How Efficient are Replay Attacks against Vote Privacy? A Formal Quantitative Analysis

David Mestel
SnT/University of Luxembourg
Luxembourg
david.mestel@uni.lu

Johannes Müller
SnT/University of Luxembourg
Luxembourg
johannes.mueller@uni.lu

Pascal Reisert
University of Stuttgart
Germany
pascal.reisert@sec.uni-stuttgart.de

Abstract—Replay attacks are among the most well-known attacks against vote privacy. Many e-voting systems have been proven vulnerable to replay attacks, including systems like Helios that are used in real practical elections.

Despite their popularity, it is commonly believed that replay attacks are inefficient but the actual threat that they pose to vote privacy has never been studied formally. Therefore, in this paper, we precisely analyze for the first time how efficient replay attacks really are.

We study this question from commonly used and complementary perspectives on vote privacy, showing as an independent contribution that a simple extension of a popular game-based privacy definition corresponds to a strong entropy-based notion.

Our results demonstrate that replay attacks can be devastating for a voter's privacy even when an adversary's resources are very limited. We illustrate our formal findings by applying them to a number of real-world elections, showing that a modest number of replays can result in significant privacy loss. Overall, our work reveals that, contrary to a common belief, replay attacks can be very efficient and must therefore be considered a serious threat.

I. INTRODUCTION

Electronic voting, or *e-voting*, is a reality. Systems for e-voting are nowadays used for political elections all over the world, for example, in Australia, Brazil, Estonia, India, Switzerland, or the US. Furthermore, in line with the general shift toward remote technologies, numerous institutions (e.g., academic organizations such as ACM, IACR, or SIAM) employ e-voting systems to mitigate physical barriers and increase voter turnout.

The two most fundamental properties for *secure* e-voting are *(end-to-end) verifiability* and *(vote) privacy*. Verifiability [17] enables external and internal observers to detect and reject falsely computed election results, even when the underlying cause is an unknown programming error or malicious behavior of some of the participants. Privacy [5] guarantees that all data published during an election (including data for proving the integrity of the final result) does not leak more information on the single voters' choices than what can be derived from the public (unbiased) election result.

Designing secure e-voting systems is very challenging, with a long and rich history going back to the 1980's [4]. Since then, numerous e-voting systems have been proposed which aim to provide both verifiability and privacy (see, e.g., [1, 13–16, 34, 36, 37]), sometimes with additional security proper-

ties such as receipt-freeness [13] or coercion-resistance [16]. Some of these e-voting systems have been and are used in practice, for example for political elections in Australia [11], Estonia [38], Switzerland [40], and the US [12], or for non-political ones, such as IACR elections [25], to name just a few.

It is crucial to protect the voters' privacy not only against passive observers but also against adversaries who control some of the protocol participants (e.g., voters or tallying authorities) and who let these corrupted participants actively deviate from their specified roles in order to undermine privacy of some or all (uncorrupted) voters. Guaranteeing privacy in the face of such active adversaries is a common standard which (most) modern e-voting systems aim to provide. However, it turned out that numerous e-voting systems fall short of this goal in their respective threat scenarios, including seminal systems like Helios (see [7, 10, 18]), Civitas (see [24, 27]), or Prêt à Voter (see [10]).

One of the most prominent classes of attacks against privacy—if not *the* most prominent one—are *replay attacks* (see, e.g., [7, 13, 18, 19, 24]) to which many e-voting systems have been proven vulnerable (e.g., [1, 8, 16, 34, 36, 37]). Roughly speaking, a replay attack works as follows. The adversary, who controls some corrupted voters, targets an uncorrupted voter whose privacy he wants to undermine. The adversary waits until the targeted voter has submitted her ballot and reads it from the public bulletin board. Then the adversary instructs (some of) the corrupted voters to submit (possibly a re-randomization of) the same ballot the targeted voter had submitted before. If, in a particular e-voting system, these replayed ballots are not discarded prior to tallying, then the targeted voter's choice is amplified in the public election result. Because the adversary now obtains more information about the targeted voter's choice than what he could derive from an unbiased election result, vote privacy is undermined.

Despite their popularity, the risk of replay attacks is often regarded as a "largely theoretical" [39] issue, even in the scientific community. In publications which unveil replay attacks against vote privacy of an e-voting system, the effect of replay attacks is typically illustrated for extreme cases only, e.g., elections with just two honest voters [10] or with many corrupted ones who all replay a single voter's ballot [19]. While such completely artificial toy examples can be useful

to explain why privacy is formally broken, they seemingly suggest that replay attacks do not pose a serious threat. It is therefore not surprising that, for instance, in response to the replay attack against Helios [1] discovered in [18], Helios Voting replied that "the risk of this attack being successfully carried out is low, as it requires "wasting" a number of votes to compromise the privacy of one voter", concluding that, most likely, replay attacks would not matter [23]. This—as we shall see, fallacious—perspective may also explain why the latest version of Helios [1] (used, e.g., for IACR elections [25]) has not yet been patched to defend against the replay attacks discovered in [7].

Indeed, at first glance, it seems necessary to replay a targeted voter's ballot many times in order to significantly amplify this voter's choice in the final result. However, somewhat surprisingly, this common conjecture has never come under close scrutiny. The study by Cortier and Smyth [18] is the only previous work which attempted (in Section III.C) to analyze how replay attacks scale, but the authors considered a "definitive mathematical analysis" as future work because their underlying model was "rather naïve" [18].

In this work, we challenge the abovementioned conjecture for the first time, both rigorously and extensively, from two established and complementary perspectives on vote privacy. We precisely measure how efficient replay attacks really are, i.e., how much the affected voter's privacy loss increases depending on the number of replays. In particular, we show that replay attacks can be devastating for a voter's privacy even when an adversary's resources are highly limited so that he can (or is willing to) replay a targeted voter's ballot only very few times. This observation disproves a common conjecture that vote privacy would only be at risk if the number of replays was high. Our novel insights are immediately relevant for the security of real elections because e-voting systems vulnerable to replay attacks have been, are, and most likely will be used in practice (e.g., the latest version of Helios for IACR elections).

A. Our contributions

a) Categorization of replay attacks (Sec. II): We begin by reviewing the scientific literature to extract all replay attacks against vote privacy that have been published to date. We categorize these attacks into different classes, depending on their specific forms. Our extensive presentation highlights that replay attacks play a central role in modern secure e-voting, which demonstrates the importance of our subsequent analysis.

b) Efficiency analysis based on the KTV vote privacy definition (Sec. IV): We first formally analyse the efficiency of replay attacks using the vote privacy definition by Küsters, Truderung, and Vogt [30], hereafter called the *KTV privacy definition* (Sec. III). The KTV privacy definition is not only established and widely used (see, e.g., [3, 9, 28, 29, 32]) but it proves particularly useful for our purposes because it allows us to *measure* the loss of vote privacy and thus the efficiency of replay attacks.

We first define an ideal functionality for an e-voting protocol which allows the adversary to replay a targeted voter's ballot n_{repl} times, and compute the KTV privacy loss of this protocol. We obtain a useful reduction from the privacy loss for a general election to that for an election with only three candidates.

This allows us to analyze how the ideal privacy loss is affected by the number of replays n_{repl}. As we shall see, even for small numbers of n_{repl}, the privacy loss can be devastating. We illustrate our abstract results with a number of realistic examples.

c) A new entropy-based vote privacy definition (Sec. V): A limitation of the KTV privacy definition [30] (observed for instance in [5]) is that it only measures privacy with respect to a specific security game, namely the adversary's ability to guess between two possible votes. In particular this means that for the ideal functionality (including replays) the privacy loss is (as we will see in Section IV) entirely determined by the two least popular candidates, with the other candidates having no effect whatsoever.

Entropy-based measures of vote privacy (e.g., [6, 35]) provide a complementary view because they consider privacy with respect to a variety of goals for the adversary. Unfortunately, as we will explain in Sec. V, they are limited in various ways which make them difficult to apply in practice to analyse concrete elections.

In Sec. V, we propose a simple extension of the KTV vote privacy definition, which we show is equivalent to a computational version of a strong entropy-based notion. This is independent of the replay attack setting, and serves to somewhat unify the KTV and strong entropy-based approaches.

We show that our novel definition can be efficiently and accurately estimated for the ideal functionality using Monte Carlo methods, and so we are able to use it to study the efficiency of replay attacks from an entropy-based perspective complementary to the game-based perspective of the KTV definition.

d) Analysis of real-world elections (Sec. VI): In order to complement our formal analysis, we study how replay attacks would scale in practical elections. We therefore apply our formal results to publicly available data of political elections in Estonia, Germany, the UK, and the USA. In this way, we can realistically simulate to which degree vote privacy would decrease if in such elections replay attacks had been executed. Our "field test" confirms the gist of our abstract results: even if the number of replays is very low, vote privacy can be undermined significantly.

B. Structure of the paper

The structure of our paper essentially follows our contributions as presented above. In Sec. II, we categorize all replay attacks described in the literature. In Sec. III, we recall the KTV privacy definition as well as the ideal privacy loss a voting protocol can achieve w.r.t. the KTV definition. In Sec. IV, we study the efficiency of replay attacks based on the KTV privacy definition. In Sec. V, we propose our new

entropy-based vote privacy definition, describe its relationship to the KTV definition, and show that it can be efficiently estimated for the ideal functionality of Sec. IV. In Sec. VI, we illustrate our theoretical results using concrete election data from political elections and discuss the consequences of our insights.

II. Categorization of Replay Attacks

We provide the first comprehensive categorization of all replay attacks against vote privacy described in the literature. We identified three different variants of replay attacks: *basic replay attacks*, *homomorphic replay attacks*, and *re-voting replay attacks*. We summarize our insights at the end of the section.

A. Basic replay attacks

In its most basic form, a replay attack works as follows. Assume that we have n_V voters and that the adversary aims to break privacy of some voter V_{obs}, the *voter under observation*. We assume that the adversary controls a number n_V^d of further voters. The adversary waits until V_{obs} has submitted her ballot b_{obs}, containing her secret choice c_{obs}, to the bulletin board. The adversary reads b_{obs} from the bulletin board and instructs all of his corrupted voters to submit V_{obs}'s ballot b_{obs} as well. If, due to the specification of the e-voting scheme invoked, all $n_V^d + 1$ identical ballots $b_{obs}, \ldots, b_{obs}$ are tallied, then the public election result contains n_V^d additional votes for V_{obs}'s choice c_{obs}. By this, V_{obs}'s choice is amplified in the final result and thus her vote privacy is undermined.

Numerous e-voting schemes have been proven vulnerable against this basic version of replay attacks. Cortier and Smyth [18] demonstrated that basic replay attacks are possible in Helios [1], in the voting scheme by Sako and Kilian [36], and in the one by Schoenmakers [37]. The basic replay attacks against these voting schemes can be prevented by rejecting (partially) duplicated ballots.

B. Homomorphic replay attacks

Even if duplicated ballots are rejected in order to protect against basic replay attacks (see above), it may be possible to exploit malleability of the underlying cryptographic primitives in order to execute (more subtle) replay attacks. In what follows, we explain the general idea of such *homomorphic replay attacks*.[1] E-voting schemes with homomorphic tallying assume that voters' ciphertexts are re-randomizable, i.e., it is possible to transform ciphertext $e = \mathsf{Enc}(pk, m; r)$ into ciphertext $e' = \mathsf{Enc}(pk, m; r')$ without knowledge of the secret key sk, plaintext m, or randomness r. In a homomorphic replay attack, the adversary re-randomizes the observed voter's ballot b_{obs} into n_V^d ballots $b_1, \ldots, b_{n_V^d}$. Because the ballots $b_{obs}, b_1, \ldots, b_{n_V^d}$ are mutually distinct (with overwhelming probability if the encryption scheme is semantically secure), they will all be tallied even if ballot duplicates are strictly removed. By this, analogously to the basic replay attack (see

[1] We restrict our attention to the ballots' ciphertexts and put further primitives (signatures etc.) aside for simplicity.

above), the observed voter's privacy is undermined because all ballots $b_{obs}, b_1, \ldots, b_{n_V^d}$ contain V_{obs}'s choice c_{obs}.

Several e-voting schemes are vulnerable to such homomorphic replay attacks, for example the one by Lee et al. [34] (pointed out by Dreier, Lafourcade and Lakhnech [19]), or the one by Blazy, Fuchsbauer, Pointcheval, and Vergnaud [8] (pointed out by Chaidos, Cortier, Fuchsbauer, and Galindo [13]) which is the predecessor of BeleniosRF [13].

In order to protect against homomorphic replay attacks, many e-voting schemes employ zero-knowledge proofs (ZKPs) of knowledge which each voter uses to prove that she *knows* the plaintexts (and randomness) in the ciphertexts of her ballot. By this, a corrupted voter can no longer re-randomize the observed voter's ballot because he is not able to come up with a (valid) proof of plaintext knowledge.

Typically, e-voting schemes employ ZKPs of knowledge which are *non-interactive*, i.e., where the voter does not communicate with the verifier while proving knowledge (and correctness) of her encrypted choice. To construct such non-interactive ZKPs, most (modern) e-voting schemes use the Fiat-Shamir transformation [20]. However, as we will recall in what follows, applying the Fiat-Shamir transformation *correctly* is non-trivial.

Bernhard, Pereira, and Warinschi [7] demonstrated that great care has to be taken when the Fiat-Shamir transformation is used. Bernhard et al. showed that the Fiat-Shamir transformation in the implementation of Helios [1] is too weak because the hash function does not take the statement to be proven as input. Therefore, a voter's ZKP in Helios [1] is in fact not a proof of knowledge, enabling an adversary to still execute homomorphic replay attacks.

