way, it is straightforward to see that: for all probabilistic polynomial-time adversaries \mathcal{B} , there exists a negligible function negl and for all security parameters κ , we have $\text{Succ}(\text{Receipt-Freeness-B}(\Gamma, V, \mathcal{B}, \kappa)) \leq \frac{1}{2} + \text{negl}(\kappa)$. It remains to consider game $\text{Receipt-Freeness-A}(\Gamma, V, \mathcal{A}, \kappa)$.

Since the coin space of algorithm VoteClient is empty, it is straightforward to see that coins need not be output by the oracle in game $\text{Receipt-Freeness-A}$. Moreover, by definition of algorithms V and VoteClient , we can replace the oracle's computation of $b \leftarrow \text{VoteClient}^{\text{VoteServer}}(pk, v_0, nc, \kappa; r)$, respectively $(b, r) \leftarrow V(pk, v_1, v_0, nc, \kappa)$, with $b \leftarrow \text{VoteServer}(pk, v_0, nc, \kappa)$, respectively $b \leftarrow \text{VoteServer}(pk, v_1, nc, \kappa)$, hence, we can replace the if-then-else branch with $b \leftarrow \text{VoteServer}(pk, v_\beta, nc, \kappa)$. Thus, in the context of this proof, we have reduced game $\text{Receipt-Freeness-A}$ to game Ballot-Secrecy , which suffices to conclude. \square

D Helios

Smyth, Frink & Clarkson [SFC17] formalise a generic construction for Helios-like election schemes (Definition 27), which is parameterised on the choice of homomorphic encryption scheme and sigma protocols for the relations introduced in the following definition.

Definition 26 (from [SFC17]). *Let $(\text{Gen}, \text{Enc}, \text{Dec})$ be a homomorphic asymmetric encryption scheme and Σ be a sigma protocol for a binary relation R .⁴³*

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

Further, suppose that (pk, sk, \mathbf{m}) is the output of $\text{Gen}(\kappa; s)$, for some security parameter κ and coins s .

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

Definition 27 (Generalised Helios [SFC17]). *Suppose $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ is an additively homomorphic asymmetric encryption scheme, Σ_1 is a sigma protocol that proves correct key generation, Σ_2 is a sigma protocol that proves plaintext knowledge in a subspace, Σ_3 is a sigma protocol that proves correct decryption,*

⁴³Given 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 .

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 election scheme generalised Helios, denoted $\text{Helios}(\Pi, \Sigma_1, \Sigma_2, \Sigma_3, \mathcal{H}) = (\text{Setup}, \text{Vote}, \text{Tally})$, as follows.⁴⁴

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

```

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

```

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

- **Tally**($sk', \mathbf{bb}, nc, \kappa$). Initialise vectors \mathbf{v} of length nc and pf of length $nc-1$. Compute **for** $1 \leq j \leq nc$ **do** $\mathbf{v}[j] \leftarrow 0$. Parse sk' as a pair (pk, sk) . Output (\mathbf{v}, pf) if parsing fails. Let $\{b_1, \dots, b_\ell\}$ be the largest subset of \mathbf{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 $2 \cdot nc - 1$ and $\bigwedge_{j=1}^{nc-1} \text{VerCiph}((pk, b_i[j], \{0, 1\}), b_i[j + nc - 1], j, \kappa) = 1 \wedge \text{VerCiph}((pk, b_i[1] \otimes \dots \otimes b_i[nc - 1], \{0, 1\}), b_i[2 \cdot nc - 1], nc, \kappa) = 1$. If $\{b_1, \dots, b_\ell\} = \emptyset$, then output (\mathbf{v}, pf) , otherwise, compute:

```

for  $1 \leq j \leq nc - 1$  do
   $c \leftarrow b_1[j] \otimes \dots \otimes b_\ell[j]$ ;
   $\mathbf{v}[j] \leftarrow \text{Dec}(sk, c)$ ;
   $pf[j] \leftarrow \text{ProveDec}((pk, c, \mathbf{v}[j]), sk, \kappa)$ ;
 $\mathbf{v}[nc] \leftarrow \ell - \sum_{j=1}^{nc-1} \mathbf{v}[j]$ ;

```

Output (\mathbf{v}, pf) .

The above algorithms assume $nc > 1$. Smyth, Frink & Clarkson define special cases of Vote and Tally when $nc = 1$. We omit those cases for brevity and, henceforth, assume nc is always greater than one.

The generic construction can be instantiated to derive Helios 2.0 and Helios'16.

⁴⁴We omit algorithm Verify for brevity.

Definition 28 (Weak Fiat-Shamir transformation [BPW12a]). *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: $chal \leftarrow \mathcal{H}(comm)$.*

Definition 29 (Helios 2.0 [SFC17]). *Let \widehat{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 $\widehat{Helios}(\Pi, \Sigma_1, \Sigma_2, \Sigma_3, \mathcal{H})$, where Π is additively homomorphic El Gamal [CGS97, §2], Σ_1 is the sigma protocol for proving knowledge of discrete logarithms by Chaum et al. [CEvP87, Protocol 2], Σ_2 is the sigma protocol for proving knowledge of disjunctive equality between discrete logarithms by Cramer et al. [CFSY96, Figure 1], Σ_3 is the sigma protocol for proving knowledge of equality between discrete logarithms by Chaum and Pedersen [CP93, §3.2], and \mathcal{H} is SHA-256 [NIS12].*

Definition 30 (Helios 3.1.4 [SFC17]). *Election scheme Helios 3.1.4 is Helios 2.0 after modifying the sigma protocols to perform the checks proposed by Chang-Fong & Essex [CE16, §4].*

Definition 31 (Helios'16 [SFC17]). *Election scheme Helios'16 is $Helios(\Pi, \Sigma_1, \Sigma_2, \Sigma_3, \mathcal{H})$, where Π is additively homomorphic El Gamal [CGS97, §2], Σ_1 is the sigma protocol for proving knowledge of discrete logarithms by Chaum et al. [CEvP87, Protocol 2], Σ_2 is the sigma protocol for proving knowledge of disjunctive equality between discrete logarithms by Cramer et al. [CFSY96, Figure 1], Σ_3 is the sigma protocol for proving knowledge of equality between discrete logarithms by Chaum & Pedersen [CP93, §3.2], \mathcal{H} is a random oracle, and the sigma protocols are modified to perform the checks proposed by Chang-Fong & Essex [CE16, §4].*

Although Helios actually uses SHA-256 [NIS12], we assume that \mathcal{H} is a random oracle to prove Theorem 7. Moreover, we assume the sigma protocols used by Helios'16 satisfy the preconditions of generalised Helios, that is, [CEvP87, Protocol 2] is a sigma protocol for proving correct key generation, [CFSY96, Figure 1] is a sigma protocol for proving plaintext knowledge in a subspace, and [CP93, §3.2] is a sigma protocol for proving decryption. We leave formally proving this assumption as future work.

D.1 Proof of Theorem 7

The construction for Helios-like schemes produces election schemes with zero-knowledge tallying proofs (Lemma 27) that satisfy universal verifiability [SFC17] and, thus, Additivity (Lemma 31). They also satisfy ballot independence (Proposition 28). Hence, they satisfy ballot secrecy too (Theorem 4). We show that Helios'16 satisfies ballot secrecy.

Henceforth, we assume Π , Σ_1 , Σ_2 and Σ_3 satisfy the preconditions of Definition 27, and \mathcal{H} is a random oracle. Let $Helios(\Pi, \Sigma_1, \Sigma_2, \Sigma_3, \mathcal{H}) = (Setup, Vote,$

Tally) and $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$. Moreover, let $\text{FS}(\Sigma_1, \mathcal{H}) = (\text{ProveKey}, \text{VerKey})$, $\text{FS}(\Sigma_2, \mathcal{H}) = (\text{ProveCiph}, \text{VerCiph})$, and $\text{FS}(\Sigma_3, \mathcal{H}) = (\text{ProveDec}, \text{VerDec})$.

Lemma 27. *If $(\text{ProveDec}, \text{VerDec})$ is zero-knowledge, then $\text{Helios}(\Pi, \Sigma_1, \Sigma_2, \Sigma_3, \mathcal{H})$ has zero-knowledge tallying proofs.*

Proof sketch. Suppose \mathcal{A} is an adversary and κ is a security parameter. Further suppose (pk, sk, mb, mc) is an output of $\text{Setup}(\kappa)$, (nc, \mathbf{bb}) is an output of $\mathcal{A}(pk, \kappa)$, and (\mathbf{v}, pf) is an output of $\text{Tally}(sk, \mathbf{bb}, nc, \kappa)$, such that $|\mathbf{bb}| \leq mb \wedge nc \leq mc$. By inspection of algorithm Tally , tallying proof pf is a vector of proofs produced by ProveDec . Thus, there trivially exists a non-interactive proof system that could compute pf , moreover, that proof system is zero-knowledge because $(\text{ProveDec}, \text{VerDec})$ is zero-knowledge, which concludes our proof. \square

Proposition 28. *Suppose Π is perfectly correct and satisfies IND-CPA. Further suppose $(\text{ProveKey}, \text{VerKey})$ and $(\text{ProveCiph}, \text{VerCiph})$ satisfy special soundness and special honest verifier zero-knowledge. We have $\text{Helios}(\Pi, \Sigma_1, \Sigma_2, \Sigma_3, \mathcal{H})$ satisfies IND-CVA.*

Proof. By Theorem 19, the proof systems have extractors and simulators. Let SimProveKey , respectively SimProveCiph , be the simulator for $(\text{ProveKey}, \text{VerKey})$, respectively $(\text{ProveCiph}, \text{VerCiph})$. And let ExtProveCiph be the extractor for $(\text{ProveCiph}, \text{VerCiph})$.

Let IND-CPA* be a variant of IND-CPA in which: 1) the adversary outputs two vectors of messages \mathbf{m}_0 and \mathbf{m}_1 such that $|\mathbf{m}_0| = |\mathbf{m}_1|$ and for all $1 \leq i \leq |\mathbf{m}_0|$ we have $|\mathbf{m}_0[i]| = |\mathbf{m}_1[i]|$ and $\mathbf{m}_0[i]$ and $\mathbf{m}_1[i]$ are from the encryption scheme's message space, and 2) the challenger computes $c_1 \leftarrow \text{Enc}(pk, \mathbf{m}_\beta[1]); \dots; c_{|\mathbf{m}_\beta|} \leftarrow \text{Enc}(pk, \mathbf{m}_\beta[|\mathbf{m}_\beta|])$ and inputs $c_1, \dots, c_{|\mathbf{m}_\beta|}$ to the adversary. We have Π satisfies IND-CPA* [KL07, §10.2.2].

