Voting: You Can't Have Privacy without Individual Verifiability

Véronique Cortier, JosephALLEMAND

► To cite this version:

Véronique Cortier, JosephALLEMAND. Voting: You Can't Have Privacy without Individual Verifiability. ACM CCS 2018 - 25th ACM Conference on Computer and Communications Security, Oct 2018, Toronto, Canada. 10.1145/3243734.3243762 . hal-01900086

HAL Id: hal-01900086

<https://hal.inria.fr/hal-01900086>

Submitted on 20 Oct 2018

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

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Voting: You Can't Have Privacy without Individual Verifiability

Véronique Cortier

CNRS, Loria

Nancy, France

veronique.cortier@loria.fr

Joseph Lallemand

Inria, Loria

Nancy, France

joseph.lallemand@loria.fr

ABSTRACT

Electronic voting typically aims at two main security goals: vote privacy and verifiability. These two goals are often seen as antagonistic and some national agencies even impose a hierarchy between them: first privacy, and then verifiability as an additional feature. Verifiability typically includes individual verifiability (a voter can check that her ballot is counted); universal verifiability (anyone can check that the result corresponds to the published ballots); and eligibility verifiability (only legitimate voters may vote).

We show that actually, privacy implies individual verifiability. In other words, systems without individual verifiability cannot achieve privacy (under the same trust assumptions). To demonstrate the generality of our result, we show this implication in two different settings, namely cryptographic and symbolic models, for standard notions of privacy and individual verifiability. Our findings also highlight limitations in existing privacy definitions in cryptographic settings.

CCS CONCEPTS

• **Security and privacy** → **Mathematical foundations of cryptography**; **Formal security models**; **Logic and verification**;

KEYWORDS

e-voting; privacy; verifiability; provable cryptography; symbolic verification

ACM Reference Format:

Véronique Cortier and Joseph Lallemand. 2018. Voting: You Can't Have Privacy without Individual Verifiability. In *2018 ACM SIGSAC Conference on Computer and Communications Security (CCS '18)*, October 15–19, 2018, Toronto, ON, Canada. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3243734.3243762>

1 INTRODUCTION

Electronic voting is often seen as a convenient way for running elections as it allows voters to vote from any place. Moreover, it eases the tally and it can therefore often be used for non trivial counting procedures such as Single Transferable Vote or Condorcet. Numerous voting systems have been proposed so far, like Helios [4],

Belenios [15], Civitas [14], Prêt-à-voter [29], or the protocols deployed in Estonia [23] or in Australia [11] to cite a few. On the other hand, many weaknesses or even attacks have been unveiled [30, 31], from voting machines [22] to Internet voting [32].

In order to carefully analyse voting systems, security requirements have been defined. The two main security properties are:

- privacy: no one should know how I voted;
- verifiability is typically described through the three following sub-properties.
 - *individual verifiability*: a voter can check that her ballot is counted;
 - *universal verifiability*: anyone can check that the results corresponds to the published ballots;
 - *eligibility verifiability*: only legitimate voters may vote.

These two main properties seem antagonistic and an impossibility result has even been established between verifiability and unconditional privacy [13], that is, a notion of privacy that is independent of the power of the attacker.

The main contribution of this paper is to establish that, in fact, (computational) *privacy implies individual verifiability*, that is, guarantees that all the honest votes will be counted. This result holds for arbitrary primitives and voting protocols without anonymous channels. To show that this implication is not due to a choice of a very particular definition, we prove this implication in two very distinct contexts, namely symbolic and cryptographic models. In symbolic models, messages are represented by terms and the attacker's behaviour is typically axiomatised through a set of logical formulas or rewrite rules. Cryptographic models are more precise. They represent messages as bitstrings and consider attackers that can be any probabilistic polynomial time Turing machines. Proofs of security are made by reduction to well accepted security assumptions such as hardness of factorisation or discrete logarithm. In both models, we consider a standard notion of privacy, already used to analyse several protocols. In both cases, we establish that privacy implies individual verifiability for a (standard) basic notion of individual verifiability, namely that the result of the election must contain the votes of all honest voters.

We now describe the main idea of the result. Actually, we show the contrapositive implication: if there is an attack against individual verifiability, then there is an attack against privacy. To explain the idea, let's consider a very simple protocol, not at all verifiable. In this simple protocol, voters simply encrypt their votes with the public key of the election. The ballot box stores the ballots and, at the end of the election, it provides the list of recorded ballots to the talliers, who detain the private key, possibly split in shares. The talliers compute and publish the result of the election. The ballot box is not public and no proof of correct decryption is provided so voters have no control over the correctness of the result. Such

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CCS '18, October 15–19, 2018, Toronto, ON, Canada

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-5693-0/18/10...\$15.00

<https://doi.org/10.1145/3243734.3243762>

a system is of course not satisfactory but it is often viewed as a “basic” system that can be used in contexts where only privacy is a concern. Indeed, it is typically believed that such a system guarantees privacy provided that the attacker does not have access to the private key of the election. In particular, the ballot box (that is, the voting server) seems powerless. This is actually *not* the case. If the ballot box aims at knowing how a particular voter, say Alice, voted, he may simply keep Alice’s ballot in the list of recorded ballots and then replace all the other ballots by encryptions of valid votes of his choice, possibly following a plausible distribution, to make the attack undetected. When the result of the election is published, the ballot box will know all the votes but Alice’s vote, and will therefore be able to deduce how Alice voted.

One may argue that such an attack is not realistic: the ballot box needs to be able to change *all* ballots but one. Note however that elections are often split in many small voting stations (sometimes as small as 20 voters in total [17]). Therefore changing a few ballots can be sufficient to learn how Alice voted. Maybe more importantly, this attack highlights the fact that it is not possible to require privacy without verifiability as sometimes specified by national agencies. For example, in France, only privacy is required [1]. In Switzerland, privacy is a pre-requisite and the level of verifiability depend on the percentage of voters that can vote electronically [2]. Our findings point out that if voters cannot trust some authorities *w.r.t.* the fact that their vote will be counted they cannot trust the same authorities *w.r.t.* their privacy, even for entities that do not have access to the secret keys. Beyond the attack explained on a simple (and naive) protocol, our proof that privacy implies individual verifiability shows that as soon as a protocol is not verifiable, then the adversary can take advantage of the fact that he may modify a vote without being detected in order to break privacy. Individual verifiability is only one part of verifiability. It does not account for universal nor eligibility verifiability. So our result cannot be used to conclude that a private voting scheme ensures all desirable verifiability properties. Instead, it demonstrates that there is no hope to design a private voting system if it does not include some degree of verifiability, namely individual verifiability at least.

Our results also emphasise issues in existing privacy definitions. Indeed, if privacy implies individual verifiability, how is it possible to prove Helios [8] or Civitas [5] without even modelling the verification aspects? How can a system that is not fully verifiable like the Neuchâtel protocol be proved private [21]? As already pointed out in [9], existing cryptographic definitions of privacy (see [7] for a survey) implicitly assume an honest voting ballot box: honest ballots are assumed to be properly stored and then tallied. Actually, we notice that the same situation occurs in symbolic models. Although the well adopted definition of privacy [20] does not specify how the ballot box should be modelled, most symbolic proofs of privacy (see e.g. [5, 17, 18, 20]) actually assume that the votes of honest voters always reach the ballot box without being modified and that they are properly tallied. The reason is that the authors were aware of the fact that if the adversary may block all ballots but Alice’s ballot, he can obviously break privacy. However, to avoid this apparently systematic attack, they make a very strong assumption: the ballot box needs to be honest. This means that previous cryptographic and symbolic privacy analyses only hold assuming an honest ballot box while the corresponding voting systems aim at privacy *without*

trusting the ballot box. This seriously weakens the security analysis and attacks may be missed, like the attack of P. Roenne [28] on Helios, for which there is no easy fix.

Why is it so hard to define vote privacy *w.r.t.* a dishonest ballot box? Intuitively, vote privacy tries to capture the idea that, no matter how voters vote, the attacker should not be able to see any difference. The key issue is that the result of the election *does* leak some information (typically the sum of the votes) and the adversary may notice a difference based on this. This particularity makes vote privacy differ from privacy in other contexts, where the adversary really should learn no information. Therefore, most definitions of vote privacy (roughly) say that, no matter how honest voters voted, provided that the aggregation of the corresponding votes remains the same, then the attacker should not see any difference. However, as soon as the ballot box is dishonest, it may discard some honest ballots and break privacy, as already discussed. The first definition of privacy *w.r.t.* a dishonest ballot box [9] weakens privacy by requiring that among the ballots that are ready to be tallied, the (sub-)tally of the honest ones does not change. This preliminary definition has two limitations. First, it assumes that the tallied ballots are exactly the same as the cast ones, which is not the case of all protocols (e.g. in ThreeBallots [27], only a part of the ballot is published; in BeleniosRF [12], ballots are re-randomised). Second, it does not model re-voting: the tally process cannot discard ballots due to some revote policy.

We propose here another approach. Instead of changing the privacy definition, we now include a model of the verification process: the ballots should be tallied only if the honest voters have successfully performed the tests specified by the protocol. We compare our definition with [9] and an original definition of privacy [6] on a selection of well-studied protocols, that have different levels of verifiability (Helios, Civitas, Belenios, Neuchâtel, and our simple - non verifiable - protocol). We show again that our notion of privacy, *w.r.t.* a dishonest ballot box, implies individual verifiability. We do not consider our new definition of privacy as final but it opens the way to a better understanding of privacy in the context of fully dishonest authorities.

Threat model. We show that privacy implies individual verifiability, *under the same trust assumptions*, that is, trusting the same group of authorities, channels, etc. In symbolic models, the privacy definition does not make prior assumptions on the threat model. Instead, the encoding of the protocol defines which parties are trusted. In particular, as already discussed, existing proofs of privacy [5, 17, 18, 20] often implicitly assume that honest ballots reach the ballot box without any modification. We show that whenever privacy holds then individual verifiability holds, for the same encoding, hence the same assumptions. In contrast, most cryptographic definitions of privacy implicitly assume an honest ballot box. Therefore, we first show that privacy implies individual verifiability, assuming an honest ballot box, considering the standard definition of privacy by Benaloh [6]. Then we show that privacy still implies individual verifiability, assuming a dishonest ballot box, considering our novel definition of privacy, that explicitly models the verification steps.

Related work. As already mentioned, [13] shows an impossibility result between universal verifiability and unconditional privacy.

We show in contrast that the commonly used (computational) definitions of privacy actually imply verifiability. The discrepancy between the two results comes from the fact that [13] considers unconditional privacy while most protocols achieve only computational privacy, that is against a polynomially bounded adversary. Interestingly, the impossibility result still holds between unconditional privacy and our notion of individual verifiability. [19] establishes a hierarchy between privacy, receipt-freeness, and coercion resistance, while in a quantitative setting, [26] shows that this hierarchy does not hold anymore. [16] recasts several definitions of verifiability in a common setting, providing a framework to compare them. Besides [13], none of these approaches relates privacy with verifiability. Many privacy definitions have been proposed as surveyed in [7]. However, they all assume an honest ballot box. To our knowledge, [9] is the only exception, as already discussed in details. [17] shows how to break privacy by replaying a ballot. If an attacker may replay Alice's ballot and cast it in his own name (or cast a related ballot), then he introduces a bias in the result, that leaks some information on Alice's vote. Note that this replay attack does not break individual verifiability: honest votes are correctly counted. We show here another breach for privacy: if an attacker may remove some honest votes, then he breaks privacy as well.