C. Re-voting replay attacks

Bursuc, Dragan, and Kremer [10] explained that, even if (partial) ballot duplicates are strictly removed and a (correct) ZKP of knowledge is used (see above), replay attacks against Helios [1] are still feasible if the ballot box is corrupted. We note that, in principle, this replay attack is not restricted to the case of Helios. In what follows, we describe the idea of this replay attack, which is due to P. B. Rønne originally (according to [10]).

If the adversary controls the ballot box (i.e., the server to which voters send their ballots), it can claim that the ballot casting of the voter under observation V_{obs} was not successful. The voter under observation may then try a second attempt with the same vote c_{obs}. This way, the adversary obtains two different ballots b_{obs}, b'_{obs}, both containing the observed voter's vote c_{obs}. Now, the adversary can submit b_{obs} on behalf of one of the corrupted voters, whereas the voter under observation V_{obs} submits b'_{obs}. Because b_{obs} and b'_{obs} do not contain identical entries (with overwhelming probability due to the semantic security of the underlying cryptographic primitives), they will both be in the input of the tallying phase. The attack can be repeated several times to obtain more ballots of V_{obs}'s vote c_{obs}. By this, analogously to the basic replay

attack (see above), the observed voter's privacy is undermined. The attack could, for example, be prevented by including each voter's ID in the statement to be proven, in particular in the hash of the Fiat-Shamir transformation.

D. Summary

Our comprehensive presentation demonstrates that replay attacks are a recurrent and often subtle issue in the construction and employment of secure e-voting systems, even when deliberately designed to protect against them. While some pitfalls making replay attacks possible are straightforward to solve (e.g., removing duplicates), others are more subtle and require very close attention (e.g., using strong Fiat-Shamir transformations). Based on our systematic literature review, we conjecture that, despite its popularity, the threat of replay attacks is a recurrent issue of e-voting. It is therefore important to precisely understand the risk that replay attacks pose to the crucial property of vote privacy. In the remainder of this paper, we provide the first formal analysis of this fundamental threat.

III. KTV Vote Privacy Definition

The first part of our formal analysis of replay attacks (Sec. IV) is based on the vote privacy definition proposed by Küsters, Truderung, and Vogt [30], hereafter called the *KTV (privacy) definition*. We explain the motivation for this privacy definition and the formal definition itself in Sec. III-B, after first recalling the underlying computational model in Sec. III-A. In Sec. III-C, we recall the best possible privacy loss an arbitrary voting protocol can achieve according to the KTV privacy definition; this *ideal privacy loss* is expressed as a parameterized formula that we will use to precisely measure the efficiency of replay attacks in Sec. IV.[2]

A. Computational model

We briefly recall the computational model of the KTV privacy definition, in particular the notions of processes, protocols, instances, and properties. We refer to [30] for full technical details.

a) Process: A *process* is a set of probabilistic polynomial-time interactive Turing machines (ITMs, also called *programs*), which are connected via named tapes (also called *channels*). We write a process π as $\pi = p_1 \| \cdots \| p_l$, where $p_1, \ldots, p_l$ are programs. If π_1 and π_2 are processes, then $\pi_1 \| \pi_2$ is a process, provided that the processes are connectible: two processes are *connectible* if common external channels have opposite directions (input/output). A process π where all programs are given the security parameter ℓ is denoted by $\pi^{(\ell)}$. The processes we consider are such that the length of a run is always polynomially bounded in ℓ. A run is uniquely determined by the random coins used by the programs in π.

b) Protocol: A *protocol* P specifies a set of agents (also called parties or protocol participants) and a set of channels these agents can communicate over. Moreover, P specifies, for every agent a, a set Π_a of all programs the agent a may run and a program $\hat{\pi}_a \in \Pi_a$, *the honest program of a*, i.e., the program that a runs if a is honest, and hence, follows the protocol.

c) Instance: Let P be a protocol with agents $a_1, \ldots, a_n$. An *instance of* P is a process of the form $\pi = (\pi_{a_1} \| \ldots \| \pi_{a_n})$ with $\pi_{a_i} \in \Pi_{a_i}$. An agent a_i is called *honest* in the instance π if and only if $\pi_{a_i} = \hat{\pi}_{a_i}$. A *run of* P (with security parameter ℓ) is a run of some instance of P (with security parameter ℓ); we consider the instance to be part of the description of the run. An agent a_i is honest in a run r, if r is a run of an instance of P with honest a_i.

d) Property: A *property* γ *of* P is a subset of the set of all runs of P. By $\neg\gamma$ we denote the complement of γ.

e) Negligible, overwhelming, δ-bounded: As usual, a function f from the natural numbers to the interval $[0, 1]$ is *negligible* if, for every $c > 0$, there exists ℓ_0 such that $f(\ell) \leq \frac{1}{\ell^c}$ for all $\ell > \ell_0$. The function f is *overwhelming* if the function $1 - f$ is negligible. A function f is δ-*bounded* if, for every $c > 0$ there exists ℓ_0 such that $f(\ell) \leq \delta + \frac{1}{\ell^c}$ for all $\ell > \ell_0$.

B. Privacy definition

The KTV privacy definition [30] formalizes privacy of an e-voting protocol as the inability of an adversary π_A to distinguish whether some voter V_{obs}, the *voter under observation* who runs her honest program, voted for choice c_j or choice $c_{j'}$. Unlike binary privacy notions according to which a voting protocol either does or does not protect privacy (see [5]), the KTV privacy definition *measures* the *privacy loss* a voting protocol provides. Being able to measure vote privacy, in particular to measure the loss of vote privacy due to attacks, is crucial for the purposes of our paper (see Sec. IV).

To be more precise, according to [30], a voting protocol provides δ-*privacy* if any adversary π_A is able to distinguish whether V_{obs} voted for c_j or $c_{j'}$ with probability at most δ; or, to phrase it differently, if any adversary's advantage is δ-bounded. To define the KTV privacy notion formally, we first introduce the following notation for an arbitrary e-voting protocol P. Given a voter V_{obs} and choice c, we consider instances of P that induce a set of processes of the form $(\hat{\pi}_{V_{obs}}(c) \| \pi^* \| \pi_A)$ where $\hat{\pi}_{V_{obs}}(c)$ is the honest program of the voter V_{obs} under observation who takes c as her choice, π^* is the composition of programs of the remaining parties in P, and π_A is the program of the adversary. Let $\Pr[(\hat{\pi}_{V_{obs}}(c) \| \pi^* \| \pi_A)^{(\ell)} \mapsto 1]$ denote the probability that the adversary writes the output 1 on some dedicated tape in a run of $(\hat{\pi}_{V_{obs}}(c) \| \pi^* \| \pi_A)$ with security parameter ℓ and some choice c, where the probability is taken over the random coins used by the parties in $(\hat{\pi}_{V_{obs}}(c) \| \pi^* \| \pi_A)$.

Now, the intuition described above is formally defined as follows.

Definition 1 (Vote Privacy [30]): Let P be a voting protocol, V_{obs} be the voter under observation, and $\delta \in [0, 1]$. Then, P

---

[2]What we call privacy *loss* in this work was in the original paper [30] called privacy *level*. Because the privacy bound δ is higher when more private information is leaked, we prefer to use the term privacy loss for δ.

achieves δ-privacy, if for all possible choices $c_j, c_{j'}$ and all adversaries π_A (implicitly on input $(c_j, c_{j'})$) the difference

$$\left| \Pr[(\hat{\pi}_{V_{obs}}(c_j) \| \pi^* \| \pi_A)^{(\ell)} \mapsto 1] - \Pr[(\hat{\pi}_{V_{obs}}(c_{j'}) \| \pi^* \| \pi_A)^{(\ell)} \mapsto 1] \right|$$

is δ-bounded as a function of the security parameter 1^ℓ.

In other words, the privacy loss δ is an upper bound of an arbitrary adversary's advantage to distinguish whether V_{obs} voted for c_j or $c_{j'}$. Clearly, $\delta = 0$ would be desirable but typically we have $\delta > 0$, even for an ideal e-voting protocol with a completely passive adversary. The reason is that in many real-world elections, there exist choices which are picked only with low probability, for example unpopular candidates or unreasonable rankings (e.g., the green party is ranked first and the coal mining party next to it). Now, in most e-voting systems, including all systems mentioned in Sec. II, the final election result consists of the number of votes for each choice. Therefore, if the voter under observation V_{obs} chooses c_j or $c_{j'}$ but all other voters vote for one (or both) of these choices with low probability only, then V_{obs}'s choice is not hidden sufficiently well—in the worst case, V_{obs}'s choice is completely revealed.

C. Ideal privacy

Since we have seen that the privacy loss δ is typically not perfect, the following questions are obvious: What is the best possible privacy loss that can be achieved in a given election? How does this ideal privacy loss depend on basic parameters, such as the number of voters or the voters' preferences? These questions have been answered precisely in [30] and we will recall the results in what follows. These results will be the foundation of our formal analysis of replay attacks in Sec. IV.

a) Ideal voting protocol: In order to have a lower bound on the privacy loss for *all* voting protocols, Küsters et al. [30] derived a formula for the privacy loss an ideal voting protocol provides.[3] Let us describe this ideal voting protocol, denoted by $\mathcal{I}_{priv}$, starting with the parameters it depends on:

- *Number of choices* n_C: The set $C = \{c_1, \ldots, c_{n_C}\}$ consists of all possible choices c_j that a voter can choose.
- *Number of honest voters* n_V^h: We denote the number of voters which cannot be corrupted, the *honest voters*, by n_V^h.[4]
- *Voting distribution* $\vec{p}$: Each honest voter V_i picks her choice according to the distribution $\vec{p}$ over C, i.e. $\vec{p}[l]$ is the probability that an honest voter chooses c_l.[5]

The ideal voting protocol $\mathcal{I}_{priv}(C, n_V^h, \vec{p})$ works as follows. For each of the n_V^h honest voters V_i, the ideal voting protocol

---

[3]The ideal privacy loss derived in [30] is formulated for result functions that reveal the complete tally, i.e., number of votes for each choice. Subsequently, a more general ideal privacy loss was derived in [28] which is formulated for arbitrary result functions, including *tally-hiding* result functions that may, for instance, only reveal the winner but nothing else. Because all e-voting systems mentioned in Sec. II employ a result function which returns the complete tally, we restrict our attention to the ideal privacy loss derived in [30].

[4]The number of dishonest voters is not relevant for result functions that reveal the full tally because an adversary can derive the "honest" election result by subtracting the dishonest voters' choices from the final election result.

[5]In slight abuse of notation we identify $\vec{p}$ and its probability mass function.

$\mathcal{I}_{priv}(C, n_V^h, \vec{p})$

Parameters:
- Finite set $C \subset \mathbb{Z}$
- Number of honest voters n_V^h
- Probability distribution $\vec{p}$ over C

On (setup) from S do:
  1) res $\leftarrow 0^{|C|}$
  2) $b_0, b_1 \leftarrow 0$
  3) Return success

On (init, honest) from S do:
  1) $\forall i \in \{1, \ldots, n_V^h\}$:
     If $(\star \xleftarrow{\vec{p}} C) = j$, set res$_j \leftarrow$ res$_j + 1$
  2) $b_0 \leftarrow 1$
  3) Return success

On (setChoice, (j, j')) from S do:
  1) If $j, j' \notin C$, return $\perp$.
  2) If $(\star \xleftarrow{\$} \{0, 1\}) = 0$, set res$_j \leftarrow$ res$_j + 1$, else set res$_{j'} \leftarrow$ res$_{j'} + 1$
  3) $b_1 \leftarrow 1$
  4) Return success

On (compute) from S do:
  1) If $b_0 \cdot b_1 = 0$, return $\perp$.
  2) Return res

Fig. 1. Protocol of ideal voting functionality.

$\mathcal{I}_{priv}$ chooses V_i's choice according to $\vec{p}$. For the voter under observation V_{obs}, the ideal voting protocol expects as input a tuple of choices $(c_j, c_{j'})$ from the adversary, and then picks one of them uniformly at random. Eventually, $\mathcal{I}_{priv}$ returns the result res $\in \mathbb{N}^{|C|}$ which contains the number of votes for each choice made by all honest voters and by the voter under observation. The protocol $\mathcal{I}_{priv}(C, n_V^h, \vec{p})$ is formally defined in Fig. 1.

b) Ideal privacy loss: We now recall from [30] how the privacy loss δ^{ideal} of the ideal voting protocol $\mathcal{I}_{priv}(C, n_V^h, \vec{p})$ can be expressed as a parameterized formula $\delta^{ideal}_{C, n_V^h, \vec{p}}$. Recall that we defined the privacy loss of a voting protocol by the (level of) inability to distinguish whether the voter under observation V_{obs} voted for choice c_j or choice $c_{j'}$ (Definition 1). Now, the intuition behind the definition of $\delta^{ideal}_{C, n_V^h, \vec{p}}$ is as follows. If the adversary, given a final election result res, wants to decide whether the observed voter voted for choice c_j or $c_{j'}$, then the best strategy of the adversary is to opt for $c_{j'}$ if and only if the output res is more likely if the voter voted for choice $c_{j'}$. In order to capture this intuition formally, we introduce the following terminology.