Suppose $\text{Helios}(\Pi, \Sigma_1, \Sigma_2, \Sigma_3, \mathcal{H})$ does not satisfy IND-CVA. Hence, there exists a probabilistic polynomial-time adversary \mathcal{A} , such that for all negligible functions negl , there exists a security parameter κ and $\frac{1}{2} + \text{negl}(\kappa) < \text{IND-CVA}(\Gamma, \mathcal{A}, \kappa)$. Since \mathcal{A} is a winning adversary, we have $\mathcal{A}(pk', \kappa)$ outputs (v_0, v_1, nc) such that $v_0 \neq v_1$ with non-negligible probability, hence, either $v_0 < v_1$ or $v_1 < v_0$. For brevity, we suppose $v_0 < v_1$. (Our proof can be adapted to consider cases such that $v_1 < v_0$, but these details provide little value, so we do not pursue them.) We construct the following adversary \mathcal{B} against IND-CPA* from \mathcal{A} :

- $\mathcal{B}(pk, \mathbf{m}, \kappa)$ outputs $((1, 0), (0, 1))$.
- $\mathcal{B}(\mathbf{c})$ proceeds as follows. First, compute:

$$\begin{aligned} \rho &\leftarrow \text{SimProveKey}((\kappa, pk, \mathbf{m}), \kappa); \\ pk' &\leftarrow (pk, \mathbf{m}, \rho); \\ (v_0, v_1, nc) &\leftarrow \mathcal{A}(pk', \kappa); \end{aligned}$$

Secondly, select coins r_1, \dots, r_{nc-1} uniformly at random and compute:

```

for  $j \in \{1, \dots, nc - 1\} \setminus \{v_0, v_1\}$  do
   $c_j \leftarrow \text{Enc}(pk, 0; r_j)$ ;
   $\sigma_j \leftarrow \text{ProveCiph}((pk, c_j, \{0, 1\}), (0, r_j), j, \kappa)$ ;
 $c_{v_0} \leftarrow \mathbf{c}[1]$ ;
 $\sigma_{v_0} \leftarrow \text{SimProveCiph}((pk, c_{v_0}, \{0, 1\}), v_0, \kappa)$ ;
if  $v_1 \neq nc$  then
   $c_{v_1} \leftarrow \mathbf{c}[2]$ ;
   $\sigma_{v_1} \leftarrow \text{SimProveCiph}((pk, c_{v_1}, \{0, 1\}), v_1, \kappa)$ ;
 $c \leftarrow c_1 \otimes \dots \otimes c_{nc-1}$ ;
 $\sigma_{nc} \leftarrow \text{SimProveCiph}((pk, c, \{0, 1\}), nc, \kappa)$ ;
 $b \leftarrow (c_1, \dots, c_{nc-1}, \sigma_1, \dots, \sigma_{nc})$ ;
 $\mathbf{bb} \leftarrow \mathcal{A}(b)$ ;

```

Thirdly, compute $\{b_1, \dots, b_\ell\}$ as the largest subset of \mathbf{bb} satisfying the conditions of algorithm Tally. Fourthly, initialise \mathbf{H} as a transcript of the random oracle's input and output, \mathbf{P} as a transcript of simulated proofs, \mathbf{Q} as a vector of length $nc - 1$, and \mathbf{v} as a zero-filled vector of length nc . Fifthly, compute:

```

 $\mathbf{Q} \leftarrow \left( \left( (pk, b_1[1], \{0, 1\}), b_1[nc] \right), \dots, \right.$ 
   $\left( (pk, b_\ell[1], \{0, 1\}), b_\ell[nc] \right), \dots,$ 
   $\left( (pk, b_1[nc - 1], \{0, 1\}), b_1[2 \cdot (nc - 1)] \right), \dots,$ 
   $\left. \left( (pk, b_\ell[nc - 1], \{0, 1\}), b_\ell[2 \cdot (nc - 1)] \right) \right)$ ;
 $\mathbf{W} \leftarrow \text{ExtProveCiph}(\mathbf{H}, \mathbf{P}, \mathbf{Q})$ ;
 $\mathbf{v} \leftarrow (\sum_{i=1}^\ell \mathbf{W}[i][1], \dots, \sum_{i=\ell \cdot (nc-2)+1}^{\ell \cdot (nc-1)} \mathbf{W}[i][1], \ell - \sum_{j=1}^{nc-1} \mathbf{v}[j])$ ;
 $g \leftarrow \mathcal{A}(\mathbf{v})$ ;

```

Finally, output g .

We prove that \mathcal{B} wins IND-CPA*.

Suppose (pk, sk, \mathbf{m}) is an output of $\text{Gen}(\kappa)$ and $(\mathbf{m}_0, \mathbf{m}_1)$ is an output of $\mathcal{B}(pk, \mathbf{m}, \kappa)$. Let $\beta \in \{0, 1\}$. Further suppose c_1 is an output of $\text{Enc}(pk, \mathbf{m}_\beta[1])$ and c_2 is an output of $\text{Enc}(pk, \mathbf{m}_\beta[2])$. Let $\mathbf{c} = (c_1, c_2)$. Moreover, suppose ρ is an output of $\text{SimProveKey}((\kappa, pk, \mathbf{m}), \kappa)$. Let $pk' = (pk, \mathbf{m}, \rho)$. Suppose (v_0, v_1, nc) is an output of $\mathcal{A}(pk', \kappa)$. Since SimProveKey is a simulator for $(\text{ProveKey}, \text{VerKey})$, we have \mathcal{B} simulates the challenger in IND-CVA to $\mathcal{A}(pk', \kappa)$. In particular, pk' is a triple containing a public key and corresponding message space generated Gen , and a (simulated) proof of correct key generation. Suppose \mathcal{B} computes b and \mathbf{bb} is an output of $\mathcal{A}(b)$. Further suppose \mathcal{B} computes \mathbf{v} , and g is an output of $\mathcal{A}(\mathbf{v})$. The following claims prove that \mathcal{B} simulates the challenger in IND-CVA to $\mathcal{A}(b)$ and $\mathcal{A}(\mathbf{v})$, hence, $g = \beta$, with at least the probability that \mathcal{A} wins IND-CVA, concluding our proof.

Claim 29. *Adversary \mathcal{B} 's computation of b is equivalent to computing b as $b \leftarrow \text{Vote}(pk', v_\beta, nc, \kappa)$.*

Proof of Claim 29. We have pk' parses as a vector (pk, \mathbf{m}, ρ) . Moreover, since (pk, sk, \mathbf{m}) is an output of $\text{Gen}(\kappa)$, there exist coins r such that $(pk, sk, \mathbf{m}) = \text{Gen}(\kappa; r)$. Hence, (sk, r) is a witness for statement (κ, pk, \mathbf{m}) . Furthermore, since SimProveKey is a simulator for $(\text{ProveKey}, \text{VerKey})$ and proofs output by ProveKey are indistinguishable from outputs of SimProveKey , we have $\text{VerKey}((\kappa, pk, \mathbf{m}), \rho, \kappa) \kappa, pk, \mathbf{m} \rho = 1$, with non-negligible probability. In addition, since \mathcal{B} is a winning adversary, we have $v_0, v_1 \in \{1, \dots, nc\}$, with non-negligible probability. It follows that $\text{Vote}(pk', v_\beta, nc, \kappa)$ does not output \perp , with non-negligible probability. Indeed, computation $b \leftarrow \text{Vote}(pk', v_\beta, nc, \kappa)$ is equivalent to the following. Select coins r_1, \dots, r_{nc-1} uniformly at random and compute:

```

for  $1 \leq j \leq nc - 1$  do
  if  $j = v_\beta$  then  $m_j \leftarrow 1$ ; else  $m_j \leftarrow 0$ ;
   $c_j \leftarrow \text{Enc}(pk, m_j; r_j)$ ;
   $\sigma_j \leftarrow \text{ProveCiph}((pk, c_j, \{0, 1\}), (m_j, r_j), j, \kappa)$ ;
 $c \leftarrow c_1 \otimes \dots \otimes c_{nc-1}$ ;
 $m \leftarrow m_1 \odot \dots \odot m_{nc-1}$ ;
 $r \leftarrow r_1 \oplus \dots \oplus r_{nc-1}$ ;
 $\sigma_{nc} \leftarrow \text{ProveCiph}((pk, c, \{0, 1\}), (m, r), nc, \kappa)$ ;
 $b \leftarrow (c_1, \dots, c_{nc-1}, \sigma_1, \dots, \sigma_{nc})$ ;

```

Since $v_\beta \in \{v_0, v_1\}$, ciphertexts computed by the above for-loop all contain plaintext 0, except (possibly) ciphertext c_{v_0} and, if defined, ciphertext c_{v_1} . (Ciphertext c_{v_1} only exists if $v_1 < nc$.) Given that $v_0 < v_1 \leq nc$, ciphertext c_{v_0} contains $1 - \beta$, i.e., if $\beta = 0$, then c_{v_0} contains 1, otherwise ($\beta = 1$), c_{v_0} contains 0. If $v_1 < nc$, then ciphertext c_{v_1} contains β . Moreover, since \odot is the addition operator in group (\mathbf{m}, \odot) and 0 is the identity element in that group, if $v_1 = nc$, then plaintext m computed by the above algorithm is $1 - \beta$, otherwise, $m = 1 - \beta \odot \beta = 1$. Hence, the above algorithm is equivalent to selecting coins r_1, \dots, r_{nc-1} uniformly at random and computing:

```

for  $j \in \{1, \dots, nc - 1\} \setminus \{v_0, v_1\}$  do
   $c_j \leftarrow \text{Enc}(pk, 0; r_j)$ ;
   $\sigma_j \leftarrow \text{ProveCiph}((pk, c_j, \{0, 1\}), (0, r_j), j, \kappa)$ ;
 $c_{v_0} \leftarrow \text{Enc}(pk, 1 - \beta; r_{v_0})$ ;
 $\sigma_{v_0} \leftarrow \text{ProveCiph}((pk, c_{v_0}, \{0, 1\}), (1 - \beta, r_{v_0}), v_0, \kappa)$ ;
if  $v_1 \neq nc$  then
   $c_{v_1} \leftarrow \text{Enc}(pk, \beta; r_{v_1})$ ;
   $\sigma_{v_1} \leftarrow \text{ProveCiph}((pk, c_{v_1}, \{0, 1\}), (\beta, r_{v_1}), v_1, \kappa)$ ;
 $c \leftarrow c_1 \otimes \dots \otimes c_{nc-1}$ ;
if  $v_1 = nc$  then  $m \leftarrow 1 - \beta$ ; else  $m \leftarrow 1$ ;
 $r \leftarrow r_1 \oplus \dots \oplus r_{nc-1}$ ;
 $\sigma_{nc} \leftarrow \text{ProveCiph}((pk, c, \{0, 1\}), (m, r), nc, \kappa)$ ;
 $b \leftarrow (c_1, \dots, c_{nc-1}, \sigma_1, \dots, \sigma_{nc})$ ;

