

Coercion-Resistant Electronic Elections

Ari Juels¹ and Dario Catalano² and Markus Jakobsson³

¹ RSA Laboratories

Bedford, MA, USA

e-mail: ajuels@rsasecurity.com

² CNRS-Ecole Normale Supérieure

75230 Paris Cedex 05 - France, France

e-mail: dario.catalano@ens.fr

³ Indiana University, School of Informatics

Bloomington, IN, USA

e-mail: markus@indiana.edu

Abstract. We introduce a model for electronic election schemes that involves a more powerful adversary than in previous work. In particular, we allow the adversary to demand of coerced voters that they vote in a particular manner, abstain from voting, or even disclose their secret keys. We define a scheme to be *coercion-resistant* if it is infeasible for the adversary to determine whether a coerced voter complies with the demands.

A first contribution of this paper is to describe and characterize a new and strengthened adversary for coercion in elections. (In doing so, we additionally present what we believe to be the first formal security definitions for electronic elections of *any* type.) A second contribution is to demonstrate a protocol that is secure against this adversary. While it is clear that a strengthening of attack models is of theoretical relevance, it is important to note that our results lie close to practicality. This is true both in that we model real-life threats (such as vote-buying and vote-cancelling), and in that our proposed protocol combines a fair degree of efficiency with an unusual lack of structural complexity. Furthermore, previous schemes have required use of an untappable channel throughout. Ours only carries the much more practical requirement of an anonymous channel during the casting of ballots, and an untappable channel during registration (potentially using postal mail).

Key words: coercion-resistance, electronic voting, mix networks, receipt-freeness

1 Introduction

Most voters participating in shareholder elections in the United States have the option of casting ballots via a Web browser [1]. Some voters near Geneva participating in recent referenda in Switzerland in 2003-4 have been able to cast binding votes electronically [19]. The UK government has enunciated plans are to allow its citizens to vote electronically “some time after 2006” [18]. These are just a few instances of a broadening trend toward Internet-based voting. While voting of this kind appears to encourage higher voter turnout [38] and make accurate accounting for votes easier, it also carries the potential of making abuse easier to perform, and easier to perform at a large scale. A number of papers in the cryptographic literature have described ways of achieving robust and verifiable *electronic* elections, i.e., elections in which ballots and processing data are posted to a publicly accessible bulletin board. For some recent examples (but not by any means an exhaustive list), see [8, 16, 21, 22, 27, 30, 34, 37, 41].

There are two other threats, however, that it is equally crucial to address in a fair and democratic election process: We speak of *voter coercion* and *vote buying*. Internet-based voting does not introduce these problems, but it does have the potential to exacerbate them by extending the reach and data collection abilities of an attacker. This is highlighted in one way by the presence of a notorious Web site that provides a forum for the auctioning of votes [2]. Seller compliance was in that case merely voluntary. Conventional Internet voting schemes, however, including those described in the literature, actually provide an attacker with ready-made tools for verifying voter behavior and thereby exerting influence or control over voters. Without careful system design, the threats of coercion and vote buying are potentially far more problematic in Internet voting schemes than in ordinary, physical voting schemes.

One commonly proposed way of achieving secure electronic voting systems is to use a cryptographic system known as a *mix network* [14]. This is a tool that enables a collection of servers to take as input a collection of ciphertexts and to output the corresponding plaintexts according to a secret permutation. A straightforward way to achieve an election system that preserves the privacy of voters, then, is to assign a private digital signing key to each voter. To cast a ballot, the voter encrypts her choice and signs it, and then posts it to a bulletin board (i.e., a publicly accessible memory space). When all ballots have been collected and the corresponding signatures have

been checked, the ciphertexts are passed through a mix network. The resulting plaintext versions of the voter choices may then be tallied. Thanks to the privacy preserving property of the mix network, an adversary cannot tell which vote was cast by which voter. This approach is frequently advocated in the mix-network literature, as in, e.g., [8, 14, 22, 27].

In an ordinary mix-based scheme of this kind, an adversary can coerce a voter straightforwardly. The adversary can simply furnish the voter with a ciphertext on a particular candidate, and then verify that the voter posted a ballot containing that ciphertext. Alternatively, the adversary can demand the private signing key of the voter and verify its correctness against the corresponding public key. An adversary attempting to buy votes can use the same means. Other types of cryptographic voting schemes, namely homomorphic schemes [5, 16] and schemes based on blind signatures [20, 37], suffer from similar vulnerabilities.

