

Measuring Vote Privacy, Revisited

David Bernhard
University of Bristol
Bristol, United Kingdom

Véronique Cortier
CNRS, Loria, UMR 7503
Vandoeuvre-lès-Nancy,
F-54500, France

Olivier Pereira
Université Catholique de
Louvain – ICTEAM
B-1348 Louvain-la-Neuve,
Belgium

Bogdan Warinschi
University of Bristol
Bristol, United Kingdom

ABSTRACT

We propose a new measure for privacy of votes. Our measure relies on computational conditional entropy, an extension of the traditional notion of entropy that incorporates both information-theoretic and computational aspects. As a result, we capture in a unified manner privacy breaches due to two orthogonal sources of insecurity: combinatorial aspects that have to do with the number of participants, the distribution of their votes and published election outcome as well as insecurity of the cryptography used in an implementation.

Our privacy measure overcomes limitations of two previous approaches to defining vote privacy and we illustrate its applicability through several case studies. We offer a generic way of applying our measure to a large class of cryptographic protocols that includes the protocols implemented in Helios. We also describe a practical application of our metric on Scantegrity audit data from a real election.

Categories and Subject Descriptors

H.1.1 [Models and Principles]: Systems and Information Theory—*Information Theory*; K.4.1 [Computers and Society]: Public Policy Issues—*Privacy*

General Terms

Measurement, Security, Theory

Keywords

Voting, Privacy, Entropy, Cryptography

1. INTRODUCTION

The design and analysis of voting systems has a long and rich history. Existing systems range from traditional paper-only ballot systems to purely electronic voting schemes where voters may vote from the privacy of their own computers (e.g., Helios [1, 2] or Civitas [11]) and also include hybrid systems that make use of

paper ballots but where computers facilitate the tally (e.g., Three-Ballot [28], Prêt-à-Voter [29] and Scantegrity [10]).

Of the many security properties that voting systems should satisfy, privacy of votes is of central concern. The development of rigorous security models for this important property started with the work of Benaloh [12, 3] but soon evolved towards the related (but seemingly much stronger) notions of receipt-freeness and coercion resistance [4, 24, 19]. Achieving these stronger notions is certainly desirable but it seems to come at the expense of efficiency and usability. It should then come as no surprise that systems in use (e.g. Helios) chose usability over coercion-resistance and aim to achieve “only” vote privacy. The study of this notion has only recently started to receive more attention with new models being developed in both symbolic models [13] and computational ones [21, 23, 5].

Models for related notions like confidential message transmission serve as a good but insufficient source of inspiration: unlike in those applications, vote privacy is not absolute but relative to specific election bylaws and voter choices. An extreme but relevant example is that of a voting system that discloses the number of votes received by each candidate. Such a system essentially reveals how each voter voted in the improbable but not impossible event that all voters vote for the same person. Yet classifying the system as insecure is clearly undesirable and one should search for a more nuanced classification.

A second difference concerns the information that the adversary tries to learn. In other scenarios where privacy is important it is usually clear what information the adversary targets (e.g. the privacy of plaintexts for the case of encryption). In contrast, adversaries against voting protocols may be interested in many possible targets ranging from how some individual voted, to complex relations between votes (e.g. have certain persons voted in the same way, or has a certain subset of voters supported a candidate more than another subset [2]). These two examples perfectly reflect the shortcomings of existing models for vote privacy which either target specific classes of protocols or are limited in the class of adversarial targets that they consider.

Contributions

In this paper we motivate and propose new privacy measures for voting schemes. We develop our definitions in two steps. The starting point is an information theoretic variant which, very roughly, declares the privacy of the information that the adversary targets to be the entropy left in the targeted information given what the adversary learns during the election process. This definition is too strong as it essentially requires security against unbounded adversaries and would declare practical systems where encryptions of votes are made public as completely insecure. In the second step,

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. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CCS'12, October 16–18, 2012, Raleigh, North Carolina, USA.
Copyright 2012 ACM 978-1-4503-1651-4/12/10 ...\$15.00.

we extend the applicability of our definition to systems where security is ensured only against computationally bounded adversaries. We do so by replacing information theoretic entropy with a conditional computational entropy which we introduce. We obtain a privacy measure that is a function of three parameters: the distribution \mathcal{D} on the votes of honest parties, the information that is the target of the adversary T , and the voting protocol π (and implicitly the tallying function which the voting protocol computes). We now discuss some of the features and shortcomings of our definition.

Simplicity. Entropy has long been identified and used as a natural measure of the uncertainty (and therefore privacy) of critical information [31]. Since our definition ultimately relies on entropy it inherits its associated intuition that has been used in many other contexts before (e.g., anonymity [8], leakage [32], information flow [32]). The move to computational entropy, although a bit technical, preserves this intuition and extends the applicability of the information theoretic approach to systems that employ cryptography.

Generality. The rather straightforward approach that we took in defining our privacy measure turns out to be quite powerful. We establish a link between a particular case of our notion and one recently proposed by Küsters, Truderung, and Vogt [23]: in the purely information-theoretic case, their notion can be obtained for a very specific distribution of the votes and by a particular choice of the parameters of our measure. We also show how a previously introduced computational notion [5] can be seen as a tool for moving from the computational to the IT version of our measure.

We keep our notion of privacy as independent as possible from the details of any fixed execution model. In particular, our definition requires and employs only an abstract notion of an adversarial view. Defining such a view is system/execution model dependent but on the one hand the details of how this is done are standard, and on the other such details are irrelevant to the understanding of our privacy measure and would only obscure our approach.

Usability. The general approach we took makes our measure of privacy applicable to a large variety of systems. The relaxation to computational entropies allows to meaningfully account for privacy breaches that are due to weak cryptography, the underlying use of entropy allows capturing insecure ways of releasing results, and the inclusion of an arbitrary distribution on the honest votes allows measuring the impact on privacy of some a priori knowledge of the votes the adversary may have. Our privacy measure could be used to identify systems that are clearly insecure (those for which the measure of privacy is always close to 0). However, we envisage its primary usage to be in comparing the level of privacy that different parameters and systems ensure. For example, consider a situation where one has to choose between different tallying systems, say one which reveals only the winner of an election, one where the number of votes each candidate obtain is revealed, or one that reveals a permutation of the individual ballots. While there is clearly a difference in the privacy offered by these options, our measure helps in understanding the different choices within the parameters of the election (e.g. the loss of privacy may strongly depend on the number of voters). Our approach can not only be applied to analyze a voting protocol but also to the data of a specific election. We discuss data from the 2009 Takoma Park election organized with the Scantegrity system [7]. The privacy of the Scantegrity protocol has been studied by Küsters et al. [22]; we do not redo such an analysis in this paper but we study privacy implications that arise from the specific results and audit data of one election:

- The audit data, which contain anonymized ballots, are enough to favor an Italian attack as soon as there are more than 2 candidates on a single question. Indeed, while almost all possi-

ble ballot fillings appear on the bulletin board for questions with two candidates, approximately 75% of the possible ballots do not appear for questions with three candidates.

- If someone obtains access to the receipts of voters, which are not expected to be kept confidential, then the privacy of those voters might be fully compromised. This is for instance the case for one specific voter in one of the six wards in this election.

We think that our observations could motivate some adaptations in the Scantegrity for future elections.

Definitional obstacles. Entropy is an appealing notion with a well-established place in defining privacy of communication. Its uses in the context of our work raises two specific obstacles that we needed to overcome. In general, computational entropies (involve an existential quantifier over an infinite class of simulators and) are difficult to compute. For systems that do not involve cryptography we show that our privacy notion collapses to an information theoretic entropy which can be computed in a similar manner as in previous work [23]. For systems that involve cryptography we provide a theorem which allows the following two-step approach. First, show cryptographic indistinguishability from an idealized system (essentially security in the sense of a computational model [5]). This technique is common in cryptography and has already been informally applied to voting protocols [21, 23]. In the second step we can simply use information-theoretic entropy.

A second difficulty is that there are many different flavors of entropy and it is not immediately clear which of the many well-established variants should form the basis of our notion. It turns out that there is no unique correct answer to this question as the different notions reflect related but different aspects of security (e.g. worst-case versus average-case insecurity). We study some of the most prominent choices and clarify the applicability of the different options through examples.

2. RELATED WORK AND LIMITATIONS

The quantification of privacy has been studied from many different angles. Our paper is concerned only with a small part of this research area: how to deal with cryptographic constructions in an analysis of privacy of voting systems. To place our paper in context, we begin by reviewing some work that is relevant to our area and point out the limitations of some existent approaches.