```


Computation $c_{v_0} \leftarrow \text{Enc}(pk, 1 - \beta; r_{v_0})$ is equivalent to $c_{v_0} \leftarrow \mathbf{c}[1]$, because if $\beta = 0$, then $\mathbf{c}[1]$ contains plaintext 1, otherwise ($\beta = 1$), $\mathbf{c}[1]$ contains plaintext 0. Similarly, if $v_1 \neq nc$, then computation $c_{v_1} \leftarrow \text{Enc}(pk, \beta; r_{v_1})$ is equivalent to $c_{v_1} \leftarrow \mathbf{c}[1]$. Moreover, proof $\text{ProveCiph}((pk, c_{v_0}, \{0, 1\}), (1 - \beta, r_{v_0}), v_0, \kappa)$, respectively $\text{ProveCiph}((pk, c_{v_1}, \{0, 1\}), (\beta, r_{v_1}), v_1, \kappa)$, can be simulated by $\text{SimProveCiph}((pk, c_{v_0}, \{0, 1\}), v_0, \kappa)$, respectively $\text{SimProveCiph}((pk, c_{v_1}, \{0, 1\}), v_1, \kappa)$. Furthermore,

```

c ← c1 ⊗ ⋯ ⊗ cnc-1;
if v1 = nc then m ← 1 - β; else m ← 1;
r ← r1 ⊕ ⋯ ⊕ rnc-1;
σnc ← ProveCiph((pk, c, {0, 1}), (m, r), nc, κ);

```

can be simulated by

```

c ← c1 ⊗ ⋯ ⊗ cnc-1;
σnc ← SimProveCiph((pk, c, {0, 1}), nc, κ);

```

Hence, we conclude the proof of this claim.

Claim 30. *Adversary \mathcal{B} 's computation of \mathbf{v} is equivalent to computing \mathbf{v} as $(\mathbf{v}, pf) \leftarrow \text{Tally}(sk', \mathbf{bb}, nc, \kappa)$, where $sk' = (pk, sk)$.*

Proof of Claim 30. Let $\{b_1, \dots, b_\ell\}$ be the largest subset of \mathbf{bb} satisfying the conditions of algorithm **Tally**. It is trivial to see that the claim holds when $\{b_1, \dots, b_\ell\} = \emptyset$, because \mathbf{v} is computed as a zero-filled vector of length nc in both cases. We prove the claim also holds when $\{b_1, \dots, b_\ell\} \neq \emptyset$.

By simulation sound extractability, for all $1 \leq i \leq \ell$ and $1 \leq j \leq nc - 1$, there exists a message $m_{i,j} \in \{0, 1\}$ and coins $r_{i,j}$ and $r_{i,j+nc-1}$ such that $b_i[j] = \text{Enc}(pk, m_{i,j}; r_{i,j})$ and $b_i[j + nc - 1] = \text{ProveCiph}((pk, b_i[j], \{0, 1\}), (m_{i,j}, r_{i,j}), j, \kappa; r_{i,j+nc-1})$, with overwhelming probability. Suppose \mathbf{Q} and \mathbf{W} are computed by \mathcal{B} . We have for all $1 \leq i \leq \ell$ and $1 \leq j \leq nc - 1$ that $\mathbf{Q}[\ell \cdot (j - 1) + i] = ((pk, b_i[j], \{0, 1\}), b_i[j + nc - 1])$ and $\mathbf{W}[\ell \cdot (j - 1) + i]$ is a witness for $(pk, b_i[j], \{0, 1\})$, i.e., $(m_{i,j}, r_{i,j})$, and $\mathbf{W}[\ell \cdot (j - 1) + i][1] = m_{i,j}$. Hence, adversary \mathcal{B} 's computation of \mathbf{v} is equivalent to computing \mathbf{v} as:

$$\mathbf{v} \leftarrow (\sum_{i=1}^{\ell} m_{i,1}, \dots, \sum_{i=1}^{\ell} m_{i,nc-1}, \ell - \sum_{j=1}^{nc-1} \mathbf{v}[j])$$

Moreover, computing \mathbf{v} as $(\mathbf{v}, pf) \leftarrow \text{Tally}(sk', \mathbf{bb}, nc, \kappa)$ is equivalent to initializing \mathbf{v} as a zero-filled vector of length nc and computing

```

for 1 ≤ j ≤ nc - 1 do
  | c ← b1[j] ⊗ ⋯ ⊗ bℓ[j];
  | v[j] ← Dec(sk, c);
v[nc] ← ℓ - ∑j=1nc-1 v[j];

```

Since Π is a homomorphic encryption scheme, we have for all $1 \leq j \leq nc - 1$ that $b_1[j] \otimes \dots \otimes b_\ell[j]$ is a ciphertext with overwhelming probability. And although ciphertext $b_1[j] \otimes \dots \otimes b_\ell[j]$ may not have been computed using coins chosen uniformly at random, we nevertheless have $\text{Dec}(sk, b_1[j] \otimes \dots \otimes b_\ell[j]) = m_{1,j} \odot \dots \odot m_{\ell,j}$ with overwhelming probability, because Π is perfectly correct. It

follows that $\mathbf{v} = (m_{1,1} \odot \cdots \odot m_{\ell,1}, \dots, m_{1,nc-1} \odot \cdots \odot m_{\ell,nc-1}, \ell - \sum_{j=1}^{nc-1} \mathbf{v}[j])$, with overwhelming probability. Let mb be the largest integer such that $\{0, \dots, mb\} \subseteq \mathbf{m}$. Since \mathcal{A} is a winning adversary, we have $\ell \leq mb$. Moreover, since $m_{1,j}, \dots, m_{\ell,j} \in \{0, 1\}$ for all $1 \leq j \leq nc - 1$ and \odot is the addition operator in group (\mathbf{m}, \odot) , we have $m_{1,j} \odot \cdots \odot m_{\ell,j} = \sum_{i=1}^{\ell} m_{i,j}$, which suffices to conclude the proof of this claim. \square

For Helios’16, encryption scheme Π is additively homomorphic El Gamal [CGS97, §2]. Moreover, $(\text{ProveKey}, \text{VerKey})$, respectively $(\text{ProveCiph}, \text{VerCiph})$ and $(\text{ProveDec}, \text{VerDec})$, is the non-interactive proof system derived by application of the Fiat-Shamir transformation [FS87] to a random oracle \mathcal{H} and the sigma protocol for proving knowledge of discrete logarithms by Chaum *et al.* [CEvP87, Protocol 2], respectively the sigma protocol for proving knowledge of disjunctive equality between discrete logarithms by Cramer *et al.* [CFSY96, Figure 1] and the sigma protocol for proving knowledge of equality between discrete logarithms by Chaum & Pedersen [CP93, §3.2].

Bernhard, Pereira & Warinschi [BPW12a, §4] remark that the sigma protocols underlying non-interactive proof systems $(\text{ProveKey}, \text{VerKey})$ and $(\text{ProveCiph}, \text{VerCiph})$ both satisfy special soundness and special honest verifier zero-knowledge, hence, Theorem 19 is applicable. Bernhard, Pereira & Warinschi also remark that the sigma protocol underlying $(\text{ProveDec}, \text{VerDec})$ satisfies special soundness and “almost special honest verifier zero-knowledge” and argue that “we could fix this[, but] it is easy to see that ... all relevant theorems [including Theorem 19] still hold.” We adopt the same position and assume that Theorem 19 is applicable.

Proof of Theorem 7. Helios’16 has zero-knowledge tallying proofs (Lemma 27), subject to the applicability of Theorem 19 to the sigma protocol underlying $(\text{ProveDec}, \text{VerDec})$. Moreover, since Helios’16 satisfies UV [SFC17], we have Helios’16 satisfies $\text{Additivity}(\Gamma, \mathcal{A}, \kappa)$ (Lemma 31). Furthermore, since El Gamal satisfies IND-CPA [TY98, KL07] and is perfectly correct, and since non-interactive proof systems $(\text{ProveKey}, \text{VerKey})$ and $(\text{ProveCiph}, \text{VerCiph})$ satisfy special soundness and special honest verifier zero-knowledge, we have Helios’16 satisfies IND-CVA (Proposition 28). Hence, Helios’16 satisfies Ballot-Secrecy too (Theorem 4). \square

E Universal verifiability implies tally soundness

We recall the definition of universal verifiability by Smyth, Frink & Clarkson [SFC17] and show that verifiable election schemes satisfy Tally-Soundness (Lemma 31). This is useful to simplify applications of Theorems 4, 14, & 33. Indeed, our ballot-secrecy proofs for Helios and Helios Mixnet make use of this result.

We extend our syntax for election schemes (Definition 1) to include a probabilistic polynomial-time algorithm Verify :

- **Verify**, denoted $s \leftarrow \text{Verify}(pk, \mathbf{bb}, nc, \mathbf{v}, pf, \kappa)$, is run to audit an election. It takes as input a public key pk , a bulletin board \mathbf{bb} , some number of candidates nc , an election outcome \mathbf{v} , a tallying proof pf , and a security parameter κ . It outputs a bit s , where 1 signifies success and 0 failure.

We previously omitted algorithm **Verify**, because we did not focus on verifiability in the main body.

For universal verifiability, anyone must be able to check whether the election outcome represents the votes used to construct ballots on the bulletin board. The formal definition of universal verifiability by Smyth, Frink & Clarkson requires algorithm **Verify** to accept if and only if the election outcome is correct. The *if* requirement is captured by completeness (Definition 32), which stipulates that election outcomes produced by algorithm **Tally** will actually be accepted by algorithm **Verify**. And the *only if* requirement is captured by soundness (Definition 34), which challenges an adversary to concoct a scenario in which algorithm **Verify** accepts, but the election outcome is not correct.