1.1 Previous work

Previous investigations of coercion-resistant voting have been confined to a property known as *receipt-freeness*. Roughly stated, receipt-freeness is the inability of a voter to prove to an attacker that she voted in a particular manner, even if the voter wishes to do so. For a more formal definition, see [37]. The property of receipt-freeness ensures that an attacker cannot determine exact voter behavior and therefore cannot coerce a voter by dictating her choice of candidate. It also protects against vote-buying by preventing a potential vote buyer from obtaining proof of the behavior of voters; voters can thereby *pretend* to sell their votes, but defraud the vote buyer. The notion of receipt-freeness first appeared in work by Benaloh and Tuinstra [5]; their scheme, based on homomorphic encryption, was shown in [25] not to possess receipt-freeness as postulated. An independent introduction of the idea appeared in Niemi and Renvall [35]. Okamoto [36] proposed a voting scheme which he himself later showed to lack the postulated receipt-freeness; a repaired version by the same author, making use of blind signatures, appears in [37]. Sako and Kilian [39] proposed a multi-authority scheme employing a mix network to conceal candidate choices, and a homomorphic encryption scheme for production of the final tally. The modelling of their scheme was clarified and refined by Michels and Horster [33]. The Sako and Kilian scheme served as a conceptual basis for the later work of Hirt and Sako [25], followed by the more efficient approach of [3]; these two are the most efficient (and correct) receipt-free voting schemes to date. A recently proposed scheme by Magkos *et al.* [32] distinguishes itself by an approach relying on tamper-resistant hardware, but is flawed.¹

All of these receipt-free voting schemes include somewhat impractical assumptions. For example, these schemes assume the availability of an *untappable channel* between the voter and the authorities, that is, a channel that provides perfect secrecy in an information-theoretic sense. (I.e., even encryption does not provide an untappable channel.) The scheme in [37] makes the even stronger assumption of an *anonymous untappable channel*. (It is also not very practical in that it requires voter interaction with the system three times in the course of an election.) Moreover, all of these schemes (excepting [37]) lose the property of coercion-resistance if the attacker is able to corrupt even one of the tallying authorities in a distributed setting. The scheme of Hirt and Sako still retains coercion-resistance when such corruption takes place, but only under the strong assumption that the voter knows *which* tallying authorities have been corrupted; the proposal of Baudron *et al.* has a similar property.

A still more serious problem with all of the receipt-free voting schemes described in the literature, however, is the fact that the property of receipt-freeness alone fails to protect an election system against several forms of serious, real-world attack, which we enumerate here:

Randomization attack: This attack was noted by Schoenmakers in 2000 [42]; he described its applicability to the scheme of Hirt and Sako. The idea is for an attacker to coerce a voter by requiring that she submit randomly composed balloting material. In this attack, the attacker (and perhaps even the voter) is unable to learn what candidate

¹ We are unaware of any mention of a break of this scheme in the literature, and therefore briefly describe one here. The Magkos *et al.* system employs an interactive honest-verifier ZK proof made by a smartcard to the voter. Presumably because of the simulability of this proof, the authors describe the proof as being “non-transferable”. This is not true. In particular, an adversary can stipulate that the voter engage in the proof using a challenge that the adversary has pre-selected. The proof then becomes transferable, yielding a means of receipt construction by the adversary. As noted in [25], this type of attack also explains why *deniable encryption* [13] does not solve the problem of coercion in a voting system.

the voter cast a ballot for. The effect of the attack, however, is to nullify the choice of the voter with a large probability. For example, an attacker favoring the Republican party in a United States election would benefit from mounting a randomization attack against voters in a heavily Democratic district.

