

Secrecy and independence for election schemes

Ben Smyth

Mathematical and Algorithmic Sciences Lab,
Huawei Technologies Co. Ltd., France

December 28, 2016

Abstract

We study ballot secrecy and ballot independence for election schemes. First, we propose a definition of ballot secrecy as an indistinguishability game in the computational model of cryptography. Our definition builds upon and strengthens earlier definitions to ensure that ballot secrecy is preserved in the presence of an adversary that controls the bulletin board and communication channel. Secondly, we propose a definition of ballot independence as an adaptation of a non-malleability definition for asymmetric encryption. We also provide a simpler, equivalent definition as an indistinguishability game. Thirdly, we prove relations between our definitions. In particular, we prove that ballot independence is necessary in election schemes satisfying ballot secrecy. And that ballot independence is sufficient for ballot secrecy in election schemes with zero-knowledge tallying proofs. Fourthly, we demonstrate the applicability of our results by analysing Helios. Our analysis identifies a new attack against Helios, which enables an adversary to determine if a voter did not vote for a candidate chosen by the adversary. The attack requires the adversary to control the bulletin board or communication channel, thus, it could not have been detected by earlier definitions of ballot secrecy. Finally, we prove that ballot secrecy is satisfied by a variant of Helios that uses non-malleable ballots.

Keywords. Elections, Helios, independence, non-malleability, privacy, provable security, secrecy, voting.

1 Introduction

An election is a decision-making procedure to choose representatives. Choices should be made freely, and this has started a movement towards voting as a secret act. The movement is championed by the United Nations [UN48, Article 21], the Organization for Security and Cooperation in Europe [OSC90, Paragraph 7.4], and the Organization of American States [OAS69, Article 23]. And

has led to the emergence of ballot secrecy¹ as a *de facto* standard requirement of voting systems.

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

Many voting systems – including systems that have been deployed in real-world, large-scale public elections – attempt to satisfy ballot secrecy by placing extensive trust in software and hardware. Unfortunately, many systems are not trustworthy, and are vulnerable to attacks that could compromise ballot secrecy [GH07, Bow07, WWH⁺10, WWIH12, SFD⁺14]. Such vulnerabilities can be avoided by formulating ballot secrecy as a rigorous and precise security definition, and proving that systems satisfy this definition. We propose such a definition in the computational model of cryptography. Our definition builds upon and strengthens earlier definitions of ballot secrecy by Bernhard *et al.* [BCP⁺11, BPW12b, SB13a, SB14, BCG⁺15b] to ensure that ballot secrecy is preserved in the presence of an adversary that controls the bulletin board and the communication channel, whereas definitions by Bernhard *et al.* only consider trusted bulletin boards and channels.

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, i.e., ballots are non-malleable.

Cortier & Smyth [CS13, CS11, SC11] attribute a class of ballot secrecy attacks to the absence of ballot independence. Their attribution caused some debate. In particular, Bulens, Giry & Pereira [BGP11, §3.2] highlight the investigation of systems which allow the submission of related votes, whilst preserving ballot secrecy, as an interesting research problem. And Desmedt & Chaidos [DC12] claim to provide a solution.² We facilitate the study of ballot independence by proposing two definitions of independence in the computational model. Our first definition is a straightforward adaptation of a non-malleability definition for asymmetric encryption. And our second definition is a straightforward adaptation of an indistinguishability game for asymmetric encryption. The former definition naturally captures ballot independence, but it is complex and proofs of non-malleability are relatively difficult. The latter definition is equivalent, yet simpler, and proofs of indistinguishability are easier.

We demonstrate relations between our definitions of secrecy and independence. In particular, we prove that ballot secrecy implies ballot independence, hence, ballot independence is necessary (assuming ballot secrecy is required).

¹*Ballot secrecy* and *privacy* occasionally appear as synonyms in the literature. We favour ballot secrecy because it avoids confusion with other privacy notions, such as receipt-freeness and coercion resistance, for example.

²Smyth & Bernhard [SB13a, §5.1] critique the results by Desmedt & Chaidos [DC12] and argue that their security results do not support their claims.

We also prove the inverse implication for a class of voting systems with zero-knowledge tallying proofs. And show that the inverse implication does not hold in general, hence, ballot secrecy is strictly stronger than ballot independence.

We employ our ballot secrecy definition to analyse Helios [AMPQ09, Per16], a web-based voting system that has been deployed in the real-world. This scheme is vulnerable to attacks against ballot secrecy [CS13, CS11, SC11]. The next Helios release, henceforth *Helios'12*, is intended to mitigate against those attacks. 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, assuming the bulletin board and communication channel are secure, despite the use of malleable ballots. Nevertheless, it follows from our results that ballot secrecy is not satisfied when this assumption is dropped. And this leads to the discovery of a new attack against Helios, whereby an adversary can determine if a voter did not vote for a candidate chosen by the adversary. Violations of ballot secrecy can be overcome using a variant of Helios that uses non-malleable ballots, and we formally prove that our definition of ballot secrecy is satisfied by that variant.

Contribution. This paper contributes to the security of voting systems by: proposing definitions of ballot secrecy (§3) and ballot independence (§4) in the computational model; proving that ballot secrecy is strictly stronger than ballot independence in general, and that secrecy and independence coincide for elections schemes with zero-knowledge tallying proofs (§5); and identifying a new attack against Helios, proposing a fix, and proving that ballot secrecy is satisfied when the fix is incorporated (§6).

2 Election schemes

We recall syntax for *election schemes*³ from Smyth, Frink & Clarkson [SFC16].⁴

Definition 1 (Election scheme [SFC16]). *An election scheme is a tuple of probabilistic polynomial-time algorithms (Setup, Vote, Tally) such that:*

Setup, denoted⁵ $(pk, sk, mb, mc) \leftarrow \text{Setup}(\kappa)$, is run by the tallier⁶. Setup 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 .

³Election schemes capture an interesting class of voting systems, which includes Helios.

⁴We omit algorithm `Verify` from our syntax and we omit the condition that election schemes must satisfy notions of completeness and injectivity, because we focus on ballot secrecy, rather than verifiability, in this paper. (Verifiability is studied elsewhere, e.g., [SFC16].)

⁵Let $A(x_1, \dots, x_n; r)$ denote the output of probabilistic algorithm A on inputs x_1, \dots, x_n and random coins r . Let $A(x_1, \dots, x_n)$ denote $A(x_1, \dots, x_n; r)$, where r is chosen uniformly at random. And let \leftarrow denote assignment.

⁶Some election schemes (e.g., Helios) permit the tallier's role to be distributed amongst several talliers. For simplicity, we consider only a single tallier in this paper. Generalising syntax and security definitions to multiple talliers is a possible direction for future work.

Vote, denoted $b \leftarrow \text{Vote}(pk, v, nc, \kappa)$, is run by voters. **Vote** takes as input a public key pk , a voter's vote v , some number of candidates nc , and a security parameter κ . A voter's vote should be selected from a sequence $1, \dots, nc$ of distinct candidates. **Vote** 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. **Tally** 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.⁷ It outputs an election outcome \mathbf{v} and a non-interactive tallying proof pf (i.e., a proof that the outcome is correct). An election outcome is a vector \mathbf{v} of length nc such that $\mathbf{v}[v]$ indicates⁸ the number of votes for candidate v .

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

$$\Pr[(pk, sk, mb, mc) \leftarrow \text{Setup}(\kappa);$$

$$\quad \text{for } 1 \leq i \leq nb \text{ do}$$

$$\quad \left[\begin{array}{l} b_i \leftarrow \text{Vote}(pk, v_i, nc, \kappa); \\ \mathbf{v}[v_i] \leftarrow \mathbf{v}[v_i] + 1; \end{array} \right.$$

$$\quad (\mathbf{v}', pf) \leftarrow \text{Tally}(sk, \{b_1, \dots, b_{nb}\}, nc, \kappa) :$$

$$\quad nb \leq mb \wedge nc \leq mc \Rightarrow \mathbf{v} = \mathbf{v}'] > 1 - \text{negl}(\kappa).$$

3 Ballot Secrecy

Our informal definition of ballot secrecy (§1) could be formulated as an indistinguishability game, similar to indistinguishability games for asymmetric encryption (e.g., IND-CPA): we could challenge the adversary to determine whether a ballot is for one of two possible votes. This formalisation is too weak, because election schemes also output the election outcome and a tallying proof, which needs to be incorporated into the game. Unfortunately, it is insufficient to simply grant the adversary access to an oracle that provides an election outcome and tallying proof corresponding to some ballots, because such a game is unsatisfiable, in particular, the adversary can use the oracle to reveal the vote encapsulated inside the challenge ballot. This reveals some limitations in our informal definition of ballot secrecy.

For simplicity, our informal definition of ballot secrecy deliberately omits some side-conditions, which are necessary for satisfiability, in particular, we did not stress that a voter's vote may be revealed in the following scenarios: unanimous election outcomes reveal how everyone voted and, more generally, election outcomes can be coupled with partial knowledge about the distribution of voters' votes to deduce voters' votes. For example, suppose Alice, Bob and Mallory

⁷Bulletin boards are modelled as sets to avoid the class of attacks against ballot secrecy that arise when duplicate ballots appear on bulletin boards [CS11, CS13].

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

vote in a referendum and the outcome is two “yes” votes and one “no” vote. 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. Accordingly, ballot secrecy must concede that election outcomes reveal partial information about voters’ votes,⁹ hence, we refine our informal definition of ballot secrecy 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 on the distribution of votes.

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.

3.1 Indistinguishability game

We formalise ballot secrecy as an indistinguishability game between an adversary and a challenger.¹⁰

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 $\text{Ballot-Secrecy}(\Gamma, \mathcal{A}, \kappa)$ be the following game.¹¹*

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

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

Predicate $\text{balanced}(\mathbf{bb}, nc, L)$ holds when: 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 oracle \mathcal{O} is defined as follows:¹²

⁹Alternative formalisations of election schemes might permit different results. For instance, voting systems which only announce the winning candidate [BY86, HK02, HK04, DK05], rather than the number of votes for each candidate (i.e., the election outcome, in our terminology), could offer stronger notions of ballot secrecy.

¹⁰Games are probabilistic algorithms that output booleans. An adversary *wins* a game by causing it to output true (\top). We denote an adversary’s *success* $\text{Succ}(\text{Exp}(\cdot))$ in a game $\text{Exp}(\cdot)$ as the probability that the adversary wins, that is, $\text{Succ}(\text{Exp}(\cdot)) = \Pr[b \leftarrow \text{Exp}(\cdot) : b = \top]$. Adversaries are assumed to be *stateful*, that is, information persists across invocations of the adversary in a single game, in particular, the adversary can access earlier assignments.