Definition 32 (Completeness [SFC17]). *An election scheme (Setup, Vote, Tally, Verify) satisfies completeness, if for all probabilistic polynomial-time adversaries \mathcal{A} , there exists a negligible function negl , such that for all security parameters κ , we have $\Pr[(pk, sk, mb, mc) \leftarrow \text{Setup}(\kappa); (\mathbf{bb}, nc) \leftarrow \mathcal{A}(pk, \kappa); (\mathbf{v}, pf) \leftarrow \text{Tally}(sk, \mathbf{bb}, nc, \kappa) : |\mathbf{bb}| \leq mb \wedge nc \leq mc \Rightarrow \text{Verify}(pk, \mathbf{bb}, nc, \mathbf{v}, pf, \kappa) = 1] > 1 - \text{negl}(\kappa)$.*

Definition 33 (Injectivity [Smy18b, SFC17]). *An election scheme (Setup, Vote, Tally, Verify) satisfies injectivity, if for all probabilistic polynomial-time adversaries \mathcal{A} , security parameters κ and computations $(pk, nc, v, v') \leftarrow \mathcal{A}(\kappa); b \leftarrow \text{Vote}(pk, v, nc, \kappa); b' \leftarrow \text{Vote}(pk, v', nc, \kappa)$ such that $v \neq v' \wedge b \neq \perp \wedge b' \neq \perp$, we have $b \neq b'$.*

Definition 34 (Soundness [SFC17]). *An election scheme $\Gamma = (\text{Setup}, \text{Vote}, \text{Tally}, \text{Verify})$ satisfies soundness, if Γ satisfies injectivity and for all probabilistic polynomial-time adversaries \mathcal{A} , there exists a negligible function negl , such that for all security parameters κ , we have $\Pr[(pk, nc, \mathbf{bb}, \mathbf{v}, pf) \leftarrow \mathcal{A}(\kappa) : \mathbf{v} \neq \text{correct-outcome}(pk, nc, \mathbf{bb}, \kappa) \wedge \text{Verify}(pk, \mathbf{bb}, nc, \mathbf{v}, pf, \kappa) = 1] \leq \text{negl}(\kappa)$.*

Definition 35 (UV [Smy18b, SFC17]). *An election scheme Γ satisfies universal verifiability (UV), if completeness, injectivity and soundness are satisfied.*

Lemma 31. *If election scheme Γ satisfies completeness and soundness, then Γ satisfies Tally-Soundness.*

Proof. Let $\Gamma = (\text{Setup}, \text{Vote}, \text{Tally}, \text{Verify})$. Suppose there exists a probabilistic polynomial-time adversary \mathcal{A} that wins Tally-Soundness against Γ . We construct an adversary \mathcal{B} against Exp-UV-Ext from \mathcal{A} . We define \mathcal{B} such that $\mathcal{B}(\kappa) = (pk, sk, mb, mc) \leftarrow \text{Setup}(\kappa); (nc, \mathbf{bb}) \leftarrow \mathcal{A}(pk, \kappa); (\mathbf{v}, pf) \leftarrow \text{Tally}(sk, \mathbf{bb}, nc, \kappa)$; **return** $(pk, nc, \mathbf{bb}, \mathbf{v}, pf)$. Suppose (pk, sk, mb, mc) is an output of $\text{Setup}(\kappa)$, (nc, \mathbf{bb}) is an output of $\mathcal{A}(pk, \kappa)$, and (\mathbf{v}, pf) is an output of $\text{Tally}(sk, \mathbf{bb}, nc, \kappa)$. Since

\mathcal{A} is a winning adversary, we have $\mathbf{v} \neq \text{correct-outcome}(pk, nc, \mathbf{bb}, \kappa) \wedge |\mathbf{bb}| \leq mb \wedge nc \leq mc$, with non-negligible probability. And, by completeness, we have $\text{Verify}(pk, \mathbf{bb}, nc, \mathbf{v}, pf, \kappa) = 1$, with overwhelming probability. Thereby concluding our proof. \square

The reverse implication of Lemma 31 does not hold: Observe that Tally-Soundness only ensures algorithm Tally tallies ballots correctly, whereas UV additionally ensures that anyone can check whether ballots are tallied correctly.

F Encryption-based voting systems

We have seen that election scheme $\text{Enc2Vote}(\Pi)$ satisfies HK-Injectivity, if Π is perfectly correct (Lemma 10). But, HK-Injectivity assumes public keys are computed using the key generation algorithm. Thus, perfect correctness is insufficient to ensure injectivity when public keys are controlled by an adversary. Nonetheless, this can be ensured using proofs of correct key generation. A subclass of schemes generated by Enc2Vote prove correct key generation. Indeed, we can consider schemes $\text{Enc2Vote}(\Pi)$ such that Gen proves correct key generation and Enc verifies such proofs, where $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$. Alternatively, we can couple Enc2Vote with proofs of correct key generation:

Definition 36 (Enc2Vote^+ [Smy18b]). *Suppose $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ is an asymmetric encryption scheme, Σ is a sigma protocol that proves correct key generation, and \mathcal{H} is a hash function. Let $\text{FS}(\Sigma, \mathcal{H}) = (\text{ProveKey}, \text{VerKey})$. We define $\text{Enc2Vote}^+(\Pi, \Sigma, \mathcal{H}) = (\text{Setup}, \text{Vote}, \text{Tally})$ such that:*

- $\text{Setup}(\kappa)$ selects coins s uniformly at random, computes $(pk, sk, \mathbf{m}) \leftarrow \text{Gen}(\kappa; s)$; $\rho \leftarrow \text{ProveKey}((\kappa, pk, \mathbf{m}), (sk, s), \kappa)$; $pk' \leftarrow (pk, \mathbf{m}, \rho)$; $sk' \leftarrow (pk, sk)$, derives mc as the largest integer such that $\{0, \dots, mc\} \subseteq \{0\} \cup \mathbf{m}$ and for all $m_0, m_1 \in \{1, \dots, mc\}$ we have $|m_0| = |m_1|$, and outputs $(pk', sk', p(\kappa), mc)$, where p is a polynomial function.
- $\text{Vote}(pk', v, nc, \kappa)$ parses pk' as vector (pk, \mathbf{m}, ρ) , outputting \perp if parsing fails or $\text{VerKey}((\kappa, pk, \mathbf{m}), \rho, \kappa) \neq 1 \vee v \notin \{1, \dots, nc\} \vee \{1, \dots, nc\} \not\subseteq \mathbf{m}$, computes $b \leftarrow \text{Enc}(pk, v)$, and outputs b .
- $\text{Tally}(sk', \mathbf{bb}, nc, \kappa)$ initialises \mathbf{v} as a zero-filled vector of length nc , parses sk' as pair (pk, sk) , outputting (\mathbf{v}, \perp) if parsing fails, computes **for** $b \in \mathbf{bb}$ **do** $v \leftarrow \text{Dec}(sk, b)$; **if** $1 \leq v \leq nc$ **then** $\mathbf{v}[v] \leftarrow \mathbf{v}[v] + 1$, and outputs (\mathbf{v}, ϵ) , where ϵ is a constant symbol.

Lemma 32. *Given an asymmetric encryption scheme Π satisfying IND-CPA, a sigma protocol Σ that proves correct key generation, and a hash function \mathcal{H} , we have $\text{Enc2Vote}^+(\Pi, \Sigma, \mathcal{H})$ is an election scheme.*

A proof of Lemma 32 follows from [Smy18b].⁴⁵

Although the set of election schemes produced by Enc2Vote^+ is not a subset of the schemes produced by Enc2Vote , there is nonetheless a straightforward mapping from the former to the latter. Thus, the results in Section 5 also hold for Enc2Vote^+ :

Theorem 33. *Let $\text{Enc2Vote}^+(\Pi, \Sigma, \mathcal{H}) = (\text{Setup}, \text{Vote}, \text{Tally})$, where Π is an asymmetric encryption scheme, Σ is a sigma protocol that proves correct key generation, and \mathcal{H} is a random oracle. Moreover, let $\Gamma = (\text{Setup}, \text{Vote}, \text{Tally}')$ for some algorithm Tally' such that Γ is an election scheme with zero-knowledge tallying proofs. Suppose Π is perfectly correct and satisfies IND-PA0 and well-definedness. Moreover, suppose Σ is perfectly complete and $\text{FS}(\Sigma, \mathcal{H})$ satisfies zero-knowledge. Further suppose Γ satisfies Tally-Soundness. We have Γ satisfies Ballot-Secrecy.*

Proof. Let $\text{FS}(\Sigma, \mathcal{H}) = (\text{ProveKey}, \text{VerKey})$ and $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$. Moreover, let asymmetric encryption scheme $\Pi' = (\text{Gen}', \text{Enc}', \text{Dec})$ such that

- $\text{Gen}'(\kappa)$ selects coins s uniformly at random, computes $(pk, sk, \mathbf{m}) \leftarrow \text{Gen}(\kappa; s)$; $\rho \leftarrow \text{ProveKey}((\kappa, pk, \mathbf{m}), (sk, s), \kappa)$; $pk' \leftarrow (pk, \mathbf{m}, \rho)$, and outputs (pk', sk, \mathbf{m}) .
- $\text{Enc}'(pk, v)$ parses pk' as a vector (pk, \mathbf{m}, ρ) , outputting \perp if parsing fails or $\text{VerKey}((\kappa, pk, \mathbf{m}), \rho, \kappa) \neq 1$, computes ciphertext $c \leftarrow \text{Enc}(pk, v)$, and outputs c .

Since Π is perfectly correct and Σ is perfectly complete, we have Π' is perfectly correct. Moreover, since Π satisfies well-definedness, we have Π' does too. Furthermore, since $\text{FS}(\Sigma, \mathcal{H})$ satisfies zero-knowledge and Π satisfies IND-PA0, we have Π' satisfies IND-PA0. It follows that $\text{Enc2Vote}(\Pi')$ satisfies Tally-Soundness and IND-CVA (Corollary 11 & Lemma 13).

We have $\text{Enc2Vote}(\Pi') = (\text{Setup}', \text{Vote}', \text{Tally})$ such that Setup' is Setup except Setup outputs public key pk' as a vector (pk, \mathbf{m}, ρ) , whereas Setup' outputs public key (pk, \mathbf{m}) . Moreover, Vote' is Vote except Vote inputs public key (pk, \mathbf{m}) whereas Vote' inputs public key (pk, \mathbf{m}, ρ) . (This blight motivated the inclusion of this appendix.) Hence, it is straightforward to see that $\text{Enc2Vote}^+(\Pi, \Sigma, \mathcal{H})$ satisfies Tally-Soundness and IND-CVA, because $\text{Enc2Vote}(\Pi')$ does. Thus, Γ satisfies IND-CVA (Proposition 9) and Ballot-Secrecy (Theorem 4 & Lemma 8). \square

G Helios Mixnet

We recall a generic construction for election schemes similar to Helios Mixnet (Definition 38). The construction is parameterised on the choice of homomor-

⁴⁵Smyth considers instantiating Enc2Vote^+ with a broad class of asymmetric encryption schemes that produce distinct ciphertexts with overwhelming probability [Smy18b], whereas we consider a strictly narrower class of schemes satisfying IND-CPA. This avoids having to recall Smyth's notion of distinct ciphertexts.

phic encryption scheme and sigma protocols for the relations introduced in Definition 26 and the following definition.