Forced-abstention attack: This is an attack related to the previous one based on randomization. In this case, the attacker coerces a voter by demanding that she refrain from voting. All of the schemes cited above are vulnerable to this simple attack. This is because the schemes authenticate voters directly in order to demonstrate that they are authorized to participate in the election. Thus, an attacker can see who has voted, and use this information to threaten and effectively bar voters from participation.²

Simulation attack: The receipt-free schemes described above assume that the attacker cannot coerce a voter by causing her to divulge her private keying material after the registration process but prior to the election process. Such an attack, however, is a real and viable one in previously proposed schemes, because these permit an attacker to verify the correctness of private keying material. For example, in [37], the voter provides a digital signature which, if correct, results in the authority furnishing a blind digital signature. In [25], the voter, when casting a ballot, proves knowledge of a private key relative to a publicly committed or published value. In general, receipt-freeness does not prevent an attacker from coercing voters into divulging private keys or buying private keys from voters and then *simulating* these voters at will, i.e., voting on their behalf.

1.2 Our contribution

Our contribution in this paper is twofold. First, we investigate a stronger and broader notion of coercive attacks than receipt-freeness. This notion, which we refer to as *coercion-resistance*, captures what we believe to be the fullest possible range of adversarial behavior in a real-world, Internet-based voting scheme. A coercion-resistant scheme offers not only receipt-freeness, but also defense against randomization, forced-abstention, and simulation attacks – all potentially in the face of corruption of a minority of tallying authorities. We propose a formal definition of coercion-freeness in the body of this paper. Two other properties are essential for any voting scheme, whether or not it is coercion-resistant. These are *correctness* and *verifiability*. As formal definitions for these properties seem to be lacking in the literature, we provide them as well in the paper appendix; we thus provide what we believe to be the first formal security framework for electronic elections in general.

To demonstrate the practical realizability of our definitions, we describe a voting scheme that possesses the strong property of coercion-resistance proposed in this paper – and also the properties of correctness and verifiability. Our scheme does not require untappable channels during the voting process, but instead assumes voter access to an anonymous channel at some point during the voting process. (Registration does involve an untappable channel, but may be achieved via, e.g., postal mail.) The assumption of anonymous channels in an election can be realized in a practical way. (Exactly how practical is subject to some debate.) For example, an adversary may be able to view communications on the Internet, but not trace IP addresses to identities effectively. Or voters may achieve anonymity by using public terminals in Internet cafés, libraries, workplaces, or physical polling places to participate in an otherwise electronic election. Popular use of mixnets would similarly help, e.g., [22, 34]; broadcast channels are another possible mechanism for anonymity. Anonymous channels are, of course, a strictly weaker assumption than untappable ones. Thus our scheme can make use of untappable channels too. The assumption of untappable channels may be reasonable in some limited cases: It is not a straightforward matter for an unsophisticated attacker to tap point-to-point communications on the Internet, for example.

Anonymous channels are in fact a minimal requirement for *any* coercion-resistant schemes: An attacker that can identify which voters have participated can obviously mount a forced-abstention attack.

A drawback of our scheme is that, even with use of asymptotically efficient mix networks as in [22, 34], the overhead for tallying authorities is quadratic in the number of voters. Thus the scheme is only practical for small

² An exception is the scheme in [37], which does not appear to be vulnerable to a forced-abstention attack. This is because the scheme seems to assume that the authority checks voter enrollment privately. In other words, the scheme does not permit public verification that participating voters are present on a published voter roll. This is potentially a problem in its own right.

elections. Our hope and belief, however, is that our proposed scheme might serve as the basis for refinements with a higher degree of practical application. We provide a security proof for our proposed scheme in the paper appendix.

1.3 Intuition behind our scheme

In a conventional voting scheme, and also in receipt-free schemes like [25], the voter V_i identifies herself at the time she casts her ballot. This may be accomplished by means of a digital signature on the ballot, or by an interactive authentication protocol. The key idea behind our scheme is for the identity of a voter to remain hidden during the election process, and for the validity of ballots instead to be checked blindly against a voter roll. When casting a ballot, a voter incorporates a concealed credential. This takes the form of a ciphertext on a secret value σ that is unique to the voter. The secret σ is a kind of *anonymous credential*, quite similar in spirit to, e.g., [9, 10]. To ensure that ballots are cast by legitimate voters, the tallying authority \mathcal{T} performs a blind comparison between hidden credentials and a list L of encrypted credentials published by an election registrar \mathcal{R} alongside the plaintext names of registered voters.