¹¹Let $x \leftarrow_R S$ denote assignment to x of an element chosen uniformly at random from set S . And let $|\mathbf{v}|$ denote the length of vector \mathbf{v} .

¹²Oracles may access game parameters, e.g., pk .

- $\mathcal{O}(v_0, v_1)$ 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 (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 $\text{Succ}(\text{Ballot-Secrecy}(\Gamma, \mathcal{A}, \kappa)) \leq 1/2 + \text{negl}(\kappa)$.

The game captures a setting in which the tallier generates a key pair using the scheme’s Setup algorithm, publishes the public key, and only uses the private key to compute the election outcome and tallying proof.

The adversary has access to a left-right oracle [BDJR97, BR05] which can compute ballots on the adversary’s behalf. Ballots can be computed by the left-right oracle in two ways, corresponding to a bit β chosen uniformly at random by the challenger. If $\beta = 0$, then, given a pair of votes v_0, v_1 , the oracle computes a ballot for v_0 and outputs the ballot to the adversary. Otherwise ($\beta = 1$), the oracle outputs a ballot for v_1 . The adversary constructs a bulletin board, which may include ballots computed by the oracle. Thus, the game captures a setting where the bulletin board is constructed by an adversary that casts ballots on behalf of a subset of voters and controls the distribution of votes cast by the remaining voters.

The challenger tallies the adversary’s bulletin board to derive an election outcome and tallying proof. The adversary is given the outcome and proof, and wins by determining whether $\beta = 0$ or $\beta = 1$. Intuitively, if the adversary wins, then there exists a strategy to distinguish ballots. On the other hand, if the adversary loses, then the adversary is unable to distinguish between a ballot for vote v_0 and a ballot for vote v_1 , therefore, voters’ votes cannot be revealed.

Our notion of ballot secrecy considers election schemes which reveal the number of votes for each candidate (i.e., the election outcome). Hence, to avoid trivial distinctions in our ballot secrecy game, we insist the game is *balanced*: “left” and “right” inputs to the left-right oracle are equivalent, when the corresponding left-right oracle’s 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.¹³ For instance, an adversary that constructs 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 bulletin board.

¹³A 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.”

3.2 Non-malleable encryption is sufficient for secrecy

To demonstrate the applicability of our definition, we recall a construction by Quaglia & Smyth [QS16] for election schemes from asymmetric encryption schemes.¹⁴

Definition 3 (Enc2Vote [QS16]). *Given an asymmetric encryption scheme $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$,¹⁵ we define $\text{Enc2Vote}(\Pi)$ as follows.*

- $\text{Setup}(\kappa)$ computes $(pk, sk, \mathbf{m}) \leftarrow \text{Gen}(\kappa)$ and outputs $(pk, sk, \text{poly}(\kappa), |\mathbf{m}|)$.
- $\text{Vote}(pk, v, nc, \kappa)$ computes $b \leftarrow \text{Enc}(pk, v)$ and outputs b , if $1 \leq v \leq nc$, and outputs \perp , otherwise.
- $\text{Tally}(sk, \mathbf{bb}, nc, \kappa)$ initialises vector \mathbf{v} of length nc , 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) .

Algorithm Setup requires poly to be a polynomial function and $\mathbf{m} = \{1, \dots, |\mathbf{m}|\}$. Algorithm Tally requires ϵ to be a constant symbol.

Lemma 1. *Given an asymmetric encryption scheme Π with perfect correctness, we have $\text{Enc2Vote}(\Pi)$ is an election scheme (i.e., $\text{Enc2Vote}(\Pi)$ satisfies correctness).*

A proof of Lemma 1 appears in [QS16, §C.4].¹⁶

Intuitively, given a non-malleable asymmetric encryption scheme Π , the construction $\text{Enc2Vote}(\Pi)$ derives ballot secrecy from Π until tallying and algorithm Tally maintains ballot secrecy by returning only the number of votes for each candidate. A formal proof of ballot secrecy follows from Quaglia & Smyth, in particular, Quaglia & Smyth show that $\text{Enc2Vote}(\Pi)$ satisfies a stronger notion of ballot secrecy [QS16, Proposition 5 & 16], hence, $\text{Enc2Vote}(\Pi)$ satisfies our notion of ballot secrecy too.

Corollary 2. *Let Π be an encryption scheme with perfect correctness. If Π satisfies IND-PA0, then election scheme $\text{Enc2Vote}(\Pi)$ satisfies Ballot-Secrecy.*

The reverse implication of Corollary 2 does not hold.

Proposition 3. *There exists an asymmetric encryption scheme Π such that election scheme $\text{Enc2Vote}(\Pi)$ satisfies Ballot-Secrecy, but Π does not satisfy IND-PA0.*

A proof of Proposition 3 and all further proofs, except where otherwise stated, appear in Appendix B.

¹⁴The construction by Quaglia & Smyth builds upon constructions by Bernhard *et al.* [SB14, SB13a, BPW12b, BCP⁺11].

¹⁵We define asymmetric encryption and an associated security property (namely, IND-PA0) in Appendix A.1.

¹⁶Quaglia & Smyth only consider asymmetric encryption schemes with perfect correctness, because they require election schemes to satisfy injectivity, and perfect correctness is required to show that $\text{Enc2Vote}(\Pi)$ satisfies injectivity. We adopt the same assumption to capitalise upon their results.

4 Ballot independence

Our informal definition of ballot independence (§1) essentially states that an adversary is unable to construct a ballot meaningfully related to a non-adversarial ballot. That is, ballots are non-malleable. Hence, we formulate ballot independence using non-malleability. The first formalisation of non-malleability is due to Dolev, Dwork & Naor [DDN91, DDN00], in the context of asymmetric encryption. Bellare & Sahai [BS99] build upon their results, and results by Bellare *et al.* [BDPR98], to introduce an alternative non-malleability definition for asymmetric encryption. We formalise non-malleability for election schemes as a straightforward adaptation of that definition.

Our formalisation of non-malleability for election schemes captures an intuitive notion of ballot independence, but the definition is complex and proofs of non-malleability are relatively difficult. Bellare & Sahai [BS99] observe similar complexities of non-malleability for encryption and show that their non-malleability definition for encryption is equivalent to a simpler, indistinguishability game for encryption. In a similar direction, we derive a simpler, equivalent definition of ballot independence as a straightforward adaptation of that indistinguishability game.

4.1 Non-malleability game

We formalise ballot independence as a non-malleability game.

Definition 4 (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 $\text{cnm-cva}(\Gamma, \mathcal{A}, \kappa)$ and $\text{cnm-cva-}\$(\Gamma, \mathcal{A}, \kappa)$ be the following games.¹⁷*

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

In the above games, we insist 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 $\text{Succ}(\text{cnm-cva}(\Gamma, \mathcal{A}, \kappa)) - \text{Succ}(\text{cnm-cva-}\$(\Gamma, \mathcal{A}, \kappa)) \leq \text{negl}(\kappa)$.

¹⁷We abbreviate $x \leftarrow_R S; x' \leftarrow_R S$ as $x, x' \leftarrow_R S$.

Similarly to game **Ballot-Secrecy**, games **cnm-cva** and **cnm-cva-\$** capture: key generation using algorithm **Setup**, publication of the public key, and only using the private key to compute the election outcome and tallying proof.

CNM-CVA is satisfied if no adversary can distinguish between games **cnm-cva** and **cnm-cva-\$**. That is, for all adversaries, we have with negligible probability that the adversary wins **cnm-cva** iff the adversary loses **cnm-cva-\$**. The first three steps of games **cnm-cva** and **cnm-cva-\$** are identical, thus, these steps cannot be distinguished. Then, game **cnm-cva-\$** performs an additional step: the challenger samples a second vote v' from vote space V . Thereafter, game **cnm-cva**($\Gamma, \mathcal{A}, \kappa$), respectively game **cnm-cva-\$**($\Gamma, \mathcal{A}, \kappa$), proceeds as follows: the challenger constructs a challenge ballot b for v , respectively v' ; the adversary is given ballot b and must compute a relation R and bulletin board \mathbf{bb} ; the challenger tallies \mathbf{bb} and outputs the election outcome \mathbf{v} ; and the challenger evaluates whether $R(v, \mathbf{v})$ holds. Hence, **CNM-CVA** is satisfied if there is no advantage of the relation constructed by an adversary given a challenge ballot for v , over the relation constructed by an 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 **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 , respectively v' , in **cnm-cva**, respectively **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 **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] demonstrate the following attack: an adversary observes a voter's ballot, casts a meaningfully related ballot, and exploits the relation to recover the voter's vote from the election outcome.

Comparing CNM-CVA and CNM-CPA. The main distinction between non-malleability for asymmetric encryption (**CNM-CPA**) and non-malleability for election schemes (**CNM-CVA**) is: **CNM-CPA** performs a parallel decryption, whereas, **CNM-CVA** performs a single tally. It follows that non-malleability for encryption reveals plaintexts corresponding to ciphertexts, whereas, non-malleability for elections reveals the number of ballots for each candidate.

4.2 Indistinguishability game

We formalise an alternative definition of ballot independence as an indistinguishability game.

Definition 5 (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 $\text{IND-CVA}(\Gamma, \mathcal{A}, \kappa)$ be the following game.*

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

We say Γ satisfies ballot independence or 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 $\text{IND-CVA}(\Gamma, \mathcal{A}, \kappa) \leq 1/2 + \text{negl}(\kappa)$.

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 constructed 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. Unsurprisingly, the distinction between indistinguishability for asymmetric encryption (IND-PA0) and indistinguishability for election schemes (IND-CVA), is similar to the distinction between non-malleability for asymmetric encryption and non-malleability for election schemes (§4.1), namely, IND-PA0 performs a parallel decryption, whereas, IND-CVA performs a single tally.

4.3 Equivalence between games

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

4.4 Non-malleable encryption is sufficient for independence

It follows naturally from our definitions that non-malleable ciphertexts are sufficient for ballot independence. Indeed, we can derive non-malleable ballots in our construction `Enc2Vote` using encryption schemes satisfying CNM-CPA.¹⁸

Corollary 5. *Let Π be an encryption scheme with perfect correctness. If Π satisfies CNM-CPA, then election scheme `Enc2Vote`(Π) satisfies CNM-CVA.*

A proof of Corollary 5 follows from Corollary 2 and Theorems 4 & 7. The reverse implication of Corollary 5 does not hold.

Corollary 6. *There exists an asymmetric encryption scheme Π such that election scheme `Enc2Vote`(Π) satisfies CNM-CVA, but Π does not satisfy CNM-CPA.*

A proof of Corollary 6 follows from Proposition 3 and Theorems 4 & 7.