The work on privacy in the context of anonymous communication (see Chaum [9] and numerous subsequent work) is mostly focused on traffic analysis techniques. While some voting protocols might rely on anonymous communication to ensure privacy, this is not necessary, and it is often possible to have elections offering private voting without anonymous channels [12]. Besides, traffic analysis is often based on getting statistics from repeated or correlated events, while the task of secure function evaluation is a one-shot procedure (one is not able to repeat an election for instance). The work on database privacy, including the differential privacy approach [15] for example, is interested in limiting whether the addition or removal of a record in a database might affect the outcome of a statistic. While, at first sight, auctions or voting can be seen as taking a statistic on a set of records, the concerns are again different: we actually expect that one single record might completely change a statistic (by changing the highest bid for instance) and still want to consider that an auction or election may offer some level of privacy even though it does not provide any differential privacy.

The cryptographic literature also contains numerous examples of definitions of what it means for a protocol to offer privacy (see, e.g., Goldreich et al. [18]). Their work however concentrates on expressing that a protocol offers as much privacy as an ideal task, but

not on giving meaningful measures of privacy loss when a privacy gap exists.

Early work on vote privacy. The same criticism applies to early definitions of privacy proposed for voting by Benaloh [12, 3]. Furthermore, this work focuses on the comparison of honest vote assignments that offer the same sub-tally, which seems too restrictive for general tallying functions (e.g. if the tallying function only announces the winner). It also does not capture the information that an adversary can learn if it knows the expected distribution of the votes in advance. Privacy is also mentioned in several papers defining receipt-freeness [4, 24] or coercion-resistance [19, 33, 21]. While privacy informally seeks to hide information on votes, receipt freeness and coercion resistance are stronger properties that require that a voter cannot prove how he/she voted even if he is willing to (or forced to) do so. However, receipt freeness and coercion-resistance may simply not be satisfied for some systems that are designed for low-coercion environments such as Helios [1, 2].

Symbolic models. Vote privacy has also been defined in the context of symbolic models, where messages are represented by terms. For example in [13], a protocol is said to preserve privacy if an attacker cannot detect when two votes are swapped. This again does not apply to all voting scenarios. For example, if voters can give a score of 0, 1, or 2, it may be the case that an attacker cannot distinguish a pair of votes (0, 2) from (2, 0) but could well distinguish (0, 2) from (1, 1). Of course, symbolic models also have the usual drawback of being too abstract, possibly missing attacks that occur with only some probability (not 1 nor negligible).

Computational privacy notions. In [5], the authors define a game-based notion of privacy that is used to analyze the privacy of Helios [2]. Their definition however is tailored to a class of voting schemes where casting a vote corresponds to submitting an encrypted ballot to some bulletin board. This is of course not the case for all voting systems (see e.g. ThreeBallot [28]).

Küsters, Truderung, and Vogt provide a privacy definition for voting which is closest to the notion we propose [23]. Very roughly, they measure the difference observed by an attacker when an (honest) voter changes his vote, while all of the other (honest) votes follow a given distribution. The authors show how to analyze the privacy offered by several voting protocols such as ThreeBallot and VAV.

Their definition is the first that can be used to capture information-theoretic aspects (e.g. how much information can be inferred from the results of the election) but is limited in several respects. First, their definition is more restricted than ours with respect to the possible distributions on the honest voters' choices: every honest voter except one "target" voter casts a vote drawn independently from the same distribution. Secondly, they focus on wondering whether an adversary can decide whether a single voter voted in one way or another. These limitations exclude, for example, a scenario with two voters who always cast opposite votes in a yes/no ballot and an adversary who tries to tell which one voted "yes"; such a scenario appears in a privacy notion using symbolic models [13]. Thirdly, the following example explains why the definition of [23] is sometimes too strong (and then non informative) for some very natural cases.

Consider an election with one million voters doing approval voting on 100 candidates, and where the outcome of the election is a shuffled version of all ballots: this is the view of the election officers doing the tally in a traditional paper-based election for instance, or the view of everyone in an election relying on verifiable mixnets for the tally. For such an election, whatever distribution of votes is followed by honest voters, there will be at least two of

the 2^{100} possible votes, say v and v' , that appear with probability at most $2^{-80} \approx 10^6/2^{100}$. But then, an adversary who looks for the votes v and v' in the tally will be able to distinguish the cases where a voter submitted a vote for v or for v' with probability almost 1: the probability that v' or v gets submitted by a honest voter is close to 2^{-80} . What this example shows is that for elections where the so called *short ballot assumption* [28] does not hold (i.e. the number of possible votes is much larger than the total number of voters), the definition of [23] would declare a complete privacy breach, whereas this may not be the case. Indeed, it is clearly reasonable to consider that voters taking part to such an election do have some level of privacy: no one is able to decide which one of the one million submitted ballots comes from a specific voter. Our privacy notion aims to address these limitations. We provide a more detailed comparison with the notion of [23] in Section 6.2.

3. ENTROPY NOTIONS

As explained in the introduction the privacy measure that we consider relies on a computational version of conditional entropy. In this section we briefly review some standard notions of entropy and explain how to transfer the definitions from an information theoretic to a computational setting. Here and throughout the paper we denote random variables by boldface capital letters and use \mathbb{P} and \mathbb{E} to denote probability and expected value. We also use boldface capital letters for ensembles of random variables.

Information-theoretic entropy

For any $k \in [0, \infty]$ the Rényi entropy [26] with parameter k of a random variable \mathbf{X} with range \mathcal{X} is defined as:

$$\mathbb{H}_k(\mathbf{X}) := \frac{1}{1-k} \log \left(\sum_{x \in \mathcal{X}} \mathbb{P}[\mathbf{X} = x]^k \right)$$

(where the values 1 and ∞ are understood as limits). Three instances of k yield particularly useful, well-known notions of entropy: the case $k = 0$ is Hartley entropy, $k = 1$ is Shannon entropy and $k = \infty$ min-entropy.

For any fixed entropy notion \mathbb{H} there are several possibilities for defining the conditional entropy of one random variable given another (and each of these variants can be used as privacy measures). The *conditional entropy of \mathbf{X} given \mathbf{Y}* is defined as the *expected value* of the entropy of \mathbf{X} , where the expectation is taken over \mathbf{Y} :

$$\mathbb{H}(\mathbf{X} | \mathbf{Y}) := \mathbb{E}_{y \in \mathcal{Y}} [\mathbb{H}(\mathbf{X} | \mathbf{Y} = y)]$$

Note that $\mathbb{H}(\mathbf{X} | \mathbf{Y} = y)$ is *not* a conditional entropy: the expression over which the entropy is taken is simply a random variable.

The *minimal entropy of \mathbf{X} given \mathbf{Y}* is defined as the min value of the entropy \mathbf{X} , where the min is over all possible values of \mathbf{Y} :

$$\mathbb{H}^\perp(\mathbf{X} | \mathbf{Y}) := \min_{y \in \mathcal{Y}} \mathbb{H}(\mathbf{X} | \mathbf{Y} = y).$$

Finally, one can consider *average entropy*, similar in spirit to the above but computed by taking the expected value of an exponential function of the entropy under consideration and applying the logarithm to this result. In general we denote average entropy by $\tilde{\mathbb{H}}(\mathbf{X} | \mathbf{Y})$. Unlike for conditional and minimal entropy, there is no universally established way to define average conditional entropies, and existing notions are usually informed by the intended applications, or intuitive appeal. Dodis et al. [14] defined average min-entropy as:

$$\tilde{\mathbb{H}}_\infty(\mathbf{X} | \mathbf{Y}) := -\log \left(\mathbb{E}_{y \in \mathcal{Y}} \left[2^{-\mathbb{H}_\infty(\mathbf{X} | \mathbf{Y} = y)} \right] \right)$$

This measures the average probability of successfully guessing the value of x given y . Later in the paper (Section 6.2) we expand the above point and explain why this way of defining average conditional min-entropy extends the notion of min-entropy as maximal guessing probability to conditional random variables.

We propose an analogue for the notion of average Hartley entropy. Recall that Hartley entropy measures the size of the range of a random variable: if \mathbf{X} has range \mathcal{X} then $\mathbb{H}_0(\mathbf{X}) = \log |\mathcal{X}|$. Our average Hartley entropy measures the average size of the range of the random variable $(\mathbf{X} | \mathbf{Y} = y)$ where the average is taken over \mathbf{Y} .

$$\tilde{\mathbb{H}}_0(\mathbf{X} | \mathbf{Y}) := \log \left(\mathbb{E}_{y \in \mathcal{Y}} \left[2^{\mathbb{H}_0(\mathbf{X} | \mathbf{Y} = y)} \right] \right)$$

Note that the formula differs in the lack of minus signs from the one for average min-entropy. This is a technical but important detail. Each of these notions measures a different aspect of the ‘‘uncertainty’’ about the random variable \mathbf{X} when the value of the random variable \mathbf{Y} is available. We discuss this point in further detail in Section 4.

Conditional Privacy Measures

Each of the information-theoretic notions defined above can serve as a basis for our computational notion of privacy. It turns out that we can offer a unified treatment of the resulting possibilities by abstracting the many variants of entropy into a single notion which we call *conditional privacy measure* (the name is due to its application in defining privacy). The following definition sets two minimal and desirable conditions that such measures should satisfy.