Roadmap. We first prove that privacy implies individual verifiability in symbolic models, in Section 3, and then in cryptographic models, in Section 4. These two parts are rather independent. In Section 6, we examine a selection of well-studied voting protocols and compare the effect of different (cryptographic) notions of privacy when the ballot box is dishonest.

The technical details and proofs omitted due to space constraints are available in the companion technical report [?].

2 PRELIMINARIES

Notations: The multiset of elements a, a, b, c is denoted $\llbracket a, a, b, c \rrbracket$. The union of two multisets S_1 and S_2 is denoted $S_1 \uplus S_2$.

In both cryptographic and symbolic models, we assume a set $\mathcal{V}$ of votes and a set $\mathcal{R}$ of possible results, equipped with an associative and commutative operator $*$ (e.g. addition of vectors). A *counting function* is a function ρ that associates a result $r \in \mathcal{R}$ to a multiset of votes. We assume that counting functions have a *partial tally* property: it is always possible to count the votes in two distinct multisets and then combine the results.

$$\forall V, V' \rho(V \uplus V') = \rho(V) * \rho(V')$$

A vote v is said to be *neutral* if $\rho(v)$ is neutral w.r.t. $*$.

Example 2.1. Consider a finite set of candidates $C = \{a_1, \dots, a_k\}$. In case voters should select between k_1 and k_2 candidates or vote blank, we can represent valid votes by vectors representing the selection of candidates

$$\mathcal{V}_{k_1, k_2} = \left\{ v \in \{0, 1\}^k \mid k_1 \leq \sum_{i=0}^k v_i \leq k_2 \right\} \cup \{v^{\text{blank}}\}$$

where v^{blank} is the null vector $(0, \dots, 0)$, representing a blank vote. In a mixnet-based tally, all the individual votes are revealed. Thus $\mathcal{R}$ is the set of multisets of votes in $\mathcal{V}_{k_1, k_2}$ and $*$ is the union of multisets. The corresponding counting function is $\rho_{\text{mix}}(V) = V$, where V is a multiset of elements of $\mathcal{V}_{k_1, k_2}$.

In an homomorphic-based tally, the votes are added together. Thus $\mathcal{R} = \mathbb{N}^k$, the set of vectors of k elements, and $*$ is the addition of vectors. The corresponding counting function is $\rho_{\text{hom}}(V) = \sum_{v \in V} v$.

Both ρ_{mix} and ρ_{hom} have the partial tally property. The vote v^{blank} is a neutral vote w.r.t. ρ_{hom} but not ρ_{mix} .

The result of the election r may have several representations. For example, a multiset may be represented by several lists (where the order changes). In symbolic models, the result will be represented by abstract terms and we wish our result to be independent of a particular choice of representation. Therefore, we will simply say that a *representation* R is a function that associates to a result $r \in \mathcal{R}$ a set of possible representations with an injectivity property:

$$\forall r \neq r'. R(r) \cap R(r') = \emptyset$$

Intuitively, a result can be associated to several representations but a given representation can correspond to at most one result.

For our proofs in a cryptographic setting, we will also assume that given an election result r and a set of votes V , one can decide efficiently (in polynomial time) whether r includes all the votes of V , that is, whether there exists V' such that $r = \rho(V \uplus V')$. This condition is satisfied by ρ_{mix} and ρ_{hom} and all standard counting functions.

3 SYMBOLIC MODEL

3.1 Model

In symbolic models, security protocols are often modelled through a process algebra, in the spirit of the applied pi-calculus [3], that offers a small, abstract language for specifying communications, where messages are represented as terms. We present here a calculus inspired from the calculus underlying the ProVerif tool [10].

3.1.1 Terms. We consider an infinite set of names $\mathcal{N}$ that model fresh values such as nonces and keys. We distinguish the set $\mathcal{FN}$ of free nonces (generated by the attacker) and the set $\mathcal{BN}$ of bound nonces (generated by the protocol agents). We also assume an infinite set of variables $\mathcal{V} = \mathcal{X} \uplus \mathcal{AX}$ where $\mathcal{X}$ contains variables used in processes (agent's memory) while $\mathcal{AX}$ contains variables used to store messages (adversary's memory). Cryptographic primitives are represented through a set of function symbols, called *signature* $\mathcal{F}$. Each function symbol has an arity, that is, the number of its arguments. We assume an infinite set $\mathcal{C} \subseteq \mathcal{F}$ of public constants, which are functions of arity 0.

Example 3.1. The standard primitives, public keys, symmetric and asymmetric encryption, concatenation, as well as addition, can be modelled by the following signature.

$$\mathcal{F}_c = \{\text{pk}/1, \text{enc}/2, \text{aenc}/2, \langle \cdot, \cdot \rangle / 2, +/2\}$$

The companion primitives (symmetric and asymmetric decryption, projections) are then represented by the following signature:

$$\mathcal{F}_d = \{\text{dec}/2, \text{adec}/2, \pi_1/1, \pi_2/1\}$$

Given a signature $\mathcal{F}$, a set of names $\mathcal{N}$, a set of variables $\mathcal{V}$, the set of *terms* $\mathcal{T}(\mathcal{F}, \mathcal{V}, \mathcal{N})$ is the set inductively defined by applying functions to variables in $\mathcal{V}$ and names in $\mathcal{N}$. The set of names resp. variables occurring in t is denoted $\text{names}(t)$ (resp. $\text{vars}(t)$). A

Processes:

$$\begin{array}{ll}
P, Q & ::= \\
& \emptyset \\
& | \nu n.P \quad \text{for } n \in \mathcal{BN} \text{ (} n \text{ bound in } P \text{)} \\
& | \text{out}(c, M).P \\
& | \text{in}(c, x).P \quad \text{for } x \in \mathcal{X} \text{ (} x \text{ bound in } P \text{)} \\
& | \text{event}(M_1, \dots, M_n).P \quad \text{for event} \in \text{Ev of arity } n \\
& | P \mid Q \\
& | \text{let } x = M \text{ in } P \quad \text{for } x \in \mathcal{X} \text{ (} x \text{ bound in } P \text{)} \\
& | \text{if } M = N \text{ then } P \text{ else } Q \\
& | !P
\end{array}$$

where $M, N, M_1, \dots, M_n$ are messages and $c \in \text{Ch}$ is a channel.

Figure 1: Syntax for processes.

term is *ground* if it does not contain any variable. The set of terms $\mathcal{T}(\mathcal{F}, \mathcal{AX}, \mathcal{FN})$ represents the *attacker terms*, that is, terms built from the messages sent on the network and stored thanks to the variables in $\mathcal{AX}$.

A *substitution* $\sigma = \{M_1/x_1, \dots, M_k/x_k\}$ maps variables $x_1, \dots, x_k \in \mathcal{V}$ to messages $M_1, \dots, M_k$. Its *domain* is denoted $\text{dom}(\sigma) = \{x_1, \dots, x_k\}$. The application of σ to a term t is denoted $t\sigma$ and is defined as usual. A substitution σ is *ground* if its messages $M_1, \dots, M_k$ are ground.

The properties of the cryptographic primitives are modelled through an *equational theory* E , which is a finite set of equations of the form $M = N$ where $M, N \in \mathcal{T}(\mathcal{F}, \mathcal{X}, \emptyset)$ are messages without names. Equality modulo E , denoted by $=_E$, is defined as the smallest equivalence relation on terms that is closed under context and substitution. We denote disequalities modulo E by $M \neq_E N$.

Example 3.2. Considering the signature $\mathcal{F}_c \cup \mathcal{F}_d \cup \mathcal{C}$ from Example 3.1, the following equational theory describes the ability to decrypt symmetrically, asymmetrically, and to project pairs. It also characterises $+$ as an associative and commutative operator.

$$\begin{aligned}
\text{dec}(\text{enc}(x, y), y) &= x \\
\text{adec}(\text{aenc}(x, \text{pk}(y)), y) &= x \\
\pi_1(\langle x, y \rangle) &= x \\
\pi_2(\langle x, y \rangle) &= y \\
x + (y + z) &= (x + y) + z \\
x + y &= y + x
\end{aligned}$$

3.1.2 Processes. The behaviour of protocol parties is described through processes. Let Ch be an infinite set of channel names, representing the channels on which the messages are exchanged. All channels will be public. We consider different channels nevertheless to model the fact that an attacker can identify the provenance of a message. We also consider a finite set Ev of event symbols, given together with their arity. Events are used to record that participants have reached a certain step, with some associated knowledge. Protocols are modelled through a process algebra, whose syntax is displayed in Figure 1.

As usual, we identify processes up to α -renaming, to avoid capture of bound names and variables.

A *configuration* of the system is a triple $(E; \mathcal{P}; \phi)$ where:

- $\mathcal{P}$ is a multiset of processes that represents the current active processes;
- $\mathcal{E}$ is a set of names, which represents the private names of the processes;
- ϕ is a substitution with $\text{dom}(\phi) \subseteq \mathcal{AX}$ that represents the messages sent on the network. We assume ϕ to be ground, that is for any $x \in \text{dom}(\phi)$, $\phi(x)$ is a ground term.

The semantics of processes is given through a transition relation $\xrightarrow{\alpha}$ provided in Figure 2, where α is the *action* associated to the transition. τ denotes a silent action. Events are recorded but will be invisible to the attacker. Intuitively, process $\nu n.P$ creates a fresh nonce, stored in $\mathcal{E}$, and behaves like P . Process $\text{out}(c, M).P$ emits M on c and behaves like P . Process $\text{in}(c, x).P$ inputs a term computed by the attacker (that is a term built from ϕ using an attacker term) on channel c and then behaves like P . Process $\text{event}(M_1, \dots, M_n).P$ triggers the event $\text{event}(M_1, \dots, M_n)$, and then behaves like P . Process $P \mid Q$ corresponds to the parallel composition of P and Q . Process $\text{let } x = M \text{ in } P$ behaves like P in which x is replaced with M . Process $\text{if } M = N \text{ then } P \text{ else } Q$ behaves like P if M and N are equal modulo E , and behaves like Q otherwise. The replicated process $!P$ behaves as an unbounded number of copies of P .

We denote by $\xrightarrow{w}_*$ the reflexive transitive closure of $\xrightarrow{\alpha}$, where w is the concatenation of all actions. We also write equality up to silent actions and events $=_\tau$.

A *trace* of a process P is any possible sequence of transitions starting from P . Traces correspond to all possible executions in the presence of an attacker that may read, forge, and send messages. Formally, the set of traces $\text{trace}(P)$ is defined as follows.

$$\text{trace}(P) = \{(w, \text{new } \mathcal{E}. \phi) | (\emptyset; \{P\}; \emptyset) \xrightarrow{w}_* (\mathcal{E}; \mathcal{P}; \phi)\}$$

A sequence of actions t is *blocking* in a process P if it cannot be executed.

$$\text{blocking}(t, P) \stackrel{\text{def}}{=} \forall \phi. (t, \phi) \notin \text{trace}(P).$$

Example 3.3. Helios [4] is a simple voting protocol used in several elections, like the election of the rector of the university of Louvain-la-Neuve. A voter simply encrypts her vote with the public key of the election. This encrypted vote forms the ballot, which is sent to the ballot box. The voter may check that her ballot is on the ballot box since the ballot box is public. There are two ways for tallying, either homomorphic tally or mixnet-based tally. We model here the two options in an abstract way: given the ballots, the talliers output the aggregation of the decryption of the ballot. This aggregation could be the addition or just the votes in a random order. For simplicity, we describe here a simple version with only two honest voters A and B , a dishonest voter C , and a voting server S . This protocol can be modelled by the following process.

$$\begin{aligned}
P_{\text{Helios}}(v_a, v_b) = & \nu k_{as}, k_{bs}, k_{cs}, k_e. \\
& (\text{out}(c, k_{cs}).\text{out}(c, \text{pk}(k_e)) \mid \\
& \text{Voter}(A, v_a, c_a, c'_a, k_{as}, k_e) \mid \text{Voter}(B, v_b, c_b, c'_b, k_{bs}, k_e) \mid \\
& \text{Tally}_{\text{Helios}}(c_a, c_b, c_c, c_s, k_{as}, k_{bs}, k_{cs}, k_e))
\end{aligned}$$

where $\text{Voter}(a, v, c, c', k, k_e)$ represents voter a willing to vote for v using the channels c and c' , the election key k_e and the credential k to authenticate to the server, while $\text{Tally}_{\text{Helios}}$ represents the voting server.