5 Relations between secrecy and independence

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 our ballot independence game, whereas, the left-right oracle produces arbitrarily many challenge ballots for the adversary in our ballot secrecy game.
2. The adversary in our ballot secrecy game has access to a tallying proof, but the adversary in our ballot independence game does not.
3. The winning condition in our ballot secrecy game requires the bulletin board to be balanced, whereas, the bulletin board must not contain the challenge ballot in our ballot independence game.

The second point distinguishes our two games and shows that ballot secrecy is stronger than ballot independence.¹⁹ Hence, non-malleable ballots are necessary in election schemes satisfying ballot secrecy.

Theorem 7 (`Ballot-Secrecy` \Rightarrow `IND-CVA`). *Given an election scheme Γ satisfying `Ballot-Secrecy`, we have Γ satisfies `IND-CVA`.*

Moreover, since tallying proofs can reveal voters' votes (e.g., a variant of `Enc2Vote` could define tallying proofs that map ballots to votes) and these proofs are available to the adversary in our ballot secrecy game, but not in our ballot independence game, it follows that ballot secrecy is strictly stronger than ballot independence.

¹⁸Bellare & Sahai [BS99, §5] show that `IND-PA0` coincides with CNM-CPA, thus it suffices to consider `IND-PA0` in Corollaries 5 & 6.

¹⁹Smyth & Bernhard explain that alternative formalisations of election schemes might permit different results [SB13a, §5.2].

Proposition 8 (IND-CVA $\not\equiv$ Ballot-Secrecy). *There exists an election scheme Γ such that Γ satisfies IND-CVA, but not Ballot-Secrecy.*

A proof of Proposition 8 follows immediately from our informal reasoning and we omit a formal proof.

Although ballot secrecy is generally stronger than ballot independence, we show that ballot independence is sufficient for ballot secrecy in the class of election schemes without tallying proofs (Definition 6), assuming a soundness condition (Definition 7), which asserts that adding a ballot for v to the bulletin board effects the election outcome by exactly vote v . (This condition is required to hold in the presence of an adversary, whereas correctness is not. We show the condition is implied by universal verifiability in Appendix C.)

Definition 6. *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 multisets \mathbf{bb} we have:*

$$\Pr[(pk, sk, mb, mc) \leftarrow \text{Setup}(\kappa); (\mathbf{v}, pf) \leftarrow \text{Tally}(sk, \mathbf{bb}, nc, \kappa) : pf = \epsilon] = 1.$$

Definition 7 (Soundness). *Let $\Gamma = (\text{Setup}, \text{Vote}, \text{Tally})$ be an election scheme, \mathcal{A} be an adversary, κ be a security parameter, and $\text{Soundness}(\Gamma, \mathcal{A}, \kappa)$ be the following game.*

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

```

(pk, sk, mb, mc)  $\leftarrow$  Setup( $\kappa$ );
(v, nc,  $\mathbf{bb}_0$ )  $\leftarrow$   $\mathcal{A}$ (pk,  $\kappa$ );
b  $\leftarrow$  Vote(pk, v, nc,  $\kappa$ );
( $\mathbf{v}_0$ ,  $pf_0$ )  $\leftarrow$  Tally(sk,  $\mathbf{bb}_0$ , nc,  $\kappa$ );
( $\mathbf{v}_1$ ,  $pf_1$ )  $\leftarrow$  Tally(sk,  $\mathbf{bb}_0 \cup \{b\}$ , nc,  $\kappa$ );
 $\mathbf{v}^* \leftarrow$  ( $\mathbf{v}_0[1], \dots, \mathbf{v}_0[v-1], \mathbf{v}_0[v] + 1, \mathbf{v}_0[v+1], \dots, \mathbf{v}_0[|\mathbf{v}_0|]$ );
return  $\mathbf{v}^* \neq \mathbf{v}_1 \wedge b \notin \mathbf{bb}_0 \wedge 1 \leq v \leq nc \leq mc \wedge |\mathbf{bb}_0 \cup \{b\}| \leq mb$ ;

```

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

Proposition 9 (Ballot-Secrecy = IND-CVA, without tallying proofs). *Let Γ be an election scheme without tallying proofs. Suppose Γ satisfies Soundness. We have Γ satisfies Ballot-Secrecy iff Γ satisfies IND-CVA.*

Our equivalence result generalises to the class of election schemes with zero-knowledge tallying proofs, that is, election schemes that construct tallying proofs using zero-knowledge non-interactive proof systems.

Definition 8 (Zero-knowledge tallying proofs). *Let $\Gamma = (\text{Setup}, \text{Vote}, \text{Tally})$ be an election scheme. We say Γ has zero-knowledge tallying proofs, if there exists a zero-knowledge non-interactive 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 $pf = \text{Prove}((pk, \mathbf{bb}, nc, \mathbf{v}), sk, \kappa; r)$, such that coins r are chosen uniformly at random by Tally .*

Theorem 10 (Ballot-Secrecy = IND-CVA, with zero-knowledge tallying proofs). *Let Γ be an election scheme with zero-knowledge tallying proofs. Suppose Γ satisfies Soundness. We have Γ satisfies Ballot-Secrecy iff Γ satisfies IND-CVA.*

6 Case Study: Helios

Helios is an open-source, web-based electronic voting system,²⁰ which has been used in the real-world: the International Association of Cryptologic Research (IACR) has used Helios annually since 2010 to elect board members [BVQ10, HBH10],²¹ the 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 [Adi09].²²

Informally, Helios can be modelled as an election scheme (Setup, Vote, Tally) such that:

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

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

Tally proceeds as follows. First, any ballots on the bulletin board for which proofs do not hold are discarded. Secondly, the ciphertexts in the remaining ballots are homomorphically combined,²³ the homomorphic combination is decrypted to reveal the election outcome, and correctness of decryption is proved in zero-knowledge. Finally, the election outcome and proof of correct decryption are output.

Helios was first implemented as *Helios 2.0*.²⁴

Helios 2.0 is vulnerable to attacks against ballot secrecy [CS13, CS11, SC11, BPW12a].²⁵ And the next Helios release (Helios'12) is intended to mitigate against those attacks. In particular, the specification [Adi14] incorporates the Fiat-Shamir heuristic (rather than the weak Fiat-Shamir heuristic [BPW12a], which does not include statements in hashes). And there are plans to omit

²⁰<https://vote.heliosvoting.org>, accessed 19 Aug 2015.

²¹<https://www.iacr.org/elections/>, accessed 1 Sep 2015.

²²<https://princeton.heliosvoting.org/>, accessed 1 Sep 2015.

²³The homomorphic combination of ciphertexts is straightforward for two-candidate elections [CF85, BY86, SK94, Ben96, HS00], since votes (e.g., “yes” or “no”) can be encoded as 1 or 0. Multi-candidate elections are also possible [BY86, Hir10, DJN10].

²⁴<https://github.com/benadida/helios/releases/tag/2.0>, released 25 Jul 2009, accessed 16 Nov 2015.

²⁵Beyond ballot secrecy, attacks against universal verifiability [BPW12a, SFC16, CE16] and eligibility [SP13, SP15, SM16] are known.

meaningfully related ballots before tallying.^{26,27} 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.²⁸ These notions assume ballots are *recorded-as-cast*, i.e., cast ballots are preserved with integrity through the ballot collection process [AN06, §2]. Unfortunately, ballot secrecy is not satisfied without this assumption, because Helios'12 uses malleable ballots,²⁹ which are incompatible with ballot secrecy (§5).

Theorem 11. *Helios'12 does not satisfy* Ballot-Secrecy.

Proof sketch. Suppose an adversary queries the left-right oracle with inputs v_0 and v_1 to derive a ballot for v_β , where β is the bit chosen by the challenger. Further suppose the adversary exploits malleability to derive a related ballot b for v_β and outputs bulletin board $\{b\}$.³⁰ The board is balanced, because it does not contain the ballot output by the left-right oracle. Suppose the adversary performs the following computation on input of the election outcome \mathbf{v} : if $\mathbf{v}[v_0] = 1$, then output 0, otherwise, output 1. Since b is a ballot for v_β , it follows by correctness that $\mathbf{v}[v_0] = 1$ iff $\beta = 0$, and $\mathbf{v}[v_1] = 1$ iff $\beta = 1$, hence, the adversary wins the game. \square

Our informal proof of Theorem 11 is straightforward. A formal proof would require a formal description of Helios'12. Such a formal description can be derived by adapting the formalisation of Helios 3.1.4 by Smyth, Frink & Clarkson [SFC16] to omit meaningfully related ballots during tallying. These details provide little value, so we do not pursue them further.

The proof sketch of Theorem 11 does not immediately give way to an attack against Helios. Nevertheless, we can derive an attack (as the following example demonstrates) by extrapolating from the proof sketch and Cortier & Smyth's *permutation attack*, which asserts: given a ballot b for vote v , we can exploit 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 voted for v . (Thus, we exclude election outcomes without any votes for candidate v , which would

²⁶Cf. <http://documentation.heliosvoting.org/attacks-and-defenses>, <https://github.com/benadida/helios-server/issues/8>, and <https://github.com/benadida/helios-server/issues/35>, accessed 9 Aug 2016.

²⁷Mechanisms to omit ballots have been proposed, e.g., [CS11, SC11, Smy12, CS13, SB13b, BCG⁺15b, BCG⁺15a], but the specification for Helios'12 does not yet define a particular mechanism.

²⁸Proofs by Bernhard, Pereira & Warinschi and Bernhard *et al.* are limited to two candidate elections.

²⁹Ballots are non-malleable for two candidate elections. (Bernhard, Pereira & Warinschi and Bernhard *et al.* are reliant on non-malleability for their proofs.) Ballots are malleable in elections with more than two candidates. (Cf. [SHM15, §6].)

³⁰The recorded-as-cast assumption is violated because the ballot output by the left-right oracle does not appear on the bulletin board.

permit Mallory to trivially convince herself that Alice did not vote for candidate v .) Let us assume that votes for v' are not expected. Mallory proceeds as follows: she intercepts ballot b_1 , exploits 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 $\{b, b_2, b_3\}$. 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] would not detect the attack, because interception is not possible when ballots are recorded-as-cast.³¹

We have seen that non-malleable ballots are necessary for ballot secrecy (§5), hence, future Helios releases should adopt non-malleable ballots. Smyth, Frink & Clarkson [SFC16] make progress in this direction by proposing Helios'16, a variant of Helios which satisfies verifiability and is intended, but not proven, to use non-malleable ballots (cf. [SHM15]). We recall their formal description in Appendix D. And, using that formalisation, we prove that Helios'16 satisfies ballot secrecy.