Let A^l_{res} denote the conditional probability that the choices made by the honest voters and by the voter under observation yield the final result res, under the condition that the voter under observation V_{obs} chooses c_l. The probability A^l_{res} can be expressed as follows using the multinomial distribution:

$$A_{\text{res}}^l = \frac{n_V^h!}{\prod_{j=1}^{n_C} \text{res}[j]!} \cdot \left(\prod_{j=1}^{n_C} \vec{p}[j]^{\text{res}[j]} \right) \cdot \frac{\text{res}[l]}{\vec{p}[l]} \qquad (1)$$

where $\text{res}[j]$ is the number of votes for choice c_j. We can now define the set of outputs res for which it is more likely that the voter voted for choice $c_{j'}$ as follows:

$$M_{j,j'}^* = \{\text{res}: A_{\text{res}}^j \leq A_{\text{res}}^{j'}\}. \qquad (2)$$

The intuition of the ideal privacy loss described above is formally captured by the following definition:

$$\delta_{C,n_V^h,\vec{p}}^{\text{ideal}} = \max_{j,j' \in \{1,\dots,n_C\}} \sum_{\text{res} \in M_{j,j'}^*} (A_{\text{res}}^{j'} - A_{\text{res}}^j). \qquad (3)$$

The following theorem (Theorem 3 of [30], proved in Appendix C of the eprint version [31]) states that the loss $\delta_{C,n_V^h,\vec{p}}^{\text{ideal}}$ is indeed optimal for the ideal voting protocol $\mathcal{I}_{\text{priv}}(C, n_V^h, \vec{p})$. As a consequence, no voting protocol can achieve a better privacy loss than $\delta_{C,n_V^h,\vec{p}}^{\text{ideal}}$.

Theorem 1 ([30]): The ideal protocol $\mathcal{I}_{\text{priv}}(C, n_V^h, \vec{p})$ achieves a privacy loss of $\delta_{C,n_V^h,\vec{p}}^{\text{ideal}}$. Moreover, it does not achieve δ'-privacy for any $\delta' < \delta_{C,n_V^h,\vec{p}}^{\text{ideal}}$.

IV. Efficiency Analysis based on the KTV Vote Privacy Definition

In this section, we formally study the efficiency of replay attacks using the KTV privacy definition. First, in Sec. IV-A we focus on capturing the *effect* of replay attacks. We define a suitable ideal functionality for a voting protocol whose only flaw is that it allows the adversary to execute a replay attack. We characterise its KTV privacy loss analogously to the characterisation for the truly ideal protocol (without replays) in [30] (Theorem 2). Because this characterisation is not computationally tractable, we then show a reduction to an election with only three candidates (Theorem 3), and obtain a tractable formula which we use to demonstrate the devastating effect of even a small number of replays on realistically-sized example elections.

Based on these insights, in Sec. IV-B we then study the *efficiency* of replay attacks in general: we analyse (Theorem 4) how the ideal privacy loss behaves asymptotically in the number of replayed ballots n_{repl} and the number of honest voters n_V^h, in particular for fairly small values of n_{repl}.

A. Ideal privacy loss

We analyse the ideal privacy loss if the adversary can replay the observed voter's choice n_{repl} times. To this end, we modify the ideal voting functionality $\mathcal{I}_{\text{priv}}(C, n_V^h, \vec{p})$ (Fig. 1) so that it adds the observed voter's choice $(1 + n_{\text{repl}})$ times to the final result, instead of only once. The resulting ideal voting functionality with n_{repl} replays $\mathcal{I}_{\text{priv}}(C, n_V^h, \vec{p}, n_{\text{repl}})$ is defined in Fig. 2; note that $\mathcal{I}_{\text{priv}}(C, n_V^h, \vec{p}) = \mathcal{I}_{\text{priv}}(C, n_V^h, \vec{p}, 0)$. By

---

$\mathcal{I}_{\text{priv}}(C, n_V^h, \vec{p}, \textcolor{red}{n_{\text{repl}}})$

Parameters:
- Finite set $C \subset \mathbb{Z}$
- Number of honest voters n_V^h
- Probability distribution $\vec{p}$ over C
- Number of replays $\textcolor{red}{n_{\text{repl}}}$

On (setup) from S do:
1) $\text{res} \leftarrow 0^{|C|}$
2) $b_0, b_1 \leftarrow 0$
3) Return success

On (init, honest) from S do:
1) $\forall i \in \{1, \dots, n_V^h\}$:
   If $(\star \xleftarrow{\vec{p}} C) = j$, set $\text{res}_j \leftarrow \text{res}_j + 1$
2) $b_0 \leftarrow 1$
3) Return success

On (setChoice, (j, j')) from S do:
1) If $j, j' \notin C$, return $\perp$.
2) If $(\star \xleftarrow{\$} \{0, 1\}) = 0$, set $\text{res}_j \leftarrow \text{res}_j + 1 + \textcolor{red}{n_{\text{repl}}}$,
   else set $\text{res}_{j'} \leftarrow \text{res}_{j'} + 1 + \textcolor{red}{n_{\text{repl}}}$
3) $b_1 \leftarrow 1$
4) Return success

On (compute) from S do:
1) If $b_0 \cdot b_1 = 0$, return $\perp$.
2) Return res

Fig. 2. Protocol of ideal voting functionality which allows for replaying the observed voter's choice n_{repl} times. The differences between the original ideal voting protocol (Fig. 1) and the one presented here are highlighted in red.

using the ideal voting functionality $\mathcal{I}_{\text{priv}}(C, n_V^h, \vec{p}, n_{\text{repl}})$, we can model exactly that the adversary is able to execute only a replay attack with n_{repl} replays but no other kind of privacy attack. This means that the privacy loss $\delta_{C,n_V^h,\vec{p},n_{\text{repl}}}^{\text{ideal}}$ provided by $\mathcal{I}_{\text{priv}}(C, n_V^h, \vec{p}, n_{\text{repl}})$ is a lower bound for the privacy loss of *any* voting protocol in which the adversary can replay the observed voter's choice n_{repl} times.

In what follows, we first derive a representation of the ideal privacy loss $\delta_{C,n_V^h,\vec{p},n_{\text{repl}}}^{\text{ideal}}$ which is conceptually similar to the ideal privacy loss without replays $\delta_{C,n_V^h,\vec{p}}^{\text{ideal}}$, as defined in Eq. 1 (Sec. III-C). We observe that $\delta_{C,n_V^h,\vec{p},n_{\text{repl}}}^{\text{ideal}}$ is indeed the privacy loss of the ideal voting functionality with replays $\mathcal{I}_{\text{priv}}(C, n_V^h, \vec{p}, n_{\text{repl}})$ (Theorem 2). We then derive an alternative representation of $\delta_{C,n_V^h,\vec{p},n_{\text{repl}}}^{\text{ideal}}$ which reduces vote privacy (under the KTV definition) from dependence on all n_C possible choices to dependence only on the two most unpopular choices (Theorem 3).

a) First representation: Analogously to A_{res}^l (Sec. III-C), let $A_{\text{res}}^{l,n_{\text{repl}}}$ be the probability of obtaining result res under the condition that the voter under observation voted for candidate l, where now her ballot is replayed n_{repl} times. Note that $A_{\text{res}}^l = A_{\text{res}}^{l,0}$. It is easy to see that we have

$$A_{\text{res}}^{l,n_{\text{repl}}} = \frac{n_V^h! \cdot p[1]^{\text{res}[1]} \cdot \dots \cdot p[l]^{\text{res}[l]-n_{\text{repl}}-1} \cdot \dots \cdot p[n_C]^{\text{res}[n_C]}}{\text{res}[1] \cdot \dots \cdot (\text{res}[l] - n_{\text{repl}} - 1) \cdot \dots \cdot \text{res}[n_C]}$$

$$= \frac{n_{\mathsf{V}}^{\mathsf{h}}!}{\prod_{j=1}^{n_C} \text{res}[j]!} \cdot \left(\prod_{j=1}^{n_C} \vec{p}[j]^{\text{res}[j]} \right) \cdot \frac{\prod_{\nu=0}^{n_{\text{repl}}} (\text{res}[l] - \nu)}{\vec{p}[l]^{n_{\text{repl}}+1}}$$

$$= \frac{\prod_{\nu=1}^{n_{\text{repl}}} (\text{res}[l] - \nu)}{\vec{p}[l]^{n_{\text{repl}}}} \cdot A_{\text{res}}^l$$

Now, analogously to $M_{j,j'}^*$ in Sec. III-C, we define

$$M_{j,j'}^{*,n_{\text{repl}}} = \left\{ \text{res} : A_{\text{res}}^{j,n_{\text{repl}}} \le A_{\text{res}}^{j',n_{\text{repl}}} \right\}$$

to be the set of all events that occur with higher likelihood under the condition that the observed voter chose $c_{j'}$ than under the condition that she chose c_j.

We thus obtain our first representation of the ideal privacy loss with n_{repl} replays, as follows:

$$\delta_{C,n_{\mathsf{V}}^{\mathsf{h}},\vec{p},n_{\text{repl}}}^{\text{ideal}} = \max_{j,j' \in \{1,\ldots,n_C\}} \sum_{\text{res} \in M_{j,j'}^{*,n_{\text{repl}}}} (A_{\text{res}}^{j',n_{\text{repl}}} - A_{\text{res}}^{j,n_{\text{repl}}}). \quad (4)$$

The following theorem states that the loss $\delta_{C,n_{\mathsf{V}}^{\mathsf{h}},\vec{p},n_{\text{repl}}}^{\text{ideal}}$ is indeed optimal for the ideal voting protocol with n_{repl} replays. This means that no voting protocol which is subject to a replay attack with n_{repl} replays can achieve a better privacy loss than $\delta_{C,n_{\mathsf{V}}^{\mathsf{h}},\vec{p},n_{\text{repl}}}^{\text{ideal}}$.

Theorem 2: The ideal protocol $\mathcal{I}_{\text{priv}}(C, n_{\mathsf{V}}^{\mathsf{h}}, \vec{p}, n_{\text{repl}})$ achieves a privacy loss of $\delta_{C,n_{\mathsf{V}}^{\mathsf{h}},\vec{p},n_{\text{repl}}}^{\text{ideal}}$. Moreover, it does not achieve δ'-privacy for any $\delta' < \delta_{C,n_{\mathsf{V}}^{\mathsf{h}},\vec{p},n_{\text{repl}}}^{\text{ideal}}$.

The proof of Theorem 2 is exactly the same as the proof of Theorem 1, with A_{res}^j replaced by $A_{\text{res}}^{j,n_{\text{repl}}}$ throughout.

Note that the formula in (4) for $\mathcal{I}_{\text{priv}}(C, n_{\mathsf{V}}^{\mathsf{h}}, \vec{p}, n_{\text{repl}})$ involves maximising over $O(n_C^2)$ terms, each of which is a sum consisting of $|M_{j,j'}^{*,n_{\text{repl}}}|$ summands. In general $|M_{j,j'}^{*,n_{\text{repl}}}|$ will have size comparable to the number of possible results, which is a multinomial coefficient of order $(n_{\mathsf{V}}^{\mathsf{h}})^{n_C-1}$. This is clearly intractable for large elections, so before we can analyse the efficiency of replay attacks for real-world elections, we will need to do some work to put (4) into a more tractable form.

b) Second representation: We now show that for the ideal functionality of Fig. 2, the definition of the KTV privacy loss can be greatly simplified. Recall that in Definition 1, we measure privacy as the adversary's maximum advantage *over all possible choices* $c_j, c_{j'}$ to successfully distinguish whether the voter under observation voted for c_j or $c_{j'}$. Our reduction states that in fact this is equal to the adversary's advantage in distinguishing only between a vote for *the least popular choice* c_j and a vote for *the second least popular choice* $c_{j'}$, i.e., those choices for which $\vec{p}[j]$ and $\vec{p}[j']$ are the two lowest probabilities. This holds both for the cases with and without replay attacks (since the latter is a special case of the former with $n_{\text{repl}} = 0$).

Theorem 3: Let j, j' be such that $\vec{p}[j] \le \vec{p}[j'] \le \vec{p}[l]$ for all $l \ne j$. Then, the ideal privacy loss is given by the following identity:

$$\delta_{C,n_{\mathsf{V}}^{\mathsf{h}},\vec{p},n_{\text{repl}}}^{\text{ideal}} = \sum_{\text{res} \in M_{j,j'}^{*,n_{\text{repl}}}} (A_{res}^{j',n_{\text{repl}}} - A_{res}^{j,n_{\text{repl}}}).$$

In order to prove Theorem 3 we will first show the technical Lemma 1. Lemma 1 states that for each choice of j and j', the corresponding term in the $\max$ of equation (4) only depends on three probabilities - $\vec{p}[j]$, $\vec{p}[j']$ and one "dummy" probability $\vec{p}[j,j'] = 1 - \vec{p}[j] - \vec{p}[j']$ that collects the probabilities of all the other choices.

The formula (5) of Lemma 1, combined with Theorem 3, gives an explicit expression for the ideal privacy loss as a sum of at most $n_{\mathsf{V}}^{\mathsf{h}} n_{\text{repl}}$ terms. This means that we are now comfortably able to analyse real-world-sized elections, as we do later in this section and in Sec. VI.