By means of mixing and blind comparison of ciphertext values, it is possible to check whether a concealed credential is in the list L or not, without revealing which voter the credential has been assigned to. In consequence, an attacker who is given a fake credential $\tilde{\sigma}$ by a coerced voter cannot tell whether or not the credential is valid. (The attacker will learn how many ballots were posted with bad credentials. Provided, however, that some spurious ones are injected by honest players, authorities, or even outsiders, the individuals associated with bad ballots will remain concealed.) Moreover, the attacker cannot mount randomization or forced-abstention attacks, since there is no feasible way to determine whether an individual voter has posted a ballot or not. In particular, after divulging fake credential $\tilde{\sigma}$, a voter can go and vote again using her real credential σ .

1.4 Organization

In section 2, we describe our setup and attack models and sketch a few of the major adversarial strategies. We provide formal definitions for the security property of coercion-resistance in section 3. We describe the particulars of our proposed scheme in section 4, prefaced by a summary of the underlying cryptographic building blocks. In the appendices to the paper, we offer formal definitions for the correctness and verifiability of election schemes, a detailed security-proof outline, and details on our choice of primitives for realizing our proposed scheme.

2 Modelling

An election system consists of several sets of entities:

1. *Registrars*: Denoted by $\mathcal{R} = \{R_1, R_2, \dots, R_{n_R}\}$, this is a set of n_R entities responsible for jointly issuing keying material, i.e., credentials to voters.
2. *Authorities (Talliers)*: Denoted by $\mathcal{T} = \{T_1, T_2, \dots, T_{n_T}\}$, authorities are responsible for processing ballots and jointly counting votes and publishing a final tally.
3. *Voters*: The set of n_V voters, denoted by $\mathcal{V} = \{V_1, V_2, \dots, V_{n_V}\}$, are the entities participating in a given election administered by \mathcal{R} . We let i be a public identifier for V_i .

We make use of a *bulletin board*, denoted by \mathcal{BB} . This is a piece of universally accessible memory to which all players have appendive-write access. In other words, any player can write data to \mathcal{BB} , but cannot overwrite or erase existing data. Moreover, voters will be able to read the contents of \mathcal{BB} once the vote casting phase has ended. For notational convenience, we assume that data are written to \mathcal{BB} in μ -bit blocks for an appropriate choice of μ . Shorter data segments may be padded appropriately. For simplicity of exposition, we assume no ordering on the contents of \mathcal{BB} .

2.1 Functions

We define a *candidate slate* C to be an ordered set of n_C distinct identifiers $\{c_1, c_2, \dots, c_{n_C}\}$, each of which corresponds to a voter choice, typically a candidate or party name. In an election, choice c_j may be identified according to its index j . Thus, for cryptographic purposes the candidate slate consists of the integers $\{1, 2, \dots, n_C\}$ and may be specified by n_C alone. We define a *tally* on an election under slate C to be a vector \mathbf{X} of n_C positive integers x_1, x_2, \dots, x_{n_C} such that x_j indicates the number of votes cast for choice c_j . The protocols composing an election system are then as follows:

- **Registering:** The function $\text{register}(SK_{\mathcal{R}}, i, k_1) \rightarrow (sk_i, pk_i)$ takes as input the private registrar key $SK_{\mathcal{R}}$, a (voter) identifier i and a security parameter k_1 , and outputs a key pair (sk_i, pk_i) . This is computed jointly by players in \mathcal{R} , possibly in interaction with voter V_i .
- **Voting:** The function $\text{vote}(sk, PK_{\mathcal{T}}, n_C, \beta, k_2) \rightarrow \text{ballot}$ takes as input a private voting key, the public key of the authorities \mathcal{T} , the candidate-slate specification n_C , a candidate selection β , and a security parameter k_2 , and yields a ballot of bit length at most μ . The form of the ballot will vary depending on the design of the election system, but is in essence a digitally signed vote choice encrypted under $PK_{\mathcal{T}}$.
- **Tallying:** The function $\text{tally}(SK_{\mathcal{T}}, \mathcal{BB}, n_C, \{pk_i\}_{i=1}^{n_V}, k_3) \rightarrow (\mathbf{X}, P)$ takes as input the private key of the authority \mathcal{T} , the full contents of the bulletin board, the candidate-slate size, all public voting keys, and a security parameter k_3 and outputs a vote tally \mathbf{X} , along with a non-interactive proof P that the tally was correctly computed.
- **Verifying:** The function $\text{verify}(PK_{\mathcal{T}}, \mathcal{BB}, n_C, \mathbf{X}, P) \rightarrow \{0, 1\}$ takes as input the public key of the authorities, the contents of the bulletin board, the candidate-slate size, the voting tally, and a non-interactive proof of correct tallying. It outputs a ‘0’ if the tally is incorrect and a ‘1’ otherwise. (We characterize the behavior of `verify` more formally in the paper appendix.)