Theorem 12. *Helios'16 satisfies Ballot-Secrecy.*

Proof sketch. We prove that Helios'16 has zero-knowledge tallying proofs. And, since Helios'16 satisfies universal verifiability [SFC16], it also satisfies Soundness (§C). Hence, by Theorem 10, it suffices to prove that Helios'16 satisfies IND-CVA. And we show that satisfying IND-CVA reduces to the security of the encryption scheme (namely, IND-CPA of El Gamal) underlying Helios'16. \square

A formal proof of Theorem 12 appears in Appendix E. The proof assumes the random oracle model [BR93]. This proof, coupled with the proof of verifiability by Smyth, Frink & Clarkson [SFC16], provides strong motivation for future Helios releases being based upon Helios'16, since it is the only variant of Helios which is known to be secure.

7 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].³² 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, SB13a, SB14, BCG⁺15b]. The definition of ballot secrecy by Bernhard, Pereira & Warinschi computes

³¹This 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.

³²Bernhard *et al.* [BCG⁺15b, BCG⁺15a] survey ballot secrecy definitions.

tallying proofs using algorithm Tally or a simulator [BPW12a], but the resulting definition is too weak [BCG⁺15b, §3.4] and some strengthening is required [BCG⁺15b, §4]. (Cortier *et al.* [CGGI13a, CGGI13b] propose a variant of the ballot secrecy definition by Bernhard, Pereira & Warinschi. That variant is also too weak [BCG⁺15b].) By comparison, the definition by Smyth & Bernhard computes tallying proofs using only algorithm Tally [SB13a], but the resulting definition is too strong [BCG⁺15b, §3.5] and a weakening is required [SB14]. The relative merits of ballot secrecy definitions due to Smyth & Bernhard [SB14, Definition 5] and 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. Benaloh [Ben96, §2.9] 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. And, more generally, Sako & Kilian [SK95, §2.4], Michels & Horster [MH96, §3] and Cortier & Smyth [CS13, CS11] discuss how malleable ballots can be exploited to compromise ballot secrecy. The first definition of ballot independence seems to be due to Smyth & Bernhard [SB13a, SB14]. Moreover, Smyth & Bernhard formally prove relations between their definitions of ballot secrecy and ballot independence.

All of the ballot secrecy definitions by Bernhard *et al.* [BCP⁺11, BPW12b, BPW12a, SB13a, SB14, BCG⁺15b] and the ballot independence definition by Smyth & Bernhard [SB13a, SB14] focus on detecting attacks by adversaries that control some voters. Attacks by adversaries that control the bulletin board or 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 [SB13a, §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, despite the use of malleable ballots. 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 due to Bernhard *et al.* by detecting more attacks.

Some of the ideas presented in this paper previously appeared in a technical report by Smyth [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 by Smyth [Smy14]. And Definition 2 is based upon the definition of ballot secrecy proposed by Smyth [Smy14, Definition 3]. The main distinction between Definition 2 and the definition by Smyth is syntax: this paper adopts syntax for election schemes from Smyth, Frink & Clarkson [SFC16], whereas, Smyth adopts syntax by Smyth & Bern-

hard [SB14, SB13a]. The change in syntax is motivated by the superiority of syntax by Smyth, Frink & Clarkson. Unfortunately, the change has a drawback: we cannot immediately prove that the definition of ballot secrecy proposed in this paper is strictly stronger than the definition proposed by Smyth & Bernhard [SB14, SB13a]. By comparison, the technical reports contain such proofs. Nevertheless, the advantages of the syntax change outweigh the disadvantages. Moreover, we can capitalise upon results by Smyth, Frink & Clarkson [SFC16] and Quaglia & Smyth [QS16].

McCarthy, Smyth & Quaglia [MSQ14] have shown that auction schemes can be constructed from election schemes, and Quaglia & Smyth [QS16] 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 election verifiability by Smyth, Frink & Clarkson [SFC16] to a definition of auction verifiability, and prove that their construction produces schemes satisfying auction verifiability.) Thus, this research has applications beyond voting.

8 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. This enabled definitions of ballot independence 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 election schemes. This prompted the study of relations between ballot secrecy and ballot independence, resulting in a proof that non-malleable ballots are necessary for ballot secrecy. And, moreover, a proof that non-malleable ballots are sufficient for ballot secrecy in election schemes with zero-knowledge tallying proofs. Furthermore, a separation result demonstrates that ballot secrecy is strictly stronger than ballot independence.

In light of the revelation that non-malleable ballots are necessary for ballot secrecy, and in the knowledge that Helios ballots are malleable, it was discovered that Helios does not satisfy ballot secrecy. Although the proof sketch of this result did not immediately uncover an attack against Helios, an extrapolation

from that proof sketch revealed an attack that allows an adversary to determine if a voter did not vote for a candidate chosen by the adversary. This naturally led to the consideration of whether definitions of ballot secrecy by Bernhard *et al.* could have detected this attack and the conclusion that they could not, because the attack requires the adversary to control the bulletin board or communication channel, which is prohibited by their definitions.

We exploited our results to prove that a variant of Helios satisfies ballot secrecy. This proof is particularly significant due to the use of Helios in the real-world. And we encourage Helios developers to base future releases on this variant, since it is the only variant of Helios which is known to be secure.

Acknowledgements. Some of the prose (in particular, the two opening paragraphs of Section 3) were prepared in collaboration with David Bernhard and I am very grateful for David’s contribution. I am also grateful to David for extensive discussions that helped improve this paper and, more generally, my knowledge of cryptography. In addition, I am grateful to Elizabeth Quaglia for her valuable feedback that also helped improve this paper. This work was performed in part at INRIA, with support from the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013) / ERC project *CRYSP* (259639).

A Cryptographic primitives

A.1 Asymmetric encryption

Definition 9 (Asymmetric encryption scheme [KL07]). *An asymmetric encryption scheme is a tuple of probabilistic polynomial-time algorithms (Gen, Enc, Dec), such that:*³³

- **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.

³³Our definition differs from Katz and Lindell’s original definition [KL07, Definition 10.1] in that we formally state the plaintext space.

Definition 10 (Homomorphic encryption [SFC16]). *An asymmetric encryption scheme $\Gamma = (\text{Gen}, \text{Enc}, \text{Dec})$ is homomorphic, with respect to ternary operators \odot , \oplus , and \otimes ,³⁴ if there exists a negligible function negl , such that for all security parameters κ , we have the following.³⁵ First, for all messages m_1 and m_2 we have $\Pr[(pk, sk, \mathbf{m}) \leftarrow \text{Gen}(\kappa); c_1 \leftarrow \text{Enc}(pk, m_1); c_2 \leftarrow \text{Enc}(pk, m_2) : m_1, m_2 \in \mathbf{m} \Rightarrow \text{Dec}(sk, c_1 \otimes_{pk} c_2) = \text{Dec}(sk, c_1) \odot_{pk} \text{Dec}(sk, c_2)] > 1 - \text{negl}(\kappa)$. 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 Γ 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 11 (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 $\text{IND-CPA}(\Pi, \mathcal{A}, \kappa)$ be the following game.³⁶*

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

In the above game, we insist $m_0, m_1 \in \mathbf{m}$ and $|m_0| = |m_1|$. We say Γ 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 $\text{Succ}(\text{IND-CPA}(\Pi, \mathcal{A}, \kappa)) \leq 1/2 + \text{negl}(\kappa)$.

Definition 12 (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 $\text{IND-PA0}(\Pi, \mathcal{A}, \kappa)$ be the following game.*

$\text{IND-PA0}(\Pi, \mathcal{A}, \kappa) =$

³⁴Henceforth, we implicitly bind ternary operators, i.e., we write Γ is a homomorphic asymmetric encryption scheme as opposed to the more verbose Γ is a homomorphic asymmetric encryption scheme, with respect to ternary operators \odot , \oplus , and \otimes .

³⁵We 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.

³⁶Our 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.

```

(pk, sk, m) ← Gen(κ);
(m0, m1) ← A(pk, m, κ);
β ←R {0, 1};
c ← Enc(pk, mβ);
c ← A(c);
m ← (Dec(sk, c[1]), ..., Dec(sk, c[|c|]));
g ← A(m);
return g = β ∧  $\bigwedge_{1 \leq i \leq |c|} c \neq c[i]$ ;

```

In the above game, we insist $m_0, m_1 \in \mathbf{m}$ and $|m_0| = |m_1|$. We say Γ 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 $\text{Succ}(\text{IND-PA0}(\Pi, \mathcal{A}, \kappa)) \leq 1/2 + \text{negl}(\kappa)$.

A.2 Proof systems

Definition 13 (Sigma protocol [SFC16, Dam10, HL10]). A sigma protocol for a relation R is a tuple $(\text{Comm}, \text{Chal}, \text{Resp}, \text{Verify})$ of probabilistic polynomial-time algorithms such that:

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

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

Definition 14 (Non-interactive proof system [SFC16]). A non-interactive proof system for a relation R is a tuple of algorithms $(\text{Prove}, \text{Verify})$, 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 $\Pr[\sigma \leftarrow \text{Prove}(s, w, \kappa) : \text{Verify}(s, \sigma, \kappa) = 1] > 1 - \text{negl}(\kappa)$.