Definition 37 (from [SFC17]). *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 that $(pk, sk) = \text{Gen}(\kappa; s)$, for some security parameter κ and coins s , and \mathbf{m} is the encryption scheme's plaintext space.*

- Σ proves plaintext knowledge if $((pk, c), (m, r)) \in R \Leftrightarrow c = \text{Enc}(pk, m; r) \wedge m \in \mathbf{m}$.
- Σ proves mixing 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{c} and \mathbf{c}' are both vectors of ciphertexts encrypted under pk , \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 38 (HeliosM [Smy18c, QS18b]). *Suppose $\Pi_0 = (\text{Gen}, \text{Enc}, \text{Dec})$ is a homomorphic asymmetric encryption algorithm, Σ_1 is a sigma protocol that proves correct key construction, Σ_2 is a sigma protocol that proves plaintext knowledge, and \mathcal{H} is a hash function. Let $\text{FS}(\Sigma_1, \mathcal{H}) = (\text{ProveKey}, \text{VerKey})$ and $\text{FS}(\Sigma_2, \mathcal{H}) = (\text{ProveCiph}, \text{VerCiph})$. Moreover, let $\pi(\Pi, \Sigma_2, \mathcal{H}) = (\text{Gen}, \text{Enc}', \text{Dec}')$ be an asymmetric encryption scheme such that:*

- $\text{Enc}'(pk, v)$ selects coins r uniformly at random, computes $c \leftarrow \text{Enc}(pk, v; r)$; $\sigma \leftarrow \text{ProveCiph}((pk, c), (v, r), \kappa)$, and outputs (c, σ) .
- $\text{Dec}'(sk, c')$ parses c' as (c, σ) , outputting \perp if parsing fails or $\text{VerCiph}((pk, c), \sigma, \kappa) \neq 1$, computes $v \leftarrow \text{Dec}(sk, c)$, and outputs v .

Let $\text{Enc2Vote}^+(\pi(\Pi, \Sigma_2, \mathcal{H}), \Sigma_1, \mathcal{H}) = (\text{Setup}, \text{Vote}, \text{Tally}')$. Suppose Σ_3 is a sigma protocol that proves correct decryption and Σ_4 is a sigma protocol that proves mixing. Let $\text{FS}(\Sigma_3, \mathcal{H}) = (\text{ProveDec}, \text{VerDec})$ and $\text{FS}(\Sigma_4, \mathcal{H}) = (\text{ProveMix}, \text{VerMix})$. We define $\text{HeliosM}(\Pi, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \mathcal{H}) = (\text{Setup}, \text{Vote}, \text{Tally})$, where algorithm Tally is defined below.⁴⁶

$\text{Tally}(sk', nc, \mathbf{bb}, \kappa)$ initialises \mathbf{v} as a zero-filled vector of length nc ; parses sk' as a pair (pk, sk) , outputting (\mathbf{v}, \perp) if parsing fails; and proceeds as follows:

1. Remove invalid ballots. Let $\{b_1, \dots, b_\ell\}$ be the largest subset of \mathbf{bb} such that for all $1 \leq i \leq \ell$ we have b_i is a pair and $\text{VerCiph}((pk, b_i[1]), b_i[2], \kappa) = 1$. If $\{b_1, \dots, b_\ell\} = \emptyset$, then output (\mathbf{v}, \perp) .
2. Mix. Select a permutation χ on $\{1, \dots, \ell\}$ uniformly at random, initialise \mathbf{bb} and \mathbf{r} as a vector of length ℓ , fill \mathbf{r} with coins chosen uniformly at random, and compute

⁴⁶We omit algorithm Verify for brevity.

for $1 \leq i \leq \ell$ **do**
 $\quad \perp \mathbf{bb}[i] \leftarrow b_{\chi(i)}[1] \otimes \text{Enc}(pk, \mathbf{e}; \mathbf{r}[i]);$
 $\quad pf_1 \leftarrow \text{ProveMix}((pk, (b_1[1], \dots, b_\ell[1]), \mathbf{bb}), (\mathbf{r}, \chi), \kappa);$

where \mathbf{e} is an identity element of Π 's message space with respect to \odot .

3. Decrypt. Initialise \mathbf{W} and pf_2 as vectors of length ℓ and compute:

for $1 \leq i \leq \ell$ **do**
 $\quad \mathbf{W}[i] \leftarrow \text{Dec}(sk, \mathbf{bb}[i]);$
 $\quad pf_2[i] \leftarrow \text{ProveDec}((pk, \mathbf{bb}[i], \mathbf{W}[i]), sk, \kappa);$
if $1 \leq \mathbf{W}[i] \leq nc$ **then**
 $\quad \perp \mathbf{v}[\mathbf{W}[i]] \leftarrow \mathbf{v}[\mathbf{W}[i]] + 1;$

Output $(\mathbf{v}, (\mathbf{bb}, pf_1, \mathbf{W}, pf_2))$.

Definition 39 (HeliosM'17). HeliosM'17 is the set of election schemes that includes every $\text{HeliosM}(\Pi_0, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \mathcal{H})$ such that $\Pi_0, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4$ and \mathcal{H} satisfy the preconditions of Definition 38, moreover, Π_0 is perfectly correct and Σ_1 and Σ_2 are perfectly complete, furthermore, Π_0 satisfies IND-CPA, $\Sigma_1, \Sigma_2, \Sigma_3$ and Σ_4 satisfy special soundness and special honest verifier zero-knowledge, \mathcal{H} is a random oracle, and $\text{HeliosM}(\Pi_0, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \mathcal{H})$ satisfies UV.

Smyth has shown that there exists an election scheme in HeliosM'17 that satisfies UV [Smy18c]. Hence, set HeliosM'17 is not empty.

G.1 Proof of Theorem 15

Let election scheme $\Gamma = \text{HeliosM}(\Pi_0, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \mathcal{H}) = (\text{Setup}, \text{Vote}, \text{Tally})$ and asymmetric encryption scheme $\Pi = \pi(\Pi_0, \Sigma_2, \mathcal{H})$. It follows that election scheme $\text{Enc2Vote}^+(\Pi, \Sigma_1, \mathcal{H}) = (\text{Setup}, \text{Vote}, \text{Tally}')$. Moreover, since Σ_1 satisfies special soundness and special honest verifier zero-knowledge, we have $\text{FS}(\Sigma_1, \mathcal{H})$ satisfies zero-knowledge (Theorem 19). We use Theorem 33 to prove that $\Gamma \in \text{HeliosM'17}$ satisfies Ballot-Secrecy.

Since Π_0 is perfectly correct and Σ_2 is perfectly complete, we have Π is a perfectly correct. Moreover, since Σ_2 satisfies special soundness and special honest verifier zero-knowledge, we have $\text{FS}(\Sigma_2, \mathcal{H})$ satisfies simulation sound extractability (Theorem 19), hence, Π satisfies CNM-CPA [BPW12a, Theorem 2] and, equivalently, IND-PA0 [BS99].

To prove $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ satisfies well-definedness, suppose \mathcal{A} is a probabilistic polynomial-time adversary, κ is a security parameter, (pk, sk, \mathbf{m}) is an output of $\text{Gen}(\kappa)$, and c is an output of $\mathcal{A}(pk, \mathbf{m}, \kappa)$ such that $\text{Dec}(sk, c) \neq \perp$. By definition of Dec, we have c is a pair (hence, $c \neq \perp$) such that $\text{FS}(\Sigma_2, \mathcal{H})$ can verify $c[2]$ with respect to pk and $c[1]$. Since $\text{FS}(\Sigma_2, \mathcal{H})$ satisfies simulation sound extractability, we have $c[2]$ is a proof computed using $\text{FS}(\Sigma_2, \mathcal{H})$ and there exists plaintext $m \in \mathbf{m}$ and coins r such that $c[1] = \text{Enc}(pk, m; r)$, with overwhelming probability. Thus, Π satisfies well-definedness.

Since Σ_3 and Σ_4 satisfy special soundness and special honest verifier zero-knowledge, we have $\text{FS}(\Sigma_3, \mathcal{H})$ and $\text{FS}(\Sigma_4, \mathcal{H})$ satisfy zero-knowledge (Theorem 19), therefore, Γ has zero-knowledge tallying proofs by reasoning similar to that given in the proof sketch of Lemma 27. Moreover, since Γ satisfies universal verifiability, we have Γ satisfies Tally-Soundness (Lemma 31).

We conclude by Theorem 33. □

H Ballot-Security is strictly stronger than IND-SEC

Smyth & Bernhard propose definitions of ballot secrecy that consider an adversary that cannot control the bulletin board nor the communication channel [SB13, SB14]. As discussed in Section 8, their original definition [SB13] is too strong [BCG⁺15b, §3.5] and they propose a revision [SB14]. We recall their syntax (Definition 40) and revised definition (Definition 41), define a transformation from their syntax to ours (Definition 42), and prove **Ballot-Security** is strictly stronger than their definition (Theorem 35).

Definition 40 (Election scheme with a trusted bulletin board). *An election scheme with a trusted bulletin board is a tuple of efficient algorithms (Setup, Vote, BB, Tally) such that:*

- **Setup**, denoted $(pk, sk, \mathbf{m}, \mathbf{bb}) \leftarrow \text{Setup}(\kappa)$, takes a security parameter κ as input and outputs a key pair pk, sk , a vote space \mathbf{m} , and a bulletin board \mathbf{bb} , where \mathbf{m} and \mathbf{bb} are both sets.
- **Vote**, denoted $b \leftarrow \text{Vote}(pk, v)$, takes a public key pk and vote v as input and outputs a ballot b . Vote v should be selected from the vote space \mathbf{m} .
- **BB**, denoted $\mathbf{bb}' \leftarrow \text{BB}(\mathbf{bb}, b)$, takes a bulletin board \mathbf{bb} and ballot b as input and outputs an updated bulletin board \mathbf{bb}' .
- **Tally**, denoted $(\mathbf{o}, pf) \leftarrow \text{Tally}(sk, \mathbf{bb})$, takes a private key sk and bulletin board \mathbf{bb} as input and outputs an election outcome \mathbf{o} and a tallying proof pf , where \mathbf{o} is a multiset of votes.

Election schemes with a trusted bulletin board must satisfy correctness, that is, for all security parameters κ , votes v and multisets \mathbf{bb} , we have: $\Pr[(pk, sk, \mathbf{m}, \mathbf{bb}_0) \leftarrow \text{Setup}(\kappa); b \leftarrow \text{Vote}(pk, v); \mathbf{bb}' \leftarrow \text{BB}(\mathbf{bb}, b); (\mathbf{o}, pf) \leftarrow \text{Tally}(sk, \mathbf{bb}); (\mathbf{o}', pf') \leftarrow \text{Tally}(sk, \mathbf{bb}'): \mathbf{bb}' = \mathbf{bb} \cup \{b\} \wedge (\mathbf{o} \neq \emptyset \Rightarrow \mathbf{o}' = \mathbf{o} \cup \{v\} \wedge |\mathbf{v}| = |\mathbf{bb}|) \wedge (\mathbf{o} = \emptyset \Rightarrow \mathbf{o}' = \emptyset)] > 1 - \text{negl}(\kappa)$.