$(\mathcal{E}; \{P_1 \mid P_2\} \cup \mathcal{P}; \phi)$	$\xrightarrow{\tau}$	$(\mathcal{E}; \{P_1, P_2\} \cup \mathcal{P}; \phi)$	PAR
$(\mathcal{E}; \{\emptyset\} \cup \mathcal{P}; \phi)$	$\xrightarrow{\tau}$	$(\mathcal{E}; \mathcal{P}; \phi)$	ZERO
$(\mathcal{E}; \{v n.P\} \cup \mathcal{P}; \phi)$	$\xrightarrow{\tau}$	$(\mathcal{E} \cup \{n\}; \{P\} \cup \mathcal{P}; \phi)$	NEW
$(\mathcal{E}; \{\text{out}(c, M).P\} \cup \mathcal{P}; \phi)$	$\xrightarrow{v ax_n.\text{out}(c, ax_n)}$	$(\mathcal{E}; \{P\} \cup \mathcal{P}; \phi \cup \{M/ax_n\})$	OUT
if M is a ground term, $ax_n \in \mathcal{AX}$ and $n = \phi + 1$			
$(\mathcal{E}; \{\text{in}(c, x).P\} \cup \mathcal{P}; \phi)$	$\xrightarrow{\text{in}(c, R)}$	$(\mathcal{E}; \{P[R\phi/x]\} \cup \mathcal{P}; \phi)$	IN
if R is an attacker term such that $\text{vars}(R) \subseteq \text{dom}(\phi)$			
$(\mathcal{E}; \{\text{event}(M_1, \dots, M_n).P\} \cup \mathcal{P}; \phi)$	$\xrightarrow{\text{event}(M_1, \dots, M_n)}$	$(\mathcal{E}; \{P\} \cup \mathcal{P}; \phi)$	EVENT
if $\forall i. M_i$ is a ground message			
$(\mathcal{E}; \{\text{let } x = M \text{ in } P\} \cup \mathcal{P}; \phi)$	$\xrightarrow{\tau}$	$(\mathcal{E}; \{P[M/x]\} \cup \mathcal{P}; \phi)$	LET-IN
if M is ground			
$(\mathcal{E}; \{\text{if } M = N \text{ then } P \text{ else } Q\} \cup \mathcal{P}; \phi)$	$\xrightarrow{\tau}$	$(\mathcal{E}; \{P\} \cup \mathcal{P}; \phi)$	IF-THEN
if M, N are ground messages such that $M =_E N$			
$(\mathcal{E}; \{\text{if } M = N \text{ then } P \text{ else } Q\} \cup \mathcal{P}; \phi)$	$\xrightarrow{\tau}$	$(\mathcal{E}; \{Q\} \cup \mathcal{P}; \phi)$	IF-ELSE
if M, N are ground messages such that $M \neq_E N$			
$(\mathcal{E}; \{!P\} \cup \mathcal{P}; \phi)$	$\xrightarrow{\tau}$	$(\mathcal{E}; \{P, !P\} \cup \mathcal{P}; \phi)$	REPL

Figure 2: Semantics

$\text{Voter}(a, v, c, c', k, k_e)$ simply sends an encrypted vote. To model the fact that voters communicate with the ballot box through an authenticated channel, we assume that a voter first sends her ballot privately to the server (using the encryption with k) and then sends the ballot on a public channel. Note that the key k is just a modelling artefact to abstract away the underlying password-based authenticated channel.

$$\begin{aligned} \text{Voter}(a, v, c, c', k, k_e) = \\ v r. \text{out}(c, \text{enc}(\text{aenc}(\langle v, r \rangle, \text{pk}(k_e)), k)). \text{Voted}(a, v). \\ \text{out}(c', \text{aenc}(\langle v, r \rangle, \text{pk}(k_e))) \end{aligned}$$

The voting server receives ballots from voters A, B , and C and then outputs the decrypted ballots, after some mixing, modelled through the $+$ operator.

$$\begin{aligned} \text{Tally}_{\text{Helios}}(c_a, c_b, c_c, c_s, k_{as}, k_{bs}, k_{cs}, k_e) = \\ \text{in}(c_a, x_1). \text{in}(c_b, x_2). \text{in}(c_c, x_3). \\ \text{let } y_1 = \text{dec}(x_1, k_{as}) \text{ in} \\ \text{let } y_2 = \text{dec}(x_2, k_{bs}) \text{ in} \\ \text{let } y_3 = \text{dec}(x_3, k_{cs}) \text{ in} \\ \text{if } x_1 \neq x_2 \wedge x_1 \neq x_3 \wedge x_2 \neq x_3 \text{ then} \\ \text{out}(c_s, \pi_1(\text{adec}(y_1, k_e)) + \pi_1(\text{adec}(y_2, k_e)) \\ + \pi_1(\text{adec}(y_3, k_e))) \end{aligned}$$

where we omit the null else-branches. $\wedge$ is syntactic sugar for a succession of tests and if $M \neq N$ then P is syntactic sugar for if $M = N$ then $\emptyset$ else P .

3.1.3 Equivalence. Sent messages are stored in a substitution ϕ while private names are stored in $\mathcal{E}$. A *frame* is simply an expression of the form $\text{new } \mathcal{E}.\phi$ where $\text{dom}(\phi) \subseteq \mathcal{AX}$. It represents the knowledge of an attacker. We define $\text{dom}(\text{new } \mathcal{E}.\phi)$ as $\text{dom}(\phi)$.

Intuitively, two sequences of messages are indistinguishable to an attacker if he cannot perform any test that could distinguish them. This is typically modelled as static equivalence [3].

Definition 3.4 (Static Equivalence). Two ground frames $\text{new } \mathcal{E}.\phi$ and $\text{new } \mathcal{E}'.\phi'$ are statically equivalent if and only if they have the same domain, and for all attacker terms R, S with variables in $\text{dom}(\phi) = \text{dom}(\phi')$, we have

$$(R\phi =_E S\phi) \iff (R\phi' =_E S\phi')$$

Two processes P and Q are in equivalence if no matter how the adversary interacts with P , a similar interaction may happen with Q , with equivalent resulting frames.

Definition 3.5 (Trace Equivalence). Let P, Q be two processes. We write $P \sqsubseteq_t Q$ if for all $(s, \psi) \in \text{trace}(P)$, there exists $(s', \psi') \in \text{trace}(Q)$ such that $s =_\tau s'$ and ψ and ψ' are statically equivalent. We say that P and Q are trace equivalent, and we write $P \approx_t Q$, if $P \sqsubseteq_t Q$ and $Q \sqsubseteq_t P$.

Note that this definition already includes the attacker's behaviour, since processes may input any message forged by the attacker.

Example 3.6. Ballot privacy is typically modelled as an equivalence property [20] that requires that an attacker cannot distinguish when Alice is voting 0 and Bob is voting 1 from the scenario where the two votes are swapped.

Continuing Example 3.3, ballot privacy of Helios can be expressed as follows:

$$P_{\text{Helios}}(0, 1) \approx_t P_{\text{Helios}}(1, 0)$$

3.2 Voting protocols

We consider two disjoint, infinite subsets of $\mathcal{C}$: a set $\mathcal{A}$ of *agent names* or *identities*, and a set $\mathcal{V}$ of *votes*. We assume given a representation R of the result.

A voting protocol is modelled as a process. It is composed of:

- processes that represent honest voters;
- a process modelling the tally;
- possibly other processes, modelling other authorities.

Formally, we define a *voting process* as follows.

Definition 3.7. A *voting process* is a process of the form

$$P = v \overrightarrow{cred}.v \overrightarrow{cred}_1 \dots v \overrightarrow{cred}_p. (\\ \text{Voter}(a_1, v_{a_1}, \overrightarrow{c}_1, \overrightarrow{cred}, \overrightarrow{cred}_1) \mid \dots \mid \\ \text{Voter}(a_n, v_{a_n}, \overrightarrow{c}_n, \overrightarrow{cred}, \overrightarrow{cred}_n) \\ \mid \text{Tally}_p(\overrightarrow{c}, \overrightarrow{cred}, \overrightarrow{cred}_1, \dots, \overrightarrow{cred}_p) \\ \mid \text{Others}_p(\overrightarrow{c}', \overrightarrow{cred}, \overrightarrow{cred}_1, \dots, \overrightarrow{cred}_p))$$

where $a_i \in \mathcal{A}$, $v_{a_i} \in \mathcal{V}$, $\overrightarrow{c}_i$, $\overrightarrow{c}$, $\overrightarrow{c}'$ are (distinct) channels, $\overrightarrow{cred}$ and $\overrightarrow{cred}_i$ are (distinct) names.

A voting process may be instantiated by various voters and vote selections. Given $A = \{b_1, \dots, b_n\} \subseteq \mathcal{A}$ a finite set of voters, and $\alpha : A \rightarrow \mathcal{V}$ that associates a vote to each voter, we define P_α by replacing a_i by b_i and v_i by $\alpha(b_i)$ in P .

Moreover, P must satisfy the following properties.

- Process $\text{Voter}(a, v_a, \overrightarrow{c}, \overrightarrow{cred}, \overrightarrow{cred})$ models an honest voter a willing to vote for v_a , using the channels $\overrightarrow{c}$, credentials $\overrightarrow{cred}$ (e.g. a signing key) and election credentials $\overrightarrow{cred}$. It is assumed to contain an event $\text{Voted}(a, v)$ that models that a has voted for v . This event is typically placed at the end of process $\text{Voter}(a, v_a, \overrightarrow{c}, \overrightarrow{cred}, \overrightarrow{cred})$. This event cannot appear in process Tally_p nor Others_p .
- Process $\text{Tally}_p(\overrightarrow{c}, \overrightarrow{cred}, \overrightarrow{cred}_1, \dots, \overrightarrow{cred}_p)$ models the tally. It is parametrised by the total number of voters p (honest and dishonest), with $p \geq n$. It is assumed to contain exactly one output action on a reserved channel c_r . The term output on this channel is assumed to represent the final result of the election.

$$\forall \alpha. \forall (tr, \phi) \in \text{trace}(P_\alpha). \text{out}(c_r, r) \in tr \Rightarrow \exists V. \phi(r) \in R(\rho(V))$$

Tally_p may of course contain input/output actions on other channels.

- Process $\text{Others}_p(\overrightarrow{c}', \overrightarrow{cred}, \overrightarrow{cred}_1, \dots, \overrightarrow{cred}_p)$ is an arbitrary process, also parametrised by p . It models the remaining of the voting protocol, for example the behaviour of other authorities. It also models the initial knowledge of the attacker by sending appropriate data (e.g. the public key of the election or dishonest credentials). We simply assume that it uses a set of channels disjoint from the channels used in Voter and Tally_p .

The channel c_r used in Tally_p to publish the result is called the *result channel* of P .

Example 3.8. The process modelling the Helios protocol, as defined in Example 3.3 is a voting process, where process Others_p consists in the output of the keys: $\text{out}(c, k_{cs}).\text{out}(c, \text{pk}(k_e))$.