DEFINITION 1. A conditional privacy measure $\mathbb{F}(\mathbf{T} | \mathbf{L})$ is a function mapping a pair of random variables to a positive real number and satisfying the following conditions. The names of the variables are mnemonics for the ‘‘target’’ of the adversary and the ‘‘leaked’’ information to which the adversary has access.

- If \mathbf{T} can be computed as a (probabilistic) function of \mathbf{L} then $\mathbb{F}(\mathbf{T} | \mathbf{L}) = 0$.
- If \mathbf{L} and \mathbf{L}' are two (probabilistic) functions such that \mathbf{L}' can be computed as a function of \mathbf{L} then for all \mathbf{T} , $\mathbb{F}(\mathbf{T} | \mathbf{L}') \geq \mathbb{F}(\mathbf{T} | \mathbf{L})$.

All of the different notions of conditional entropy satisfy these two conditions.

Computational Entropy

Information theoretic entropy quantifies uncertainty of random variables in face of unbounded adversaries. Computational entropy is the analogue notion for the case of efficient adversaries. Intuitively, a random variable has computational entropy if it is computationally close to one that has information theoretic entropy. In the rest of the section we abuse notation and refer to ensembles of random variables as random variables.

DEFINITION 2. Two random variables ensembles $\mathbf{X} = (\mathbf{X}_i)_{i \in \mathbb{N}}$ and $\mathbf{Y} = (\mathbf{Y}_i)_{i \in \mathbb{N}}$ are computationally close, written $\mathbf{X} \stackrel{c}{\approx} \mathbf{Y}$, if the quantity $|\mathbb{P}[A(\mathbf{X}_i) = 1] - \mathbb{P}[A(\mathbf{Y}_i) = 1]|$ is negligible as a function of i , for all polynomial-time algorithms A .

We propose a computational notion of conditional entropy, based on previous work by Reyzin et al. [27] and Gentry et al. [17]. Our definition is in the setting of asymptotic security as opposed to concrete and/or non-uniform adversaries. Furthermore, the application

to voting motivates a variation from the typical way of extending information theoretic entropies to computational versions. We discuss this variation after we give the definition.

DEFINITION 3. Let \mathbf{T} , \mathbf{R} and \mathbf{L} be ensembles of random variables to which we refer to as target, result, and leakage functions. Let \mathbb{F} be a conditional privacy measure.¹ We say that \mathbf{T} has at least r bits of computational conditional privacy given \mathbf{R} and \mathbf{L} (for which we write $\mathbb{F}^c(\mathbf{T} | \mathbf{R}, \mathbf{L}) \geq r$) iff $\exists \mathbf{S} = (\mathbf{S}_i)_{i \in \mathbb{N}}$ s.t.

- $(\mathbf{T}, \mathbf{R}, \mathbf{L}) \stackrel{c}{\approx} (\mathbf{T}, \mathbf{R}, \mathbf{S})$
- $(\forall i \in \mathbb{N}) \mathbb{F}(\mathbf{T}_i | \mathbf{R}_i, \mathbf{S}_i) \geq r$

To understand the above definition, it helps to think of \mathbf{T} as some sensitive information, \mathbf{R} as some information about \mathbf{T} that is certainly leaked, and \mathbf{L} as some information about \mathbf{T} that is cryptographically hidden. Informally, the above definition says that target \mathbf{T} has at least r bits of computational entropy given result \mathbf{R} and leakage \mathbf{L} if there is a distribution \mathbf{S} such that \mathbf{L} is computationally close to \mathbf{S} , even when \mathbf{T} and \mathbf{R} are known, and for any security parameter i the information theoretic entropy in \mathbf{T} given \mathbf{R} and \mathbf{S} is at least r .

The following example should help understanding the intuition behind computational conditional entropy and how we employ it later in the paper to measure vote privacy. Consider the encryption $Enc_{pk}(M)$ of some message M under a public key pk , and assume that M is selected from a distribution with non-zero entropy (say just 1 bit). Imagine that an adversary obtains in the execution of some system the encryption $Enc_{pk}(M)$ and some side information on M , say the XOR of its bits, $\oplus_i M_i$. In this situation, the information theoretic entropy left in M is 0 (as an unbounded adversary can decrypt the ciphertext and recover M). However, since for an efficient adversary the ciphertext looks like an encryption of a random message (assuming that the encryption scheme is secure) we would like to conclude that the loss of entropy in M is only due to revealing $\oplus_i M_i$. The definition above captures this intuition: the computational entropy in M given $\oplus_i M_i$ and $Enc_{pk}(M)$ is the (information theoretic) entropy of M given $\oplus_i M_i$ and the encryption $Enc_{pk}(R)$ of a random message (independent of M). This latter encryption plays the role of \mathbf{S} in our definition above: $(M, \oplus_i M_i, Enc_{pk}(M)) \stackrel{c}{\approx} (M, \oplus_i M_i, Enc_{pk}(R))$.

4. VOTE PRIVACY

In this section we introduce our measure of privacy. We start by fixing some necessary details regarding the execution model but leave others unspecified. For example, we do not enforce a particular communication infrastructure nor do we assume a particular communication model. We even abstract away many details of the adversarial model. The result is a flexible framework focused on those aspects that are essential for defining privacy. We then introduce our notion of privacy based on the details that we fix. The definition can then be easily instantiated for particular execution models/communication infrastructures etc.

Execution

We assume that voting involves a set of parties, some of which are under the control of the adversary. We do not make a distinction between voters and authorities. We write \mathcal{P} for the set of all parties and \mathcal{H} for the set of honest parties. Throughout the paper we let $n_{\mathcal{P}}$ be the total number of voters and $n_{\mathcal{H}}$ the number of honest voters. We write \mathcal{V} for the set of possible votes (including abstention).

¹Think about \mathbb{F} as some fixed information theoretic entropy notion.

A voting protocol is given by a set of interactive programs (processes), one for each party involved. Each program may use some secret information (e.g. signing keys that users use to authenticate, decryption keys that tallying authorities use to decrypt the result of the election) and the information that is publicly available. The programs for the voters also take as input a vote in \mathcal{V} , and we assume that these votes are selected according to a joint distribution \mathcal{D} on set $\mathcal{V}^{n\iota}$. We write π for a generic voting protocol, i.e. a description of the programs for the parties involved.

As discussed above we do not require a particular execution model as our definition and results do not depend on the model. We thus only assume that whatever the model one can formally define the information that an adversary obtains during the execution through its *view* of the execution. The view of the adversary is a standard cryptographic notion.

We define the view of the adversary as the (distribution of) his output; an adversary may output anything he chooses including the entire state and history of his execution and his random choices.

Clearly the view depends on the details of the execution model which formally specifies which channels are public/private what parts of the system can be corrupted, if corruption is static or adaptive, how does the adversary access the bulletin board (if any), how it accesses the result of the election, etc. We use the notation $\text{View}(A, \pi(\mathcal{D}))$ for the random variable that defines the view of the adversary A that interacts with the voting protocol π , when the votes of the honest participants are selected according to the distribution \mathcal{D} . By abusing notation we also write $\text{View}(A, \pi(\mathcal{D}))$ for the ensemble of random variables that define the view of the adversary for the different security parameters, and sometimes we simply write View when the various parameters are clear from the context. In a voting process that aims to compute a function ρ of the votes, this view includes the result of the election, which is the random variable (ensemble) $\mathbf{R}_{\mathcal{D}, v_A, \pi}$, where v_A is the distribution (ensemble) that represents the votes cast by corrupt parties.

Our Privacy Measure

In this section we motivate and define a measure for the privacy of votes in an election. We model privacy with respect to a *target function* T that models the information that the adversary is interested in. This is an important parameter of our definition as it allows modelling multiple scenarios of interest. Examples of potential adversarial targets include the vote of one, some, or even all voters. More complex information, e.g. whether two particular voters voted for the same candidate or whether a particular subset of voters supported a candidate more than others is also covered.

In addition to T , we aim to measure privacy along two dimensions: the distribution of the votes \mathcal{D} and the election protocol π . Extreme situations where \mathcal{D} contains no entropy or π simply reveals the vote of each participant entail no privacy. As soon as there is some uncertainty on how honest voters vote and these votes are somehow protected (e.g. cryptographically), then vote privacy clearly increases.

The intuition behind our definition is simple: we capture the privacy of the information targeted by the adversary $T(\mathcal{D})$ as the entropy left in the target given what the adversary learns from the voting process. The relation between the information in the target and the view of the adversary can be looked at from various angles. So we use the abstract notion of conditional privacy measure to encompass, succinctly, the different variants. We define two versions of privacy, against bounded and unbounded adversaries.