We define an election scheme ES as the collection of these functions. Thus $\text{ES} = \{\text{register}, \text{vote}, \text{tally}, \text{verify}\}$.

Remark: There are many election models in use throughout the world. The model we propose here excludes important variants. In some systems, for example, voters are asked to rank candidate choices, rather than just listing those they favor. Many systems permit the use of *write-in* votes, i.e., the casting of a ballot in favor of a candidate not listed on the slate for the election. We exclude write-in voting from our model because it undermines the possibility of coercion resistance in any scheme where an observer can see a complete election tally including write-in votes. An attacker may, for example, require coerced voters to cast write-in ballots for candidate names consisting of random strings pre-specified by the attacker. This way, the attacker can: (1) Verify that coerced voters complied with instructions, by looking for the random strings the attacker furnished, and (2) Ensure that the votes of coerced voters are not counted, since random strings will most likely not correspond to real election choices. (Thus, this would combine the forced abstention attack and the randomization attack.)

2.2 Summary of the attack model

We consider the process for a single election as proceeding in these phases, corresponding largely with the functions enumerated in section 2.1:

1. **Setup:** If not already available, key pairs are generated for or by \mathcal{R} and \mathcal{T} . The candidate slate C for the election is published by \mathcal{R} with appropriate integrity protection.
2. **Registration:** The identities and eligibility of would-be voters are verified by \mathcal{R} . Given successful verification, an individual becomes a registered voter, receiving from \mathcal{R} a credential permitting participation in the election. Previously registered voters may be able to re-use their credentials. \mathcal{R} publishes a voter roll L .
3. **Voting:** Referring to the candidate slate C , registered voters use their credentials to cast ballots.
4. **Tallying:** The authority \mathcal{T} processes the contents of the bulletin board \mathcal{BB} so as to produce a tally vector \mathbf{X} specifying the outcome of the election, along with a proof of correctness P of the tally.
5. **Verification:** Any player, whether or not a participant in the election, can refer to \mathcal{BB} , P and L to verify the correctness of the tally produced by \mathcal{T} in the previous phase.

Assumptions in setup phase: Our security definitions permit the possibility of static, active corruption by the adversary of a minority of players in \mathcal{T} in the setup phase. The security of our construction then relies on generation of the key pair $(SK_{\mathcal{T}}, PK_{\mathcal{T}})$ by a trusted third party, or, alternatively, on an interactive, computationally secure key-generation protocol such as [24] between the players in \mathcal{T} .

Assumptions prior to registration: The adversary may coerce a voter prior to the registration phase in the sense of requesting in advance that the voter retain transcripts of the registration process, or by providing data in an attempt to dictate voter interaction with the registrar.

Assumptions in registration phase: We assume that this phase is trustworthy. Strong integrity of the registrar \mathcal{R} is of course critical for any secure election system. To evade coercion, a voter must be able to receive a credential without adversarial interference. An adversary capable of corrupting and seizing the credentials of a voter in this initial phase can mount a full simulation attack. An adversary capable of preventing a voter from registering can mount a forced-abstention attack.

We assume therefore that the voter receives her credential from \mathcal{R} via an untappable channel. We are helped here by the fact that registration is generally an offline procedure. For example, a voter might receive her registration through the postal service. We must then assume that mail in the postal system is not subject to compromise. Alternatively, registration could be performed by voters in person.

It is further necessary to assume that an attacker cannot obtain registration transcript data after the fact. Where registration is electronic, for example, erasure of data from voter interaction with \mathcal{R} must be compulsory by voters (e.g., enforced by smartcards). In a postal system, we must assume that the adversary cannot feasibly gather and validate registration letters from targeted voters (at least not many of them).

Finally, we must assume the integrity of the registration procedure on the inside, i.e., we must assume that \mathcal{R} is trustworthy and does not leak credentials to the adversary.³

It may or may not be entirely feasible to achieve all of these assumptions around the registration process in real-world systems. We note, however, that if the registration process does *not* have the integrity we require here, then coercion-resistance is the least of our problems: Votes can then be tampered with.