We will introduce a new variable $T_{t,j,j'}$ which will be used to state Lemma 1 below. Let $j, j' \in C$. For each $t \in \mathbb{N}$, let $T_{t,j,j'}$ be a natural number that satisfies

$$\prod_{\nu=0}^{n_{\text{repl}}} \frac{r - \nu}{\vec{p}[j]^{n_{\text{repl}}+1}} \le \prod_{\nu=0}^{n_{\text{repl}}} \frac{t - r - \nu}{\vec{p}[j']^{n_{\text{repl}}+1}}$$

for all $r \le T_{t,j,j'}$ and

$$\prod_{\nu=0}^{n_{\text{repl}}} \frac{r - \nu}{\vec{p}[j]^{n_{\text{repl}}+1}} \ge \prod_{\nu=0}^{n_{\text{repl}}} \frac{t - r - \nu}{\vec{p}[j']^{n_{\text{repl}}+1}}$$

for $r > T_{t,j,j'}$. Note that $0 \le T_{t,j,j'} \le t$ certainly exists since $\prod_{\nu=0}^{n_{\text{repl}}} \frac{r-\nu}{\vec{p}[j]^{n_{\text{repl}}+1}}$ is monotonic in r. To simplify the notation we will sometimes omit the j, j' and just write T_t.

$T_{t,j,j'}$ describes the point at which the sign of $A_{\text{res}}^{j',n_{\text{repl}}} - A_{\text{res}}^{j,n_{\text{repl}}}$ switches. It allows us to replace the summation over $M_{j,j'}^{*,n_{\text{repl}}}$ by a summation over *all* possible outcomes. We can then use standard properties of probability distributions to remove all but two probabilites.

Lemma 1: Let $j, j \in C$ and T_t as before. Then for $\vec{p}[j,j'] := 1 - \vec{p}[j] - \vec{p}[j']$

$$\sum_{res \in M_{j,j'}^{*,n_{\text{repl}}}} (A_{res}^{j',n_{\text{repl}}} - A_{res}^{j,n_{\text{repl}}}) = \quad (5)$$

$$\sum_{t=0}^{n_{\mathsf{V}}^{\mathsf{h}}} \vec{p}[j,j']^{n_{\mathsf{V}}^{\mathsf{h}}-t} \binom{n_{\mathsf{V}}^{\mathsf{h}}}{t} \sum_{r=\max\{T_t+n_{\text{repl}}+1-n_{\text{repl}},0\}}^{\min\{T_t+n_{\text{repl}}+1,t\}} \binom{t}{r} \vec{p}[j]^r \vec{p}[j']^{t-r}$$

Proof. We give only a short sketch, with many details relegated to the Appendix in Lemma 2. To simplify notation we write $\delta_{jj'}$ for $\sum_{\text{res} \in M_{j,j'}^{*,n_{\text{repl}}}} (A_{\text{res}}^{j'} - A_{\text{res}}^j)$. First observe that $\delta_{jj'} = \delta_{j'j}$ since $M_{j,j'}^{*,n_{\text{repl}}}$ and $M_{j,j'}^{*,n_{\text{repl}}}$ are complimentary up to a trivial intersection that does not contribute to $\delta_{jj'}$ or $\delta_{j'j}$. More precisely, writing $\text{MN}_{n_{\mathsf{V}}^{\mathsf{t}}}^{res}$ for the multinomial probability density function

$$\text{MN}_{n_{\mathsf{V}}^{\mathsf{t}}}^{res} = \frac{n_{\mathsf{V}}^{\mathsf{t}}!}{\prod_{j=1}^{n_C} \text{res}[j]!} \left(\prod_{j=1}^{n_C} \vec{p}[j]^{res[j]} \right),$$

we have that

$$\frac{n_{\mathsf{V}}^{\mathsf{t}}! \delta_{jj'}}{n_{\mathsf{V}}^{\mathsf{h}}!} = \sum_{\text{res} \in M_{j,j'}^{*,n_{\text{repl}}}} \frac{\prod_{\nu=0}^{n_{\text{repl}}} (\text{res} j - \nu)}{\vec{p}[j]^{n_{\text{repl}}+1}} \text{MN}_{n_{\mathsf{V}}^{\mathsf{t}}}^{res}$$

Fig. 3. KTV privacy loss δ for the ideal protocol with 10 candidates and the uniform vote distribution. Note that the y-axis is on a log scale.

implies that $\frac{n_V^t!}{n_V^h!}(\delta_{jj'} - \delta_{j'j})$ is equal to

$$\sum_{res} \left(\frac{\prod_{\nu=0}^{n_{repl}}(res\,j - \nu)}{\vec{p}[j]^{n_{repl}+1}} - \frac{\prod_{\nu=0}^{n_{repl}}(res\,j' - \nu)}{\vec{p}[j']^{n_{repl}+1}} \right) MN_{n_V^t}^{res} = 0$$

where the sum is over all possible results (without abstention), i.e. $\sum_{l=0}^{n_C} res[l] = n_V^h + n_{repl} + 1 = n_V^t$. We also used that the two conditional probability distribution (one w.r.t. to the choice j, one w.r.t. to the choice j') are each normed. Now $\delta_{jj'} = \delta_{j'j}$ and $\delta_{jj'} = \frac{\delta_{jj'} + \delta_{j'j}}{2}$ lead to the more symmetric representation

$$\delta_{jj'} = \frac{n_V^h}{2 \cdot n_V^t!} \sum_{res} \left| \Delta_{j,j'}^{res,n_{repl}} \right| \cdot MN_{n_V^t}^{res}.$$

with $\Delta_{j,j'}^{res,n_{repl}} = \frac{\prod_{\nu=0}^{n_{repl}}(res[j]-\nu)}{\vec{p}[j]^{n_{repl}+1}} - \frac{\prod_{\nu=0}^{n_{repl}}(res[j']-\nu)}{\vec{p}[j']^{n_{repl}+1}}$. By moving the $\vec{p}[j], \vec{p}[j']$ out of the product we can sum over all $res[l]$ for $j \neq l \neq j'$ under the restriction that $\sum_{j\neq l\neq j'} res[l] = n_V^t - t$ for $t = res[j] + res[j']$. Furthermore, we can then consider the multinomial distribution with $n_V^t - t$ trials and $(n_C - 2)$ probabilities $\vec{q} := (\vec{p}[j,j]^{-1}\vec{p}[l])_{j\neq l\neq j'}$. Using the norm 1 property of this probability distribution, we get a term that depends only on j and j'.

$$\delta_{jj'} = \sum_{t=0}^{n_V^t} \sum_{res[j]+res[j']=t} \left| \Delta_{j,j'}^{res,n_{repl}} \right| \cdot \frac{n_V^h! \vec{p}[j]^{res[j]} \vec{p}[j']^{res[j']}}{2 \cdot (n_V^t - t)! res[j]! res[j']!}$$

$$\cdot \sum_{\sum_{j\neq l\neq j'} res[l]=n_V^t-t} \frac{\vec{p}[j,j']^{n_V^t-t}(n_V^t-t)!}{\prod_{j\neq l\neq j'} res[l]!} \prod_{j\neq l\neq j'} \vec{q}[l]^{res[l]}$$

$$= \sum_{t=0}^{n_V^t} \sum_{r=0}^{t} \left| \Delta_{j,j'}^{res,n_{repl}} \right| \cdot \frac{n_V^h! \vec{p}[j,j']^{n_V^t-t} \vec{p}[j]^r \vec{p}[j']^{r-t}}{2 \cdot (n_V^t - t)! r! (t-r)!}$$

where $\Delta_{j,j'}^{res,n_{repl}}$ is defined as before with $res[j] = r, res[j'] = t - r$. Finally the definition of T_t allows us to replace the absolute value to retrieve (5). For more details see Lemma 2.

Proof of Theorem 3. By definition $\delta_{jj'}$ is smooth in $\vec{p}[j] + \vec{p}[j']$. Hence its representation in (5) is smooth and it is enough to compute the differential on intervals where $T_{t,j,j'}$ is constant. After differentiating (5) w.r.t. $\vec{p}[j] + \vec{p}[j']$ we see that $\delta_{jj'}$ decreases, i.e. becomes maximal if $\vec{p}[j] + \vec{p}[j']$ is minimal. For a more detailed version, see Appendix A.

In Fig. 3, we give some concrete values of the ideal privacy loss with replays $\delta_{C,n_V^h,\vec{p},n_{repl}}^{ideal}$ for different numbers of honest

voters n_V^h and replays n_{repl}. We model an election in which the votes of the honest voters are uniformly distributed; this is the distribution that minimises δ, and for any other distribution the privacy loss would be even greater. Observe from Fig. 3 that even if the adversary replays the observed voter's choice only very few times in relation to the total number of voters, the observed voter's privacy can be reduced dramatically (corresponding to a dramatically higher value of δ). As we will prove in the remainder of this section, this is no coincidence. In fact, we will show that replay attacks are very efficient in general.

B. Asymptotics

In the first part of our analysis, we focused on the *effect* of replay attacks: our results in Sec. IV-A state which privacy loss can be achieved ideally if an adversary replays the observed voter's choice n_{repl} times. Based on these results, we now precisely analyze the *efficiency* of replay attacks, i.e., how vote privacy decreases asymptotically depending on the number of replays n_{repl}.

Our main result on the efficiency of replay attacks is Theorem 4. Its proof is based on the explicit representation of the ideal privacy loss from Theorem 1. In what follows, we will first deduce Theorem 4 and eventually illustrate it for specific settings.

We use the terminology introduced in Sec. IV-A. We remark first, that one can choose

$$\left\lfloor \frac{t\vec{p}[j]}{\vec{p}[j] + \vec{p}[j']} \right\rfloor \leq T_t \leq \left\lfloor \frac{t\vec{p}[j]}{\vec{p}[j] + \vec{p}[j']} \right\rfloor + n_{repl} + 1$$

for $\vec{p}[j] \leq \vec{p}[j']$. In particular, the coefficients in the inner sum surround the expected value $\mathbb{E}(X_j) := \frac{t\vec{p}[j]}{\vec{p}[j] + \vec{p}[j']}$. Now we can use the Integral Limit DeMoivre-Laplace theorem for multinomial distributions (see, e.g., [22]) to represent the asymptotic behaviour in terms of the multivariate Gaussian over the hypersurface given by the condition $\sum_{i=1}^{3} r_i = n_V^h$ for $r_1 = r, r_2 = t - r, r_3 - n_V^h - t$ as follows:

$$\sum_{t=0}^{n_V^h} \vec{p}[j,j']^{n_V^h-t} \binom{n_V^h}{t} \sum_{r=\max\{T_{t+n_{repl}+1}-n_{repl},0\}}^{\min\{T_{t+n_{repl}+1},t\}} MN_t^{r,t-r}$$

$$\approx \frac{\sqrt{\vec{q}[j]\vec{q}[j']\vec{q}[j,j']}}{2\pi n_V^h \sqrt{\vec{p}[j]\vec{p}[j']\vec{p}[j,j']}} \int_0^{n_V^h} \int_{\mathbb{E}(X_j)-\frac{n_{repl}+1}{2}}^{\mathbb{E}(X_j)+\frac{n_{repl}+1}{2}} e^{-\sum_{i=1}^{3} x_i^2(r,t)} dr dt$$

where we used

$$x_1(r,t) = \frac{r - n_V^h \vec{p}[j]}{\sqrt{n_V^h \vec{p}[j]\vec{q}[j]}}, \qquad x_2(r,t) = \frac{t - r - n_V^h \vec{p}[j']}{\sqrt{n_V^h \vec{p}[j']\vec{q}[j']}},$$

$$x_3(r,t) = \frac{n_V^h - t - n_V^h \vec{p}[j,j']}{\sqrt{n_V^h \vec{p}[j,j']\vec{q}[j,j']}}, \qquad \vec{q} = 1 - \vec{p}.$$

We can isolate the terms for r and t to get

$$\sqrt{\frac{\vec{q}[j,j']}{2\pi n_V^h \vec{p}[j]\vec{p}[j']}} \int_{-\frac{n_{repl}+1}{2}}^{\frac{n_{repl}+1}{2}} e^{-r^2 \frac{\vec{q}[j,j']}{2\pi n_V^h \vec{p}[j]\vec{p}[j']}} dr$$

$$\approx (n_{\text{repl}} + 1)\sqrt{\frac{\bar{q}[j, j']}{2\pi n_{\text{V}}^{\text{h}} \bar{p}[j]\bar{p}[j']}} + O\left(\left(\frac{n_{\text{repl}} + 1}{\sqrt{n_{\text{V}}^{\text{h}}}}\right)^3\right) \quad (6)$$

Since the remaining term

$$\frac{1}{\sqrt{2\pi n_{\text{V}}^{\text{h}} \bar{p}[j, j']\bar{q}[j, j']}} \int_{-n_{\text{V}}^{\text{h}}\bar{q}[j,j']}^{n_{\text{V}}^{\text{h}}\bar{p}[j,j']} e^{-t^2/(2\pi n_{\text{V}}^{\text{h}}\bar{p}[j,j']\bar{q}[j,j'])} dt$$

converges to 1 for $n_{\text{V}}^{\text{h}} \to \infty$, our approximation (6) describes the asymptotics of the whole term.[6]

From what we have shown above, we obtain the following central result.