DEFINITION 4. *Let π be a voting protocol for result function ρ , \mathcal{D} a distribution on the honest votes, and T a target function. Let \mathcal{A} be the class of efficient adversaries, \mathcal{I} the class of unbounded*

adversaries, and \mathbb{F} be a computational privacy measure. The computational privacy $M(\mathcal{D}, T, \pi)$ is defined by

$$\inf_{A \in \mathcal{A}} \mathbb{F}^c(\mathbf{T}(\mathcal{D}) \mid \mathbf{R}_{\mathcal{D}, v_A, \pi}, \text{View}(A, \pi(\mathcal{D})))$$

Information theoretic privacy of $M^I(\mathcal{D}, T, \pi)$ is defined as

$$\inf_{A \in \mathcal{I}} \mathbb{F}(\mathbf{T}(\mathcal{D}) \mid \mathbf{R}_{\mathcal{D}, v_A, \pi}, \text{View}(A, \pi(\mathcal{D})))$$

Notice that the above definition in fact introduces a family of privacy measures $M_{\mathbb{F}}$, one for each fixed conditional privacy measure \mathbb{F} . Our intention is to let \mathbb{F} vary over the different existing notions of entropy and rely on their associated intuition to understand the guarantees entailed for the privacy of votes by the resulting measures. Indeed, we are convinced that evaluating our privacy measure for different entropy notions gives answers to different natural questions that voters might have about the confidentiality of their vote.

The Choice of Entropy Notions

We discuss here different variants based on different types of Rényi entropies (min, Hartley and Shannon) and different forms of conditional entropies (average, minimal and conditional).

For the sake of our discussion, we consider the following election example:

- The ballot takes the form of one question asking for approval on 100 choices (it can therefore be filled in 2^{100} ways).
- The distribution of the votes by the honest voters is uniform, except for a couple voters P_1 and P_2 , who vote as follows: with probability $1/2$, they agree on their choices before voting and vote exactly in the same way (one single uniform choice for both) but, if they disagree, then their choices are simply independent (uniform distribution on all pairs of distinct votes).
- The tallying function ρ reveals the vote of P_1 and P_2 if they are equal and reveals nothing otherwise.
- The target is the vote of P_1 .

Let us now consider the privacy of P_1 with for the vote distribution and the ρ function above.

Min-entropy based notions. A first natural question for P_1 is: “What is the probability that an observer will be able to guess my vote?” The answer to this question is given by using min-entropy, which provides a measure of the success probability of the best guess that an observer can make. If the election outcome ρ is not empty, which happens with probability $1/2$, the observer can make a correct guess with probability 1. Otherwise, the probability of success is 2^{-100} . Using average min-entropy as our measure, we get a measure of the success probability of approximately one bit: $\mathbb{H}_{\infty}(v_1 \mid \rho) = -\log(\frac{1}{2} \cdot 1 + \frac{1}{2} \cdot 2^{-100}) \approx 1$. This single bit of entropy is in line with the behaviour of our system: on average, an observer will be able to make a correct guess with probability just slightly higher than $\frac{1}{2}$. Now, if we use min-min-entropy as our measure, we get a measure of the success probability given the worst possible outcome from a privacy point of view, which happens when ρ is not empty: $\mathbb{H}_{\infty}^{\perp}(v_1 \mid \rho) = -\log(1) = 0$. This absence of entropy reflects that there is an election outcome for which v_1 has no privacy at all. Besides these two natural questions, the conditional min-entropy measures the average of the min-entropy on all possible outcomes: $\mathbb{H}_{\infty}(v_1 \mid \rho) = \frac{1}{2} \cdot 0 + \frac{1}{2} \cdot 100 = 50$. This last measure seems much less useful however.

Hartley entropy based notions. A second natural question for P_1 is: “In how many different ways can I pretend that I have voted?” The answer to this question is given using Hartley entropy and, in particular, by the min-Hartley entropy which gives a measure

of the minimum number of ways P_1 is guaranteed to be able to pretend that he has voted: $\mathbb{H}_0^\perp(v_1 | \rho) = \log(1) = 0$. This reflects the case where ρ reveals P_1 's choice. The average Hartley entropy gives a measure of the number of ways P_1 can expect to be able to pretend he voted: $\mathbb{H}_0(v_1 | \rho) = \log\left(\frac{1}{2} \cdot 1 + \frac{1}{2} \cdot 2^{100}\right) \approx 99$. This information of an average equivocation level of 2^{99} seems however less useful: while being correct, it hides the fact that, in half of the cases, there is no ambiguity left on the vote. Conditional Hartley entropy seems even less useful.

Another observation can be made of the min-Hartley entropy: just as the classic Hartley entropy, this notion does not require any probabilities: it only reflects the size of the smallest set in a set of sets. So, this measure can be used meaningfully even when one ignores the probability distribution of the honest votes, but only know which votes are possible.

Shannon entropy based notions. Shannon entropy measures the average number of extra bits of information that an observer would need to determine a vote (or any other target) given the election outcome. We should however keep in mind that Shannon entropy remains an average notion: the measure we get here is the average number of bits needed for the worst possible outcome. So, even if the minimal Shannon entropy is very high, it remains possible that some votes could be identified with a single extra bit of information.

The minimal Shannon entropy gives that measure for the worst possible outcome from a privacy point of view: $\mathbb{H}^\perp(v_1 | \rho) = 0$, which reflects the case where ρ is not empty. The conditional Shannon entropy gives the average case: $\mathbb{H}(v_1 | \rho) = \frac{1}{2} \cdot 0 + \frac{1}{2} \cdot 100 = 50$. This measure reflects that, in half of the cases, the choice of P_1 keeps 100 bits of Shannon entropy and 0 bits in the remaining cases. It however seems even less informative than the previous one, as it hides the worst case which happens with high probability. Average Shannon entropy seems also quite unnatural in our context.

Conclusion Our analysis of the example above shows that at least three entropy notions provide useful information:

- Average min-entropy measures, for a given distribution of the votes, the probability that an observer guesses the target function of the votes.
- Min min-entropy measures the probability that an observer guesses the target function of the votes for the worst possible election outcome that is in the support of the vote distribution.
- Min Hartley entropy measures the minimum number of values that the target function can take for any assignment of the votes.

This last notion can be particularly convenient when the distribution of the votes is unknown and as an indication of the level of deniability that voters can expect.

Other entropy notions might be of interest in specific cases.

5. DISCUSSION AND EXAMPLES

In this section, we test the robustness and meaningfulness of our privacy measures in several different ways.

We first demonstrate that, for ideal elections where all that the adversary sees are the choices made by the corrupted voters and the election outcome, our computational privacy measures coincide with the corresponding purely information theoretic measures and identify the worst adversary from a privacy point of view. We believe that this is an important sanity check. Furthermore we will also show, in the next section, the benefits of this property for the

modular analysis of voting systems that rely on computationally secure cryptographic primitives.

We next illustrate, through simple case studies, the impact on the privacy of the votes of several important parameters: the result function, which may correspond to the election outcome but also to the content of various audit data, and the ballot format and filling rules.

Finally, we perform an analysis of the audit data provided as part of the audit trail of the 2009 Takoma Park election and discuss some practical lessons that can be taken from this analysis.

For readability, from now on we use the following notational convention. We write M_k for a computational privacy measure when the underlying conditional privacy function is $\mathbb{H}_k(\cdot | \cdot)$ for some k . Furthermore, for average versions of the entropy functions, i.e. $\mathbb{H}_k(\cdot | \cdot)$, we write \tilde{M}_k for the resulting computational privacy measure. For example $\tilde{M}_\infty(\cdot, \cdot, \cdot)$ is the computational privacy measure we obtain if \mathbb{F} is the average conditional min-entropy \mathbb{H}_∞ .

Since \mathbb{F} captures the loss of privacy that revealing the results of the election entails, it is sufficient to understand the relation to the context of an idealized system where a trusted third party gathers the vote and publishes the results.

5.1 Ideal Protocols

Assume that ρ is an arbitrary function on \mathcal{V}^* . We define \mathcal{I}_ρ , the ideal process that computes ρ as follows. The process samples the honest votes according to \mathcal{D} and allows the adversary to cast votes on behalf of the remaining voters. The adversary then signals that it wants to receive the result and obtains $\rho(\vec{v})$, where \vec{v} is the list of all cast votes. Notice that we are tacitly assuming that ρ is such that the order in which the votes are cast does not matter. [23] analyzes the privacy offered by several protocols of this type. Similar idealizations are possible for complex result functions; one simply needs to specify more carefully how and when the votes are cast by the honest parties and the adversary.

The following theorem says that for ideal protocols, the privacy with respect to $M_\mathbb{F}$ is essentially the information-theoretic privacy of the function ρ . This is obtained by having the adversary cast votes from a distribution v_C^* that minimizes the information-theoretic entropy left in \mathcal{D} , given the result of the vote:

THEOREM 5. *Let ρ be an order independent function on \mathcal{V}^* . Let $v_C^* = \underset{v_A}{\operatorname{argmin}} \mathbb{F}(\mathbf{T}(\mathcal{D}) | \mathbf{R}_{\mathcal{D}, v_A, \mathcal{I}_\rho})$. Then,*

$$M_\mathbb{F}(\mathcal{D}, T, \mathcal{I}_\rho) = \mathbb{F}(\mathbf{T}(\mathcal{D}) | \mathbf{R}_{\mathcal{D}, v_C^*})$$

In other words, the conditional computational privacy of an ideal protocol is the minimum information theoretic entropy that can be obtained by setting the votes of the adversary. The proof of the theorem is in the full version of our paper.