Definition 41. *Let $\Gamma = (\text{Setup}, \text{Vote}, \text{BB}, \text{Tally})$ be an election scheme with a trusted bulletin board, \mathcal{A} be an adversary, κ be a security parameter, and IND-SEC be the following game.*

IND-SEC($\Gamma, \mathcal{A}, \kappa$) =

```

(pk, sk, m, bb0) ← Setup( $\kappa$ );
bb1 ← bb0;  $\beta \leftarrow_R \{0, 1\}$ ;
L0 ←  $\emptyset$ ; L1 ←  $\emptyset$ ;
x ←  $\mathcal{A}^{\mathcal{O}}$ (pk, m);
if L0 = L1 then
  | (o, pf) ← Tally(sk, bb $\beta$ );
else
  | (o, pf) ← Tally(sk, bb0);
  | pf ←  $\perp$ ;
g ←  $\mathcal{A}$ (o, pf);
return g =  $\beta$ ;

```

In the above game, L_0 and L_1 are multisets and oracle \mathcal{O} is defined as follows:

- $\mathcal{O}(v_0, v_1)$ computes $L_0 \leftarrow L_0 \cup \{v_0\}$; $L_1 \leftarrow L_1 \cup \{v_1\}$; $b_0 \leftarrow \text{Vote}(pk, v_0)$; $bb_0 \leftarrow \text{BB}(bb_0, b_0)$; $b_1 \leftarrow \text{Vote}(pk, v_1)$; $bb_1 \leftarrow \text{BB}(bb_1, b_1)$, where $v_0, v_1 \in \mathbf{m}$.
- $\mathcal{O}(b)$ computes $bb' \leftarrow bb_\beta$; $bb_\beta \leftarrow \text{BB}(bb_\beta, b)$; **if** $bb_\beta \neq bb'$ **then** $bb_{1-\beta} \leftarrow \text{BB}(bb_{1-\beta}, b)$.
- $\mathcal{O}()$ outputs bb_β .

We say Γ satisfies IND-SEC, if for all probabilistic polynomial-time adversaries \mathcal{A} , there exists a negligible function negl , such that for all security parameters κ , we have $\text{Succ}(\text{IND-SEC}(\Gamma, \mathcal{A}, \kappa)) \leq \frac{1}{2} + \text{negl}(\kappa)$.

Observe that game IND-SEC uses bulletin boards bb_β and bb_0 , hence, bb_1 is unused when $\beta = 0$. It follows that game IND-SEC is equivalent to a variant that redefines the following oracle calls:

- $\mathcal{O}(v_0, v_1)$ computes $L_0 \leftarrow L_0 \cup \{v_0\}$; $L_1 \leftarrow L_1 \cup \{v_1\}$; $b_\beta \leftarrow \text{Vote}(pk, v_\beta)$; $bb_\beta \leftarrow \text{BB}(bb_\beta, b_\beta)$; **if** $\beta = 1$ **then** $b_0 \leftarrow \text{Vote}(pk, v_0)$; $bb_0 \leftarrow \text{BB}(bb_0, b_0)$, where $v_0, v_1 \in \mathbf{m}$.
- $\mathcal{O}(b)$ computes $bb' \leftarrow bb_\beta$; $bb_\beta \leftarrow \text{BB}(bb_\beta, b)$; **if** $bb_\beta \neq bb' \wedge \beta = 1$ **then** $bb_0 \leftarrow \text{BB}(bb_0, b)$.

Let that variant be IND-SEC*.

Lemma 34. *An election scheme with a trusted bulletin board Γ satisfies IND-SEC iff Γ satisfies IND-SEC*.*

Lemma 34 follows from our informal reasoning and we omit a formal proof.

Definition 42. *Given an election scheme with a trusted bulletin board $\Gamma = (\text{Setup}, \text{Vote}, \text{BB}, \text{Tally})$ such that for all security parameters κ and computations $(pk, sk, m, bb) \leftarrow \text{Setup}(\kappa)$ we have $\mathbf{m} = \{1, \dots, mc\}$ for some integer mc , we define $\gamma(\Gamma) = (\text{Setup}', \text{Vote}', \text{Tally}')$ such that*

- $\text{Setup}'(\kappa)$ computes $(pk, sk, \mathbf{m}, \mathbf{bb}) \leftarrow \text{Setup}(\kappa)$; $mc \leftarrow |\mathbf{m}|$; $pk' \leftarrow (pk, mc)$ and outputs $(pk', sk, mc, p(\kappa))$, where p is a polynomial function.
- $\text{Vote}'(pk', v, nc, \kappa)$ parses pk' as (pk, mc) , aborting if parsing fails; computes $b \leftarrow \text{Vote}(pk, v)$; and outputs b .
- $\text{Tally}'(sk, \mathbf{bb}, nc, \kappa)$ computes $(\mathbf{o}, pf) \leftarrow \text{Tally}(sk, \mathbf{bb})$, outputting an empty vector if \mathbf{o} is empty; assigns the largest integer in \mathbf{o} to nc ; initialises \mathbf{v} as a zero-filled vector of length nc ; computes **while** $v \in \mathbf{o}$ **do** $\mathbf{v}[v] \leftarrow \mathbf{v}[v] + 1$; $\mathbf{o} \leftarrow \mathbf{o} \setminus \{v\}$; and outputs (\mathbf{v}, pf) .

Theorem 35. *Let Γ be an election scheme with a trusted bulletin board. If $\gamma(\Gamma)$ is an election scheme satisfying Ballot-Security, then Γ satisfies IND-SEC.*

Proof sketch. Let $\Gamma = (\text{Setup}, \text{Vote}, \text{BB}, \text{Tally})$ and $\gamma(\Gamma) = (\text{Setup}', \text{Vote}', \text{Tally}')$. Suppose $\gamma(\Gamma)$ is an election scheme satisfying Ballot-Security. Moreover, suppose Γ does not satisfy IND-SEC. Hence, there exists a probabilistic polynomial-time adversary \mathcal{A} , such that for negligible functions negl , there exists a security parameter κ and $\text{Succ}(\text{IND-SEC}(\Gamma, \mathcal{A}, \kappa)) > \frac{1}{2} + \text{negl}(\kappa)$. We construct the following adversary \mathcal{B} against Ballot-Security from \mathcal{A} , where $\mathcal{O}_{\mathcal{A}}$ denotes \mathcal{A} 's oracle and $\mathcal{O}_{\mathcal{B}}$ denotes \mathcal{B} 's oracle:

- $\mathcal{B}(pk', \kappa)$ parses pk' as (pk, mc) , aborting if parsing fails, and outputs mc .
- $\mathcal{B}()$ initialises \mathbf{bb} and \mathbf{bb}_0 as empty sets and L_0 and L_1 as empty multisets; computes $\mathbf{m} \leftarrow \{1, \dots, mc\}$; $x \leftarrow \mathcal{A}(pk, \mathbf{m})$, handling \mathcal{A} 's oracle calls as follows, namely,
 - $\mathcal{O}_{\mathcal{A}}(v_0, v_1)$ computes $L_0 \leftarrow L_0 \cup \{v_0\}$; $L_1 \leftarrow L_1 \cup \{v_1\}$; $b \leftarrow \mathcal{O}_{\mathcal{B}}(v_0, v_1)$; $\mathbf{bb} \leftarrow \text{BB}(\mathbf{bb}, b)$; $b_0 \leftarrow \text{Vote}(pk, v_0)$; $\mathbf{bb}_0 \leftarrow \text{BB}(\mathbf{bb}, b_0)$,
 - $\mathcal{O}_{\mathcal{A}}(b)$ computes $\mathbf{bb}' \leftarrow \mathbf{bb}$; $\mathbf{bb} \leftarrow \text{BB}(\mathbf{bb}, b)$; **if** $\mathbf{bb} \neq \mathbf{bb}'$ **then** $bb_0 \leftarrow \text{BB}(\mathbf{bb}_0, b)$, and
 - $\mathcal{O}_{\mathcal{A}}()$ outputs \mathbf{bb} ,

and outputs \mathbf{bb} if $L_0 = L_1$ and \mathbf{bb}_0 otherwise.

- $\mathcal{B}(\mathbf{v}, pf)$ computes $\mathbf{o} \leftarrow \{v^{[v]} \mid 1 \leq v \leq mc\}$; **if** $L_0 \neq L_1$ **then** $pf \leftarrow \perp$; $g \leftarrow \mathcal{A}(\mathbf{o}, pf)$ and outputs g .

We prove that \mathcal{B} wins Ballot-Security.

Suppose (pk', sk, mb, mc) is an output of $\text{Setup}'(\kappa)$. By definition of algorithm Setup' , we have pk' is a pair (pk, mc) , where $(pk, sk, \mathbf{m}, \mathbf{bb})$ is an output of $\text{Setup}(\kappa)$ and $mc = |\mathbf{m}|$. Hence, $\mathcal{B}(pk', \kappa)$ outputs mc . Let β be a bit and suppose \mathbf{bb} is an output of $\mathcal{B}()$. It is trivial to see that $\mathcal{B}()$ simulates \mathcal{A} 's challenger to \mathcal{A} . Moreover, by Lemma 34 it is straightforward to see that \mathcal{B} simulates \mathcal{A} 's oracle too. Indeed, adversary \mathcal{B} maintains bulletin board \mathbf{bb} such that $\mathcal{O}_{\mathcal{A}}(v_0, v_1)$ constructs a ballot b for v_β using \mathcal{B} 's oracle, hence, the ballot is constructed by algorithm Vote by way of algorithm Vote' , and adds that

ballot to the bulletin board using algorithm **BB**. Suppose (\mathbf{v}, pf) is an output of $\text{Tally}(sk, \mathbf{bb}, nc, \kappa)$ and g is an output of $\mathcal{B}(\mathbf{v}, pf)$. It is straightforward to see that $\mathcal{B}(\mathbf{v}, pf)$ simulates \mathcal{A} 's challenger to \mathcal{A} , thus, $g = \beta$, with at least the probability that \mathcal{A} wins IND-SEC, concluding our proof. \square

References

- [ABR12] 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.
- [Adi08] Ben Adida. Helios: Web-based Open-Audit Voting. In *USENIX Security'08: 17th USENIX Security Symposium*, pages 335–348. USENIX Association, 2008.
- [AH10] R. Michael Alvarez and Thad E. Hall. *Electronic Elections: The Perils and Promises of Digital Democracy*. Princeton University Press, 2010.
- [AMPQ09] 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.
- [AN06] Ben Adida and C. Andrew Neff. Ballot casting assurance. In *EVT'06: Electronic Voting Technology Workshop*. USENIX Association, 2006.
- [BBH⁺17] Matthew Bernhard, Josh Benaloh, J. Alex Halderman, Ronald L. Rivest, Peter Y. A. Ryan, Philip B. Stark, Vanessa Teague, Poorvi L. Vora, and Dan S. Wallach. Public evidence from secret ballots. In *E-Vote-ID'17: 10th International Conference for Electronic Voting*, *LNCS*, pages 84–109. Springer, 2017.
- [BBP07] Romain Bertrand, Jean-Louis Briquet, and Peter Pels. Introduction: Towards a Historical Ethnography of Voting. In *The Hidden History of the Secret Ballot*. Indiana University Press, 2007.
- [BCG⁺15a] 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.
- [BCG⁺15b] David Bernhard, Véronique Cortier, David Galindo, Olivier Pereira, and Bogdan Warinschi. SoK: A comprehensive analysis