We can read which voters voted from a trace. Formally, given a sequence tr of actions, the set of voters $\text{Voters}(tr)$ who did vote in tr is defined as follows.

$$\text{Voters}(tr) = \{a \in \mathcal{A} \mid \exists v \in \mathcal{V}. \text{Voted}(a, v) \in tr\}.$$

The result of the election is emitted on a special channel c_r . It should correspond to the tally of a multiset of votes. Formally, given a trace (t, ϕ) and a multiset of votes V , the predicate $\text{result}(t, \phi, V)$

holds if the election result in (t, ϕ) corresponds to V .

$$\text{result}(t, \phi, V) \stackrel{\text{def}}{=} \exists x, t'. t = t'.\text{out}(c_r, x) \wedge \phi(x) \in R(\rho(V)).$$

3.3 Security properties

Several definitions of verifiability have been proposed. In the lines of [15, 25], we consider a very basic notion, that says that the result should at least contain the votes from honest voters.

Definition 3.9 (symbolic individual verifiability). Let P be a voting process with result channel c_r . P satisfies *symbolic individual verifiability* if, for any trace $(t, \phi) \in \text{trace}(P_\alpha)$ of the form $t'.\text{out}(c_r, x)$, there exists V_c such that the result in t corresponds to $V_a \uplus V_c$, that is $\text{result}(t, \phi, V_a \uplus V_c)$, where

$$V_a = \{v \mid \exists a. \text{Voted}(a, v) \in t\}$$

Individual verifiability typically guarantees that voters can check that their ballot will be counted. Our notion of individual verifiability goes one step further, ensuring that the corresponding votes will appear in the result, even if the tally is dishonest. One of the first definitions of verifiability was given in [24], distinguishing between individual, universal, and eligibility verifiability. Intuitively, our own notion of individual verifiability sits somewhere between individual verifiability and individual plus universal verifiability as defined in [24]. A precise comparison is difficult as individual and universal verifiability are strongly tight together in [24]. Moreover, [24] only considers the case where all voters are honest and they all vote.

We consider the privacy definition proposed in [20] and widely adopted in symbolic models: an attacker cannot distinguish when Alice is voting v_1 and Bob is voting v_1 from the scenario where the two votes are swapped.

Definition 3.10 (Privacy [20]). Let P be a voting process. P satisfies *privacy* if, for any substitution α from voters to votes, for any two voters $a, b \in \mathcal{A} \setminus \text{dom}(\alpha)$ and any two votes $v_1, v_2 \in \mathcal{V}$, we have

$$P_{\alpha \cup \{a \mapsto v_1, b \mapsto v_2\}} \approx P_{\alpha \cup \{a \mapsto v_2, b \mapsto v_1\}}$$

3.4 Privacy implies verifiability

We show that privacy implies verifiability under a couple of assumptions, typically satisfied in practice.

First, we assume a light form of determinacy: two traces with the same observable actions yield the same election result. This excludes for example cases for voters chose non deterministically how they vote. Formally, we say that a voting process P with election channel c_r is *election determinate* if, for any substitution α from voters to votes, for any two traces t, t' such that $t =_\tau t'$, $(t.\text{out}(c_r, x), \phi) \in \text{trace}(P_\alpha)$, and $(t'.\text{out}(c_r, x), \phi') \in \text{trace}(P_\alpha)$, then

$$\phi(x) \in R(\rho(V)) \Rightarrow \phi'(x) \in R(\rho(V))$$

This assumption still supports some form of non determinism but may not hold for example in the case where voters use anonymous channels that even hide who participated in the election.

Second, we assume that it is always possible for a new voter to vote (before the tally started) without modifying the behaviour of the protocol.

Formally, a voting process P is *voting friendly* if for all voter $a \in \mathcal{A}$, there exists t'' (the honest voting trace) such that for all α satisfying $a \notin \text{dom}(\alpha)$,

- for all $(t, \phi) \in \text{trace}(P_\alpha)$, such that $t = t'.\text{out}(c_r, x)$ for some t' , x , for all v , there exists tr, ψ such that $tr =_\tau t''$, $\text{Voted}(a, v) \in tr$, $(t'.\text{tr.out}(c_r, x), \psi) \in \text{trace}(P_{\alpha \cup \{a \mapsto v\}})$, and $\forall V. \phi(x) \in R(\rho(V)) \Rightarrow \psi(x) \in R(\rho(V \cup \{v\}))$. Intuitively, if a votes normally, her vote will be counted as expected, no matter how the adversary interfered with the other voters.
- for all t', x such that $\text{blocking}(t'.\text{out}(c_r, x), P_\alpha)$, for all v , tr, ψ such that $tr =_\tau t''$, we have $\text{blocking}(t'.\text{tr.out}(c_r, x), P_{\alpha \cup \{a \mapsto v\}})$. Intuitively, the fact that a voted does not suddenly unlock the tally.

In practice, most voting systems are voting friendly since voters vote independently. In particular, process P_{Helios} modelling Helios, as defined in Example 3.3, is voting friendly (assuming an honest tally). The voting friendly property prevents a fully dishonest tally since the first item requires that unmodified honest ballots are correctly counted. However, we can still consider a partially dishonest tally that, for example, discards or modifies ballots that have been flagged by the attacker.

Moreover, we assume that there exists a neutral vote, which is often the case in practice. Actually, this is a simplified (sufficient) condition. Our result also holds as soon as there is a vote that can be counted separately from the other votes (as formally defined in a companion technical report [?]).

THEOREM 3.11 (PRIVACY IMPLIES INDIVIDUAL VERIFIABILITY). *Let P be a voting friendly, election determinate voting process. If P satisfies privacy then P satisfies individual verifiability.*

The proof of this result intuitively relies on the fact that in order to satisfy privacy w.r.t. two voters Alice and Bob, a voting process has to guarantee that the vote of Alice is, if not correctly counted, at least taken into account to some extent. Indeed, if an attacker, trying to distinguish whether Alice voted for 0 and Bob for 1, or Alice voted for 1 and Bob for 0, is able to make the tally ignore completely the vote of Alice, the result of the election is then Bob's choice. Hence the attacker learns how Bob voted, which breaks privacy.

Therefore, we first we prove that if a protocol satisfies privacy, then if we compare an execution (i.e. a trace) where Alice votes 0 with the corresponding execution where Alice votes 1, the resulting election results must differ by exactly a vote for 0 and a vote for 1. Formally, we show the following property.

LEMMA 3.12. *If a voting friendly, election determinate voting process P satisfies privacy, then it satisfies*

$$\begin{aligned} [t =_\tau t' \wedge (t, \phi) \in \text{trace}(P_{\alpha \cup \{a \mapsto v_1\}}) \wedge \\ (t', \phi') \in \text{trace}(P_{\alpha \cup \{a \mapsto v_2\}}) \wedge \\ \text{result}(t, \phi, V) \wedge \text{result}(t', \phi', V')] \implies \\ \rho(V' \uplus \{v_1\}) = \rho(V \uplus \{v_2\}). \end{aligned}$$

This lemma is used as a central property to prove the theorem. Intuitively, we apply this lemma repeatedly, changing one by one all the votes from honest voters into neutral votes. Let r denote the result before this operation, and r' the result after. Let V_a denote the multiset of honest votes, and V_b the multiset containing the

same number of neutral votes. Thanks to Lemma 3.12, we can show that $r * \rho(V_b) = r' * \rho(V_a)$. Since V_b only contains neutral votes, we have $r = r' * \rho(V_a)$. This means that r contains all honest votes, hence the voting process satisfies individual verifiability.

The detailed proof of this theorem can be found in a companion technical report [?].

4 COMPUTATIONAL MODEL

Computational models define protocols and adversaries as probabilistic polynomial-time algorithms.

Notation: We may write $(id, *) \in L$ as a shorthand, meaning that there exists an element of the form (id, x) in L . If V is a multiset of elements of the form (id, v) , we define $\rho(V) = \rho(\{v \mid (id, v) \in V\})$.

4.1 Voting system

We assume that the ballot box displays a board BB , that is a list of ballots. The nature of the ballots depend on the protocol we consider.

Definition 4.1. *A voting scheme consists in six algorithms*

(Setup, Credential, Vote, VerifVoter, Tally, Valid)

- **Setup**(1^λ), given a security parameter λ , returns a pair of *election keys* (pk, sk) .
- **Credential**($1^\lambda, id$) creates a *credential cred* for voter id , for example a signing key. The credential may be empty as well. Registered voters are stored in a list U .
- **Vote**($id, cred, pk, v$) constructs a ballot containing the vote v for voter id with credential $cred$, using the election public key pk .
- **VerifVoter**($id, cred, L, BB$) checks whether the local knowledge L of voter id is consistent with the board BB . For example, a voter may check that her (last) ballot appears on the bulletin board.
- **Tally**(BB, sk, U) computes the *tally* of the ballots on the board BB , using the election secret key sk , assuming a list of registered voter identities and credentials U . The Tally algorithm first runs some test **ValidTally**(BB, sk, U) that typically checks that the ballots of BB are valid. Tally may return $\perp$ if the tally procedure fails (invalid board or decryption failure for example). If **Tally**(BB, sk, U) $\neq \perp$ then it must correspond to a valid result, that is, there exists V such that **Tally**(BB, sk, U) = $\rho(V)$.
- **Valid**(id, b, BB, pk) checks that a ballot b cast by a voter id is valid with respect to the board BB using the election public key pk . For example, the ballot b should have a valid signature or valid proofs of knowledge. The ballot b will be added to BB only if **Valid**(id, b, BB, pk) succeeds.

We will always assume a *correct* voting scheme, that is, tallying honestly generated ballots yields the expected result. Formally, for all distinct identities $U = id_1, \dots, id_n$, and credentials $cred_1, \dots, cred_n$, for all votes $v_1, \dots, v_n$, for all election keys (pk, sk) , if $BB = [\text{Vote}(id_i, cred_i, pk, v_i) \mid i \in [1, n]]$, then

$$\text{Tally}(BB, sk, U) = \rho(\{v_1, \dots, v_n\})$$

The tally algorithm typically applies a revote policy. Indeed, if voters may vote several times, the revote policy states which vote

should be counted. The two main standard revote policies are 1. the last vote counts or 2. the first vote counts (typically when revote is forbidden). In what follows, our definitions are written assuming the last ballot revote policy. However, they can easily be adapted to the first ballot revote policy and all our results hold in both cases (as shown in a companion technical report [?]).

The revote policy is either based on the identities or the credentials. We say that a voting system is *id-based* if there exists a function open_{id} which, given a ballot b , retrieves the associated identity. Formally, for any $id, cred, pk, v$,

$$\text{open}_{id}(\text{Vote}(id, cred, pk, v)) = id$$

Similarly, we say that a voting system is *cred-based* if there exists a function open_{cred} which, given a ballot b , the election secret key sk , and a list U of registered voters and credentials, retrieves the credential $cred$ used by the voter to create the ballot. Formally, for any $id, cred, sk, pk, v$,

$$\text{open}_{cred}(\text{Vote}(id, cred, pk, v), sk, U) = cred$$

Note that some schemes are neither id-based nor cred-based, in particular when the ballots contain no identifier. Such schemes typically assume that voters do not revote since there is no means to identify whether two ballots originate from the same voter.

4.2 Security properties

As usual, an *adversary* is any probabilistic polynomial time Turing machine (PPTM). We define verifiability and privacy through game-based properties.