Definition 15 (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 tuple $(\text{Prove}, \text{Verify})$ of algorithms, defined as follows:

```

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 ( $\text{comm}, \text{resp}$ );

```

```

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

```

Definition 16 (Zero-knowledge [QS16]). 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 the sigma protocol. Moreover, let \mathcal{S} be an algorithm, \mathcal{A} be an adversary, κ be a security parameter, and $\text{ZK}(\Delta, \mathcal{A}, \mathcal{H}, \mathcal{S}, \kappa)$ be the following game.

```

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} .³⁷ 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 $(\text{Prove}, \text{Verify})$.

Definition 17 (Simulation sound extractability [SFC16, BPW12a, Gro06]). Suppose Σ is a sigma protocol for relation R , \mathcal{H} is a random oracle, and $(\text{Prove}, \text{Verify})$ is a non-interactive proof system, such that $\text{FS}(\Sigma, \mathcal{H}) = (\text{Prove}, \text{Verify})$. Further suppose \mathcal{S} is a simulator for $(\text{Prove}, \text{Verify})$ and \mathcal{H} can be patched by

³⁷Random 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.

\mathcal{S} . Proof system $(\text{Prove}, \text{Verify})$ 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:³⁸

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

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 $\sigma \leftarrow \mathcal{S}(s); \mathbf{P} \leftarrow (\mathbf{P}[1], \dots, \mathbf{P}[|\mathbf{P}|], (s, \sigma))$ and outputs σ .

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

Theorem 13 (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 can be generalised to include an optional string in the hashes produced by functions Prove and Verify (§A). Simulators can be generalised to include an optional string m too. We write $\mathcal{S}(s, m, \kappa)$ for invocations of simulator \mathcal{S} which include an optional string. Theorem 13 can be extended to this generalisation.

The Fiat-Shamir transformation can be generalised to include an optional string m in the hashes produced by functions Prove and Verify . We write $\text{Prove}(s, w, m, \kappa)$ and $\text{Verify}(s, (\text{comm}, \text{resp}), m, k)$ for invocations of Prove and Verify which include an optional string. When m is provided, it is included in the hashes in both algorithms. That is, given $\text{FS}(\Sigma, \mathcal{H}) = (\text{Prove}, \text{Verify})$, the hashes are computed as follows in both algorithms: $\text{chal} \leftarrow \mathcal{H}(\text{comm}, s, m)$. Simulators can be generalised to include an optional string m too. We write $\mathcal{S}(s, m, \kappa)$ for invocations of simulator \mathcal{S} which include an optional string. Theorem 13 can be extended to this generalisation.

B Proofs

B.1 Proof of Proposition 3

We present a construction (Definition 18) for encryption schemes (Lemma 14) which are clearly not secure (Lemma 15). Nevertheless, the construction pro-

³⁸We extend set membership notation to vectors: we write $x \in \mathbf{x}$ if x is an element of the set $\{\mathbf{x}[i] : 1 \leq i \leq |\mathbf{x}|\}$.

duces encryption schemes that are sufficient for ballot secrecy (Lemma 16). The proof of Proposition 3 follows from Lemmata 14–16.

Definition 18. *Given an asymmetric encryption scheme $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ and a constant symbol ω , let $\text{Leak}(\Pi, \omega) = (\text{Gen}, \text{Enc}, \text{Dec}')$, such that $\text{Dec}'(sk, c)$ proceeds as follows: if $c = \omega$, then output sk , otherwise, compute $m \leftarrow \text{Dec}(sk, c)$ and output m .*

Lemma 14. *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 15. *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 16. *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 $\text{Enc2Vote}(\text{Leak}(\Pi, \omega)) = (\overline{\text{Setup}}, \overline{\text{Vote}}, \overline{\text{Tally}})$. By definition of Enc2Vote and Leak , we have $\text{Setup} = \overline{\text{Setup}}$ and $\text{Vote} = \overline{\text{Vote}}$. 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

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

because Enc2Vote ensures that $\overline{\mathbf{v}}$ 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 3. 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 14) such that $\text{Enc2Vote}(\text{Leak}(\Pi, \omega))$ satisfies Ballot-Secrecy (Lemma 16), but $\text{Leak}(\Pi, \omega)$ does not satisfy IND-PA0 (Lemma 15), concluding our proof. \square

B.2 Proof of Theorem 4

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 adversary \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 $\text{IND-CVA}(\Gamma, \mathcal{B}, \kappa)$, then adversary \mathcal{B} simulates \mathcal{A} 's challenger to \mathcal{A} in $\text{cnm-cva}(\Gamma, \mathcal{A}, \kappa)$ 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 $\text{cnm-cva-}\$(\Gamma, \mathcal{A}, \kappa)$ 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 $\text{Succ}(\text{IND-CVA}(\Gamma, \mathcal{A}, \kappa)) = 1/2 \cdot (\text{Succ}(\text{cnm-cva}(\Gamma, \mathcal{A}, \kappa)) + 1 - \text{Succ}(\text{cnm-cva-}\$(\Gamma, \mathcal{A}, \kappa)))$ and, therefore, $2 \cdot \text{Succ}(\text{IND-CVA}(\Gamma, \mathcal{A}, \kappa)) - 1 = \text{Succ}(\text{cnm-cva}(\Gamma, \mathcal{A}, \kappa)) - \text{Succ}(\text{cnm-cva-}\$(\Gamma, \mathcal{A}, \kappa))$. Since Γ does not satisfy CNM-CVA and a function that doubles the output of a negligible function is a negligible function, we have $\text{Succ}(\text{cnm-cva}(\Gamma, \mathcal{A}, \kappa)) - \text{Succ}(\text{cnm-cva-}\$(\Gamma, \mathcal{A}, \kappa)) > 2 \cdot \text{negl}(\kappa)$. It follows that $2 \cdot \text{Succ}(\text{IND-CVA}(\Gamma, \mathcal{A}, \kappa)) - 1 > 2 \cdot \text{negl}(\kappa)$, hence, $\text{Succ}(\text{IND-CVA}(\Gamma, \mathcal{A}, \kappa)) > 1/2 + \text{negl}(\kappa)$, concluding our proof.

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 $\text{Succ}(\text{IND-CVA}(\Gamma, \mathcal{A}, \kappa)) > 1/2 + \text{negl}(\kappa)$. We construct an adversary \mathcal{B} against CNM-CVA from adversary \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 $\text{IND-CVA}(\Gamma, \mathcal{A}, \kappa)$. 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 one time in the game. Let us consider adversary \mathcal{B} 's success in $\text{cnm-cva}(\Gamma, \mathcal{B}, \kappa)$ and $\text{cnm-cva-}\$(\Gamma, \mathcal{B}, \kappa)$.

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

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

B.3 Proof of Theorem 7

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 $\text{Succ}(\text{IND-CVA}(\Gamma, \mathcal{A}, \kappa)) > 1/2 + \text{negl}(\kappa)$. We construct a ballot secrecy adversary \mathcal{B} from the ballot independence adversary \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, $\text{Succ}(\text{Ballot-Secrecy}(\Gamma, \mathcal{B}, \kappa)) > 1/2 + \text{negl}(\kappa)$, concluding our proof. \square

B.4 Proof of Proposition 9

In essence, the proof follows from Theorem 10. 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 constructed by a zero-knowledge non-interactive proof system, but an election scheme without tallying proofs need not construct 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

B.5 Proof of Theorem 10

Let BS-0, respectively BS-1, be the game derived from **Ballot-Secrecy** by replacing $\beta \leftarrow_R \{0, 1\}$ with $\beta \leftarrow 0$, respectively $\beta \leftarrow 1$. These games are trivially related to **Ballot-Secrecy**, namely, $\text{Succ}(\text{Ballot-Secrecy}(\Gamma, \mathcal{A}, \kappa)) = \frac{1}{2} \cdot \text{Succ}(\text{BS-0}(\Gamma, \mathcal{A}, \kappa)) + \frac{1}{2} \cdot \text{Succ}(\text{BS-1}(\Gamma, \mathcal{A}, \kappa))$. Moreover, let BS-1:0 be the game derived from BS-1 by replacing $g = \beta$ with $g = 0$. We relate game BS-1:0 to BS-1, and games BS-0 and BS-1:0 to the hybrid games G_0, G_1, \dots introduced in Definition 19. We prove Theorem 10 using these relations.

Lemma 17. *Given an adversary \mathcal{A} that wins game **Ballot-Secrecy** against election scheme Γ , we have $\text{Succ}(\text{BS-1}(\Gamma, \mathcal{A}, \kappa)) = 1 - \text{Succ}(\text{BS-1:0}(\Gamma, \mathcal{A}, \kappa))$ for all security parameters κ .*

Definition 19. *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 zero-knowledge non-interactive proof system used by algorithm **Tally** to construct tallying proofs. We introduce games G_0, G_1, \dots , defined as follows.*

```

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

```

Oracle \mathcal{O} is defined 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 and tallying procedures. In particular, the first j left-right oracle queries in G_j compute ballots for the oracle's "left" input and any remaining queries compute ballots for the oracle's "right" input, whereas the left-right oracle in BS-0, respectively BS-1:0, always computes ballots for the oracle's "left," respectively "right," input. Moreover, the tallying procedure in G_j computes

the outcome by tallying the ballots on the bulletin board that were constructed by the adversary and by simulating the tallying of any remaining ballots (i.e., ballots constructed by the oracle). And the tallying proof is simulated in G_j . By comparison, the outcome and tallying proof are computed by tallying all the ballots on the bulletin board in both BS-0 and BS-1:0.

Lemma 18. *Let Γ be an election scheme, \mathcal{A} be an adversary, and κ be a security parameter. If Γ satisfies Soundness, then $\text{Succ}(\text{BS-0}(\Gamma, \mathcal{A}, \kappa)) = \text{Succ}(G_0(\Gamma, \mathcal{A}, \kappa))$ and $\text{Succ}(\text{BS-1:0}(\Gamma, \mathcal{A}, \kappa)) = \text{Succ}(G_q(\Gamma, \mathcal{A}, \kappa))$, where q is an upper-bound on \mathcal{A} 's left-right oracle queries.*

Proof. The challengers in games BS-0 and G_0 , respectively BS-1:0 and G_q , both construct 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.

Left-right oracles queries $\mathcal{O}(v_0, v_1)$ in games BS-0 and G_0 output ballots for vote v_0 , hence, the bulletin boards are equivalent in both games. The bulletin boards 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 . Thus, the bulletin board output by the second adversary call is equivalent in BS-0 and G_0 , respectively BS-1:0 and G_q .

It follows that $1 \leq nc \leq mc \wedge |\mathbf{bb}| \leq mb$ in BS-0 iff $1 \leq nc \leq mc \wedge |\mathbf{bb}| \leq mb$ in G_0 , and similarly for BS-1:0 and G_q . Moreover, predicate *balanced* is satisfied in BS-0 iff it is satisfied in G_0 , and similarly for BS-1:0 and G_q . Hence, if $1 \leq nc \leq mc \wedge |\mathbf{bb}| \leq mb$ is not satisfied or predicate *balanced* is not satisfied, then $\text{Succ}(\text{BS-0}(\Gamma, \mathcal{A}, \kappa)) = \text{Succ}(G_0(\Gamma, \mathcal{A}, \kappa))$ and $\text{Succ}(\text{BS-1:0}(\Gamma, \mathcal{A}, \kappa)) = \text{Succ}(G_q(\Gamma, \mathcal{A}, \kappa))$, concluding our proof. Otherwise, it suffices to show that the outcome and tallying proof are equivalently computed in BS-0 and G_0 , respectively BS-1:0 and G_q , since this ensures the inputs to the third adversary call are equivalent, thus the corresponding outputs are equivalent too, which suffices to conclude.

In BS-0, respectively BS-1:0, the outcome is computed by tallying the bulletin board. By comparison, in G_0 , respectively G_q , 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 such that $nc \leq mc$. Since Γ satisfies Soundness, 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