5.2 The Role of the Result Function

We now give some examples on how the result function can influence the privacy measure. As we would intuitively expect, the more the result function reveals about the votes, the lower the level of privacy.

Consider a poll in which each voter may cast a single yes/no vote by submitting 1 or 0 to the trusted party in an ideal protocol. Fix the distribution \mathcal{D} to be as follows: let there be three voters and let every voter pick his vote uniformly at random. Let T be the vote of the first voter. We compare the privacy of the following result functions where $|\vec{v}|_0$ is the number of 0-votes in the vector \vec{v} of all votes cast.

$$\begin{array}{l|l} \rho_1 & c \text{ (const.)} \\ \rho_2 & 1 \text{ if } |\vec{v}|_1 \geq |\vec{v}|_0 \text{ else } 0 \\ \rho_3 & (|\vec{v}|_0, |\vec{v}|_1) \\ \rho_4 & \vec{v} \end{array}$$

Each of these result functions contains strictly more information than the previous one. The first is just a constant, the second is the majority vote (with 1 in case of a tie), the third is the number of 0- and 1-votes submitted and the fourth reveals all votes.

We consider the privacy measures based on the three entropy notions highlighted the previous section, namely the average min-entropy $\tilde{\mathbb{H}}_\infty$, the min-min-entropy M_∞^\perp , and the min-Hartley entropy M_0^\perp .

ρ	\tilde{M}_∞	M_∞^\perp	M_0^\perp
ρ_1	1	1	1
ρ_2	≈ 0.415	≈ 0.415	1
ρ_3	≈ 0.415	0	0
ρ_4	0	0	0

We give some interpretation of these results.

- Rows 1 and 4 of this table are obvious: if nothing about the target vote (that starts with one bit of entropy) is revealed, the entropy remains at 1 bit; if everything is revealed the conditional privacy drops to zero.
- $M_0^\perp(\rho_2) = 1$ occurs because for any of the two possible outcomes (majority vote is 0 or 1) the target voter could have cast a 0-vote or a 1-vote. In other words, the conditional probabilities $\mathbb{P}[\mathbf{T} = t \mid \mathbf{R} = r]$ (where \mathbf{R} is the random variable for the result) are nonzero for any (t, r) pair in their respective domains. We may observe that $M_0(\rho_2)$ has the same value.
- However, in the case of average-min and min-min entropies, the level of privacy decreases. This reflects the fact that, given any outcome, the probability that P_1 supported the winning candidate is $3/4$, resulting in an entropy of $0.415 \approx -\log 3/4$.
- In the case of ρ_3 , the min-min and min-Hartley entropies fall to 0. This is because the worst-case outcomes $(0, 3)$ and $(3, 0)$ determine the choice of P_1 .
- The average min-entropy does not decrease, though, which might be surprising. We can however observe that the probability of the adversary guessing the choice of P_1 remains equal to $3/4$: with probability $1/4$, the three candidates voted in the same way, giving a probability of guessing the choice of P_1 equal to 1, and with probability $3/4$, two candidates voted in one way and one voted in the other, giving a probability of $2/3$ of doing a correct guess given the outcome. We eventually observe that $1 \cdot 1/4 + 1/3 \cdot 3/4 = 3/4$.

All these results hence are in accordance with the intuition that we can get from those simple examples.

5.3 The Role of the Ballot Format

Ballot formats differ largely among elections: ballots can contain from one or two candidates up to a few hundred candidates and even offer the possibility to voters to nominate people who are not listed, they can require the voter to pick a single candidate, to pick up to a fixed number of candidates, to rank the candidates, and so on. The tallying rules can be very diverse as well.

One common approach for running verifiable elections, especially when ballots are complex, is to use mixnets [9, 30, 16, 7, 6]. All voters encrypt their vote, the encrypted ballots are shuffled by a series of mixers and then decrypted. This allows any observer to verify the result (as they can recompute it themselves from the

revealed votes) but the random order does not allow anyone to link votes to voters as long as at least one of the mixers is honest. However, mixnet-based tallying can destroy entropy if the number of choices a voter is presented with greatly exceeds the number of voters. We give an example.

Consider a poll in which each of 2^{10} voters is asked to answer n yes/no questions, whose answers are encrypted in a single ciphertext to be shuffled. The main factor that distinguishes mixnet-based tallying from the other main approach to cryptographic voting, namely homomorphic tallying, is that the ballots themselves can reveal relations between the answers to the various questions: for example, it may be possible to observe that the second question was more often answered with “yes” by voters who also said “yes” to the first question. This kind of information would not be deducible from the number of “yes” answers given to each question individually.

We consider the level of privacy offered by such an election as a function of the number n of questions. The level of privacy can clearly not exceed n : there are just 2^n choices. It might however be lower than n when the number of voters becomes smaller than 2^n : there will not be enough voters to make all possible choices, which will allow an observer to eliminate some of them. So, in our election, the level of privacy is also upper bounded by 10, since no more than 2^{10} choices can be made.

Assuming that the voters make their choices uniformly at random, we estimate our privacy measure based on average min-entropy for different values of n .

n	2	6	10	14	18	22
M_∞	1.9	5.3	7.5	8.7	9.1	9.9

This table shows that, for 2 and 6 choices, the measure of privacy is fairly close to optimal: an observer will not be able to guess the choice of a specific voter much better than by a random guess. When the number of choices increases, the measure of privacy progressively tends to 10. This corresponds to the fact that, when seeing the decrypted ballots, an observer can be convinced that any voter made one of the $\approx 2^{10}$ choices that appear after decryption instead of the 2^n choices that were initially possible. These $\approx 2^{10}$ choices still give 10 bits of privacy, a non-negligible measure, which is acknowledged by our measure.

Our measures of min-min-entropy and min-Hartley entropy will all provide an outcome equal to 0. There is indeed a non-zero probability that the 2^{10} voters make exactly the same choice, in which case they completely lose their privacy.

5.4 Analysis of the 2009 Takoma Park Election Data

The Scantegrity [10] voting system has been used in two public elections in Takoma Park, in 2009 [7] and in 2011. Scantegrity is a paper-based universally verifiable voting system. As part of the voting process, each voter is invited to fill in his paper ballot with a special pen: every time a voter marks a candidate, a confirmation code is unveiled, which the voter is invited to write on his receipt. When coming back home the voter has then the possibility to check the presence of the receipt codes he noted on a public bulletin board. The design of the system is expected to guarantee that it is not possible to link any specific code to a particular voter choice.

The tallying procedure then proceeds to a verifiable shuffle of the ballots, and the shuffled ballots are made available on the bulletin board as well, for audit purposes. We note that this mixnet-based tallying procedure provides voters with more information than usual: all voters can now see all shuffled ballots, while they are

usually only able to see the final election outcome only (expressed in terms of number of votes for each candidate), unless they are part of the tallying officers.

The Takoma Park elections have another specifics: they are based on instant-runoff voting (IRV), that is, voters are invited to rank any number of candidates on their ballot (including a write-in position). A first side-effect of IRV is that there is a very large number of ways to fill in a ballot, even for a relatively small number of candidates: for n candidates, there are $\sum_{i=0}^n \frac{n!}{(n-i)!}$ valid ways of filling a ballot. Furthermore, if we take into account the fact that voters are not forced to respect the ranking rule (e.g., they can produce ballots with several candidates ranked first, or they can skip positions in their ranking), this number grows to $(n+1)^n$. So, even for relatively small values of n , we can expect that specific choices that are possible in theory, given the election outcome, will be ruled out just by looking at the bulletin board. A second side-effect of IRV is that not all voters will rank the same number of candidates. As a result, a voter’s receipt actually leaks some information on the choices of that voter: the number of choices he made.

So, there seem to be three natural levels of privacy to compare in such elections:

- View 1 The amount of privacy that voters have when the only information revealed is the election outcome, that is, the privacy obtained with respect to a non-tallying voter if the Scantegrity audit trail is not available.
- View 2 The amount of privacy that voters have when they also see the bulletin board, which is the actual view of the voters for the Takoma Park elections using Scantegrity.
- View 3 The amount of privacy that voters keep if they show their receipt to someone else. We stress that receipts are not supposed to be kept secret.