4.2.1 Verifiability. For verifiability, we propose a simple definition, inspired from [15, 25]. Intuitively, we require that the election result contains at least the votes of all honest voters. This notion was called weak verifiability in [15] but we will call it individual verifiability to match the terminology used in symbolic settings. More sophisticated and demanding definitions have been proposed, for example controlling how many dishonest votes can be inserted [15] or tolerating some variations in the result [25]. The main missing part (in terms of security) is that our definition does not control ballot stuffing: arbitrarily many dishonest votes may be added to the result. The reason is that ballot stuffing seems unrelated to privacy. Moreover, our definition assumes an honest tally, and thus does not capture universal verifiability aspects. The main reason is that existing privacy definitions in computational settings assume an honest tally and we compare the two notions under the same trust assumptions. We leave as future work to determine how to extend these two definitions to a dishonest tally, and whether the implication still holds.

Verifiability is defined through the game $\text{Exp}_{\mathcal{A}}^{\text{verif}}(\lambda)$ displayed on Figure 4. In a first step, the adversary may use oracles $O_{\text{reg}}(id)$ and $O_{\text{corr}}(id)$ (defined on Figure 3) to respectively register a voter and get her credential (in this case, the voter is said to be corrupted). Then the adversary may ask an honest voter id to vote for a given vote v through oracle $O_{\text{vote}}^v(id, v)$. In this case, the adversary sees the corresponding ballot and the fact that id voted for v is registered in the list Voted. The adversary may also cast an arbitrary ballot b in the name of a dishonest voter id through oracle $O_{\text{cast}}(id, b)$. Finally, the adversary wins if the election result does not contain

$O_{\text{reg}}(id)$	$O_{\text{corr}}(id)$
if $(id, *) \in U$ then	if $(id, *) \notin U \vee (id, *) \in CU$ then
stop	stop
else	else
$cred_{id} \leftarrow \text{Credential}(1^\lambda, id)$	$CU \leftarrow CU \parallel (id, cred_{id})$
$U \leftarrow U \parallel (id, cred_{id})$	return $cred_{id}$
	where $(id, cred_{id}) \in U$

Figure 3: Registration and corruption oracles

all the honest votes registered in Voted (where only the last vote is counted).

Definition 4.2 (Individual verifiability). A voting system is *individually verifiable* if for any adversary $\mathcal{A}$,

$$\mathbb{P} \left[\text{Exp}_{\mathcal{A}}^{\text{verif}}(\lambda) = 1 \right] \text{ is negligible.}$$

As mentioned in introduction, [13] shows an impossibility result between (unconditional) privacy and verifiability. [13] considers another aspect of verifiability, namely universal verifiability, that is, the guarantee that the result corresponds to the content of the ballot, even in presence of a dishonest tally. Interestingly, the same incompatibility result holds between individual verifiability and unconditional privacy, for the same reasons. Exactly like in [13], a powerful adversary (*i.e.* not polynomial) could tally BB and BB' where BB' is the ballot box from which Alice's ballot has been removed and infer Alice's vote by difference. More generally, unconditional privacy is lost as soon as there exists a tally function that is meaningfully related to the result, which is implied by individual verifiability.

4.2.2 Privacy. For privacy, we consider the old, well established definition of Josh Benaloh [6]. More sophisticated definitions have been proposed later (see [7] for a survey and a unifying definition). They aim in particular at getting rid of the partial tally assumption (needed in [6]). Note however that they all assume an honest ballot box. Since we also assume partial tally, the original Benaloh definition is sufficient for our needs. In particular, we do not know if privacy implies verifiability for counting functions that do not have the partial tally property. This is left as future work.

Intuitively, a voting system is private if, no matter how honest voters vote, the adversary cannot see any difference. However, the adversary always sees the election result, that leaks how the group of honest voters voted (altogether). Therefore, the election result *w.r.t.* the honest voters has to remain the same. More formally, in a first step, the adversary uses oracles $O_{\text{reg}}(id)$ and $O_{\text{corr}}(id)$ to respectively register a voter and get her credential. Then the adversary may request an honest voter id to vote either for v_0 or v_1 through oracle $O_{\text{vote}}^p(id, v_0, v_1)$. Voter id will vote v_β depending on the bit β . The adversary may also cast an arbitrary ballot b in the name of a dishonest voter id through oracle $O_{\text{cast}}(id, b)$. The election will be tallied, only if the set V_0 of votes v_0 yields the same result than the set V_1 of votes v_1 (where only the last vote is counted). Finally, the adversary wins if he correctly guesses β . Formally, privacy is defined through the game $\text{Exp}_{\mathcal{A}}^{\text{priv}, \beta}(\lambda)$ displayed on Figure 5.

$\text{Exp}_{\mathcal{A}}^{\text{verif}}(\lambda)$ <hr/> $(pk, sk) \leftarrow \text{Setup}(1^\lambda)$ $U, CU \leftarrow []$ $state \leftarrow \mathcal{A}_1^{O_{\text{reg}}, O_{\text{corr}}}(\text{pk})$ $BB, \text{Voted} \leftarrow []$ $\mathcal{A}_2^{O_{\text{vote}}, O_{\text{cast}}}(state, \text{pk})$ $r \leftarrow \text{Tally}(BB, sk, U)$ if $r \neq \perp \wedge \forall V_c \text{ (finite) } . r \neq \rho(\{v_i\}_{1 \leq i \leq k} \uplus V_c)$ then $\text{return } 1$ where $\text{Voted} = \{(id_1, v_1), \dots, (id_k, v_k)\}$	$O_{\text{vote}}^v(id, v)$ <hr/> if $(id, *) \in U \setminus CU$ then $b \leftarrow \text{Vote}(id, cred_{id}, \text{pk}, v)$ $BB \leftarrow BB \parallel b$ $\text{Voted} \leftarrow \text{Voted}' \parallel (id, cred_{id}, v)$ return b where $(id, cred_{id}) \in U$ and Voted' is obtained from Voted by removing all previous instances of $(id, *)$	$O_{\text{cast}}(id, b)$ <hr/> if $(id, *) \in CU \wedge$ $\text{Valid}(id, b, BB, \text{pk})$ then $BB \leftarrow BB \parallel (id, b)$
---	--	---

Figure 4: Verifiability

Definition 4.3 (Privacy [6]). A voting system is private if for any adversary $\mathcal{A}$,

$$\left| \mathbb{P} \left[\text{Exp}_{\mathcal{A}}^{\text{priv}, 0}(\lambda) = 1 \right] - \mathbb{P} \left[\text{Exp}_{\mathcal{A}}^{\text{priv}, 1}(\lambda) = 1 \right] \right| \text{ is negligible.}$$

4.3 Privacy implies individual verifiability

We show that privacy implies individual verifiability and we first list here our assumptions. As for the symbolic case, we assume the existence of a neutral vote. We also require that the tally can be performed piecewise, that is, informally, as soon as two boards BB_1, BB_2 are independent then $\text{Tally}(BB_1 \uplus BB_2) = \text{Tally}(BB_1) * \text{Tally}(BB_2)$. This property is satisfied by most voting schemes. Formally, we characterize this notion of “independence” depending on whether a scheme is id-based or cred-based.

An id-based voting scheme has the *piecewise tally property* if for any two boards BB_1 and BB_2 that contain ballots registered for different agents and such that $BB_1 \uplus BB_2$ is valid, that is, if

$$\text{ValidTally}(BB_1 \uplus BB_2, sk, U) \wedge \\ \forall b \in BB_1. \forall b' \in BB_2. \text{open}_{id}(b) \neq \text{open}_{id}(b'),$$

then their tally can be computed separately:

$$\text{Tally}(BB_1 \uplus BB_2, sk, U) = \text{Tally}(BB_1, sk, U) * \text{Tally}(BB_2, sk, U). \quad (*)$$

We also assume that the tally only counts ballots cast with registered *ids*, i.e. $\forall BB, sk, U. \text{Tally}(BB, sk, U) = \text{Tally}(BB', sk, U)$ where $BB' = [b \in BB \mid (\text{open}_{id}(b), *) \in U]$; and that registering more voters does not change the tally: if U, U' have no *id* in common and $\forall b \in BB. (\text{open}_{id}(b), *) \notin U'$, then $\text{Tally}(BB, sk, U) = \text{Tally}(BB, sk, U \cup U')$.

Similarly, a cred-based voting scheme has the *piecewise tally property* if for any two boards BB_1 and BB_2 that contain ballots associated to different credentials, that is

$$\forall b \in BB_1. \forall b' \in BB_2. \text{open}_{cred}(b, sk, U) \neq \text{open}_{cred}(b', sk, U)$$

then their tally can be computed separately (Property (*)).

We also assume that registering more voters does not change the tally: if U, U' share no credentials and $\forall b \in BB. (*, \text{open}_{cred}(b, sk, U \cup U')) \notin U'$, then $\text{Tally}(BB, sk, U) = \text{Tally}(BB, sk, U \cup U')$.

We say that a (id-based) voting scheme is *strongly correct* if whatever valid board the adversary may produce, adding a honestly generated ballot still yields a valid board. This property is formally

defined through the game $\text{Exp}_{\mathcal{A}}^{\text{ValidTally}}(\lambda)$ displayed in Figure 6. A similar assumption was introduced in [7]. For example, Helios is strongly correct.

A voter credential typically includes a private part used to generate a signing key for example. It should not be possible for an adversary to forge a ballot with an honest credential. Formally, we say that a voting scheme has *non-malleable credentials*, if for any adversary $\mathcal{A}$,

$$\mathbb{P} \left[\text{Exp}_{\mathcal{A}}^{\text{NM}}(\lambda) = 1 \right] \text{ is negligible}$$

where $\text{Exp}_{\mathcal{A}}^{\text{NM}}(\lambda)$ is defined on Figure 7. For example, Belenios and Civitas have non-malleable credentials.

THEOREM 4.4 (PRIVACY IMPLIES INDIVIDUAL VERIFIABILITY). *Let V be an id-based, strongly correct, voting scheme that has the piecewise tally property. If V is private, then V is individually verifiable.*

Similarly, let V be a cred-based voting scheme that has the piecewise tally property and non-malleable credentials. If V is private, then V is individually verifiable.

The proof of this theorem is inspired by the same intuition as in the symbolic case: if an attacker manages to break verifiability, that is, to obtain that not all votes from the honest voters are counted correctly, then there also exists an attack against privacy. Indeed, consider a scenario with additional, new voters, whose votes should compensate those cast by the initial voters. By performing the attack on verifiability for the initial voters, the attacker reaches a state where, in the result of the election, they are no longer compensated by the new votes. This allows the attacker to break privacy.

More precisely, the general idea of the proof is as follows. Consider an attacker $\mathcal{A}$ that breaks individual verifiability, i.e. wins the game $\text{Exp}_{\mathcal{A}}^{\text{verif}}$ with non negligible probability. We construct an attacker $\mathcal{B}$ that breaks privacy, i.e. wins $\text{Exp}_{\mathcal{A}}^{\text{priv}, \beta}$. $\mathcal{B}$ starts by registering, and corrupting, the same voters as $\mathcal{A}$, using oracles O_{reg} and O_{corr} . Let $id_1, \dots, id_n$ be this first set of voters. $\mathcal{B}$ then registers another set of n voters $id'_1, \dots, id'_n$, where the id'_i are fresh identities, that $\mathcal{A}$ does not use.