Theorem 4 (Asymptotics): Let $C, n_{\text{V}}^{\text{h}}, \vec{p}, n_{\text{repl}}$ be as above. Let $n_{\text{repl}} = o\left(\sqrt{n_{\text{V}}^{\text{h}}}\right)$.[7] Then, we have that

$$\delta_{C,n_{\text{V}}^{\text{h}},\vec{p}}^{\text{ideal}} \sim \frac{n_{\text{repl}} + 1}{\sqrt{n_{\text{V}}^{\text{h}}}}. \quad (7)$$

Intuitively, Theorem 4 essentially states that the KTV privacy loss of an election with n_{V}^{h} honest voters and n_{repl} replays is the same as that of an election with $\frac{n_{\text{V}}^{\text{h}}}{(n_{\text{repl}}+1)^2}$ voters but no replays. Loosely speaking, by replaying the targeted voter's choice n_{repl} times, it is as though the adversary could "reduce" the number of honest voters from n_{V}^{h} down to $\frac{n_{\text{V}}^{\text{h}}}{(n_{\text{repl}}+1)^2}$. This is perhaps a more intuitive way to evaluate privacy loss than a change in a numerical measure which may be difficult to interpret without context (although we should emphasise that it is equally dependent on the threat model embedded in the KTV definition).

To give some concrete examples, again with 10 candidates and a uniform distribution: in an election with 10 honest and no corrupted voters (and hence no replays possible) the adversary wins the privacy game with probability $\delta = 0.533$. In an election with 100 voters, the adversary needs to control as few as 3 out of 100 voters (and submit replays on their behalf) in order to have a similar advantage $> \frac{1}{2}$ in the privacy game. In an election with 1000 voters, as few as 9 out of 1000 voters suffice for the same purpose. At the same time, in the last two elections with 100 respectively 1000 voters, vote privacy is mostly preserved without replays ($\delta = 0.177$ and $\delta = 0.056$). Altogether, we can conclude that replay attacks can be devastating even if the adversary controls only a tiny fraction of all voters.

V. Strong vote privacy

An important limitation of the KTV privacy definition is that it only considers vulnerability with respect to a very specific goal, namely for the adversary to guess between two possible votes (which earns it a rating of 'too limited' in the survey article [5]), and a similar limitation applies to most

---

[6]Our series approximation becomes weak for $n_{\text{repl}}^2 \geq n_{\text{V}}^{\text{h}}$. Obviously $\delta_{jj'}$ is bounded by 1.

[7]Observe that $n_{\text{repl}} \ll \sqrt{n_{\text{V}}^{\text{h}}}$ covers the interesting cases because the privacy loss is obviously close to 1 for large n_{repl}.

game-based definitions. Various works have sought to address this by considering *entropy-based* privacy definitions, most notably [6] and [35] (we note that the game-based vote privacy definitions in the line of [5], which are often used to formally analyse vote privacy—see, e.g., [5, 10, 13, 26]—reduce to the entropy-based approach, as proven in [5]). In this section we will show that a simple extension of the KTV definition, which we term 'strong vote privacy', is equivalent to a computational version of a strong entropy-based definition.

A. Related work

In [6] at CCS 2012, the authors propose a family of entropy-based privacy definitions, parametrised by a number of modeling choices that must be made: firstly, we fix a distribution on the votes, both of the observed voter and the innocent third parties; secondly, we fix a 'target function' that the adversary is interested in learning, for example the observed voter's vote; thirdly, we must choose an 'entropy notion' to measure the adversary's success in learning the target information, for example the 'average min-entropy', which measures the adversary's ability to guess the value of the target function with a single guess. The privacy measure is then the posterior vulnerability of the target function with respect to the chosen entropy notion (for the examples just mentioned, that is the probability that the adversary will correctly guess the observed voter's vote after seeing the election). The setting of a computationally bounded adversary is dealt with indirectly, by saying that the vulnerability of an output distribution is the minimum vulnerability among distributions computationally indistinguishable from the true distribution.

The assumption that we know in advance the voting behaviour of the innocent third-party voters may well be reasonable—we can estimate this based on opinion polls and the results of previous elections—and the same assumption is made both in the KTV definition and in our definition below. The same assumption for the observed voter is more problematic, since the adversary may well choose to attack a highly atypical voter about whom they possess side-information (for example, someone who is known or suspected to belong to an opposition group). Note, however, that some such assumption seems inevitable in any measure focusing exclusively on posterior vulnerability, since in the case that the observed voter follows a point distribution (i.e. the adversary has total knowledge of her vote) any target function will be totally compromised.

The choice of target function may also be far from straightforward: while the observed voter's vote is certainly very reasonable, it may not be the only thing the adversary could care about. For example, if candidates are grouped into parties the adversary may only care about which party a voter voted for, or in the case of ranked choice voting the adversary may care about who the voter ranked first. Again some assumption of this kind seems necessary in any posterior vulnerability measure, since for instance a target which is a constant function will trivially be compromised.

Finally, the choice of entropy notion may not be entirely clear, but we agree with the use of average min-entropy because of its clear operational interpretation as the adversary's best single guess for the target.

The idea of the adversary's target function is further developed in [35], which considers the voting system as a communication channel from the voters to the result, and applies ideas from the theory of Quantified information flow, in particular the notion of *g-leakage*. The idea is to define a set $\mathcal{W}$ of possible guesses for the adversary, and then a *gain function* $g : \mathcal{W} \times \mathcal{C} \to [0, 1]$ quantifying the adversary's reward for making guess $w \in \mathcal{W}$ where the true choice was $c \in \mathcal{C}$. We can then define vulnerability with respect to g as the maximum possible expected payoff for a single guess by the adversary, and the *g-leakage* as the ratio of the vulnerabilities before and after seeing the tally. The authors show how to represent natural targets for the adversary (such as a specific voter's vote, or the number of voters whose vote the adversary can guess) by suitable gain functions and illustrate this for small toy example elections.

However, the framework of [35] also has a number of important limitations. It does not attempt at all to consider either adversaries who may interact during the protocol, or computationally bounded adversaries. It is still parametrised by choices of both gain function and vote distribution. Indeed, this is some sense inevitable in that we cannot hope to quantify loss of privacy by the overall capacity of the channel (which would be the supremum over possible gain functions and vote distributions), because the tally *does* reveal substantial information about the joint distribution of the voters' votes—this is the whole point of running an election!

B. Strong vote privacy

The key idea for our definition is to think of the voting system not as a channel from the votes of *all* the voters, but rather as a *noisy* channel from the vote of the observed voter, with noise coming from the random votes of the innocent third-party voters. This means that we can measure the loss of vote privacy as the min-entropy capacity of this channel, and so we can consider the maximum over *all* possible gain functions and *all* possible priors on the observed voter.

In order to allow for interactive and computationally bounded adversaries, we will phrase our definition not in the language of channels and information flow, but directly using the operational interpretation of min-entropy, in terms of the maximum advantage that can be gained by a computationally bounded interactive adversary—essentially an interactive and computational version of *g-leakage* as discussed above.

Definition 2 (Strong vote privacy): Let P be a voting protocol, $\mathsf{V}_{\mathrm{obs}}$ be the voter under observation, and $\delta \in [0, 1]$. Then, P *achieves strong δ-privacy*, if for all possible probability distributions π over the set of choices $\mathcal{C}$, all finite sets $\mathcal{W}$ and 'gain functions' $g : \mathcal{C} \times \mathcal{W} \to [0, 1]$ and all adversaries π_{A} the ratio

$$\frac{\sum_{i,w} \pi(c_i) \Pr[(\hat{\pi}_{\mathsf{V}_{\mathrm{obs}}}(c_i)\|\pi^*\|\pi_{\mathsf{A}})^{(\ell)} \mapsto w] g(c_i, w)}{\max_w \sum_i \pi(c_i) g(c_i, w)}$$

is $(1 + \delta)$-bounded as a function of the security parameter 1^ℓ.

Remarkably, it turns out that this is equivalent to a simple strengthening of the KTV privacy definition:

Definition 3 (Strong vote privacy, II): Let P be a voting protocol, $\mathsf{V}_{\mathrm{obs}}$ be the voter under observation, and $\delta \in [0, 1]$. Then, P *achieves strong δ-privacy*, if for all adversaries π_{A} the sum

$$\sum_i \Pr[(\hat{\pi}_{\mathsf{V}_{\mathrm{obs}}}(c_i)\|\pi^*\|\pi_{\mathsf{A}})^{(\ell)} \mapsto c_i]$$

is $(1 + \delta)$-bounded as a function of the security parameter 1^ℓ.

The proof of this equivalence, which will be the main theorem of this section, is essentially a computational and interactive version of the 'miracle' theorem of QIF, Theorem 5.1 of [2]. We first note that Definition 3 is indeed an extension of the KTV definition:

Proposition 1: KTV vote privacy (Definition 1) is equivalent to Definition 3 with the adversary π_{A} restricted to two possible outputs.

Proof. Let π_{A} have outputs $\{c_j, c_{j'}\}$. Then

$$\sum_i \Pr[(\hat{\pi}_{\mathsf{V}_{\mathrm{obs}}}(c_i)\|\pi^*\|\pi_{\mathsf{A}})^{(\ell)} \mapsto c_i]$$
$$= \Pr[(\hat{\pi}_{\mathsf{V}_{\mathrm{obs}}}(c_j)\|\pi^*\|\pi_{\mathsf{A}})^{(\ell)} \mapsto c_j]$$
$$\quad + \Pr[(\hat{\pi}_{\mathsf{V}_{\mathrm{obs}}}(c_{j'})\|\pi^*\|\pi_{\mathsf{A}})^{(\ell)} \mapsto c_{j'}]$$
$$= \Pr[(\hat{\pi}_{\mathsf{V}_{\mathrm{obs}}}(c_j)\|\pi^*\|\pi_{\mathsf{A}})^{(\ell)} \mapsto c_j]$$
$$\quad + \left(1 - \Pr[(\hat{\pi}_{\mathsf{V}_{\mathrm{obs}}}(c_{j'})\|\pi^*\|\pi_{\mathsf{A}})^{(\ell)} \mapsto c_j]\right)$$
$$= \Pr[(\hat{\pi}_{\mathsf{V}_{\mathrm{obs}}}(c_j)\|\pi^*\|\pi_{\mathsf{A}})^{(\ell)} \mapsto c_j]$$
$$\quad - \Pr[(\hat{\pi}_{\mathsf{V}_{\mathrm{obs}}}(c_{j'})\|\pi^*\|\pi_{\mathsf{A}})^{(\ell)} \mapsto c_j] + 1,$$

as required.

It trivially follows that the KTV privacy loss is a lower bound for the strong privacy loss, and that for two-candidate elections the definitions are equivalent.

We now establish the main theorem:

Theorem 5: Definitions 2 and 3 are equivalent.

Proof. Trivially Definition 2 implies Definition 3 (take $\mathcal{W} = \mathcal{C}$, π the uniform distribution and $g(c, c') = 1$ if $c' = c$ and 0 otherwise).

To prove the converse implication, let $\mathcal{W}$ and g be fixed, and let π_{A} be an adversary for which the quantity in Definition 2 exceeds $1 + \delta' > 1 + \delta$ infinitely often. Write $A^{(\ell)}(c_i, w) = \Pr[(\hat{\pi}_{\mathsf{V}_{\mathrm{obs}}}(c_i)\|\pi^*\|\pi_{\mathsf{A}})^{(\ell)} \mapsto w]$.

Our task is now to construct an adversary $\pi_{\mathsf{A}'}$ for which the quantity in Definition 3 exceeds $1 + \delta''$ infinitely often, for some $\delta'' > \delta$. The general idea is for $\pi_{\mathsf{A}'}$ to imitate π_{A}, except that at the end when π_{A} would output $w \in \mathcal{W}$, we will have $\pi_{\mathsf{A}'}$ output the c_i which maximises $A^{(\ell)}(c_i, w)$.

Note, however, that the values of $A^{(\ell)}(c_i, w)$ are an infite family of data (parametrised by ℓ), and so it is not possible to 'hard-code' them into the finite specification of the adversary $\pi_{\mathsf{A}'}$. It is therefore necessary for $\pi_{\mathsf{A}'}$ to estimate them on-the-fly, by simulating the behaviour of $(\hat{\pi}_{\mathsf{V}_{\mathrm{obs}}}(c_i)\|\pi^*\|\pi_{\mathsf{A}})^{(\ell)}$ for each c_i.

Define the adversary $\pi_{\mathsf{A}'}$ as follows: first simulate $(\hat{\pi}_{\mathsf{V}_{\mathrm{obs}}}(c_i)\|\pi^*\|\pi_{\mathsf{A}})^{(\ell)}$ with $\ell^2|\mathcal{W}|$ trials for each c_i to obtain estimates $\widetilde{A}(c_i,w)$ for $A^{(\ell)}(c_i,w)$. By the Chernoff bound on the sample mean we have that $|\widetilde{A}(c_i,w)-A^{(\ell)}(c_i,w)| < 1/|\mathcal{W}|\sqrt{\ell}$ with probability at least $1-2^{-\ell}$. Then $\pi_{\mathsf{A}'}$ behaves as π_{A} (in the real run of the protocol), and when π_{A} would output w, $\pi_{\mathsf{A}'}$ outputs the c_i which maximises $\widetilde{A}(c_i,w)$.

By the definition of π_{A} we have that infinitely often

$$(1+\delta')\max_w \sum_i \pi(c_i)g(c_i,w) \leq \sum_{i,w} \pi(c_i)A^{(\ell)}(c_i,w)g(c_i,w)$$

$$\leq \sum_w \left[\left(\max_i A^{(\ell)}(c_i,w)\right)\sum_i \pi(c_i)g(c_i,w)\right]$$

$$\leq \sum_w \left[\left(\max_i A^{(\ell)}(c_i,w)\right)\left(\max_w \sum_i \pi(c_i)g(c_i,w)\right)\right]$$

$$= \left(\sum_w \max_i A^{(\ell)}(c_i,w)\right)\left(\max_w \sum_i \pi(c_i)g(c_i,w)\right),$$

and hence we have that $\sum_w \max_i A^{(\ell)}(c_i,w) \geq 1 + \delta'$ for infinitely many ℓ.