We select our measure $M(\mathcal{D}, T, \pi)$ with privacy measure \mathbb{F} as follows:

- Our target T will be the exact vote of an individual voter. Other choices could have been made as well, like the first choice of any individual voter, but we keep targeting the full vote as it is a more challenging target.
- Since we do not have any a priori knowledge about the distribution of the votes, we consider any distribution whose support contains all possible vote assignments that is consistent with the view of the adversary, and use Hartley entropy as our privacy measure.
- To perform our analysis, we use the audit data provided for the 2009 Takoma Park election.² Since the view of the adversary is fully determined by these audit data, the choice between average, conditional and min Hartley entropy is irrelevant.

The result of our analysis appears in Table 1. Privacy measures appear for the answers to the two questions (0 and 1) that were submitted to the voters in each of the 6 wards of this election. We see that we have a fairly high level of entropy when the adversary only sees the election outcome: 6 bits for a question with 3 choices, 3.17 bits for a question with 2 choices. The level is just a bit lower than ideal for Question 0 in Ward 4, as one of the candidates was not ranked first by anyone there.

The level of privacy substantially decreases in many cases when the bulletin board becomes available: viewing the submitted ballot allows ruling out a lot of potential combinations. While this does not seem alarming from a pure privacy point of view, this decrease of privacy might be sufficient to force a voter to submit a ballot with an unusual ranking of the candidates, and this voter might

legitimately fear that nobody else will submit a ballot with the same choices if he does not do it. This is a well-known issue in elections in which ballots can be filled in many ways, traditionally called the Italian attack. Our measure shows, considering the actual votes, that this attack could be fairly effective as soon as a question has 3 candidates. This quantifies the bound related to the short ballot assumption [28] used in similar contexts.

Now, if the adversary sees the receipt of specific voters, the level of privacy substantially decreases again. Here, we display the minimum level of privacy depending on the number of codes appearing on the receipt. In particular, we see that we do not have any entropy left for at least one voter of Ward 5. Checking the audit data, it indeed appears that only one voter in that ward ranked all 3 candidates. This voter ranked the write-in candidate first, Candidate 1 second and Candidate 2 as third.

We believe that this potential loss of privacy indicates that voters should be encouraged to keep their receipt secret and suggests that a mechanism should be put in place to prevent information from leaking from the receipts. One possibility would be to have dummy codes available on the ballots, that would be offered for completion by the voters, allowing all voters to obtain the same number of unveiled codes on their receipts.

We stress that, even though it uses data from a real election, our analysis may not completely reflect the actual votes, and there might be in reality more than one voter who ranked all 3 candidates in Ward 5. Indeed, the audit data do not exactly match the published election results (but the differences are small and both tallies agree on the identity of the winners.)

6. RELATION WITH TWO PREVIOUS PRIVACY DEFINITIONS

We show that our framework is sufficiently general to capture two existing and quite different notions of privacy based on cryptographic indistinguishability. For lack of better names, we refer to the notion of Bernhard et al. [5] as the *game-based* notion and to the notion of Küsters, Truderung, and Vogt [23] as δ -privacy (although both these notions involve cryptographic games); we call our own notion that we introduce in this paper *entropy-based privacy*.

6.1 Comparison with the game-based notion

Game-based privacy [5] (our terminology) is a privacy notion for a class of cryptographic voting protocols inspired by cryptographic security for encryption. Bernhard et al. [5] analyze the Helios [1] voting protocol and show that with an encryption scheme meeting certain requirements, Helios meets their security notion. In this section we introduce their model and state a theorem that any scheme secure in the game-based model is as good as the ideal protocol for the same election parameters according to our privacy measure, for any choice of conditional privacy measure.

As a corollary, we get that Helios offers the same level of privacy (using our measure) as an ideal protocol. We stress that we do not need to perform a detailed analysis of Helios to get this result, as we can build on the result that Helios offers game-based privacy.

Single-Pass Voting

In single-pass voting, we consider a set \mathcal{P} of voters, a subset $\mathcal{H} \subseteq \mathcal{P}$ of honest voters, a set of administrators and a bulletin board B . All parties may post messages to the bulletin board at any time; the board decides to accept or reject a message based on its current state and a public algorithm π_B . A single-pass protocol π_ρ for a result function ρ executes in three phases. The board stores all

²See <https://scantegrity.org/svn/data/>.

Ward	1		2		3		4		5		6	
#Ballots	470		277		481		212		85		198	
Question	0	1	0	1	0	1	0	1	0	1	0	1
Entropy for View 1	6	3.17	6	3.17	6	3.17	5.95	3.17	6	3.17	6	6
Entropy for View 2	4.32	3	4.32	2.81	4.09	3	3.17	3	3.17	2.58	3.91	3.58
Entropy for View 3	1.58	1.58	2	1.58	2	1.58	1	1.58	0	1	2	1.58

Table 1: Privacy measures for the Takoma Park 2009 election

accepted messages and any party may read the board at any time to obtain the current phase and the list of all accepted messages and their senders, in the order that they were posted. The board starts out empty and in the *setup* phase.

The name *single-pass* comes from the observation that voters only have to post a single message to the bulletin board to cast their vote and need take no further part in the election.

1. In the *setup* phase, the board expects one message from each administrator (in any order) after which it may either transition to the *voting* phase or abort.
2. In the *voting* phase, the board accepts one message from each party in \mathcal{P} , in any order. After receiving these messages, the board transitions to the *result* phase.
3. In the *result* phase, the board expects one message from each administrator; after receiving these it halts (accepts no more messages but can still be read).

Correctness. We need the following correctness assumption on the execution of a single-pass protocol which intuitively says that if everyone acts correctly, the protocol indeed computes ρ on the votes submitted.

A protocol π_ρ is correct for result function ρ if there is an efficient algorithm that computes either a (claimed) result r in the range of ρ or a symbol \perp to denote failure from a board that has halted successfully. We define the *result of a board* as being the output of this algorithm on the board and say the board is *valid* if the result is not \perp . Furthermore, if all parties execute the protocol correctly then the result r of the board is the correct election result, i.e. $r = \rho((v_P)_{P \in \mathcal{P}})$ where v_P is the vote cast by voter P .

Extractibility. In addition to correctness, we require that all administrators together can extract votes from individual ballots. This is a technical point that is required in some security proofs. Although slightly stronger than the original model of [5], extractibility is satisfied by all voting protocols that we know of, in particular Helios [1].³

Game-based privacy

Consider a game between a challenger C who controls the administrators, bulletin board and honest voters and an adversary A who controls the set of corrupt voters. The adversary may in addition choose the votes of the honest voters, even adaptively:⁴ during the voting phase the adversary may both submit ballots on behalf of

³Helios has voters encrypt their votes under a key shared among the administrators. In a real execution of the Helios protocol, administrators never decrypt individual votes and as long as at least one administrator is honest, no-one can extract votes. Extractibility does not weaken vote privacy.

⁴In the original game [5] the adversary can even choose adaptively which voters are honest or corrupt. Security w.r.t. the original notion implies our version, which is sufficient to show the relationship with entropy.

his own, corrupt voters and inform the challenger that he wishes a certain honest voter to vote in a way of his choice.

The challenger chooses a random bit β at the beginning of the protocol. If $\beta = 0$, the challenger and adversary just run π together. If $\beta = 1$, when the adversary asks a honest voter to vote, the challenger notes this vote but has the voter submit a ballot for a fixed value v_ϵ instead. In the result phase, the challenger always announces to the adversary the correct result based on the ballots of the dishonest voters and the votes that the adversary chose for the honest voters.⁵

The adversary’s goal is to guess the challenger’s bit β . A protocol has game-based privacy if no efficient adversary can guess β with probability non-negligibly better than $1/2$.

Relation between entropy-based and game-based privacy

Security in the sense of the game described above leads to the following intuitive argument: since the adversary cannot distinguish ballots containing the true votes of the honest users — even if the adversary can choose these votes — from ballots containing a fixed, constant vote then the adversary cannot extract any information about the votes from ballots. Therefore, we expect that the conditional computational entropy of the honest votes given the bulletin board (which together with the random coins of the adversary form its view of the execution) to be equal to the computational entropy that the votes have on their own (the adversary can see the result of the result function on the votes in both cases). This intuition is formalized by the following theorem.

THEOREM 6. *Suppose that π is a correct and secure single-pass voting protocol for a set \mathcal{P} of voters and a result function ρ . Fix any set $\mathcal{H} \subseteq \mathcal{P}$, any efficiently samplable distribution \mathcal{D} on \mathcal{H} and any target function T . Then*

$$M(\mathcal{D}, T, \pi) = M(\mathcal{D}, T, \mathcal{I}_\rho) .$$

The proof is in the full version of our paper.

Discussion

We have shown that, informally speaking, a single-pass protocol for computing some tally function ρ which is secure in the game-based sense achieves the same level of privacy as the ideal protocol (see Section 6.1) for ρ , in the sense of entropy-based privacy. One immediate interpretation of this result is that from a protocol secure in the game based sense the adversary can extract as much information as it can extract from only seeing the result and nothing more. This is a very desirable and intuitively appealing property as it reduces understanding the level of privacy of a protocol to that of understanding the level of privacy of a corresponding ideal protocol. Furthermore proofs using the game-based definition are

⁵If $\beta = 1$ then the announced result will not match the ballots on the board. It is part of the game-based definition that the challenger can produce such a “fake” result without the adversary noticing whereas in a real election, faking a result should be infeasible.

more standard and easier to construct than those using the computational conditional entropy notions. Whenever possible, proving game-based privacy is therefore desirable.