$\mathcal{B}$ then simulates $\mathcal{A}$, using the oracle O_{vote}^p to simulate $\mathcal{A}$'s calls to O_{vote}^v . Specifically, when $\mathcal{A}$ calls $O_{\text{vote}}^v(id, v)$, $\mathcal{B}$ calls the oracle $O_{\text{vote}}^p(id, v, v^{\text{blank}})$, where v^{blank} is a neutral vote. Once $\mathcal{B}$ is done simulating $\mathcal{A}$, it triggers the new voters id'_i to vote, by calling

$\text{Exp}_{\mathcal{A}}^{\text{priv}, \beta}(\lambda)$	$\mathcal{O}_{\text{vote}}^p(id, v_0, v_1)$	$\mathcal{O}_{\text{cast}}(id, b)$
$(pk, sk) \leftarrow \text{Setup}(1^\lambda)$	if $(id, *) \in U \setminus CU$ then	if $(id, *) \in CU \wedge$
$U, CU \leftarrow []$	$b \leftarrow \text{Vote}(id, cred_{id}, pk, v_\beta)$	$\text{Valid}(id, b, BB, pk)$
$state_1 \leftarrow \mathcal{A}_1^{O_{\text{reg}}, O_{\text{corr}}}(\text{pk})$	$BB \leftarrow BB \parallel b$	then
$BB, V_0, V_1 \leftarrow []$	$V_0 \leftarrow V'_0 \parallel (id, v_0)$	$BB \leftarrow BB \parallel (id, b)$
$state_2 \leftarrow \mathcal{A}_2^{O_{\text{vote}}, O_{\text{cast}}}(state_1, pk)$	$V_1 \leftarrow V'_1 \parallel (id, v_1)$	
if $\rho(V_0) = \rho(V_1)$ then	return b	
$r \leftarrow \text{Tally}(BB, sk, U)$	where $(id, cred_{id}) \in U$	
$\beta' \leftarrow \mathcal{A}_3(state_2, pk, r)$	and V'_0 (resp. V'_1) is obtained from V_0 (resp. V_1)	
return β'	by removing all instances of $(id, *)$	

Figure 5: Privacy

$\text{Exp}_{\mathcal{A}}^{\text{ValidTally}}(\lambda)$	$\mathcal{O}_{\text{vote}}^{vt}(id, v)$
$(pk, sk) \leftarrow \text{Setup}(1^\lambda)$	if $\exists cred_{id}.(id, cred_{id}) \in U \setminus CU$ then
$U, CU \leftarrow []$	return $\text{Vote}(id, cred_{id}, pk, v)$
$state \leftarrow \mathcal{A}_1^{O_{\text{reg}}, O_{\text{corr}}}(\text{pk})$	
$(BB, id, v) \leftarrow \mathcal{A}_2^{O_{\text{vote}}^{vt}}(state, pk)$	
$b \leftarrow \text{Vote}(id, cred_{id}, pk, v)$	
where $(id, cred_{id}) \in U$	
if $(id, *) \in U \setminus CU \wedge$	
$(\forall b' \in BB. \text{open}_{id}(b') \neq id) \wedge$	
$\text{ValidTally}(BB, sk, U) \wedge$	
$\neg \text{ValidTally}(BB \parallel b, sk, U)$ then	
return 1	

Figure 6: ValidTally game

$\text{Exp}_{\mathcal{A}}^{\text{NM}}(\lambda)$	$\mathcal{O}_c(id, v)$
$(pk, sk) \leftarrow \text{Setup}(1^\lambda)$	if $(id, *) \in U \setminus CU$ then
$U, CU \leftarrow []$	$b \leftarrow \text{Vote}(id, cred_{id}, pk, v)$
$state \leftarrow \mathcal{A}_1^{O_{\text{reg}}, O_{\text{corr}}}(\text{pk})$	$L \leftarrow L \parallel b$
$L \leftarrow []$	return b
$b \leftarrow \mathcal{A}_2^{O_c}(state, pk)$	where $(id, cred_{id}) \in U$
if $b \notin L \wedge$	
$\exists (id, cred_{id}) \in U \setminus CU.$	
$\text{open}_{cred}(b, sk, U) = cred_{id}$ then	
return 1	

Figure 7: Credential non-malleability

the oracle $\mathcal{O}_{\text{vote}}^p(id'_i, v_i^{\text{blank}}, v_i)$, where v_i is the (last) vote cast by id_i . The vote of each id'_i compensates the vote of id_i , so that the condition $\rho(V_0) = \rho(V_1)$ from Exp^{priv} holds. $\mathcal{B}$ then obtains the result r of the election, which is equal to $r_1 * r_2$, where r_1 is the tally of the ballots cast by $\mathcal{A}$, and r_2 the tally of the additional ballots cast by $\mathcal{B}$. Then:

- if $\beta = 0$: then all the votes cast by the id'_i were v_i^{blank} , and the result is thus $r = r_1$. Since $\mathcal{A}$ breaks individual verifiability, r_1 does not contain the honest votes, *i.e.*, for all multiset V_c of votes, $r \neq \rho(v_1, \dots, v_n) * \rho(V_c)$.
- if $\beta = 1$ however, the votes cast by the id'_i were the v_i , and the partial tally r_2 is therefore $r_2 = \rho(v_1, \dots, v_n)$. Hence, the result r does contain the honest votes.

Therefore, by observing whether the final result of the election contains the honest votes, $\mathcal{B}$ is able to guess β , and wins Exp^{priv} .

5 PRIVACY WITH A DISHONEST BOARD

Our main theorem states that privacy implies individual verifiability. However, the privacy definition introduced by Benaloh assumes an honest ballot box, as most existing privacy definitions of the literature [7]. Therefore, our main theorem shows that whenever a voting scheme is private *w.r.t.* an honest ballot box, then it is also individually verifiable *w.r.t.* an honest ballot box, which is of course a rather weak property. However, intuitively, our proof technique does not rely on the trust assumptions.

As pointed out in introduction, extending cryptographic privacy definitions to a dishonest ballot box is difficult. Consider the natural extension of privacy as displayed in Figure 8: the game is the same than $\text{Exp}_{\mathcal{A}}^{\text{priv}, \beta}(\lambda)$ except that the adversary arbitrarily controls the ballot box. Unfortunately, an adversary can always win this new game. Indeed, he may simply query $\mathcal{O}_{\text{vote}}^d(id_1, 0, 1)$ and $\mathcal{O}_{\text{vote}}^d(id_2, 1, 0)$, yielding respectively ballots b_{id_1} and b_{id_2} . Then the adversary choses $BB = b_{id_1}$. The tally will return β , hence the adversary wins. This corresponds to the fact that an adversary may always isolate a voter and break her privacy.

5.1 Privacy with careful voters

To solve this issue, we choose another approach, which consists in explicitly modelling the verification steps made by voters: the tally will be performed only if honest voters have successfully run their checks (*e.g.* checking that their ballot belongs to the bulletin board). Therefore, we extend the privacy game as follows. The adversary arbitrarily controls the ballot box and may request honest voters to vote through $\mathcal{O}_{\text{vote}}^{p, c}(id, v_0, v_1)$ as before. Note that there is no need for the $\mathcal{O}_{\text{cast}}$ oracle since the adversary may add directly his own

$\text{Exp}_{\mathcal{A}}^{\text{dis}, \beta}(\lambda)$ <hr/> $(\text{pk}, \text{sk}) \leftarrow \text{Setup}(1^\lambda)$ $\text{U}, \text{CU} \leftarrow []$ $\text{state}_1 \leftarrow \mathcal{A}_1^{O_{\text{reg}}, O_{\text{corr}}}(\text{pk})$ $\text{V}_0, \text{V}_1 \leftarrow []$ $\text{state}_2, \text{BB} \leftarrow \mathcal{A}_2^{O_{\text{vote}}^d}(\text{state}_1, \text{pk})$ if $\rho(\text{V}_0) = \rho(\text{V}_1)$ then $r \leftarrow \text{Tally}(\text{BB}, \text{sk}, \text{U})$ $\beta' \leftarrow \mathcal{A}_3(\text{state}_2, \text{pk}, r)$ return β'	$O_{\text{vote}}^d(id, v_0, v_1)$ <hr/> if $(id, *) \in \text{U} \setminus \text{CU}$ then $b \leftarrow \text{Vote}(id, \text{cred}_{id}, \text{pk}, v_\beta)$ $\text{V}_0 \leftarrow \text{V}'_0 \parallel (id, v_0)$ $\text{V}_1 \leftarrow \text{V}'_1 \parallel (id, v_1)$ return b <div style="margin-left: 20px;"> where $(id, \text{cred}_{id}) \in \text{U}$ and V'_0 (resp. V'_1) is obtained from V_0 (resp. V_1) by removing all instances of $(id, *)$ </div>
--	--

Figure 8: Privacy against a dishonest board (PrivDis-Naive)

ballots in the ballot box. He triggers voters to run their verification tests through the oracle $O_{\text{happyBB}}(id)$. To run her verification test (using algorithm VerifVoter), the voter has access to the ballot box BB forged by the adversary as well as her local state L_{id} that contains in particular her previously generated ballots. The tally is performed only if all honest voters have successfully performed their test and if, as previously, the set V_0 of left votes v_0 yields the same result than the set V_1 of right votes v_1 . Formally, privacy with careful voters is defined through the game $\text{Exp}_{\mathcal{A}}^{\text{priv-careful}}$ displayed on Figure 9.

Definition 5.1 (Privacy with careful voters). A voting system is *private against a dishonest board with careful voters* if for any adversary $\mathcal{A}$,

$$\left| \mathbb{P} \left[\text{Exp}_{\mathcal{A}}^{\text{priv-careful}, 0}(\lambda) = 1 \right] - \mathbb{P} \left[\text{Exp}_{\mathcal{A}}^{\text{priv-careful}, 1}(\lambda) = 1 \right] \right|$$

is negligible.

While this definition models a dishonest ballot box, it implicitly assumes that all voters see the *same* (possibly dishonest) ballot box. This is a very common assumption in voting, that needs to be achieved by external means.

Similarly, we extend individual verifiability to *individual verifiability against a dishonest board* as expected, assuming that the tally is performed only if all honest voters have successfully performed their test. The formal definition of individual verifiability against a dishonest board can be found in a companion technical report [?].

5.2 Privacy implies individual verifiability against a dishonest box too

We need to assume that the verification test run by honest voters (VerifVoter) is consistent with how the voter voted. Namely, if the voter's intended ballot is the one that is selected from the board by the revoke policy (e.g. appears last w.r.t. this voter), then this voter must be satisfied with the board (that is, VerifVoter passes). Conversely, if the test VerifVoter fails for voter id then adding ballots unrelated to id (or her credential) will not change this fact (VerifVoter will still fail). These assumptions are formally stated in a companion technical report [?].