Writing $\phi(w)$ for the c_i which maximises $\widetilde{A}(c_i,w)$, we have for these ℓ

$$\sum_i \Pr[(\hat{\pi}_{\mathsf{V}_{\mathrm{obs}}}(c_i)\|\pi^*\|\pi_{\mathsf{A}'})^{(\ell)} \mapsto c_i]$$

$$= \sum_i \sum_w A^{(\ell)}(c_i,w)1_{c_i=\phi(w)}$$

$$= \sum_w A^{(\ell)}(\phi(w),w)$$

$$\geq \sum_w \left(\widetilde{A}(\phi(w),w) - 2^{-\ell} - 1/|W|\sqrt{\ell}\right)$$

$$= \sum_w \max_i \widetilde{A}(c_i,w) - \left(|\mathcal{W}|2^{-\ell} + 1/\sqrt{\ell}\right)$$

$$\geq \sum_w \max_i A^{(\ell)}(c_i,w) - 2\left(|\mathcal{W}|2^{-\ell} + 1/\sqrt{\ell}\right)$$

$$\geq 1 + \delta' - 2\left(|\mathcal{W}|2^{-\ell} + 1/\sqrt{\ell}\right) \xrightarrow{\ell\to\infty} 1 + \delta'.$$

C. Monte Carlo estimation

Precise analysis of strong vote privacy for the ideal functionality discussed above is rather less straightforward than for the KTV definition (partly since unlike the latter it depends on the entire distribution of the honest voters, rather than only with respect to the two least popular candidates). However, for the ideal replay attack functionality in which the adversary's only action is to make a guess based on the output tally it is possible to obtain fairly accurate estimates by Monte Carlo methods, as we now show. The first observation is that the optimal output is one which can be easily simulated.

Proposition 2: For a protocol for which the adversary's output is a function on some finite set of tallies T, the sum in Definition 3 is maximised by the 'maximum likelihood adversary', which on tally t outputs the vote c_i which maximises $p_{T|C}(t|c_i)$ (breaking ties arbitrarily).
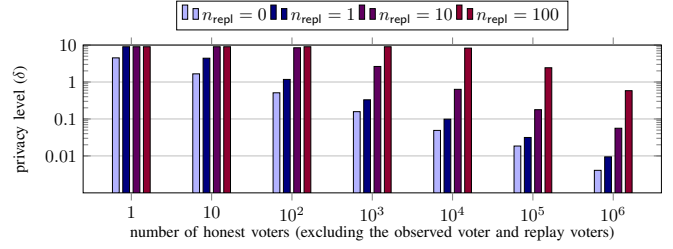

Fig. 4. Strong privacy loss δ for the ideal protocol with 10 candidates and the uniform vote distribution.

Proof. We have

$$\sum_i \Pr[(\hat{\pi}_{\mathsf{V}_{\mathrm{obs}}}(c_i)\|\pi^*\|\pi_{\mathsf{A}})^{(\ell)} \mapsto c_i] = \sum_i \sum_t p(t|c_i)1_{\pi_{\mathsf{A}}(t)=c_i}$$

$$= \sum_t p(t|\pi_{\mathsf{A}}(t)) \leq \sum_t \max_c p(t|c)$$

$$= \sum_i \Pr[(\hat{\pi}_{\mathsf{V}_{\mathrm{obs}}}(c_i)\|\pi^*\|\pi_{\widetilde{\mathsf{A}}})^{(\ell)} \mapsto c_i]$$

as required, where $\pi_{\widetilde{\mathsf{A}}}$ is the maximum likelihood adversary.

Note that the maximum likelihood adversary can be easily implemented for the ideal protocol of the previous section: for an election with k candidates, n_{repl} replay voters, and honest voters who cast their votes with probabilities $(p_1,\ldots,p_k)$, for a tally $t = (m_1,\ldots,m_k)$ we have

$$p(t|c_i) = \binom{m_1 + \ldots + m_k - (n_{\mathsf{repl}}+1)}{m_1,\ldots,m_{i-1},m_i - (n_{\mathsf{repl}}+1),m_{i+1},\ldots,m_k}$$
$$p_1^{m_1}\cdots p_{i-1}^{m_{i-1}}p_i^{m_i-(n_{\mathsf{repl}}+1)}p_{i+1}^{m_{i+1}}\cdots p_k^{m_k}.$$

By computing this for each c_i, we can find $\pi_{\widetilde{\mathsf{A}}}(t)$ for given t.

Now observe that if c_i is uniformly distributed, and t is drawn according to $p(\cdot|c_i)$ then we have that $1_{\pi_{\widetilde{\mathsf{A}}}(t)=c_i} \sim$ Bernoulli$((1+\delta)/k)$, where δ is the privacy loss of Def. 3.

Then to estimate δ we repeatedly sample c_i uniformly at random and simulate a tally t with the observed voter (hence also the replay voters) voting for c_i, and then check whether $\pi_{\widetilde{\mathsf{A}}}(t) = c_i$. If this occurs with frequency ρ then $k\rho - 1$ is an unbiased estimator for δ with standard error at most $(1+\delta)/\sqrt{n}$ (where n is the number of trials).

Figure 4 shows the privacy loss of the ideal functionality for an election with 10 candidates and various numbers of honest and replay voters (with the honest voters voting according to the uniform distribution). Figure 5 shows a direct comparison of our definition with the KTV definition, for elections with between 2 and 10 candidates. All estimates in this section and in Section VI are with 10,000,000 trials, so standard error < 0.0005.

VI. ANALYSIS OF REAL-WORLD ELECTIONS

In order to complement our formal analysis, we study how replay attacks would scale in practical elections. We therefore apply our formal results to publicly available data of political elections in Estonia, Germany, the UK, and the USA. In this way, we can realistically simulate to which degree vote privacy would decrease if in such elections replay attacks had been

Fig. 5. Comparison of privacy losses for strong privacy and KTV privacy, for an election with 10000 honest voters, 10 replays and 2–10 candidates, uniform distribution.

executed. Our "field test" confirms the gist of our abstract results: even if the number of replays is very low, vote privacy can be undermined significantly.

In the remainder of this section, we first discuss the modeling assumptions used for our analysis (Sec. VI-A), then describe our real-world examples and the results of our simulations (Sec. VI-B). Finally we discuss these results, explain why they confirm our theoretical analysis, and elaborate on the consequences of our new insights (Sec. VI-C).

A. Modeling assumptions

Throughout all of our analysis we, like Küsters, Truderung and Vogt in [30], assume that the adversary's knowledge about the actions of the honest voters is represented by a vector of vote probabilities $\vec{p}$, with the individual voters' votes drawn independently according to this distribution (other prior works in the literature [6, 35] base their examples on the simple but unrealistic assumption of a uniform prior). In applying this analysis to real-world elections, there are two similar-seeming but conceptually distinct issues to address: whether it is reasonable to think that *the adversary* knows this probability distribution, and how *we as analysts* estimate this distribution in order to perform the privacy analysis.

An important feature of all the elections we consider is that they are national elections for which results are published at the constituency or even per-polling station level. This means that if voters were nationally homogeneous then the adversary could easily discover the vote distribution $\vec{p}$ by just averaging the results of the constituencies other than the one in which he executed the replay attack.

Of course, in reality electorates are almost never nationally homogeneous, and there will be systematic variation between North and South, East and West, rich and poor regions, and so on. The adversary's goal, therefore, is to estimate the distribution $\vec{p}$ for the specific constituency (or polling district) he has attacked, using only the results from other constituencies (and perhaps also results from previous elections). Fortunately for him, there is a well-established technique in political science, called *Multilevel Regression and Poststratification (MRP)* [21], to predict the local result based on a combination of national polls, local demographic factors and previous results. It is difficult for us to make quantitative statements about the accuracy that could be achieved, since political scientists are generally interested in making predictions *before* the election using opinion polls rather than with access to actual results

outside the target constituency (and a full implementation would be far beyond the scope of this paper). Howerver, recent examples (e.g. [33]) are able to obtain local-level predictions with typical error comparable to the sampling uncertainty of national opinion polls, so it seems reasonable to expect (or at least fear) that with access to the actual national results (rather than only polls) fairly precise estimates could be obtained of the local underlying vote distribution.

The second task, for us as analysts to estimate the vote probabilities in order to perform the analysis, is considerably simpler. Unlike the adversary we have access to the actual results in the relevant area (unpolluted by replay attacks), and so we can use the proportion of votes cast for each candidate as an unbiased estimator of the underlying vote probabilities.

B. Examples

We use public data from political elections in Estonia, Germany, the UK, and the USA, to simulate the potential privacy loss if these elections had been conducted using an e-voting scheme vulnerable to replay attacks.[8]

In each of these elections, the partial election result of each polling station/area was published. We use these partial results to analyse the efficiency of replay attacks because it is reasonable to assume that an adversary knows in which partial result a targeted voter's choice is included. For each election, we chose a polling station/area where the number of votes was close to the overall average of votes per polling station/area. Our results are summarized in Fig. 6.

	Estonia		Germany		UK		US	
	KTV	SP	KTV	SP	KTV	SP	KTV	SP
0	0.266	0.485	0.103	0.187	0.003	0.003	0.222	0.331
1	0.502	1.076	0.204	0.389	0.006	0.006	0.426	0.695
5	0.950	3.811	0.560	1.354	0.017	0.017	0.875	2.103
10	0.999	6.235	0.839	2.698	0.031	0.032	0.992	3.034

Fig. 6. Ideal privacy losses with $n_{\text{repl}} = 0, 1, 5, 10$ replays based on real election data from Examples 1–4. "KTV" denotes KTV privacy definition and "SP" denotes strong privacy definition.

Example 1 (Estonia, Riigikogu Election 2019): In the Riigikogu (parliamentary) elections in 2019, 561,141 votes were cast in total. The number of polling stations was 451, which results in 1,244 votes per polling stations on average. In this example, we choose polling station S53P in Mustamäe linnaosa where 1,404 valid votes were cast.[9] The public partial election result at this polling station was $(14, 233, 22, 82, 9, 210, 31, 702, 5, 92, 4)$.[10]

Example 2 (Germany, Landtag Election 2021): In the Landtag (parliamentary) election in the state of Rhineland-Palatinate in 2021, 1,922,579 votes were cast in total. The number of polling stations was (roughly) 2,300, which results

---

[8]We published our implementation at http://hdl.handle.net/10993/51209.

[9]See https://rk2019.valimised.ee/en/voting-result/local-municipality-0482-voting-result.html (accessed 11.04.2022).

[10]The public election result is even more fine-grained because the number of votes per candidate on each party list is revealed. We aggregated the number of votes for each party list and consider adversaries who merely know the aggregated result; this makes our overall argument only stronger.

in < 836 votes per polling station on average. In this example, we choose polling station Pluwig where 855 votes were cast.[11] The public partial election result at this polling station was $(291, 253, 34, 35, 141, 27, 74)$.

Example 3 (UK, EU Referendum 2016): In the EU referendum in the UK 2016, 33,551,983 votes were cast in total. The number of areas was 382, which results in 87,832 votes per area on average. In this example, we choose the area of Kingston-upon-Thames (London) where 85,270 votes were cast.[12] The public partial election result at this polling station was $(52533, 32737)$.

Example 4 (USA, Presidential Election 2020): For the US presidential election in 2020, we were not able to determine the number of polling stations nationwide, so we focused on the results of one state, namely Massachusetts. Here, the average number of votes per polling station was (roughly) 1,500. In this example, we choose the polling station for Precinct 13 of Ward 1 in Boston, where 1,430 votes were cast.[13] The public partial election result at this polling station was $(995, 404, 12, 5, 9)$.

C. Discussion

Observe that for the Estonian, German and US elections, which have many candidates, even a very small number of replays would have a devastating impact on vote privacy. For example, in the Estonian election even a single replay already has a substantial effect, and with only 5 replays privacy is almost completely lost (similarly for Germany and the US 1–5 replays compromise privacy, and 5–10 destroy it completely).

On the other hand, in the UK Brexit referendum, which has only two 'candidates' and far more votes at the most granular reporting level, we see that the effect of up to 10 replays is far less. This example also illustrates most clearly the result of Theorem 4 that the KTV privacy loss scales approximately proportionately to $n_{\mathsf{repl}} + 1$, which we also see in the small-δ regions of the other examples (we also see the consequence of Proposition 1 that for a two-candidate election the KTV and strong privacy losses agree, apart from stochastic sampling error).

The referendum results should not make us too complacent, however, because in fact the number of replays required to obtain a KTV privacy loss $\delta > 1/2$ is just 196—equivalent to just 0.2% of the total number of votes.

It is interesting to compare the results for the US election with the UK results, because these are both elections for which the vast majority of votes went to just two candidates, but the privacy loss for the US is much greater than for the UK (by a factor of approximately 70). The number of votes in the UK example exceeds that in the US by a factor of approximately 60; Theorem 4 tells us that the KTV privacy loss δ scales as

$1/\sqrt{n_{\mathsf{V}}^{\mathsf{h}}}$, and so the discrepancy in total votes predicts a ratio of approximately only $\sqrt{60} \approx 8$, leaving around a factor of 10 still to explain. The reason for this remaining difference is that both the KTV and strong privacy definitions are heavily (in the case of strong privacy) or entirely (in the case of KTV privacy) influenced by the least popular candidates, and so the fact that the US election has a few *very* unpopular candidates has a large effect.