There are several downsides to the game-based notion. The notion does not account for the loss of privacy due to the result function ρ . A voting protocol where ρ is the identity function and which is implemented by letting the tallying authorities to simply decrypt the ballots and output the votes in clear would be secure in the game-based sense. Clearly however privacy of the votes is actually lost. The theorem we gave in this section allows for the following methodology: prove game-based security for a voting protocol and then analyze the privacy of the ideal protocol for ρ to understand the entropy-based privacy of the overall protocol.

This route is not possible for protocols that are not in the single-pass class, or for protocols that are single-pass but where some “little” (potentially useless parts) information about the honest votes is revealed. In both cases our entropy-based notion may still apply and in the latter case, it allows for giving a more refined quantification of the loss of privacy, than simply declaring the protocols insecure.

6.2 Relation with δ -privacy

In this section we establish a relation between the notion of privacy introduced by Küsters, Truderung, and Vogt [23] at IEEE S&P 2011 and the privacy measure that we introduce in this paper.

Definition of δ -privacy

In the model for δ -privacy, the adversary’s target is the vote of a single “voter under observation”. This vote is restricted to being one of two votes. All other voters cast votes according to some fixed distribution.⁶ The adversary interacts with the protocol by controlling a set of dishonest parties (voters and/or authorities). At the end of the execution, the adversary outputs a bit to indicate which of the two possible target votes he thinks the voter under observation has cast. The system offers δ -privacy if the adversary can distinguish between the two situations with probability no better than δ (so the smaller δ the better).

The adversary is quantified over an abstract set of “observer processes” which is either the set of bounded computations or that of unbounded computations. This dichotomy gives rise to two distinct privacy notions, one computational and the other one information theoretic. Below, we call the set over which the adversaries is quantified *admissible adversaries*.

We restate the definition of privacy definition of [23] but use slightly different notation. Without loss of generality we assume that the voter under observation is the first voter. We write \mathcal{D}_j for the distribution where the vote of the first voter is set to j and the votes of all other parties are selected according to \mathcal{D} . We write $A^\pi(\mathcal{D})$ for the random variable (ensemble) that describes the output of an adversary A interacting with protocol π when the honest votes are selected according to \mathcal{D} . δ -privacy of π requires that no adversary can tell if the first voter votes for i or votes for j except with probability δ (no matter what i and j are). We refer to an adversary for this experiment as a distinguishing adversary.

DEFINITION 7 ([23]). *A protocol π achieves δ -privacy if for any admissible distinguishing adversary A and any two distinct*

⁶In fact, the assumption in [23] is that everyone else votes independently and identically distributed according to some probability vector \vec{p} on the possible votes but this can be easily generalised. In particular, one may consider a joint distribution \mathcal{D}' on all other honest voters.

votes i and j there exists a negligible function ν such that

$$\mathbf{Adv}_{A,i,j}^{\text{dist}}(k) = \mathbb{P}[A^\pi(\mathcal{D}_i) = 1] - \mathbb{P}[A^\pi(\mathcal{D}_j) = 1] \leq \delta + \nu(k).$$

A protocol is exactly δ -private if it achieves δ -privacy and does not achieve δ' -privacy for any $\delta' < \delta$. We write $\delta(\pi, \mathcal{D})$ for the exact level of privacy achieved by protocol π when the honest votes are selected according to \mathcal{D} . By a slight abuse of notation, we write $\delta(\pi, \mathcal{D}_{ij})$ for the maximum level of privacy achieved if i and j are fixed and only the adversary is allowed to vary.

Relation between entropy-based privacy and δ -privacy

First, we link δ -privacy with the ability of any adversary to correctly guess the vote of the first voter, if this vote is either i or j . Specifically, let the distribution \mathcal{D}_{ij} be such that the vote of the first voter is selected uniformly at random from the set $\{i, j\}$, and the votes of the remaining honest voters are selected according to \mathcal{D} . Consider a modified version of the experiment that defines δ -privacy. In this modified version the votes of the honest voters are distributed according to \mathcal{D}_{ij} and the goal of the adversary is to output a guess $g \in \{i, j\}$ as to what the vote of the first voter is. We refer to such an adversary as a guessing adversary. The adversary wins if it guesses correctly, i.e. we define its guessing advantage as $\mathbf{Adv}_{A,i,j}^{\text{guess}} = \mathbb{P}[A^\pi(\mathcal{D}_{ij}) = v_1]$, where $v_1 \in \{i, j\}$ is the vote cast by the first voter in the execution. The following lemma establishes a well-known relation between winning a “distinguishing game” and guessing the value to be distinguished.

LEMMA 8. *Let i, j be fixed. For any admissible (distinguishing) adversary D there exists an admissible (guessing) adversary G_D such that:*

$$\mathbf{Adv}_{G_D,i,j}^{\text{guess}}(k) \geq \frac{1}{2} \cdot (\mathbf{Adv}_{D,i,j}^{\text{dist}}(k) + 1).$$

Conversely, for any admissible (guessing) adversary G there exists an admissible (distinguishing) adversary D_G such that:

$$\mathbf{Adv}_{D_G,i,j}^{\text{dist}}(k) \geq 2 \cdot \mathbf{Adv}_{G,i,j}^{\text{guess}}(k) - 1.$$

One implication of the above lemma is that for protocols which achieve δ -privacy only for large δ there are adversaries that are successful in guessing the vote of the first party. This relation between δ -privacy and guessing abilities suggests a link between δ -privacy and a privacy measure that captures the ability of the adversary to guess a target vote.

We make this intuition formal by instantiating our privacy measure in a particular way. We set the distribution on the votes to be \mathcal{D}_{ij} (the vote of the first voter is i or j with probability a half), we set the target function to T_1 , the function that returns the vote of the first voter, and set the underlying entropy measure to be $\tilde{\mathbb{H}}_\infty$ the average conditional min-entropy (that measures precisely the ability of the adversary to guess the target). We thus relate δ -privacy with $\tilde{M}_\infty(\mathcal{D}_{ij}, T_1, \pi)$, where i, j are the votes of the first voter that yield that largest possible δ . The relation between guessing probability and average min-entropy is the following. Consider \mathbf{T} and \mathbf{L} two (possibly correlated) random variables. Then for any adversary A , the probability that A guesses \mathbf{T} given that it sees \mathbf{L} is $\mathbb{P}[A(\mathbf{L}(\mathbf{T})) = \mathbf{T}]$, which is the same as:

$$\sum_l \mathbb{P}[\mathbf{L} = l] \cdot \mathbb{P}[A(\mathbf{L}(\mathbf{T})) = \mathbf{T} \mid \mathbf{L} = l] = \mathbb{E} \left(2^{-\tilde{\mathbb{H}}_\infty(A(\mathbf{L}(\mathbf{T}))=\mathbf{T}|\mathbf{L}=l)} \right) = 2^{-\tilde{\mathbb{H}}_\infty(\mathbf{T}|\mathbf{L})}$$

We use this connection to relate δ -privacy with $\tilde{M}_\infty(\mathcal{D}, T, \pi)$ for the case of bounded adversaries, and with $\tilde{M}_\infty^I(\mathcal{D}, T, \pi)$ for unbounded ones. For unbounded adversaries the connection is made

precise by the following theorem which says that δ -privacy in the sense of [23] is equivalent to privacy as captured by $\tilde{M}_\infty^I(\mathcal{D}_{ij}, T_1, \pi)$. (The proof is in the full version of our paper.)

THEOREM 9. *Let i, j arbitrary votes. For unbounded adversaries $\delta(\pi, \mathcal{D}_{ij}) = 2^{1-\tilde{M}_\infty^I(\mathcal{D}_{ij}, T_1, \pi)} - 1$*

Perhaps unsurprisingly, for bounded adversaries we were able to prove a connection only in one direction. Specifically, we argue that if a protocol is not δ -private (i.e. there exists i, j votes for the first voter, and a distinguishing adversary with advantage larger than δ), then our computational privacy measure $\tilde{M}_\infty(\mathcal{D}_{ij}, T_1, \pi)$ is also upperbounded appropriately. We omit the proof due to space limitations.

THEOREM 10. *Let π be an arbitrary protocol, T_1 the target function that returns the vote of the first voter and \mathcal{D} a distribution on the honest votes. Then for any i, j :*

$$\delta(\pi, \mathcal{D}_{ij}) \leq 2^{1-\tilde{M}_\infty(\mathcal{D}_{ij}, T_1, \pi)} - 1.$$

The following corollary (obtained by setting i and j appropriately) makes the relation between δ -privacy and \tilde{M}_∞ precise.