THEOREM 5.2 (PRIVACY IMPLIES INDIVIDUAL VERIFIABILITY AGAINST A DISHONEST BOARD). *Let V be an id -based, strongly correct voting*

$\text{Exp}_{\mathcal{A}}^{\text{priv-careful}, \beta}(\lambda)$ <hr/> $(\text{pk}, \text{sk}) \leftarrow \text{Setup}(1^\lambda)$ $\text{U}, \text{CU} \leftarrow []$ $\text{state}_1 \leftarrow \mathcal{A}_1^{O_{\text{reg}}, O_{\text{corr}}}(\text{pk})$ $\text{V}_0, \text{V}_1, L_{id} \text{ (for all } id \text{ in U)} \leftarrow []$ $\text{state}_2, \text{BB} \leftarrow \mathcal{A}_2^{O_{\text{vote}}^{p,c}}(\text{state}_1, \text{pk})$ $\text{H} \leftarrow []$ $\text{state}_3 \leftarrow \mathcal{A}_3^{O_{\text{happyBB}}}(\text{state}_2)$ if $\forall id. (id, *) \in \text{V}_0, \text{V}_1 \Rightarrow id \in \text{H}$ and $\rho(\text{V}_0) = \rho(\text{V}_1)$ then $r \leftarrow \text{Tally}(\text{BB}, \text{sk}, \text{U})$ else $r \leftarrow \perp$ $\beta' \leftarrow \mathcal{A}_4(\text{state}_3, \text{pk}, r)$ return β'	$O_{\text{vote}}^{p,c}(id, v_0, v_1)$ <hr/> if $(id, *) \in \text{U} \setminus \text{CU}$ then $b \leftarrow \text{Vote}(id, \text{cred}_{id}, \text{pk}, v_\beta)$ $\text{V}_0 \leftarrow \text{V}'_0 \parallel (id, v_0)$ $\text{V}_1 \leftarrow \text{V}'_1 \parallel (id, v_1)$ $L_{id} \leftarrow L_{id} \parallel (b, v_\beta)$ return b <div style="margin-left: 20px;"> where $(id, \text{cred}_{id}) \in \text{U}$ and V'_0 (resp. V'_1) is obtained from V_0 (resp. V_1) by removing all instances of $(id, *)$ </div>
--	--

$O_{\text{happyBB}}(id)$ <hr/> if $(id, \text{cred}_{id}) \in \text{U} \setminus \text{CU}$ then if VerifVoter($id, \text{cred}_{id}, L_{id}, \text{BB}$) then $\text{H} \leftarrow \text{H} \parallel id$

Figure 9: Privacy game against a dishonest board with careful voters (Priv-careful)

scheme that has the piecewise tally property. If V is private against a dishonest board with careful voters, then V is individually verifiable against a dishonest board with careful voters.

Similarly, let V be a cred-based voting scheme that has the piecewise tally property and non-malleable credentials. If V is private against a dishonest board with careful voters, then V is individually verifiable against a dishonest board with careful voters.

6 COMPARING PRIVACY

We compare different notions of privacy, with and without an honest ballot box, on four standard protocols (Helios, Belenios, Civitas, and Neuchâtel) as well as on our simple protocol, sketched in introduction.

To our knowledge, the only other definition of privacy with a dishonest ballot box is the privacy notion introduced by Bernhard and Smyth [9]. We first start by discussing this definition.

6.1 PrivacyBS

The privacy notion introduced by Bernhard and Smyth [9] is recalled in Figure 10 (PrivacyBS). The adversary may request a voter id to vote for v_0 or v_1 (depending on the bit β) through the oracle $O_{\text{vote}}^{bs}(id, v_0, v_1)$. He produces an arbitrary ballot box BB and the tally will be performed provided that, looking at honest ballots that appear in BB, counting the corresponding left and right votes yields the same result.

The main interest of [9] is to highlight the fact that previous definitions implicitly assume an honest ballot box. The first attempt at defining privacy w.r.t. a dishonest ballot box (PrivacyBS) has several limitations. First, it strongly assumes that the ballots that appear in the ballot box are exactly the same than the cast ballots. This is not

$\text{Exp}_{\mathcal{A}}^{\text{BS}, \beta}(\lambda)$	$\mathcal{O}_{\text{vote}}^{\text{bs}}(id, v_0, v_1)$
$(pk, sk) \leftarrow \text{Setup}(1^\lambda)$	if $(id, *) \in U \setminus CU$ then
$U, CU \leftarrow []$	$b \leftarrow \text{Vote}(id, cred_{id}, pk, v_\beta)$
$state_1 \leftarrow \mathcal{A}_1^{\mathcal{O}_{\text{reg}}, \mathcal{O}_{\text{corr}}}(\text{pk})$	$L \leftarrow L \parallel (b, v_0, v_1)$
$L \leftarrow []$	return b
$state_2, BB \leftarrow \mathcal{A}_2^{\mathcal{O}_{\text{vote}}^{\text{bs}}}(\text{state}_1, \text{pk})$	where $(id, cred_{id}) \in U$
if $\forall v.$	
$ \{b \mid b \in BB \wedge \exists v'. (b, v, v') \in L\} =$	
$ \{b \mid b \in BB \wedge \exists v'. (b, v', v) \in L\} $	
then	
$r \leftarrow \text{Tally}(BB, sk, U)$	
$\beta' \leftarrow \mathcal{A}_3(\text{state}_2, \text{pk}, r)$	
return β'	

Figure 10: PrivacyBS [9]

the case for example of the ThreeBallots protocol [27] where the ballot box only contains two shares (out of three) of the original ballot. It is not applicable either to a protocol like BeleniosRF [12] where ballots are re-randomised before their publication. Second, it requires ballots to be non-malleable [9]. This means that, as soon as a ballot includes a malleable part (for example the voter's id like in Helios, or a timestamp), privacy cannot be satisfied. This severely restricts the class of protocols that can be considered. Third, PrivacyBS does not account for a revote policy. As soon as revote is allowed (for example in Helios), then PrivacyBS is broken since some ballots may not be counted. Indeed, an attacker may call $\mathcal{O}_{\text{vote}}^{\text{bs}}(id_1, 1, 0)$, followed by $\mathcal{O}_{\text{vote}}^{\text{bs}}(id_1, 0, 1)$, obtaining ballots b_1, b'_1 , and return the board $BB = [b_1, b'_1]$. The equality condition on the number of ballots in BB produced by $\mathcal{O}_{\text{vote}}^{\text{bs}}$ holds, since for $v = 0, 1$:

$$|\{b \in BB \mid \exists v'. (b, v, v') \in L\}| = |\{b \in BB \mid \exists v'. (b, v', v) \in L\}| = 1$$

where $L = [(b_1, 1, 0), (b'_1, 0, 1)]$. Hence the tally is computed. According to the revote policy, only b'_1 is counted, and the result is β , which lets the attacker win Exp^{BS} .

6.2 Protocols

We consider four standard protocols (Helios, Belenios, Civitas, and Neuchâtel) as well as our simple protocol, presented in introduction. We briefly explain each of them in this section. In what follows $\mathcal{E} = (\text{gen}, \text{enc}, \text{dec})$ denotes an encryption algorithm.

Simple. We detail the simple protocol sketched in introduction. Recall that voters simply send their encrypted votes to the ballot box, and, at the end of the voting phase, the tally computes and publishes the result of the election. No revote is allowed, and the voters do not have any means of verifying that their vote is taken into account. Identities and credentials are not used in this protocol. The corresponding algorithms of this protocol are:

- $\text{Vote}(id, cred, pk, v) = \text{enc}(pk, v)$
- $\text{VerifVoter}(id, cred, L, BB) = \text{true}$ (voters do not make any checks)

- $\text{Tally}(BB, sk, U)$ checks that all the ballots in BB are distinct, and returns $\perp$ if not. The tally performs a random permutation of the ballots, decrypts all of them and returns the multiset of the votes they contain.
- $\text{Valid}(id, b, BB, pk)$ checks that b does not already occur in BB.

Helios [4] is similar to Simple, except that revote is allowed, and the last vote cast by each id is counted. To make this revote policy possible, the ballots contain the id of the voter: $\text{Vote}(id, cred, pk, v) = (id, \text{enc}(pk, v))$. $\text{enc}(pk, v)$ here also includes a proof that v is a valid vote. Credentials are unused. The tally computes the result of the election similarly to Simple except that it also features an homomorphic mode, where the tally homomorphically computes the sum of the ballots in BB, decrypts the resulting ciphertext and returns the result. Moreover, the tally returns a proof of correct decryption. In addition, the board which will be tallied is made public, allowing the voters to check that their last ballot is indeed the last ballot with their id on the board:

$$\begin{aligned} \text{VerifVoter}(id, cred, L_{id}, BB) = \\ \text{the last element in } L_{id} \text{ is the last ballot} \\ \text{registered for } id \text{ in BB.} \end{aligned}$$

Similarly to Simple, the Valid function checks that there is no duplicated ciphertext and also checks that the ballot is submitted under the right id .

$$\text{Valid}(id, (id', c), BB, pk) = (id = id') \wedge c \text{ does not occur in BB}$$

This models an authenticated channel between the ballot box and each voter: a voter id may not cast a vote in the name of id' .

Belenios [15] is similar to Helios, except that voters sign their encrypted vote thanks to their credential:

$$\text{Vote}(id, k, pk, v) = (id, \text{signElGamal}(v, pk, k))$$

where $\text{signElGamal}(\cdot, \cdot, \cdot)$ denotes the combination of the (ElGamal) encryption and the signature. As for Helios, it also includes a proof that v is a valid vote. Tally checks that there exists a bijection between the ids and the credentials in the final board, i.e. that the same id is always associated with the same signature, and vice versa. The revote policy counts the last ballot corresponding to a given credential. Voters can verify that their last ballot is indeed the last one signed by their key on the board.

Civitas [14]. In Civitas, voters privately receive a credential, that is published encrypted on the bulletin board. To cast a vote, a voter encrypts her vote, also encrypts her credential, and produces a proof π of well-formedness that links the two ciphertexts together. The corresponding ballot is of the form

$$\text{Vote}(id, cred, pk, v) = (\text{enc}(pk, cred), \text{enc}(pk, v), \pi).$$

The voters can verify that their vote will be taken into account by checking that it is present on the board that will be tallied.

$$\text{VerifVoter}(id, cred, L_{id}, BB) = b \in BB$$

where b is the ballot in L_{id} . In theory, revote is allowed. However, we unveil a small discrepancy in how revote should be performed. Assume for example that the last ballot should be counted. Since an adversary may recast old ballots generated by an honest voter, a voter should memorise all the ballots he generated and check that they appear in the right order on the ballot box. Such a check seems

Protocol	Honest board [6]	PrivDis-Naive	PrivacyBS [9]	Priv-careful
Simple (no revote)	✓	✗	✓	✗
Helios	✓	✗	✗	✗
Belenios	✓	✗	✗	✓
Civitas (no revote)	✓	✗	✓	✓
Neuchâtel (no revote)	✓	✗	✓	✗

Figure 11: Comparison of several privacy definitions
(✓: the protocol is private, ✗: there exists an attack on privacy)

highly cumbersome for an average voter and we could not find its description in [14]. Therefore, we simply assume here that honest voters do not revoke.

Neuchâtel [21]. Voters privately receive a code sheet, where each candidate is associated to a (short) code. To cast a vote, voters send their encrypted votes to the ballot box, similarly to Simple or Helios. The ballot box then provides a return code allowing the voter to check that the ballot has been received and that it encrypts their candidate, as intended. This offers a protection against a dishonest voting client (e.g. if the voter’s computer is corrupted). No revoke is allowed. Since the bulletin board is not published, voters cannot check that their ballots really belong to the final board (used for tally), which we model by $\text{VerifVoter}(id, cred, L, BB) = \text{true}$. Voters have to trust the voting server (or other internal components) on this aspect.

6.3 Attacks

Simple. As described in introduction, a dishonest ballot box may break ballot privacy of any voter by simply replacing the other votes by votes of its choice. In other words, even if the ballot box does not detain any decryption key, it can learn how Alice’s voted.

Neuchâtel. Exactly like the Simple protocol, a dishonest ballot box may break ballot privacy of any voter by simply replacing the other votes. This is due to the fact that voters have no control over the ballots that are actually tallied. Note that the Neuchâtel protocol actually includes internal mechanisms that render such an attack difficult. However, from the point of view of a voter, if the ballot box is compromised, her privacy is no longer guaranteed.