```

(v, pf) ← Tally(sk, bb \ {b | (b, v_0, v_1) ∈ L}, nc, κ);
for b ∈ bb ∧ (b, v_0, v_1) ∈ L do
  (v', pf') ← Tally(sk, {b}, nc, κ);
  v ← v + v';

```

Thus, to prove the outcome is computed equivalently in BS-0 and G_0 , respectively BS-1:0 and G_q , it suffices to prove that the simulations are valid, i.e., computing the above is equivalent to computing

```

(v, pf) ← Tally(sk, bb \ {b | (b, v_0, v_1) ∈ L}, nc, κ);
for b ∈ bb ∧ (b, v_0, v_1) ∈ L do
  v[v_0] ← v[v_0] + 1;

```

In G_0 , respectively G_q , we have for all $(b, v_0, v_1) ∈ L$ that b is an output of $\text{Vote}(pk, v_0, nc, \kappa)$, respectively $\text{Vote}(pk, v_1, nc, \kappa)$, such that $v_0, v_1 ∈ \{1, \dots, nc\}$. Moreover, by correctness of Γ , we have $\text{Tally}(sk, \{b\}, nc, \kappa)$ outputs (v', pf') such that v' is a zero-filled vector, except for index v_0 , respectively v_1 , which contains one. Hence, the simulation is valid in G_0 . Furthermore, since predicate *balanced* holds in G_q , we have for all $v ∈ \{1, \dots, nc\}$ that $|\{b | b ∈ bb ∧ \exists v_1 . (b, v, v_1) ∈ L\}| = |\{b | b ∈ bb ∧ \exists v_0 . (b, v_0, v) ∈ L\}|$. Hence, in G_q , computing

```

for b ∈ bb ∧ (b, v_0, v_1) ∈ L do v[v_0] ← v[v_0] + 1;

```

is equivalent to computing

```

for b ∈ bb ∧ (b, v_0, v_1) ∈ L do v[v_1] ← v[v_1] + 1;

```

Thus, the simulation is valid in G_q too.

In BS-0, respectively BS-1:0, the tallying proof is computed by tallying the bulletin board. By comparison, in G_0 , respectively G_q , 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 (v, pf) output by $\text{Tally}(sk, bb, nc, \kappa)$, we have $pf = \text{Prove}((pk, bb, nc, v), sk, \kappa; r)$, such that coins r are chosen uniformly at random by Tally . Moreover, since \mathcal{S} is a simulator for (Prove, Verify), proofs output by $\text{Prove}((pk, nc, bb, v), w, \kappa)$ are indistinguishable from outputs of $\mathcal{S}((pk, nc, bb, v), \kappa)$. Thus, tallying proofs are equivalently computed in BS-0 and G_0 , respectively BS-1:0 and G_q , thereby concluding our proof. \square

Proof of Theorem 10. By Theorem 7, 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

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

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

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

And, by Lemma 17, we have

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

Let q be an upper-bound on \mathcal{A} 's left-right oracle queries. Hence, by Lemma 18, we have

$$= \frac{1}{2} + \frac{1}{2} \cdot (\text{Succ}(\text{G}_0(\Gamma, \mathcal{A}, \kappa)) - \text{Succ}(\text{G}_q(\Gamma, \mathcal{A}, \kappa)))$$

which can be rewritten as the telescoping series

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

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

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

Thus,

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

From \mathcal{A} , we construct an adversary \mathcal{B} against IND-CVA whose success is at least $\frac{1}{2} + \frac{1}{2} \cdot (\text{Succ}(\text{G}_{j-1}(\Gamma, \mathcal{A}, \kappa)) - \text{Succ}(\text{G}_j(\Gamma, \mathcal{A}, \kappa)))$.

Let $\Gamma = (\text{Setup}, \text{Vote}, \text{Tally})$. Since Γ has zero-knowledge tallying proofs, tallying proofs output by Tally are constructed by a zero-knowledge non-interactive proof system. Let algorithm \mathcal{S} be the simulator for that proof system. We define \mathcal{B} as follows.

- $\mathcal{B}(pk, \kappa)$ 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** $\mathbf{v}[v_0] \leftarrow \mathbf{v}[v_0] + 1$, and $pf \leftarrow \mathcal{S}((pk, nc, \mathbf{bb}, \mathbf{v}), \kappa)$; $g \leftarrow \mathcal{A}(\mathbf{v}, pf)$, and outputs g .

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

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 \mathbb{G}_{j-1} and \mathbb{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 \mathbb{G}_{j-1} and \mathbb{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 $\text{Succ}(\mathbb{G}_{j-1}(\Gamma, \mathcal{A}, \kappa)) - \text{Succ}(\mathbb{G}_j(\Gamma, \mathcal{A}, \kappa))$ 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 \mathbb{G}_{j-1} to \mathcal{A} , otherwise, $\mathcal{B}(b)$ simulates the oracle in \mathbb{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 \mathbb{G}_{j-1} when $\beta = 0$, respectively \mathbb{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 \mathbb{G}_{j-1} to \mathcal{A} , thus, $g = \beta$ with at least the probability that \mathcal{A} wins \mathbb{G}_{j-1} ; or
2. $\beta = 1$ and \mathcal{B} simulates \mathbb{G}_j to \mathcal{A} , thus, $g \neq 0$ with at least the probability that \mathcal{B} loses \mathbb{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 $\frac{1}{2} \cdot \text{Succ}(\mathbb{G}_{j-1}(\Gamma, \mathcal{A}, \kappa)) + \frac{1}{2} \cdot (1 - \text{Succ}(\mathbb{G}_j(\Gamma, \mathcal{A}, \kappa)))$, thus we conclude our proof. \square

C Universal verifiability implies soundness

We recall the full syntax for election schemes in Definition 20. (The syntax for election schemes presented in Section 2 omitted algorithm **Verify** and the condition that election schemes must satisfy notions of completeness and injectivity, because we did not consider verifiability in the main body.)

Definition 20 (Election scheme [SFC16]). *An election scheme is a tuple of probabilistic polynomial-time algorithms (**Setup**, **Vote**, **Tally**, **Verify**) such that (**Setup**, **Vote**, **Tally**) is an election scheme and*

- **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 , some number of candidates nc , a bulletin board \mathbf{bb} , an election outcome \mathbf{v} , a proof pf , and a security parameter κ . It outputs a bit s , which is 1 if the election verifies successfully or 0 otherwise.

Election schemes must satisfy Completeness: there exists a negligible function negl , such that for all security parameters κ , bulletin boards \mathbf{bb} , and integers nc ,

we have

$$\Pr[(pk, sk, mb, mc) \leftarrow \text{Setup}(\kappa); (\mathbf{v}, pf) \leftarrow \text{Tally}(sk, \mathbf{bb}, nc, \kappa) \\ : |\mathbf{bb}| \leq mb \wedge nc \leq mc \Rightarrow \text{Verify}(pk, \mathbf{bb}, nc, \mathbf{v}, pf, \kappa) = 1] > 1 - \text{negl}(\kappa).$$

Election schemes must also satisfy Injectivity: for all security parameters κ , integers nc , and votes v and v' , such that $v \neq v'$, we have

$$\Pr[(pk, sk, mb, mc) \leftarrow \text{Setup}(\kappa); b \leftarrow \text{Vote}(pk, v, nc, \kappa); \\ b' \leftarrow \text{Vote}(pk, v', nc, \kappa) : b \neq \perp \wedge b' \neq \perp \Rightarrow b \neq b'] = 1.$$

Universal verifiability (Definition 21) challenges the adversary to concoct a scenario in which `Verify` accepts, but the election outcome is not correct. Formally, we capture the correct outcome using function *correct-outcome*,³⁹ defined such that for all $pk, nc, \mathbf{bb}, \kappa, \ell$, and $v \in \{1, \dots, nc\}$, we have:

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

That is, component v of vector *correct-outcome*($pk, \mathbf{bb}, nc, \kappa$) equals ℓ iff there exist ℓ ballots on the bulletin board that are votes for candidate v . The vector produced by *correct-outcome* must be of length nc .

Definition 21 (Exp-UV-Ext [SFC16]). *Let $\Gamma = (\text{Setup}, \text{Vote}, \text{Tally}, \text{Verify})$ be an election scheme, \mathcal{A} be an adversary, κ be a security parameter, and Exp-UV-Ext($\Gamma, \mathcal{A}, \kappa$) be the following game.*

$$\text{Exp-UV-Ext}(\Gamma, \mathcal{A}, \kappa) = \\ (pk, nc, \mathbf{bb}, \mathbf{v}, pf) \leftarrow \mathcal{A}(\kappa); \\ \text{return} \\ \mathbf{v} \neq \text{correct-outcome}(pk, nc, \mathbf{bb}, \kappa) \wedge \text{Verify}(pk, \mathbf{bb}, nc, \mathbf{v}, pf, \kappa) = 1;$$

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

We show that universally verifiable election schemes satisfy Soundness (Proposition 20).

Lemma 19. *Given an election scheme (Setup, Vote, Tally), there exists a negligible function negl , such that for all security parameters κ , integers nc , and votes $v \in \{1, \dots, nc\}$, we have*

$$\Pr[(pk, sk, mb, mc) \leftarrow \text{Setup}(\kappa); b \leftarrow \text{Vote}(pk, v, nc, \kappa) \\ : 1 \leq mb \wedge nc \leq mc \Rightarrow b \neq \perp] > 1 - \text{negl}(\kappa).$$

³⁹Function *correct-outcome* uses a counting quantifier [Sch05] denoted \exists^{ℓ} . Predicate $(\exists^{\ell} x : P(x))$ holds exactly when there are ℓ distinct values for x such that $P(x)$ is satisfied. Variable x is bound by the quantifier, whereas ℓ is free.

Proof. Suppose κ is a security parameter and nc and v are integers, such that $v \in \{1, \dots, nc\}$. Further suppose (pk, sk, mb, mc) is an output of $\text{Setup}(\kappa)$, b is an output of $\text{Vote}(pk, v, nc, \kappa)$, and (\mathbf{v}, pf) is an output of $\text{Tally}(sk, \{b\}, nc, \kappa)$, such that $1 \leq mb \wedge nc \leq mc$. By correctness, we have \mathbf{v} is a zero-filled vector of length nc , except for index v which contains integer 1, with overwhelming probability. Given that $\text{Tally}(sk, \{b\}, nc, \kappa)$ and $\text{Tally}(sk, \{b, b\}, nc, \kappa)$ input the same set $\{b\}$, correctness ensures the probability of $\text{Vote}(pk, v, nc, \kappa)$ outputting two identical ballots is upper-bounded by a negligible function. It follows that the probability of $\text{Vote}(pk, v, nc, \kappa)$ outputting error symbol \perp twice is upper-bounded by a negligible function too. Moreover, the probability of $\text{Vote}(pk, v, nc, \kappa)$ outputting error symbol \perp is also upper-bounded by a negligible function, thereby concluding our proof. \square

Proposition 20 (Exp-UV-Ext \Rightarrow Soundness). *If election scheme Γ satisfies Exp-UV-Ext, then Γ satisfies Soundness.*

Proof. Let $\Gamma = (\text{Setup}, \text{Vote}, \text{Tally}, \text{Verify})$. Suppose Γ does not satisfy 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 $\text{negl}(\kappa) < \text{Succ}(\text{Soundness}(\Gamma, \mathcal{A}, \kappa))$. We construct an adversary \mathcal{B} against Exp-UV-Ext from \mathcal{A} . We define \mathcal{B} as follows.