Assumptions on voting, tallying and verification phases: Subsequent to the registration phase, we assume that the adversary may seize control of a minority of players in \mathcal{T} and any number of voters in a static, active manner. (Since \mathcal{R} does not participate in the process subsequent to registration, we need not consider corruption of \mathcal{R} at this point.) The adversary may also attempt to coerce voters outside its control by requesting that they divulge private keying material⁴ or behave in a prescribed manner in voting. Voters are assumed to be able to cast their ballots via anonymous channels, i.e., channels such that an attacker cannot determine whether or not a given voter cast a ballot. We noted some practical underpinnings for this assumption above. As also noted above, this assumption is a requirement for any election scheme to be fully coercion-resistant: If an attacker can tell whether or not a given voter cast a ballot, then the attacker can easily mount a forced-abstention attack. We assume that either sessions are integrity protected (on a per-session basis, since they are anonymous), or that the adversary has only passive interaction with the network.

Not considered here: We do not treat the problem of denial-of-service attacks. Mechanisms like [29] might help, but are beyond the scope of our research. We also do not treat the problem of enabling voters to verify that their votes have been counted. This is in principle possible whilst retaining coercion resistance. Honest voters could submit their credentials via anonymous channels for verification of proper processing by authorities. Authorities would concoct appropriate false replies for invalid credentials. We regard this as future work.

³ Some relaxation of this assumption is possible. If \mathcal{R} comprises multiple players, computations by \mathcal{R} are distributed, and communication between players in \mathcal{R} are made directly with voters, then even if a minority of players in \mathcal{R} is corrupted, coercion-resistance is still possible. The catch is that voters must know *which* players in \mathcal{R} have been corrupted. For details on this idea, see [25].

⁴ We assume that the coercion takes place remotely. For example, the adversary may not continuously watch over the shoulder of a voter, monitor her hard-drive, etc. Our proposed protocol does potentially defend against some shoulder-surfing, however, by permitting voters to use fake keys and/or re-vote.

3 Formal definitions

We now turn our attention to formal security definitions of the essential properties of *correctness*, *verifiability*, and *coercion-resistance*, respectively abbreviated *corr*, *ver*, and *c-resist*. Our definitions hinge on a set of experiments involving an adversary \mathcal{A} in interaction with components of the election system ES. This adversary is assumed to retain state throughout the duration of an experiment. We formulate our experiments such that in all cases, the aim of the adversary is to cause an output value of ‘1’. Thus, for experiment $\text{Exp}_{\text{ES},\mathcal{A}}^E(\cdot)$ on property $E \in (\text{ver}, \text{corr}, \text{c-resist})$, we define $\text{Succ}_{\text{ES},\mathcal{A}}^E(\cdot) = \Pr[\text{Exp}_{\text{ES},\mathcal{A}}^E(\cdot) = \text{‘1’}]$.

According to the standard definition, we say that a quantity $f(k)$ is *negligible* in k if for every positive integer c there is some l_c such that $f(k) < k^{-c}$ for $k > l_c$. In most cases, we use the term negligible alone to mean negligible with respect to the full set of relevant security parameters. Similarly, in saying that an algorithm has *polynomial running time*, we mean that its running time is asymptotically bounded by some polynomial in the relevant security parameters. As the properties of correctness and verifiability are of less relevance to our work than coercion-resistance, we relegate the first two definitions to appendices A and B.

Coercion resistance: Coercion resistance may be regarded as an extension of the basic property of privacy. Privacy in an election system is defined in terms of an adversary that cannot interact with voters during the election process. In particular, we say that an election is private if such an adversary cannot guess the vote of any voter better than an adversarial algorithm whose only input is the election tally. (Note, for example, in an election where all voters vote Republican, the system may have the property of privacy, even though the adversary knows how all voters cast their ballots in that election.)

Coercion resistance is a strong form of privacy in which it is assumed that the adversary may interact with voters. In particular, the adversary may instruct targeted voters to divulge their private keys subsequent to registration, or may specify that these voters cast ballots of a particular form. If the adversary can determine whether or not voters behaved as instructed, then the adversary is capable of blackmail or otherwise exercising undue influence over the election process. Hence a coercion-resistant voting system is one in which the user can deceive the adversary into thinking that she has behaved as instructed, when the voter has in fact cast a ballot according to her own intentions.