It may seem odd or even undesirable that the measured loss of privacy should be so heavily influenced by the small number of voters who support minority parties. However, we would argue that ballot privacy must mean privacy for *all* voters. Indeed, it may well be supporters of unpopular candidates who are at most risk of stigmatisation or reprisals; note that the adversary does not have to choose the targeted voter at random, but can choose to target someone they already suspect of supporting a minority position.

REFERENCES

[1] Ben Adida. Helios: Web-based Open-Audit Voting. In Paul C. van Oorschot, editor, *Proceedings of the 17th USENIX Security Symposium, July 28-August 1, 2008, San Jose, CA, USA*, pages 335–348. USENIX Association, 2008.

[2] Mário S. Alvim, Kostas Chatzikokolakis, Catuscia Palamidessi, and Geoffrey Smith. Measuring information leakage using generalized gain functions. In *Proceedings of the 2012 IEEE 25th Computer Security Foundations Symposium*, CSF '12, pages 265–279, USA, 2012. IEEE Computer Society.

[3] Gergei Bana, Marco Biroli, Megi Dervishi, Fatima-Ezzahra El Orche, Rémi Géraud-Stewart, David Naccache, Peter B. Rønne, Peter Y. A. Ryan, and Hugo Waltsburger. Time, Privacy, Robustness, Accuracy: Trade Offs for the Open Vote Network Protocol. *IACR Cryptol. ePrint Arch.*, page 1065, 2021.

[4] Josh Daniel Cohen Benaloh. *Verifiable Secret-Ballot Elections*. PhD thesis, 1987.

[5] David Bernhard, Véronique Cortier, David Galindo, Olivier Pereira, and Bogdan Warinschi. SoK: A Comprehensive Analysis of Game-Based Ballot Privacy Definitions. In *2015 IEEE Symposium on Security and Privacy,*

---

[11] See https://www.wahlen.rlp.de/de/ltw/wahlen/2021/ergebnisse/2242350410700.html (accessed 11.04.2022).

[12] See https://www.electoralcommission.org.uk/who-we-are-and-what-we-do/elections-and-referendums/past-elections-and-referendums/eu-referendum (accessed 11.04.2022).

[13] See https://electionstats.state.ma.us/elections/view/140751/filter_by_county:Suffolk (accessed 11.04.2022).

SP 2015, San Jose, CA, USA, May 17-21, 2015, pages 499–516, 2015.

[6] David Bernhard, Véronique Cortier, Olivier Pereira, and Bogdan Warinschi. Measuring vote privacy, revisited. In Ting Yu, George Danezis, and Virgil D. Gligor, editors, *ACM Conference on Computer and Communications Security (CCS 2012)*, pages 941–952. ACM, 2012.

[7] David Bernhard, Olivier Pereira, and Bogdan Warinschi. How Not to Prove Yourself: Pitfalls of the Fiat-Shamir Heuristic and Applications to Helios. In Xiaoyun Wang and Kazue Sako, editors, *Advances in Cryptology - ASIACRYPT 2012 - 18th International Conference on the Theory and Application of Cryptology and Information Security, Beijing, China, December 2-6, 2012. Proceedings*, volume 7658 of *Lecture Notes in Computer Science*, pages 626–643. Springer, 2012.

[8] Olivier Blazy, Georg Fuchsbauer, David Pointcheval, and Damien Vergnaud. Signatures on Randomizable Ciphertexts. In Dario Catalano, Nelly Fazio, Rosario Gennaro, and Antonio Nicolosi, editors, *Public Key Cryptography - PKC 2011 - 14th International Conference on Practice and Theory in Public Key Cryptography, Taormina, Italy, March 6-9, 2011. Proceedings*, volume 6571 of *Lecture Notes in Computer Science*, pages 403–422. Springer, 2011.

[9] Xavier Boyen, Thomas Haines, and Johannes Müller. Epoque: Practical End-to-End Verifiable Post-Quantum-Secure E-Voting. In *IEEE European Symposium on Security and Privacy, EuroS&P 2021, Vienna, Austria, September 6-10, 2021*, pages 272–291. IEEE, 2021.

[10] Sergiu Bursuc, Constantin Catalin Dragan, and Steve Kremer. Private Votes on Untrusted Platforms: Models, Attacks and Provable Scheme. In *IEEE European Symposium on Security and Privacy, EuroS&P 2019, Stockholm, Sweden, June 17-19, 2019*, pages 606–620. IEEE, 2019.

[11] Craig Burton, Chris Culnane, James Heather, Thea Peacock, Peter Y. A. Ryan, Steve Schneider, Vanessa Teague, Roland Wen, Zhe Xia, and Sriramkrishnan Srinivasan. Using Prêt à Voter in Victoria State Elections. In J. Alex Halderman and Olivier Pereira, editors, *2012 Electronic Voting Technology Workshop / Workshop on Trustworthy Elections, EVT/WOTE '12, Bellevue, WA, USA, August 6-7, 2012*. USENIX Association, 2012.

[12] Richard Carback, David Chaum, Jeremy Clark, John Conway, Aleksander Essex, Paul S. Herrnson, Travis Mayberry, Stefan Popoveniuc, Ronald L. Rivest, Emily Shen, Alan T. Sherman, and Poorvi L. Vora. Scantegrity II Municipal Election at Takoma Park: The First E2E Binding Governmental Election with Ballot Privacy. In *19th USENIX Security Symposium, Washington, DC, USA, August 11-13, 2010, Proceedings*, pages 291–306. USENIX Association, 2010.

[13] Pyrros Chaidos, Véronique Cortier, Georg Fuchsbauer, and David Galindo. BeleniosRF: A Non-interactive Receipt-Free Electronic Voting Scheme. In Edgar R.

Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, Vienna, Austria, October 24-28, 2016*, pages 1614–1625. ACM, 2016.

[14] David Chaum, Aleksander Essex, Richard Carback, Jeremy Clark, Stefan Popoveniuc, Alan T. Sherman, and Poorvi L. Vora. Scantegrity: End-to-End Voter-Verifiable Optical-Scan Voting. *IEEE Secur. Priv.*, 6(3):40–46, 2008.

[15] David Chaum, Peter Y. A. Ryan, and Steve A. Schneider. A Practical Voter-Verifiable Election Scheme. In Sabrina De Capitani di Vimercati, Paul F. Syverson, and Dieter Gollmann, editors, *Computer Security - ESORICS 2005, 10th European Symposium on Research in Computer Security, Milan, Italy, September 12-14, 2005, Proceedings*, volume 3679 of *Lecture Notes in Computer Science*, pages 118–139. Springer, 2005.

[16] Michael R. Clarkson, Stephen Chong, and Andrew C. Myers. Civitas: Toward a Secure Voting System. In *2008 IEEE Symposium on Security and Privacy (S&P 2008), 18-21 May 2008, Oakland, California, USA*, pages 354–368. IEEE Computer Society, 2008.

[17] Véronique Cortier, David Galindo, Ralf Küsters, Johannes Müller, and Tomasz Truderung. SoK: Verifiability Notions for E-Voting Protocols. In *IEEE Symposium on Security and Privacy, SP 2016, San Jose, CA, USA, May 22-26, 2016*, pages 779–798, 2016.

[18] Véronique Cortier and Ben Smyth. Attacking and Fixing Helios: An Analysis of Ballot Secrecy. In *Proceedings of the 24th IEEE Computer Security Foundations Symposium, CSF 2011, Cernay-la-Ville, France, 27-29 June, 2011*, pages 297–311. IEEE Computer Society, 2011.

[19] Jannik Dreier, Pascal Lafourcade, and Yassine Lakhnech. Vote-Independence: A Powerful Privacy Notion for Voting Protocols. In Joaquín García-Alfaro and Pascal Lafourcade, editors, *Foundations and Practice of Security - 4th Canada-France MITACS Workshop, FPS 2011, Paris, France, May 12-13, 2011, Revised Selected Papers*, volume 6888 of *Lecture Notes in Computer Science*, pages 164–180. Springer, 2011.

[20] Amos Fiat and Adi Shamir. How to Prove Yourself: Practical Solutions to Identification and Signature Problems. In Andrew M. Odlyzko, editor, *Advances in Cryptology - CRYPTO '86, Santa Barbara, California, USA, 1986, Proceedings*, volume 263 of *Lecture Notes in Computer Science*, pages 186–194. Springer, 1986.

[21] Andrew Gelman and Thomas C Little. Poststratification into many categories using hierarchical logistic regression. *Survey Methodology*, 46(1), 1997.

[22] B.V. Gnedenko. *Theory of Probability*. Taylor & Francis, 6 edition, 1998.

[23] Helios Voting. Attacks and Defenses. https://documentation.heliosvoting.org/attacks-and-defenses (accessed 11.04.2022).

[24] Lucca Hirschi, Lara Schmid, and David A. Basin. Fixing

the Achilles Heel of E-Voting: The Bulletin Board. In *34th IEEE Computer Security Foundations Symposium, CSF 2021, Dubrovnik, Croatia, June 21-25, 2021*, pages 1–17. IEEE, 2021.

[25] IACR. IACR Elections, 2020. https://www.iacr.org/elections/ (accessed 11.04.2022).

[26] Vincenzo Iovino, Alfredo Rial, Peter B. Rønne, and Peter Y. A. Ryan. Universal Unconditional Verifiability in E-Voting without Trusted Parties. In *33rd IEEE Computer Security Foundations Symposium, CSF 2020, Boston, MA, USA, June 22-26, 2020*, pages 33–48. IEEE, 2020.

[27] Shahram Khazaei and Douglas Wikström. Randomized Partial Checking Revisited. In Ed Dawson, editor, *Topics in Cryptology - CT-RSA 2013 - The Cryptographers' Track at the RSA Conference 2013, San Francisco,CA, USA, February 25-March 1, 2013. Proceedings*, volume 7779 of *Lecture Notes in Computer Science*, pages 115–128. Springer, 2013.

[28] Ralf Küsters, Julian Liedtke, Johannes Müller, Daniel Rausch, and Andreas Vogt. Ordinos: A Verifiable Tally-Hiding E-Voting System. In *IEEE European Symposium on Security and Privacy, EuroS&P 2020, Genoa, Italy, September 7-11, 2020*, pages 216–235. IEEE, 2020.

[29] Ralf Küsters, Johannes Müller, Enrico Scapin, and Tomasz Truderung. sElect: A Lightweight Verifiable Remote Voting System. In *IEEE 29th Computer Security Foundations Symposium, CSF 2016, Lisbon, Portugal, June 27 - July 1, 2016*, pages 341–354, 2016.

[30] Ralf Küsters, Tomasz Truderung, and Andreas Vogt. Verifiability, Privacy, and Coercion-Resistance: New Insights from a Case Study. In *32nd IEEE Symposium on Security and Privacy, S&P 2011, 22-25 May 2011, Berkeley, California, USA*, pages 538–553, 2011.

[31] Ralf Küsters, Tomasz Truderung, and Andreas Vogt. Verifiability, Privacy, and Coercion-Resistance: New Insights from a Case Study. *IACR Cryptol. ePrint Arch.*, 2011:517, 2011.

[32] Ralf Küsters, Tomasz Truderung, and Andreas Vogt. Formal Analysis of Chaumian Mix Nets with Randomized Partial Checking. In *2014 IEEE Symposium on Security and Privacy, SP 2014, Berkeley, CA, USA, May 18-21, 2014*, pages 343–358, 2014.

[33] Benjamin E Lauderdale, Delia Bailey, Jack Blumenau, and Douglas Rivers. Model-based pre-election polling for national and sub-national outcomes in the us and uk. *International Journal of Forecasting*, 36(2):399–413, 2020.

[34] Byoungcheon Lee, Colin Boyd, Ed Dawson, Kwangjo Kim, Jeongmo Yang, and Seungjae Yoo. Providing Receipt-Freeness in Mixnet-Based Voting Protocols. In Jong In Lim and Dong Hoon Lee, editors, *Information Security and Cryptology - ICISC 2003, 6th International Conference, Seoul, Korea, November 27-28, 2003, Revised Papers*, volume 2971 of *Lecture Notes in Computer Science*, pages 245–258. Springer, 2003.

[35] Annabelle McIver, Tahiry Rabehaja, Roland Wen, and Carroll Morgan. Privacy in elections: How small is "small"? *J. Inf. Secur. Appl.*, 36(C):112–126, October 2017.

[36] Kazue Sako and Joe Kilian. Secure Voting Using Partially Compatible Homomorphisms. In Yvo Desmedt, editor, *Advances in Cryptology - CRYPTO '94, 14th Annual International Cryptology Conference, Santa Barbara, California, USA, August 21-25, 1994, Proceedings*, volume 839 of *Lecture Notes in Computer Science*, pages 411–424. Springer, 1994.

[37] Berry Schoenmakers. A Simple Publicly Verifiable Secret Sharing Scheme and Its Application to Electronic Voting. In Michael J. Wiener, editor, *Advances in Cryptology - CRYPTO '99, 19th Annual International Cryptology Conference, Santa Barbara, California, USA, August 15-19, 1999, Proceedings*, volume 1666 of *Lecture Notes in Computer Science*, pages 148–164. Springer, 1999.

[38] Smartmatic. Online Voting Successfully Solving the Challenges, 2021. https://www.smartmatic.com/fileadmin/user_upload/Whitepaper_Online_Voting_Challenge_Considerations_TIVI.pdf (accessed 11.04.2022).