- of game-based ballot privacy definitions. In *SEP'15: 36th Security and Privacy Symposium*, pages 499–516. IEEE Computer Society, 2015.
- [BCP⁺11] David Bernhard, Véronique Cortier, Olivier Pereira, Ben Smyth, and Bogdan Warinschi. Adapting Helios for provable ballot privacy. In *ESORICS'11: 16th European Symposium on Research in Computer Security*, volume 6879 of *LNCS*, pages 335–354. Springer, 2011.
- [BDJR97] Mihir Bellare, Anand Desai, E. Jorjipii, and Phillip Rogaway. A Concrete Security Treatment of Symmetric Encryption. In *FOCS'97: 38th Annual Symposium on Foundations of Computer Science*, pages 394–403. IEEE Computer Society, 1997.
- [BDPR98] 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.
- [Ben96] Josh Benaloh. *Verifiable Secret-Ballot Elections*. PhD thesis, Department of Computer Science, Yale University, 1996.
- [Ber14] David Bernhard. *Zero-Knowledge Proofs in Theory and Practice*. PhD thesis, Department of Computer Science, University of Bristol, 2014.
- [BGP11] Philippe Bulens, Damien Giry, and Olivier Pereira. Running Mixnet-Based Elections with Helios. In *EVT/WOTE'11: Electronic Voting Technology Workshop/Workshop on Trustworthy Elections*. USENIX Association, 2011.
- [Bjo04] Eric C. Bjornlund. *Beyond Free and Fair: Monitoring Elections and Building Democracy*. Woodrow Wilson Center Press / Johns Hopkins University Press, 2004.
- [Bow07] 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, August 2007.
- [BPW12a] 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.
- [BPW12b] 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.

- [BR93] 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. ACM, 1993.
- [BR05] Mihir Bellare and Phillip Rogaway. Symmetric Encryption. In *Introduction to Modern Cryptography*, chapter 4. 2005. <http://cseweb.ucsd.edu/~mihir/cse207/w-se.pdf>.
- [Bre06] Peter Brent. The Australian ballot: Not the secret ballot. *Australian Journal of Political Science*, 41(1):39–50, 2006.
- [BS99] 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.
- [BS15] David Bernhard and Ben Smyth. Ballot secrecy with malicious bulletin boards. Cryptology ePrint Archive, Report 2014/822 (version 20150413:170300), 2015.
- [BS16] Bruno Blanchet and Ben Smyth. Automated reasoning for equivalences in the applied pi calculus with barriers. In *CSF'16: 29th Computer Security Foundations Symposium*, pages 310–324. IEEE Computer Society, 2016.
- [BS18] Bruno Blanchet and Ben Smyth. Automated reasoning for equivalences in the applied pi calculus with barriers. *Journal of Computer Security*, 26(3):367–422, 2018.
- [BSCS16] Bruno Blanchet, Ben Smyth, Vincent Cheval, and Marc Sylvestre. *ProVerif 1.96: Automatic Cryptographic Protocol Verifier, User Manual and Tutorial*, 2016.
- [BT94] Josh Cohen Benaloh and Dwight Tuinstra. Receipt-free secret-ballot elections. In *STOC'94: 26th Theory of computing Symposium*, pages 544–553. ACM Press, 1994.
- [BVQ10] 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, Sept 2010.
- [BW14] 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.

- [BY86] 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.
- [CCC⁺08] 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*. USENIX Association, 2008.
- [CCFG16] Pyrros Chaidos, Véronique Cortier, Georg Fuschbauer, and David Galindo. BeleniosRF: A Non-interactive Receipt-Free Electronic Voting Scheme. In *CCS'16: 23rd ACM Conference on Computer and Communications Security*, pages 1614–1625. ACM Press, 2016.
- [CE16] 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.
- [CEvP87] 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.
- [CFSY96] 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.
- [CGGI13a] Veronique Cortier, David Galindo, Stephane Glondu, and Malika Izabachene. A generic construction for voting correctness at minimum cost - Application to Helios. *Cryptology ePrint Archive*, Report 2013/177 (version 20130521:145727), 2013.
- [CGGI13b] Véronique Cortier, David Galindo, Stéphane Glondu, and Malika Izabachene. Distributed elgamal à la pedersen: Application to helios. In *WPES'13: Workshop on Privacy in the Electronic Society*, pages 131–142. ACM Press, 2013.
- [CGGI14] Véronique Cortier, David Galindo, Stephane 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.

- [CGMA85] Benny Chor, Shafi Goldwasser, Silvio Micali, and Baruch Awerbuch. Verifiable Secret Sharing and Achieving Simultaneity in the Presence of Faults. In *FOCS'85: 26th Foundations of Computer Science Symposium*, pages 383–395. IEEE Computer Society, 1985.
- [CGS97] 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.
- [CH17] Cas Cremers and Lucca Hirschi. Improving Automated Symbolic Analysis for E-voting Protocols: A Method Based on Sufficient Conditions for Ballot Secrecy. arXiv, Report 1709.00194, 2017.
- [Cha81] David L. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24:84–90, 1981.
- [CHI⁺] Michael Clarkson, Brian Hay, Meador Inge, abhi shelat, David Wagner, and Alec Yasinsac. Software review and security analysis of Scytl remote voting software. <http://election.dos.state.fl.us/voting-systems/pdf/FinalReportSept19.pdf>.
- [CP93] 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.
- [CR87] Benny Chor and Michael O. Rabin. Achieving Independence in Logarithmic Number of Rounds. In *PODC'87: 6th Principles of Distributed Computing Symposium*, pages 260–268. ACM Press, 1987.
- [CRS05] 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.
- [CS11] 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.
- [CS13] Véronique Cortier and Ben Smyth. Attacking and fixing Helios: An analysis of ballot secrecy. *Journal of Computer Security*, 21(1):89–148, 2013.
- [CSD⁺17] Véronique Cortier, Benedikt Schmidt, Constantin Cătălin Drăgan, Pierre-Yves Strub, Francois Dupressoir, and Bogdan Warinschi.

- Machine-Checked Proofs of Privacy for Electronic Voting Protocols. In *S&P'17: 37th IEEE Symposium on Security and Privacy*. IEEE Computer Society, 2017.
- [DC12] Yvo Desmedt and Pyrros Chaidos. Applying Divertibility to Blind Ballot Copying in the Helios Internet Voting System. In *ESORICS'12: 17th European Symposium on Research in Computer Security*, volume 7459 of *LNCS*, pages 433–450. Springer, 2012.
- [DDN91] Danny Dolev, Cynthia Dwork, and Moni Naor. Non-Malleable Cryptography. In *STOC'91: 23rd Theory of computing Symposium*, pages 542–552. ACM Press, 1991.
- [DDN00] Danny Dolev, Cynthia Dwork, and Moni Naor. Nonmalleable Cryptography. *Journal on Computing*, 30(2):391–437, 2000.
- [DK05] Yvo Desmedt and Kaoru Kurosawa. Electronic Voting: Starting Over? In *ISC'05: International Conference on Information Security*, volume 3650 of *LNCS*, pages 329–343. Springer, 2005.
- [DKR06] Stéphanie Delaune, Steve Kremer, and Mark Ryan. Coercion-Resistance and Receipt-Freeness in Electronic Voting. In *CSFW'06: 19th Computer Security Foundations Workshop*, pages 28–42. IEEE Computer Society, 2006.
- [DKR09] 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.
- [DKRS11] 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.
- [DRS08] Stéphanie Delaune, Mark D. Ryan, and Ben Smyth. Automatic verification of privacy properties in the applied pi-calculus. In *IFIPTM'08: 2nd Joint iTrust and PST Conferences on Privacy, Trust Management and Security*, volume 263 of *International Federation for Information Processing (IFIP)*, pages 263–278. Springer, 2008.
- [FS87] 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.
- [Gen95] Rosario Gennaro. Achieving independence efficiently and securely. In *PODC'95: 14th Principles of Distributed Computing Symposium*, pages 130–136. ACM Press, 1995.

- [Gen00] Rosario Gennaro. A Protocol to Achieve Independence in Constant Rounds. *IEEE Transactions on Parallel and Distributed Systems*, 11(7):636–647, 2000.
- [Ger09] Bundesverfassungsgericht (Germany’s Federal Constitutional Court). *Use of voting computers in 2005 Bundestag election unconstitutional*, March 2009. Press release 19/2009.
- [GGR09] Ryan W. Gardner, Sujata Garera, and Aviel D. Rubin. Coercion Resistant End-to-end Voting. In *FC’09: 13th International Conference on Financial Cryptography and Data Security*, volume 5628 of *LNCS*, pages 344–361. Springer, 2009.
- [GH07] Rop Gonggrijp and Willem-Jan Hengeveld. Studying the Nedap/Groenendaal ES3B Voting Computer: A Computer Security Perspective. In *EVT’07: Electronic Voting Technology Workshop*. USENIX Association, 2007.
- [Gro04] Jens Groth. Efficient maximal privacy in boardroom voting and anonymous broadcast. In *FC’04: 8th International Conference on Financial Cryptography*, volume 3110 of *LNCS*, pages 90–104. Springer, 2004.
- [Gro06] 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.
- [Gum05] Andrew Gumbel. *Steal This Vote: Dirty Elections and the Rotten History of Democracy in America*. Nation Books, 2005.
- [HBH10] 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>, May 2010.
- [HK02] Alejandro Hevia and Marcos A. Kiwi. Electronic Jury Voting Protocols. In *LATIN’02: Theoretical Informatics*, volume 2286 of *LNCS*, pages 415–429. Springer, 2002.
- [HK04] Alejandro Hevia and Marcos A. Kiwi. Electronic jury voting protocols. *Theoretical Computer Science*, 321(1):73–94, 2004.
- [HRZ10] Fao Hao, Peter Y. A. Ryan, and Piotr Zieliński. Anonymous voting by two-round public discussion. *Journal of Information Security*, 4(2):62 – 67, 2010.