Our definition of coercion resistance requires addition of a new function to voting system ES:

- The function $\text{fakekey}(PK_{\mathcal{T}}, sk, pk) \rightarrow \tilde{sk}$ takes as input the public key of the authorities, and the private/public key pair of the voter. It outputs a spurious key \tilde{sk} .

Of course, for the function fakekey to enable coercion resistance, the key \tilde{sk} must be indistinguishable by the adversary \mathcal{A} from a valid key, and only distinguishable by a majority of talliers \mathcal{T} . This property is captured in our experiment characterizing coercion resistance. To simplify the formulation of the experiment, we assume implicitly that tally is computed by an oracle (with knowledge of $SK_{\mathcal{T}}$). It suffices, however, for \mathcal{T} to be computed via a protocol that achieves correct output and is computationally simulable by the adversary \mathcal{A} (who, it will be recalled, may corrupt a minority of \mathcal{T}).

Our definition of coercion resistance centers on a kind of game between the adversary \mathcal{A} and a voter targeted by the adversary for coercive attack. A coin is flipped; the outcome is represented by a bit b . If $b = 0$, then the voter casts a ballot with a particular choice β , and provides the adversary with a false voting key \tilde{sk} ; in other words, the voter attempts to evade adversarial coercion. If $b = 1$, on the other hand, then the voter submits to the coercion of the adversary; she simply furnishes the adversary with her valid voting key sk , and does not cast a ballot. The task of the adversary is to guess the value of the coin b , that is, to determine whether or not the targeted voter in fact cast a ballot. We permit the adversary in this definitional game to specify the ballot value β . While it is somewhat unnatural for the adversary thus to specify the intention of the voter, this is a necessary (but not sufficient) condition to achieve the strongest possible security definition. The adversary could potentially use control of this ballot to improve its probability of coercion.

We characterize the voting pattern of honest voters in terms of a probability distribution D_{n,n_C} . This distribution models the state of knowledge of the adversary about the intentions of these voters. Of course, the adversary generally does not have perfect knowledge of the voting intentions of honest voters. (Indeed, as we explain below, imperfect

adversarial knowledge is actually a requirement for meaningful coercion-resistance.) For a collection of n voters outside the control of the adversary – i.e., voters not subject to coercion – we characterize the view of the adversary in terms of a probability distribution D_{n,n_C} . We let ϕ be a symbol denoting a null ballot, i.e., an abstention, and let λ denote a ballot cast with an invalid credential. Then D_{n,n_C} is a distribution over vectors $(\beta_1, \beta_2, \dots, \beta_n) \in (n_C \cup \phi \cup \lambda)^n$, i.e., over the set of possible ballot choices for an election plus abstentions and invalid ballots. For a set of n voting credentials $\{sk_i\}$, we let $\text{vote}(\{sk_i\}, PK_{\mathcal{T}}, n_C, D_{n,n_C}, k_2)$ denote the casting of ballots according to distribution D_{n,n_C} . In other words, a vector $(\beta_1, \beta_2, \dots, \beta_n)$ is drawn from D_{n,n_C} and vote β_i is cast using credential sk_i .

We are now ready to present an experiment *c-resist* that defines the game described above between an adversary and a voter targeted for coercion. Recall that k_1, k_2 , and k_3 are security parameters defined above, n_V is the total number of eligible voters for the election, and n_C is the number of candidates, i.e., the size of the candidate slate. We let n_A denote the number of voters that may be completely controlled, i.e., corrupted by the adversary. We define $n_U = n_V - n_A - 1$. In other words, the number of uncertain votes n_U equals the total number of honest votes, i.e., the number of possible votes, minus those coming from voters controlled by the attacker, minus the vote coming from the voter the attacker is trying to coerce (in the experiment).

We consider a static adversary, i.e., one that selects voters to corrupt prior to protocol execution. We assume that the adversary has a list of “voter names,” i.e., a roll of potential participating voters.

We let \leftarrow denote assignment and \Leftarrow denote the append operation, while % denotes the beginning of an annotative comment on the experiment. Our experiment treats the case in which the adversary seeks to coerce a single voter; extension of the definition to coercion of multiple voters is straightforward. The experiments defined here halt when an output value is produced.