COROLLARY 11. *Let π be an arbitrary protocol that is not δ -private. Then there exists i and j such that*

$$\tilde{M}_\infty(\mathcal{D}_{ij}, T, \pi) \leq 1 - \log(1 + \delta)$$

Discussion

The results of this section show that security in the sense captured by one instantiation of our privacy notion implies security in the sense defined by δ -privacy. Furthermore, for information theoretically secure protocols (e.g., ideal ones) and for specific distributions the two notions coincide. We think that the relations exhibited in this section between our entropy-based notion, the cryptographic notion of [5] and δ -privacy [23] support all three notions and their respective approaches to privacy. Our privacy measure allows us to make statements about privacy in cryptographic voting protocols that would be much harder to establish using the game-based model of [5] directly. Our measure also applies to a more general class of protocols, vote distributions and targets than those that have been studied previously using δ -privacy.

7. CONCLUSION

Entropy is a natural choice to measure privacy in an information-theoretic setting and we demonstrate how different formulations of conditional entropy answer different intuitive questions about vote privacy. Through an appropriate notion of computational conditional entropy we have extended the reach of this idea to the computational setting and have established a theorem that enables accurate analysis of privacy offered by complex cryptographic voting protocols while simply disregarding the details of their implementation. Furthermore, the underlying entropy-based approach makes our measure applicable to non-cryptographic protocols and we have shown through the Takoma Park example how to obtain meaningful results for a real election. We completed the investigation of our notion by establishing powerful connections with two existing privacy notions for votes [23, 5].

As our definition does not concentrate on any specific election rule or on any specific target function, in future work we plan to explore if and how it can be applied to related problems, e.g., sealed-bid auctions [25, 20] or more generally any secure function evaluation problem.

Acknowledgments

The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement no 258865, project ProSecure, under the ICT-2007-216676 European Network of Excellence in Cryptology II, under the HOME/2010/ISEC/AG/INT-011 project B-CENTRE, and under the SCOOP Action de Recherche Concertées. Olivier Pereira is a Research Associate of the F.R.S.-FNRS.

8. REFERENCES

- [1] B. Adida. Helios: Web-based Open-Audit Voting. In *17th USENIX Security Symposium*, pages 335–348, 2008. Helios website: <http://heliosvoting.org>.
- [2] B. Adida, O. de Marneffe, O. Pereira, and J.-J. Quisquater. Electing a University President Using Open-Audit Voting: Analysis of Real-World Use of Helios. In *Electronic Voting Technology Workshop/Workshop on Trustworthy Elections*. Usenix, Aug. 2009.
- [3] J. Benaloh. Verifiable secret-ballot elections. Technical Report 561, Yale University Department of Computer Science, September 1987.
- [4] J. Benaloh and D. Tuinstra. Receipt-free secret-ballot elections. In *26th ACM Symposium on Theory of Computing*, pages 544–553, 1994.
- [5] D. Bernhard, V. Cortier, O. Pereira, B. Smyth, and B. Warinschi. Adapting helios for provable ballot secrecy. In Springer, editor, *16th European Symposium on Research in Computer Security (ESORICS'11)*, volume 6879 of LNCS, 2011.
- [6] P. Bulens, D. Giry, and O. Pereira. Running mixnet-based elections with Helios. In *Electronic Voting Technology Workshop/Workshop on Trustworthy Elections*. Usenix, 2011.
- [7] R. Carback, D. Chaum, J. Clark, J. Conway, A. Essex, P. S. Herrnson, T. Mayberry, S. Popoveniuc, R. L. Rivest, E. Shen, A. T. Sherman, and P. L. Vora. Scantegrity II Municipal Election at Takoma Park: The First E2E Binding Governmental Election with Ballot Privacy. In *USENIX Security Symposium*, pages 291–306. USENIX Association, 2010.
- [8] K. Chatzikokolakis, C. Palamidessi, and P. Panangaden. Anonymity protocols as noisy channels. *Information and Computation*, 2-4(206):378–401, 2008.
- [9] D. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2):84–88, February 1981.
- [10] D. Chaum, A. Essex, R. Carback, J. Clark, S. Popoveniuc, A. Sherman, and P. Vora. Scantegrity: End-to-End Voter-Verifiable Optical-Scan Voting. *IEEE Security and Privacy*, 6(3):40–46, 2008.
- [11] M. R. Clarkson, S. Chong, and A. C. Myers. Civitas: Toward a Secure Voting System. In *29th Security and Privacy Symposium (S&P'08)*. IEEE, 2008.
- [12] J. Cohen (Benaloh) and M. Fischer. A robust and verifiable cryptographically secure election scheme. In *26th Symposium on Foundations of Computer Science.*, pages 372–382, Portland, OR, 1985. IEEE.
- [13] S. Delaune, S. Kremer, and M. D. Ryan. Verifying privacy-type properties of electronic voting protocols. *Journal of Computer Security*, 17(4):435–487, 2009.

- [14] Y. Dodis, R. Ostrovsky, L. Reyzin, and A. Smith. Fuzzy Extractors: How to Generate Strong Keys from Biometrics and Other Noisy Data. *SIAM Journal of Computing*, 38(1):97–139, 2008.
- [15] C. Dwork. Differential privacy. In *Automata, Languages and Programming, 33rd International Colloquium, ICALP 2006*, volume 4052 of *LNCS*, pages 1–12. Springer, 2006.
- [16] J. Furukawa, K. Mori, and K. Sako. An implementation of a mix-net based network voting scheme and its use in a private organization. In *Towards Trustworthy Elections*, volume 6000 of *LNCS*, pages 141–154. Springer, 2010.
- [17] C. Gentry and D. Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In *43rd ACM Symposium on Theory of Computing*, pages 99–108, 2011.
- [18] O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game: A completeness theorem for protocols with honest majority. In *19th Annual ACM Symposium on the Theory of Computing (STOC)*, pages 218–229. ACM Press, 1987.
- [19] A. Juels, D. Catalano, and M. Jakobsson. Coercion-Resistant Electronic Elections. In *4th Workshop on Privacy in the Electronic Society (WPES 2005)*, pages 61–70. ACM, 2005.
- [20] A. Juels and M. Szydlo. A two-server, sealed-bid auction protocol. In *6th international conference on Financial cryptography (FC'02)*, pages 72–86. Springer, 2003.
- [21] R. Küsters, T. Truderung, and A. Vogt. A Game-Based Definition of Coercion-Resistance and its Applications. In *23rd IEEE Computer Security Foundations Symposium (CSF'10)*, pages 122–136. IEEE, 2010.
- [22] R. Küsters, T. Truderung, and A. Vogt. Proving Coercion-Resistance of Scantegrity II. In *12th International Conference on Information and Communications Security (ICICS 2010)*, volume 6476 of *LNCS*, pages 281–295, 2010.
- [23] R. Küsters, T. Truderung, and A. Vogt. Verifiability, Privacy, and Coercion-Resistance: New Insights from a Case Study. In *IEEE Symposium on Security and Privacy (S&P 2011)*, pages 538–553. IEEE Computer Society, 2011.
- [24] T. Moran and M. Naor. Receipt-Free Universally-Verifiable Voting with Everlasting Privacy. In *26th International Cryptology Conference (CRYPTO'06)*, volume 4117 of *LNCS*, pages 373–392. Springer, 2006.
- [25] M. Naor, B. Pinkas, and R. Sumner. Privacy preserving auctions and mechanism design. In *1st ACM conf. on Electronic Commerce*, 1999.
- [26] A. Rényi. On measures of information and entropy. In *4th Berkeley Symposium on Mathematics, Statistics and Probability*, pages 547–561, 1960.
- [27] L. Reyzin. Some notions of entropy for cryptography - (invited talk). In *Information Theoretic Security – ICITS*, pages 138–142, 2011.
- [28] R. L. Rivest and W. D. Smith. ThreeVotingProtocols: ThreeBallot, VAV, and Twin. In *Electronic Voting Technology Workshop (EVT 2007)*, 2007.
- [29] P. Ryan, D. Bismark, J. Heather, S. Schneider, and Z. Xia. The prêt à voter verifiable election system. *IEEE Transactions on Information Forensics and Security*, 4:662–673, 2009.
- [30] K. Sako and J. Kilian. Receipt-free mix-type voting scheme - a practical solution to the implementation of a voting booth. In *Advances in Cryptology - EUROCRYPT '95*, volume 921 of *LNCS*, pages 393–403. Springer, 1995.
- [31] C. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, pages 379–423 and 623–656, 1948.
- [32] G. Smith. Quantifying information flow using min-entropy. In *8th International Conference on Quantitative Evaluation of Systems (QEST'11), invited paper*, pages 159–167, 2011.
- [33] D. Unruh and J. Müller-Quade. Universally Composable Incoercibility. In *30th International Cryptology Conference (CRYPTO'10)*, volume 6223 of *LNCS*, pages 411–428. Springer, 2010.