```

 $\mathcal{B}(\kappa) =$ 
   $(pk, sk, mb, mc) \leftarrow \text{Setup}(\kappa);$ 
   $(v, nc, \mathbf{bb}_0) \leftarrow \mathcal{A}(pk, \kappa);$ 
   $(\mathbf{v}_0, pf_0) \leftarrow \text{Tally}(sk, \mathbf{bb}_0, nc, \kappa);$ 
   $\beta \leftarrow_R \{0, 1\};$ 
  if  $\beta = 1$  then
     $b \leftarrow \text{Vote}(pk, v, nc, \kappa);$ 
     $\mathbf{bb}_1 \leftarrow \mathbf{bb}_0 \cup \{b\};$ 
     $(\mathbf{v}_1, pf_1) \leftarrow \text{Tally}(sk, \mathbf{bb}_1, nc, \kappa);$ 
  return  $(pk, nc, \mathbf{bb}_\beta, \mathbf{v}_\beta, pf_\beta);$ 

```

We prove that \mathcal{B} wins Exp-UV-Ext with non-negligible probability.

Suppose (pk, sk, mb, mc) is an output of $\text{Setup}(\kappa)$, (v, nc, \mathbf{bb}_0) is an output of $\mathcal{A}(pk, \kappa)$, b is an output of $\text{Vote}(pk, v, nc, \kappa)$, (\mathbf{v}_0, pf_0) is an output of $\text{Tally}(sk, \mathbf{bb}_0, nc, \kappa)$, and (\mathbf{v}_1, pf_1) is an output of $\text{Tally}(sk, \mathbf{bb}_1, nc, \kappa)$, where $\mathbf{bb}_1 = \mathbf{bb}_0 \cup \{b\}$. Let $\mathbf{v}^* \leftarrow (\mathbf{v}_0[1], \dots, \mathbf{v}_0[v-1], \mathbf{v}_0[v] + 1, \mathbf{v}_0[v+1], \dots, \mathbf{v}_0[|v_0|])$. Since \mathcal{A} is a winning adversary, we have $\mathbf{v}^* \neq \mathbf{v}_1 \wedge b \notin \mathbf{bb}_0 \wedge 1 \leq v \leq nc \leq mc \wedge |\mathbf{bb}_0 \cup \{b\}| \leq mb$, with probability greater than $\text{negl}(\kappa)$.

Suppose β is a bit chosen uniformly at random. By Completeness, we have $\text{Verify}(pk, \mathbf{bb}_\beta, nc, \mathbf{v}_\beta, pf_\beta, \kappa) = 1$, with overwhelming probability. Hence, it suffices to prove that $\mathbf{v}_\beta \neq \text{correct-outcome}(pk, nc, \mathbf{bb}_\beta, \kappa)$, with non-negligible probability. Let δ_0 , respectively δ_1 , be the probability that $\mathbf{v}_0 \neq \text{correct-outcome}(pk, nc, \mathbf{bb}_0, \kappa)$, respectively $\mathbf{v}_1 \neq \text{correct-outcome}(pk, nc, \mathbf{bb}_1, \kappa)$. It follows that $\text{Succ}(\text{Exp-UV-Ext}(\Gamma, \mathcal{B}, \kappa)) = \frac{1}{2} \cdot \delta_0 + \frac{1}{2} \cdot \delta_1$ and it remains to show that $\delta_0 + \delta_1$

is non-negligible. It suffices to prove that $\mathbf{v}_0 = \text{correct-outcome}(pk, nc, \mathbf{bb}_0, \kappa) \wedge \mathbf{v}_1 = \text{correct-outcome}(pk, nc, \mathbf{bb}_1, \kappa)$ is false with overwhelming probability.

Suppose $\mathbf{v}_0 = \text{correct-outcome}(pk, nc, \mathbf{bb}_0, \kappa)$. By definition of function *correct-outcome*, we have $\exists^{=v_0[v]} b' \in \mathbf{bb}_0 \setminus \{\perp\} : \exists r : b' = \text{Vote}(pk, v, nc, \kappa; r)$. Since $1 \leq |\mathbf{bb}_0 \cup \{b\}| \leq mb$, we have $b \neq \perp$ by Lemma 19, with overwhelming probability. Given that b is an output of $\text{Vote}(pk, v, nc, \kappa)$, $b \notin \mathbf{bb}_0$, and $\mathbf{v}^*[v] = v_0[v] + 1$, it follows that $\exists^{=v^*[v]} b' \in \mathbf{bb}_0 \cup \{b\} \setminus \{\perp\} : \exists r : b' = \text{Vote}(pk, v, nc, \kappa; r)$. Moreover, by Injectivity, b is not an output of $\text{Vote}(pk, v', nc, \kappa)$ for all $v' \in \{1, \dots, |\mathbf{v}^*| \setminus \{v\}\}$. Thus, for all $v' \in \{1, \dots, |\mathbf{v}^*| \setminus \{v\}\}$ we have $\exists^{=v^*[v']]} b' \in \mathbf{bb}_0 \cup \{b\} \setminus \{\perp\} : \exists r : b' = \text{Vote}(pk, v', nc, \kappa; r)$. Given that $\mathbf{bb}_1 = \mathbf{bb}_0 \cup \{b\}$, we have $\mathbf{v}^* = \text{correct-outcome}(pk, nc, \mathbf{bb}_1, \kappa)$. Moreover, given that $\mathbf{v}^* \neq \mathbf{v}_1$, we have $\mathbf{v}_1 \neq \text{correct-outcome}(pk, nc, \mathbf{bb}_1, \kappa)$ with overwhelming probability, which suffices to conclude our proof. \square

D Helios

Smyth, Frink & Clarkson [SFC16] formalise a generic construction for Helios-like election schemes (Definition 23), which is parameterised on the choice of homomorphic encryption scheme and sigma protocols for the relations introduced in the following definition.

Definition 22 (from [SFC16]). *Let $(\text{Gen}, \text{Enc}, \text{Dec})$ be a homomorphic asymmetric encryption scheme and Σ be a sigma protocol for a binary relation R .⁴⁰*

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

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

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

Definition 23 (Generalised Helios [SFC16]). *Suppose $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ is an additively homomorphic asymmetric encryption scheme with a message space that, for sufficiently large security parameters, includes $\{0, 1\}$, Σ_1 proves correct key construction, Σ_2 proves plaintext knowledge in a subspace, Σ_3 proves correct decryption, and \mathcal{H} is a hash function. Let $\text{FS}(\Sigma_1, \mathcal{H}) = (\text{ProveKey}, \text{VerKey})$, $\text{FS}(\Sigma_2, \mathcal{H}) = (\text{ProveCiph}, \text{VerCiph})$, and $\text{FS}(\Sigma_3, \mathcal{H}) = (\text{ProveDec}, \text{VerDec})$. We define election scheme generalised Helios, denoted $\text{Helios}(\Pi, \Sigma_1, \Sigma_2, \Sigma_3, \mathcal{H}) = (\text{Setup}, \text{Vote}, \text{Tally}, \text{Verify})$, as follows.*

⁴⁰Given 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 .

Setup(κ). Select coins s , compute $(pk, sk, \mathbf{m}) \leftarrow \text{Gen}(\kappa; s)$; $\rho \leftarrow \text{ProveKey}((\kappa, pk, \mathbf{m}), (sk, s), \kappa)$; $PK_{\mathcal{T}} \leftarrow (pk, \mathbf{m}, \rho)$; $SK_{\mathcal{T}} \leftarrow (pk, sk)$, let m be the largest integer such that $\{0, \dots, m\} \subseteq \mathbf{m}$, and output $(PK_{\mathcal{T}}, SK_{\mathcal{T}}, m, m)$.

Vote($PK_{\mathcal{T}}, v, nc, \kappa$). Parse $PK_{\mathcal{T}}$ 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} 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($SK_{\mathcal{T}}, \mathbf{bb}, nc, \kappa$). 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 $SK_{\mathcal{T}}$ as a vector (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) .

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

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

Instantiations of generalised Helios work as follows [SFC16].

- **Setup** generates the tallier’s key pair. The public key includes a non-interactive proof demonstrating that the key pair is correctly constructed.
- **Vote** takes a vote $v \in \{1, \dots, nc\}$ and outputs ciphertexts c_1, \dots, c_{nc-1} such that if $v < nc$, then ciphertext c_v contains plaintext 1 and the remaining ciphertexts contain plaintext 0, otherwise, all ciphertexts contain plaintext 0. **Vote** also outputs proofs $\sigma_1, \dots, \sigma_{nc}$ so that this can be verified. In particular, proof σ_j demonstrates ciphertext c_j contains 0 or 1, for all $1 \leq j \leq nc - 1$. And proof σ_{nc} demonstrates that the homomorphic combination of ciphertexts $c_1 \otimes \dots \otimes c_{nc-1}$ contains 0 or 1. (It follows that the voter’s ballot contains a vote for exactly one candidate.)
- **Tally** homomorphically combines ciphertexts representing votes for a particular candidate and decrypts the homomorphic combinations. The number of votes for a candidate $v \in \{1, \dots, nc - 1\}$ is simply the homomorphic combination of ciphertexts representing votes for that candidate. The number of votes for candidate nc is equal to the number of votes for all other candidates subtracted from the total number of valid ballots on the bulletin board.
- **Verify** checks that each of the above steps has been performed correctly.

The generic construction can be instantiated to derive Helios’16.

Definition 24 (Helios’16 [SFC16]). *Election scheme Helios’16 is $\text{Helios}(\Pi, \Sigma_1, \Sigma_2, \Sigma_3, \mathcal{H})$, 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. [CEGP87, 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], and \mathcal{H} is a random oracle.*

Although Helios actually uses SHA-256 [NIS12], we assume that \mathcal{H} is a random oracle to prove Theorem 12. Moreover, we assume the sigma protocols used by Helios’16 satisfy the preconditions of generalised Helios, that is, [CEGP87, Protocol 2] is a sigma protocol for proving correct key construction, [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.

E Helios satisfies ballot secrecy

The construction for Helios-like schemes produces election schemes with zero-knowledge tallying proofs (Lemma 21) that satisfy universal verifiability [SFC16] and, thus, soundness (Proposition 20). They also satisfy ballot independence (Proposition 22). Hence, they satisfy ballot secrecy too (Theorem 10). Moreover, Helios’16 satisfies ballot secrecy.

Henceforth, we assume Π , Σ_1 , Σ_2 and Σ_3 satisfy the preconditions of Definition 23, and \mathcal{H} is a random oracle. Let $\text{Helios}(\Pi, \Sigma_1, \Sigma_2, \Sigma_3, \mathcal{H}) = (\text{Setup}, \text{Vote}, \text{Tally}, \text{Verify})$ 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 21. *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 construct pf , moreover, that proof system is zero-knowledge because $(\text{ProveDec}, \text{VerDec})$ is zero-knowledge, which concludes our proof. \square

Proposition 22. *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 13, the proof systems have extractors and simulators. Let SimProveKey be the simulator for $(\text{ProveKey}, \text{VerKey})$. 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 $1/2 + \text{negl}(\kappa) < \text{IND-CVA}(\Gamma, \mathcal{A}, \kappa)$. Since \mathcal{A} is a winning adversary, we have $\mathcal{A}(PK_{\mathcal{T}}, \kappa)$ 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 $((1, 0), (0, 1))$.
- $\mathcal{B}(c)$ proceeds as follows. First, compute:

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

Secondly, select coins r_1, \dots, r_{nc-1} 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:

```

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

```

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 $PK_{\mathcal{T}} = (pk, \mathbf{m}, \rho)$. Suppose (v_0, v_1, nc) is an output of $\mathcal{A}(PK_{\mathcal{T}}, \kappa)$. Since SimProveKey is a simulator for $(\text{ProveKey}, \text{VerKey})$, we have \mathcal{B} simulates the challenger in IND-CVA to $\mathcal{A}(PK_{\mathcal{T}}, \kappa)$. In particular, $PK_{\mathcal{T}}$ is a triple containing a public key and corresponding message space generated Gen , and a (simulated) proof of correct construction. 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 23. *Adversary \mathcal{B} 's computation of b is equivalent to computing b as $b \leftarrow \text{Vote}(PK_{\mathcal{T}}, v_\beta, nc, \kappa)$.*

Proof of Claim 23. We have $PK_{\mathcal{T}}$ 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 $\text{VerKey}((\kappa, pk, \mathbf{m}), \rho, \kappa) = 1$, 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 $\text{Vote}(PK_{\mathcal{T}}, v_{\beta}, nc, \kappa)$ does not output \perp , with non-negligible probability. Indeed, computation $b \leftarrow \text{Vote}(PK_{\mathcal{T}}, v_{\beta}, nc, \kappa)$ is equivalent to the following. Select coins r_1, \dots, r_{nc-1} 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} 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,

```

 $c \leftarrow c_1 \otimes \cdots \otimes c_{nc-1};$ 
if  $v_1 = nc$  then  $m \leftarrow 1 - \beta$ ; else  $m \leftarrow 1$ ;
 $r \leftarrow r_1 \oplus \cdots \oplus r_{nc-1};$ 
 $\sigma_{nc} \leftarrow \text{ProveCiph}((pk, c, \{0, 1\}), (m, r), nc, \kappa);$ 

```

can be simulated by

```

 $c \leftarrow c_1 \otimes \cdots \otimes c_{nc-1};$ 
 $\sigma_{nc} \leftarrow \text{SimProveCiph}((pk, c, \{0, 1\}), nc, \kappa);$ 

```

Hence, we conclude the proof of this claim.

Claim 24. *Adversary \mathcal{B} 's computation of \mathbf{v} is equivalent to computing \mathbf{v} as $(\mathbf{v}, pf) \leftarrow \text{Tally}(SK_{\mathcal{T}}, \mathbf{bb}, nc, \kappa)$, where $SK_{\mathcal{T}} = (pk, sk)$.*

Proof of Claim 24. 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:

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

$$b_i[j + nc - 1] = \text{ProveCiph}((pk, b_i[j], \{0, 1\}), (m_{i,j}, r_{i,j}), j, \kappa; r_{i,j+nc-1})$$

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 $(\mathbf{v}, pf) \leftarrow \text{Tally}(SK_{\mathcal{T}}, \mathbf{bb}, nc, \kappa)$ is equivalent to initialising \mathbf{v} as a zero-filled vector of length nc and computing

```

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

```

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

$\cdots \odot m_{\ell,j}$ 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.* [CEGP87, 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 13 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 13] still hold.” We adopt the same position and assume that Theorem 13 is applicable.

Proof of Theorem 12. Helios'16 has zero-knowledge tallying proofs (Lemma 21), subject to the applicability of Theorem 13 to the sigma protocol underlying $(\text{ProveDec}, \text{VerDec})$. Moreover, since Helios'16 satisfies Exp-UV-Ext [SFC16], we have Helios'16 satisfies Soundness (Proposition 20). 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 22). Hence, Helios'16 satisfies Ballot-Secrecy too (Theorem 10). \square

References

- [Adi09] Ben Adida. Helios deployed at Princeton. <http://heliosvoting.wordpress.com/2009/10/13/helios-deployed-at-princeton/> (accessed 7 May 2014), 2009.

- [Adi14] Ben Adida. Helios v4 Verification Specs. Helios documentation, <http://documentation.heliosvoting.org/verification-specs/helios-v4> (accessed 2 May 2014), 2014. A snapshot of the specification on 18 Oct 2013 is available from <https://web.archive.org/web/20131018033747/http://documentation.heliosvoting.org/verification-specs/helios-v4..>
- [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.
- [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 *SE&P'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.
- [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 http://admin.cdn.sos.ca.gov/press-releases/prior/2007/DB07_111.pdf (accessed 1 September 2015), August 2007. A snapshot of the press release on 6 February 2008 is available from https://web.archive.org/web/20080206210142/http://www.sos.ca.gov/elections/voting_systems/ttbr/db07_042_ttbr_system_decisions_release.pdf.
- [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>. A snapshot of the chapter on 21 Mar 2015 is available from <https://web.archive.org/web/20150321170845/http://cseweb.ucsd.edu/~mihir/cse207/w-se.pdf>.
- [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.

- [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.
- [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.
- [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.
- [CEGP87] 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.
- [CF85] Josh Daniel Cohen and Michael J. Fischer. A Robust and Verifiable Cryptographically Secure Election Scheme. In *FOCS'85: 26th Symposium on Foundations of Computer Science*, pages 372–382. IEEE Computer Society, 1985.
- [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.
- [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.
- [Cha81] David L. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24:84–90, 1981.
- [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.
- [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.
- [Dam10] Ivan Damgård. On Σ -protocols, 2010. Available from <http://www.daimi.au.dk/~ivan/Sigma.pdf>.
- [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.
- [DJN10] Ivan Damgård, Mads Jurik, and Jesper Buus Nielsen. A Generalization of Paillier's Public-Key System with Applications to Electronic Voting. *International Journal of Information Security*, 9(6):371–385, 2010.
- [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.

- [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.
- [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.
- [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.
- [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.
- [Hir10] Martin Hirt. Receipt-Free K -out-of- L Voting Based on ElGamal Encryption. In David Chaum, Markus Jakobsson, Ronald L. Rivest, and Peter Y. A. Ryan, editors, *Towards Trustworthy Elections: New Directions in Electronic Voting*, volume 6000 of *LNCS*, pages 64–82. Springer, 2010.
- [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.
- [HL10] Carmit Hazay and Yehuda Lindell. Sigma protocols and efficient zero-knowledge. In *Efficient Secure Two-Party Protocols*, Information Security and Cryptography, chapter 6, pages 147–175. Springer Berlin Heidelberg, 2010.
- [HS00] Martin Hirt and Kazue Sako. Efficient Receipt-Free Voting Based on Homomorphic Encryption. In *EUROCRYPT'06: 25th International Conference on the Theory and Applications of Cryptographic Techniques*, volume 1807 of *LNCS*, pages 539–556. Springer, 2000.
- [KL07] Jonathan Katz and Yehuda Lindell. *Introduction to Modern Cryptography*. Chapman & Hall/CRC, 2007.

- [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.
- [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.
- [NIS12] NIST. Secure Hash Standard (SHS). FIPS PUB 180-4, Information Technology Laboratory, National Institute of Standards and Technology, March 2012.
- [OAS69] American Convention on Human Rights, “Pact of San Jose, Costa Rica”, 1969.
- [OSC90] Document of the Copenhagen Meeting of the Conference on the Human Dimension of the CSCE, 1990.
- [Per16] Olivier Pereira. Internet Voting with Helios. In *Real-World Electronic Voting: Design, Analysis and Deployment*, volume 8604, chapter 11. CRC Press, 2016.
- [QS16] Elizabeth A. Quaglia and Ben Smyth. Constructing secret, verifiable auction schemes from election schemes. Cryptology ePrint Archive, Report 2015/1204 (version 20160524:130412), 2016.
- [SB13a] 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.
- [SB13b] Ben Smyth and David Bernhard. Ballot secrecy and ballot independence coincide. Cryptology ePrint Archive, Report 2013/235 (version 20130618:102144), 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. <http://hal.inria.fr/inria-00599182/>.

- [Sch05] Nicole Schweikardt. Arithmetic, first-order logic, and counting quantifiers. *ACM Transactions on Computational Logic*, 6(3):634–671, July 2005.
- [SFC16] 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 20160707:162850), 2016.
- [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.
- [SHM15] Ben Smyth, Yoshikazu Hanatani, and Hirofumi Muratani. NM-CPA secure encryption with proofs of plaintext knowledge. In *IWSEC'15: 10th International Workshop on Security*, volume 9241 of *LNCS*, pages 115–134. Springer, 2015.
- [SK94] Kazue Sako and Joe Kilian. Secure Voting Using Partially Compatible Homomorphisms. In *CRYPTO'94: 14th International Cryptology Conference*, volume 839 of *LNCS*, pages 411–424. Springer, 1994.
- [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.
- [SM16] Ben Smyth and Maxime Meyer. An attack against the helios election system that violates eligibility. arXiv, Report 1612.04099, 2016.
- [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.
- [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.
- [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.
- [UN48] Universal Declaration of Human Rights, 1948.
- [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.