Helios. Helios is also vulnerable to an attack when the ballot box is compromised. This attack is due to P. Roenne [28]. It involves two honest voters id_1 , id_2 , and a dishonest voter id_3 . The attacker may call $O_{\text{vote}}^{p,c}(id_1, 0, 1)$ twice and $O_{\text{vote}}^{p,c}(id_2, 1, 0)$ once, obtaining ballots (id_1, b_1) , (id_1, b'_1) , (id_2, b_2) . The adversary then returns the board $[(id_1, b'_1), (id_2, b_2), (id_3, b_1)]$. All ballots are different, hence no weeding is needed. The result of the tally is then $\rho(\{0, 1, 0\})$ if $\beta = 0$ and $\rho(\{1, 0, 1\})$ if $\beta = 1$. The attacker can therefore observe the difference in the result, which breaks privacy.

Belenios and *Civitas* remain private against a dishonest board as long as voters perform their verification checks. We formally prove privacy according to our definition Priv-careful.

6.4 Comparison

We summarise our findings in Figure 11. As explained in Section 5, the naive extension of the privacy definition to a dishonest board (PrivDis-Naive) is immediately false for any protocol.

All of our five protocols satisfy privacy against an honest ballot box. We rely here on previous results of the literature, except for Civitas (and of course the Simple protocol). Indeed, Civitas has been proved to be coercion-resistant [14] in a rather different setting. Therefore we show here that it satisfies the Benaloh definition.

PrivacyBS fails to detect the attack on the Neuchâtel protocol and the Simple protocol since it requires that the tally of the honest ballots present on the final board does not leak information. Conversely, it cannot prove Belenios private as it does not properly handle revoting as explained in Section 6.1.

7 CONCLUSION

We show a subtle relation between privacy and verifiability, namely that privacy implies individual verifiability, which is rather counter-intuitive. Our result holds in a cryptographic as well as a symbolic setting, for various trust assumptions. In contrast, privacy does not seem to imply universal verifiability nor eligibility verifiability. To show that there is indeed no implication, we plan to exhibit counter-examples, as simple as possible.

Our result assumes counting functions that have the partial tally property. Our proof technique does not extend immediately to more complex counting functions such as STV or Condorcet. We plan to study how privacy and individual verifiability are related in this context. Also, our results implicitly discard anonymous channels: computational models do not account for anonymous channels while our election determinism assumption discards at least some use of anonymous channels. Intuitively, in presence of anonymous channels, an attacker may be able to modify a ballot without being able to tell which one, hence breaking verifiability without breaking privacy. It would be interesting to identify which kind of anonymous channels and more generally, which form of non determinism, can still be tolerated.

Our findings also highlight a crucial need for a ballot privacy definition in the context of a dishonest ballot box, in a cryptographic setting. So far, privacy has only been proved assuming an honest ballot box, which forms a very strong trust assumption that was probably never made clear to voters nor election authorities.

We propose a first attempt at modelling privacy against a dishonest board, assuming that honest voters check their ballots as expected by the voting protocol. We do not see our definition as final. In particular, assuming that *all* voters check their vote is highly

unrealistic. In a realistic setting, it is more likely that a (small) fraction of honest voters perform the required tests while the others stop after casting their vote. We plan to explore how to adapt our definition to a quantitative setting, in the lines of [26].

ACKNOWLEDGMENTS

The authors would like to thank Esfandiar Mohammadi and the anonymous reviewers for their helpful comments and suggestions that greatly contributed to clarify the notions presented in the paper.

This work has been partially supported by the European Research Council (ERC) under the European Union's Horizon 2020 research (grant agreement No 645865-SPOOC).

REFERENCES

- [1] 2010. Délibération n° 2010-371 du 21 octobre 2010 portant adoption d'une recommandation relative à la sécurité des systèmes de vote électronique. French National recommendation on e-voting.
- [2] 2013. Ordonnance de la ChF sur le vote électronique (OVotE) du 13 décembre 2013 (Etat le 15 janvier 2014). Chancellerie fédérale ChF. Swiss recommendation on e-voting.
- [3] Martin Abadi and Cédric Fournet. 2001. Mobile Values, New Names, and Secure Communication. In *28th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL'01)*. ACM, 104–115. <https://doi.org/10.1145/360204.360213>
- [4] Ben Adida. 2008. Helios: Web-based Open-Audit Voting. In *17th USENIX Security Symposium (Usenix'08)*. 335–348.
- [5] Michael Backes, Catalin Hritcu, and Matteo Maffei. 2008. Automated Verification of Remote Electronic Voting Protocols in the Applied Pi-Calculus. In *21st IEEE Computer Security Foundations Symposium, (CSF 2008)*. 195–209.
- [6] J. Benaloh. 1987. *Verifiable secret-ballot elections*. Ph.D. Dissertation. Yale University.
- [7] David Bernhard, Veronique Cortier, David Galindo, Olivier Pereira, and Bogdan Warinschi. 2015. A comprehensive analysis of game-based ballot privacy definitions. In *Proceedings of the 36th IEEE Symposium on Security and Privacy (S&P'15)*. IEEE Computer Society Press, 499–516.
- [8] David Bernhard, Olivier Pereira, and Bogdan Warinschi. 2012. How Not to Prove Yourself: Pitfalls of the Fiat-Shamir Heuristic and Applications to Helios. In *Advances in Cryptology - ASIACRYPT 2012 (LNCS)*, Vol. 7658. Springer, 626–643.
- [9] David Bernhard and Ben Smyth. 2014. Ballot secrecy with malicious bulletin boards. *Cryptology ePrint Archive*, Report 2014/822.
- [10] Bruno Blanchet. 2016. Modeling and Verifying Security Protocols with the Applied Pi Calculus and ProVerif. *Foundations and Trends in Privacy and Security* 1, 1–2 (Oct. 2016), 1–135.
- [11] Ian Brightwell, Jordi Cucurull, David Galindo, and Sandra Guasch. 2015. An overview of the iVote 2015 voting system. Available at <https://www.elections.nsw.gov.au>.
- [12] Pyrros Chaidos, Veronique Cortier, Georg Fuchsbauer, and David Galindo. 2016. BeleniosRF: A Non-interactive Receipt-Free Electronic Voting Scheme. In *23rd ACM Conference on Computer and Communications Security (CCS'16)*. Vienna, Austria, 1614–1625.
- [13] Benoît Chevallier-Mames, Pierre-Alain Fouque, David Pointcheval, Julien Stern, and Jacques Traoré. 2010. On Some Incompatible Properties of Voting Schemes. In *Towards Trustworthy Elections 2010*. 191–199.
- [14] M. R. Clarkson, S. Chong, and A. C. Myers. 2008. Civitas: Toward a Secure Voting System. In *IEEE Symposium on Security and Privacy (S&P'08)*. IEEE Computer Society, 354–368.
- [15] Veronique Cortier, David Galindo, Stéphane Glondou, and Malika Izabachene. 2014. Election Verifiability for Helios under Weaker Trust Assumptions. In *19th European Symposium on Research in Computer Security (ESORICS'14) (LNCS)*, Vol. 8713. Springer, 327–344.
- [16] Veronique Cortier, David Galindo, Ralf Küsters, Johannes Müller, and Tomasz Truderung. 2016. SoK: Verifiability Notions for E-Voting Protocols. In *36th IEEE Symposium on Security and Privacy (S&P'16)*. San Jose, USA, 779–798.
- [] Veronique Cortier, Niklas Grimm, Joseph Lallemand, and Matteo Maffei. 2017. A type system for privacy properties. In *24th ACM Conference on Computer and Communications Security (CCS'17)*. ACM, Dallas, USA, 409–423.
- [] Veronique Cortier and Joseph Lallemand. 2018. *Voting: You Can't Have Privacy without Individual Verifiability*. Research Report. CNRS, Inria, LORIA. <https://hal.inria.fr/hal-01858034>
- [17] Veronique Cortier and Ben Smyth. 2013. Attacking and fixing Helios: An analysis of ballot secrecy. *Journal of Computer Security* 21, 1 (2013), 89–148.
- [18] Veronique Cortier and Cyrille Wiedling. 2012. A formal analysis of the Norwegian E-voting protocol. In *Proceedings of the 1st International Conference on Principles of Security and Trust (POST'12) (Lecture Notes in Computer Science)*, Vol. 7215. Springer, 109–128.
- [19] Stéphanie Delaune, Steve Kremer, and Mark D. Ryan. 2006. Coercion-Resistance and Receipt-Freeness in Electronic Voting. In *19th IEEE Computer Security Foundations Workshop (CSFW'06)*. IEEE Computer Society Press, Venice, Italy, 28–39.
- [20] Stéphanie Delaune, Steve Kremer, and Mark D. Ryan. 2009. Verifying Privacy-type Properties of Electronic Voting Protocols. *Journal of Computer Security* 17, 4 (2009), 435–487. <https://doi.org/10.3233/JCS-2009-0340>
- [21] David Galindo, Sandra Guasch, and Jordi Puiggali. 2015. 2015 Neuchâtel's Cast-as-Intended Verification Mechanism. In *5th International Conference on E-Voting and Identity (VoteID'15)*. 3–18.
- [22] Rop Gonggrijp and Willem-Jan Hengeveld. 2007. Studying the Nedap/Groenendaal ES3B Voting Computer: A Computer Security Perspective. In *USENIX Workshop on Accurate Electronic Voting Technology (EVT'07)*.
- [23] Sven Heiberg, Tarvi Martens, Priit Vinkel, and Jan Willemson. 2017. Improving the Verifiability of the Estonian Internet Voting Scheme. In *E-Vote-ID 2016 (LNCS)*, Vol. 10141. Springer, 92–107.
- [24] Steve Kremer, Mark D. Ryan, and Ben Smyth. 2010. Election verifiability in electronic voting protocols. In *15th European Symposium on Research in Computer Security (ESORICS'10) (LNCS)*, Vol. 6345. Springer. https://doi.org/10.1007/978-3-642-15497-3_24
- [25] Ralf Küsters, Tomasz Truderung, and Andreas Vogt. 2010. Accountability: Definition and Relationship to Verifiability. In *17th ACM Conference on Computer and Communications Security (CCS'10)*. 526–535.
- [26] Ralf Küsters, Tomasz Truderung, and Andreas Vogt. 2011. Verifiability, Privacy, and Coercion-Resistance: New Insights from a Case Study. In *32nd IEEE Symposium on Security and Privacy (S&P 2011)*. IEEE Computer Society, 538–553.
- [27] R. L. Rivest and W. D. Smith. 2007. Three Voting Protocols: ThreeBallot, VAV and Twin. In *USENIX/ACCURATE Electronic Voting Technology (EVT 2007)*.
- [28] Peter Roenne. [n. d.]. Private communication. ([n. d.]). The attack has been discovered by Peter Roenne and then described in the paper [?].
- [29] Peter Ryan. 2008. Prêt à Voter with Paillier encryption. *Mathematical and Computer Modelling* 48, 9–10 (2008), 1646–1662.
- [30] Drew Springall, Travis Finkenauer, Zakir Durumeric, Jason Kitcat, Harri Hursti, Margaret MacAlpine, and J. Alex Halderman. 2014. Security Analysis of the Estonian Internet Voting System. In *2014 ACM SIGSAC Conference on Computer and Communications Security*, Gail-Joon Ahn, Moti Yung, and Ninghui Li (Eds.). ACM, 703–715.
- [31] Scott Wolchok, Eric Wustrow, J. Alex Halderman, Hari K. Prasad, Arun Kankipati, Sai Krishna Sakhamuri, Vasavya Yagati, and Rop Gonggrijp. 2010. Security Analysis of India's Electronic Voting Machines. In *17th ACM Conference on Computer and Communications Security (CCS'10)*. Chicago, IL.
- [32] Scott Wolchok, Eric Wustrow, Dawn Isabel, and J. Alex Halderman. 2012. Attacking the Washington, D.C. Internet Voting System. In *Financial Cryptography and Data Security (FC'12)*.