- [HS12] James Heather and Steve Schneider. A formal framework for modelling coercion resistance and receipt freeness. In *FM'12: 18th International Symposium on Formal Methods*, number 7436 in LNCS, pages 217–231. Springer, 2012.
- [JCJ05] 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.
- [JCJ10] 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.
- [JS12] 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.
- [Kel12] Judith G. Kelley. *Monitoring Democracy: When International Election Observation Works, and Why It Often Fails*. Princeton University Press, 2012.
- [KL07] Jonathan Katz and Yehuda Lindell. *Introduction to Modern Cryptography*. Chapman & Hall/CRC, 2007.
- [KRA18] Shahram Khazaei and Mehri Rezaei-Aliabadi. A rigorous security analysis of a decentralized electronic voting protocol in the universal composability framework. *Journal of Information Security and Applications*, 43:99–109, 2018.
- [KSR10] Petr Klus, Ben Smyth, and Mark D. Ryan. ProSwapper: Improved equivalence verifier for ProVerif. <http://www.bensmyth.com/proswapper.php>, 2010.
- [KSRH12] Dalia Khader, Ben Smyth, Peter Y. A. Ryan, and Feng Hao. A Fair and Robust Voting System by Broadcast. In *EVOTE'12: 5th International Conference on Electronic Voting*, volume 205 of *Lecture Notes in Informatics*, pages 285–299. Gesellschaft für Informatik, 2012.
- [KSRW04] Tadayoshi Kohno, Adam Stubblefield, Aviel D. Rubin, and Dan S. Wallach. Analysis of an Electronic Voting System. In *SP'04: 25th Security and Privacy Symposium*, pages 27–40. IEEE Computer Society, 2004.
- [KTV12a] 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.

- [KTV12b] Ralf Küsters, Tomasz Truderung, and Andreas Vogt. Clash Attacks on the Verifiability of E-Voting Systems. In *SEP'12: 33rd IEEE Symposium on Security and Privacy*, pages 395–409. IEEE Computer Society, 2012.
- [KY02] Aggelos Kiayias and Moti Yung. Self-tallying elections and perfect ballot secrecy. In *PKC'01: 3rd International Workshop on Practice and Theory in Public Key Cryptography*, volume 2274 of *LNCS*, pages 141–158. Springer, 2002.
- [KZZ15] 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.
- [Lep08] Jill Lepore. Rock, Paper, Scissors: How we used to vote. *Annals of Democracy, The New Yorker*, October 2008.
- [LG84] Arend Lijphart and Bernard Grofman. *Choosing an electoral system: Issues and Alternatives*. Praeger, 1984.
- [MH96] Markus Michels and Patrick Horster. Some Remarks on a Receipt-Free and Universally Verifiable Mix-Type Voting Scheme. In *ASIACRYPT'96: International Conference on the Theory and Application of Cryptology and Information Security*, volume 1163 of *LNCS*, pages 125–132. Springer, 1996.
- [Mil30] James Mill. The Ballot. In *The Westminster Review*, volume 13. Robert Heward, 1830.
- [MN06] 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.
- [MS19] Maxime Meyer and Ben Smyth. Exploiting re-voting in the helios election system. *Information Processing Letters*, (143):14–19, 2019.
- [MSQ14] 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.
- [NA03] C. Andrew Neff and Jim Adler. Verifiable e-Voting: Indisputable electronic elections at polling places. Technical report, VoteHere, 2003.

- [Nef04] C. Andrew Neff. Practical high certainty intent verification for encrypted votes. Technical report, VoteHere, 2004.
- [NIS12] NIST. Secure Hash Standard (SHS). FIPS PUB 180-4, Information Technology Laboratory, National Institute of Standards and Technology, March 2012.
- [Nor15] Pippa Norris. *Why Elections Fail*. Cambridge University Press, 2015.
- [OAS69] Organization of American States. *American Convention on Human Rights, "Pact of San Jose, Costa Rica"*, 1969.
- [OSC90] Organization for Security and Co-operation in Europe. *Document of the Copenhagen Meeting of the Conference on the Human Dimension of the CSCE*, 1990.
- [PB12] 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.
- [Pfi94] Birgit Pfitzmann. Breaking Efficient Anonymous Channel. In *EUROCRYPT'94: 11th International Conference on the Theory and Applications of Cryptographic Techniques*, volume 950 of *LNCS*, pages 332–340. Springer, 1994.
- [PP89] Birgit Pfitzmann and Andreas Pfitzmann. How to Break the Direct RSA-Implementation of Mixes. In *EUROCRYPT'89: 6th International Conference on the Theory and Applications of Cryptographic Techniques*, volume 434 of *LNCS*, pages 373–381. Springer, 1989.
- [QS18a] Elizabeth A. Quaglia and Ben Smyth. Authentication with weaker trust assumptions for voting systems. In *AFRICACRYPT'18: 10th International Conference on Cryptology in Africa*, LNCS. Springer, 2018.
- [QS18b] Elizabeth A Quaglia and Ben Smyth. Secret, verifiable auctions from elections. *Theoretical Computer Science*, 730:44–92, 2018.
- [QS18c] Elizabeth A. Quaglia and Ben Smyth. A short introduction to secrecy and verifiability for elections. arXiv, Report 1702.03168, 2018.
- [Saa95] Thomas Saalfeld. On Dogs and Whips: Recorded Votes. In Herbert Döring, editor, *Parliaments and Majority Rule in Western Europe*, chapter 16. St. Martin's Press, 1995.

- [SB13] 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.
- [SB14] Ben Smyth and David Bernhard. Ballot secrecy and ballot independence: definitions and relations. Cryptology ePrint Archive, Report 2013/235 (version 20141010:082554), 2014.
- [SC11] Ben Smyth and Véronique Cortier. A note on replay attacks that violate privacy in electronic voting schemes. Technical Report RR-7643, INRIA, June 2011.
- [Sch99] Berry Schoenmakers. A simple publicly verifiable secret sharing scheme and its application to electronic voting. In *CRYPTO'99: 19th International Cryptology Conference*, volume 1666 of *LNCS*, pages 148–164. Springer, 1999.
- [Sch05] Nicole Schweikardt. Arithmetic, first-order logic, and counting quantifiers. *ACM Transactions on Computational Logic*, 6(3):634–671, July 2005.
- [SFC17] Ben Smyth, Steven Frink, and Michael R. Clarkson. Election Verifiability: Cryptographic Definitions and an Analysis of Helios and JCJ. Cryptology ePrint Archive, Report 2015/233 (version 20170213:132559), 2017.
- [SFD⁺14] Drew Springall, Travis Finkenauer, Zakir Durumeric, Jason Kitcat, Harri Hursti, Margaret MacAlpine, and J. Alex Halderman. Security Analysis of the Estonian Internet Voting System. In *CCS'14: 21st ACM Conference on Computer and Communications Security*, pages 703–715. ACM Press, 2014.
- [SK95] Kazue Sako and Joe Kilian. Receipt-Free Mix-Type Voting Scheme: A practical solution to the implementation of a voting booth. In *EUROCRYPT'95: 12th International Conference on the Theory and Applications of Cryptographic Techniques*, volume 921 of *LNCS*, pages 393–403. Springer, 1995.
- [Smy11] Ben Smyth. *Formal verification of cryptographic protocols with automated reasoning*. PhD thesis, School of Computer Science, University of Birmingham, 2011.
- [Smy12] Ben Smyth. Replay attacks that violate ballot secrecy in Helios. Cryptology ePrint Archive, Report 2012/185, 2012.
- [Smy14] Ben Smyth. Ballot secrecy with malicious bulletin boards. Cryptology ePrint Archive, Report 2014/822 (version 20141012:004943), 2014.

- [Smy15] Ben Smyth. Secrecy and independence for election schemes. Cryptology ePrint Archive, Report 2015/942 (version 20150928:195428), 2015.
- [Smy16] Ben Smyth. Secrecy and independence for election schemes. Cryptology ePrint Archive, Report 2015/942 (version 20160713:142934), 2016.
- [Smy17] Ben Smyth. First-past-the-post suffices for ranked voting. <https://bensmyth.com/publications/2017-FPTP-suffices-for-ranked-voting/>, 2017.
- [Smy18a] Ben Smyth. A foundation for secret, verifiable elections. Cryptology ePrint Archive, Report 2018/225 (version 20180301:164045), 2018.
- [Smy18b] Ben Smyth. Verifiability of Helios Mixnet. In *Voting'18: 3rd Workshop on Advances in Secure Electronic Voting*, LNCS. Springer, 2018.
- [Smy18c] Ben Smyth. Verifiability of Helios Mixnet. Cryptology ePrint Archive, Report 2018/017, 2018.
- [SP13] 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.).
- [SP15] Ben Smyth and Alfredo Pironti. Truncating TLS Connections to Violate Beliefs in Web Applications. Technical Report hal-01102013, INRIA, 2015.
- [Sta14] CACM Staff. ACM's 2014 General Election: Please Take This Opportunity to Vote. *Communications of the ACM*, 57(5):9–17, May 2014.
- [TPLT13] Georgios Tsoukalas, Kostas Papadimitriou, Panos Louridas, and Panayiotis Tsanakas. From Helios to Zeus. *Journal of Election Technology and Systems*, 1(1), 2013.
- [TY98] Yiannis Tsiounis and Moti Yung. On the Security of ElGamal Based Encryption. In *PKC'98: First International Workshop on Practice and Theory in Public Key Cryptography*, volume 1431 of LNCS, pages 117–134. Springer, 1998.
- [UK07] UK Electoral Commission. *Key issues and conclusions: May 2007 electoral pilot schemes*, May 2007.
- [UM10] Dominique Unruh and Jörn Müller-Quade. Universally Composable Incoercibility. In *CRYPTO'10: 30th International Cryptology Conference*, volume 6223 of LNCS, pages 411–428. Springer, 2010.

- [UN48] United Nations. *Universal Declaration of Human Rights*, 1948.
- [Wik06] Douglas Wikström. Simplified Submission of Inputs to Protocols. Cryptology ePrint Archive, Report 2006/259, 2006.
- [Wik08] Douglas Wikström. Simplified Submission of Inputs to Protocols. In *SCN'08: 6th International Conference on Security and Cryptography for Networks*, volume 5229 of *LNCS*, pages 293–308. Springer, 2008.
- [Wik16] Douglas Wikström. *Verificatum: How to Implement a Stand-alone Verifier for the Verificatum Mix-Net (VMN Version 3.0.2)*, 2016. <http://www.verificatum.com/files/vmnum-3.0.2.pdf>.
- [WWH⁺10] Scott Wolchok, Eric Wustrow, J. Alex Halderman, Hari K. Prasad, Arun Kankipati, Sai Krishna Sakhamuri, Vasavya Yagati, and Rop Gonggrijp. Security Analysis of India's Electronic Voting Machines. In *CCS'10: 17th ACM Conference on Computer and Communications Security*, pages 1–14. ACM Press, 2010.
- [WWIH12] Scott Wolchok, Eric Wustrow, Dawn Isabel, and J. Alex Halderman. Attacking the Washington, D.C. Internet Voting System. In *FC'12: 16th International Conference on Financial Cryptography and Data Security*, volume 7397 of *LNCS*, pages 114–128. Springer, 2012.