[39] Ben Smyth. Replay attacks that violate ballot secrecy in Helios. *IACR Cryptol. ePrint Arch.*, 2012:185, 2012.

[40] Swiss Post. E-voting and security, 2021. https://www.post.ch/en/business-solutions/e-voting/security-given-top-priority (accessed 11.04.2022).

APPENDIX A
PROOFS

Lemma 2: Let $j, j \in C$. For each $t \in \{0, \ldots, n_{\mathsf{V}}^{\mathsf{h}} + n_{\mathsf{repl}} + 1\}$ let $T_{t,j,j'} \in \mathbb{N}$ be the largest number such that $\prod_{\nu=0}^{n_{\mathsf{repl}}} \frac{T_{t,j,j'} - \nu}{\vec{p}[j]^{n_{\mathsf{repl}}+1}} \leq \prod_{\nu=0}^{n_{\mathsf{repl}}} \frac{t - T_{t,j,j'} - \nu}{\vec{p}[j']^{n_{\mathsf{repl}}+1}}$. Then for $\vec{p}[j, j'] := 1 - \vec{p}[j] - \vec{p}[j']$

$$\sum_{t=0}^{n_{\mathsf{V}}^{\mathsf{t}}} \sum_{r=0}^{t} \left| \Delta_{j,j'}^{res,n_{\mathsf{repl}}} \right| \cdot \frac{n_{\mathsf{V}}^{\mathsf{h}}! \vec{p}[j, j']^{n_{\mathsf{V}}^{\mathsf{t}}-t} \vec{p}[j]^{r} \vec{p}[j']^{r-t}}{2 \cdot (n_{\mathsf{V}}^{\mathsf{t}} - t)! r! (t-r)!} \quad (8)$$

$$= \sum_{t=0}^{n_{\mathsf{V}}^{\mathsf{h}}} \vec{p}[j, j']^{n_{\mathsf{V}}^{\mathsf{h}}-t} \binom{n_{\mathsf{V}}^{\mathsf{h}}}{t} \sum_{r=\max\{T_{t+n_{\mathsf{repl}}+1}-n_{\mathsf{repl}}, 0\}}^{\min\{T_{t+n_{\mathsf{repl}}+1}, t\}} \binom{t}{r} \vec{p}[j]^{r} \vec{p}[j']^{t-r}$$

with $\Delta_{j,j'}^{res,n_{\mathsf{repl}}} = \frac{\prod_{\nu=0}^{n_{\mathsf{repl}}}(r-\nu)}{\vec{p}[j]^{n_{\mathsf{repl}}+1}} - \frac{\prod_{\nu=0}^{n_{\mathsf{repl}}}(t-r-\nu)}{\vec{p}[j']^{n_{\mathsf{repl}}+1}}$.

Proof. We first remove the absolute value to get

$$\sum_{r=0}^{t} \left| \Delta_{j,j'}^{res,n_{\mathsf{repl}}} \right| \cdot \frac{\vec{p}[j]^{r} \vec{p}[j']^{t-r}}{r!(t-r)!} \quad (9)$$

$$= \sum_{r=0}^{T_t} \Delta_{j,j'}^{res,n_{\mathsf{repl}}} \cdot \frac{\vec{p}[j]^{r} \vec{p}[j']^{t-r}}{r!(t-r)!} - \sum_{r=T_t+1}^{t} \Delta_{j,j'}^{res,n_{\mathsf{repl}}} \cdot \frac{\vec{p}[j]^{r} \vec{p}[j']^{t-r}}{r!(t-r)!}$$

Check that the term is 0 for $T_t \leq t < n_{\mathsf{repl}} + 1$ and $r \leq T_t, t - r \leq t$. By index shifting we get for the different summands

of $\Delta_{j,j'}^{\mathrm{res},n_{\mathrm{repl}}}$ and $-\Delta_{j,j'}^{\mathrm{res},n_{\mathrm{repl}}}$ for $t \geq n_{\mathrm{repl}} + 1$:

$$\sum_{r=0}^{T_t} \frac{\prod_{\nu=0}^{n_{\mathrm{repl}}}(t-r-\nu)}{\vec{p}[j']^{n_{\mathrm{repl}}+1}/t!} \, \mathrm{MN}_t^{r,t-r}$$
$$= \sum_{r=0}^{\min\{T_t,t-n_{\mathrm{repl}}-1\}} \frac{\mathrm{MN}_{t-n_{\mathrm{repl}}-1}^{r,t-r-n_{\mathrm{repl}}-1}}{(t-n_{\mathrm{repl}}-1)!}$$

$$\sum_{r=0}^{T_t} \frac{\prod_{\nu=0}^{n_{\mathrm{repl}}}(r-\nu)}{\vec{p}[j]^{n_{\mathrm{repl}}+1}/t!} \, \mathrm{MN}_t^{r,t-r}$$
$$= \sum_{r=0}^{T_t-n_{\mathrm{repl}}-1} \frac{\mathrm{MN}_{t-n_{\mathrm{repl}}-1}^{r,t-r-n_{\mathrm{repl}}-1}}{(t-n_{\mathrm{repl}}-1)!}$$

$$\sum_{r=T_t+1}^{t} \frac{\prod_{\nu=0}^{n_{\mathrm{repl}}}(t-r-\nu)}{\vec{p}[j']^{n_{\mathrm{repl}}+1}/t!} \, \mathrm{MN}_t^{r,t-r}$$
$$= \sum_{r=T_t+1}^{t-n_{\mathrm{repl}}-1} \frac{\mathrm{MN}_{t-n_{\mathrm{repl}}-1}^{r,t-r-n_{\mathrm{repl}}-1}}{(t-n_{\mathrm{repl}}-1)!}$$

$$\sum_{r=T_t+1}^{t} \frac{\prod_{\nu=0}^{n_{\mathrm{repl}}}(r-\nu)}{\vec{p}[j]^{n_{\mathrm{repl}}+1}/t!} \, \mathrm{MN}_t^{r,t-r}$$
$$= \sum_{r=\max\{T_t-n_{\mathrm{repl}},0\}}^{t-n_{\mathrm{repl}}-1} \frac{\mathrm{MN}_{t-n_{\mathrm{repl}}-1}^{r,t-r-n_{\mathrm{repl}}-1}}{(t-n_{\mathrm{repl}}-1)!}$$

Combining these terms shows that (9) is equal to

$$2 \sum_{r=\max\{T_t-n_{\mathrm{repl}},0\}}^{\min\{T_t,t-n_{\mathrm{repl}}-1\}} \frac{\mathrm{MN}_{t-n_{\mathrm{repl}}-1}^{r,t-r-n_{\mathrm{repl}}-1}}{(t-n_{\mathrm{repl}}-1)!}$$

Thus (8) becomes

$$\sum_{t=n_{\mathrm{repl}}+1}^{n_{\mathsf{V}}^{\mathsf{t}}} \frac{n_{\mathsf{V}}^{\mathsf{h}}!\vec{p}[j,j']^{n_{\mathsf{V}}^{\mathsf{t}}-t}}{(n_{\mathsf{V}}^{\mathsf{t}}-t)!} \sum_{r=\max\{T_t-n_{\mathrm{repl}},0\}}^{\min\{T_t,t-n_{\mathrm{repl}}-1\}} \frac{\mathrm{MN}_{t-n_{\mathrm{repl}}-1}^{r,t-r-n_{\mathrm{repl}}-1}}{(t-n_{\mathrm{repl}}-1)!}$$

Another index shift shows the equality in (8):

$$\sum_{t=0}^{n_{\mathsf{V}}^{\mathsf{h}}} \frac{n_{\mathsf{V}}^{\mathsf{h}}!\vec{p}[j,j']^{n_{\mathsf{V}}^{\mathsf{h}}-t}}{(n_{\mathsf{V}}^{\mathsf{h}}-t)!} \sum_{r=\max\{T_{t+n_{\mathrm{repl}}+1}-n_{\mathrm{repl}},0\}}^{\min\{T_{t+n_{\mathrm{repl}}+1},t\}} \frac{\mathrm{MN}_t^{r,t-r}}{t!}$$

Proof of Corollary 3. Note that by definition $\delta_{jj'}$ is smooth in $\vec{p}[j] + \vec{p}[j']$. Hence its representation in (5) is also smooth and it is enough to compute the differential on intervals where T_t is constant. We will differentiate (5) w.r.t. $\vec{p}[j] + \vec{p}[j']$ to see that $\delta_{jj'}$ decreases, i.e. becomes maximal if $\vec{p}[j] + \vec{p}[j']$ is minimal. We get $\frac{\partial}{\partial(\vec{p}[j]+\vec{p}[j'])} f(\vec{p}[j], \vec{p}[j'])$ with $f(\vec{p}[j], \vec{p}[j']) = (1-\vec{p}[j]-\vec{p}[j'])^{n_{\mathsf{V}}^{\mathsf{h}}-t}(\vec{p}[j])^r (\vec{p}[j'])^{t-r}$ to

$$\left(\frac{r}{2\vec{p}[j]} + \frac{t-r}{2\vec{p}[j]} - \frac{n_{\mathsf{V}}^{\mathsf{h}}-t}{\vec{p}[j,j']}\right) f(\vec{p}[j], \vec{p}[j'])$$

Hence we get the three terms:

$$-\sum_{t=0}^{n_{\mathsf{V}}^{\mathsf{h}}-1} \binom{n_{\mathsf{V}}^{\mathsf{h}}}{t} \frac{\vec{p}[j,j']^{n_{\mathsf{V}}^{\mathsf{h}}-t-1}}{(n_{\mathsf{V}}^{\mathsf{h}}-t)^{-1}} \sum_{r=\max\{T_{t+n_{\mathrm{repl}}+1}-n_{\mathrm{repl}},0\}}^{\min\{T_{t+n_{\mathrm{repl}}+1},t\}} \mathrm{MN}_t^{r,t-r}$$

$$\frac{1}{2} \sum_{r=\max\{T_{t+n_{\mathrm{repl}}+1}-n_{\mathrm{repl}}-1,0\}}^{\min\{T_{t+n_{\mathrm{repl}}+1},t\}-1} \binom{t}{r}(t-r)\cdot\vec{p}[j']^{-1}\,\mathrm{MN}_t^{r,t-r}$$

$$\frac{1}{2} \sum_{r=\max\{T_{t+n_{\mathrm{repl}}+1}-n_{\mathrm{repl}},0\}}^{\min\{T_{t+n_{\mathrm{repl}}+1},t-1\}} \binom{t}{r}(t-r)\cdot\vec{p}[j']^{-1}\,\mathrm{MN}_t^{r,t-r}$$

The last two terms are 0 for $t = 0$ and we can shift them by $t \mapsto t+1$. Adding all resulting terms almost all summands cancel out. If $T_{t+n_{\mathrm{repl}}+1} = T_{t+n_{\mathrm{repl}}+2}$ we get remaining terms (up to a positive factor)

$$-\binom{t}{r_{t,\uparrow}} \vec{p}[j]^{r_{t,\uparrow}}\vec{p}[j']^{t-r_{t,\uparrow}} + \binom{t}{r_{t,\downarrow}-1}\vec{p}[j]^{r_{t,\downarrow}-1}\vec{p}[j']^{t-r_{t,\downarrow}+1}$$

where $r_{t,\uparrow} = T_{t+n_{\mathrm{repl}}+1}$ and $r_{t,\downarrow} = T_{t+n_{\mathrm{repl}}+1} - n_{\mathrm{repl}}$.[14] The term is negative if

$$\frac{\prod_{\nu=0}^{n_{\mathrm{repl}}}(t+n_{\mathrm{repl}}+1+T_{t+n_{\mathrm{repl}}+1}-\nu)}{\vec{p}[j']^{n_{\mathrm{repl}}+1}}$$
$$-\frac{\prod_{\nu=0}^{n_{\mathrm{repl}}}(T_{t+n_{\mathrm{repl}}+1}-\nu)}{\vec{p}[j]^{n_{\mathrm{repl}}+1}} \geq 0$$

which is true by definition. Analogously for $T_{t+n_{\mathrm{repl}}+1}+1 = T_{t+n_{\mathrm{repl}}+2}$:[15]

$$\binom{t}{r_{t,\uparrow}} \vec{p}[j]^{r_{t,\uparrow}+1}\vec{p}[j']^{t-r_{t,\uparrow}-1} - \binom{t}{r_{t,\downarrow}}\vec{p}[j]^{r_{t,\downarrow}}\vec{p}[j']^{t-r_{t,\downarrow}}$$

which is again negative if

$$\frac{\prod_{\nu=0}^{n_{\mathrm{repl}}}(t+n_{\mathrm{repl}}+1+T_{t+n_{\mathrm{repl}}+1}-\nu)}{\vec{p}[j']^{n_{\mathrm{repl}}+1}}$$
$$-\frac{\prod_{\nu=0}^{n_{\mathrm{repl}}}(T_{t+n_{\mathrm{repl}}+1}-\nu)}{\vec{p}[j]^{n_{\mathrm{repl}}+1}} \geq 0.$$

Hence the derivative is non-positive which shows that $\delta_{jj'}$ becomes maximal if $\vec{p}[j], \vec{p}[j'] \leq \vec{p}[k]$ for all $k \in C$.

[14]Note that we still get the same term for $r_{t,\uparrow} = t$. For $r_{t,\downarrow} = 0$ the positive term drops off.

[15]For $r_{t,\downarrow} = 0$ we get the same term, for $r_{t,\uparrow} = t$ the positive term drops off.