Experiment $\text{Exp}_{\text{ES}, \mathcal{A}, H}^{c\text{-resist}}(k_1, k_2, k_3, n_V, n_A, n_C)$

```

 $V \leftarrow \mathcal{A}(\text{voter names, “control voters”});$            %  $\mathcal{A}$  corrupts voters
 $\{(sk_i, pk_i) \leftarrow \text{register}(SK_{\mathcal{R}}, i, k_2)\}_{i=1}^{n_V};$    % voters are registered
 $(j, \beta) \leftarrow \mathcal{A}(\{sk_i\}_{i \in V}, \text{“set target voter and vote”});$  %  $\mathcal{A}$  sets coercive target
if  $|V| \neq n_A$  or  $j \notin \{1, 2, \dots, n_V\} - V$  or
   $\beta \notin \{1, 2, \dots, n_C\} \cup \phi$  then
  output ‘0’;                                           % outputs of  $\mathcal{A}$  checked for validity
 $b \in_U \{0, 1\};$                                        % coin is flipped
if  $b = 0$  then
   $\tilde{sk} \leftarrow \text{fakekey}(PK_{\mathcal{T}}, sk_j, pk_j);$ 
   $\mathcal{BB} \Leftarrow \text{vote}(sk_j, PK_{\mathcal{T}}, n_C, \beta, k_2);$ 
else
   $\tilde{sk} \leftarrow sk_j;$ 
   $\mathcal{BB} \Leftarrow \text{vote}(\{sk_i\}_{i \neq j, i \notin V}, PK_{\mathcal{T}}, n_C, D_{n_U, n_C}, k_2);$  % ballots posted for honest voters
 $\mathcal{BB} \Leftarrow \mathcal{A}(\tilde{sk}, \mathcal{BB}, \text{“cast ballots”});$            %  $\mathcal{A}$  posts to  $\mathcal{BB}$ 
 $(\mathbf{X}, P) \leftarrow \text{tally}(SK_{\mathcal{T}}, \mathcal{BB}, n_C, \{pk_i\}_{i=1}^{n_V}, k_3);$  % election results are tallied
 $b' \leftarrow \mathcal{A}(\mathbf{X}, P, \text{“guess } b\text{”});$                %  $\mathcal{A}$  guesses coin flip
if  $b' = b$  then
  output ‘1’;
else
  output ‘0’;
```

The adversary \mathcal{A} in the above experiment is quite powerful, being capable (when $b = 1$) of complete coercion of the targeted voter. In order to characterize the success of \mathcal{A} , we must compare \mathcal{A} with a second adversary \mathcal{A}' . \mathcal{A}' is capable of coercion only within the framework of an ideal voting experiment *c-resist-ideal*. In other words, \mathcal{A}' characterizes the type of security against coercion that we would like to achieve in ES.

The main feature we are aiming for in our ideal experiment *c-resist-ideal* is for \mathcal{A}' to learn nothing from the private keys she acquires from corrupted players and from the coerced player. In particular, \mathcal{A}' cannot use private

keys to perform active attacks. We cause \mathcal{A}' to express voting choices in a direct, ideal process; \mathcal{A}' cannot cast ballots, but merely enumerates the choices of players in her control. Additionally, \mathcal{A} cannot use private keys to learn information about the voting behavior of honest players or the coerced player. The *only* information that \mathcal{A}' gets is the grand total \mathbf{X} of votes in the election.

One feature of our experiment is counterintuitive. Because this is an ideal experiment, \mathcal{A}' is *always* given \tilde{sk} as the key of the coerced player. This is because \mathcal{A}' should be unable to determine, on the basis of keying material, from the situation in which coercion is successful or unsuccessful.

We require a function for the definition. We include here an ideal function *ideal-tally* that tallies the ballots posted to \mathcal{BB} in a special way. The function *ideal-tally* tallies in a normal manner all of the ballots cast by honest voters, i.e., prior to adversarial posting. The ballots cast by \mathcal{A}' , however, are treated specially. In particular, *ideal-tally* determines for each ballot B what the underlying private key sk_i is. If $i \notin V$, i.e., if the private key is not one assigned to one of the corrupted players, then the corresponding vote is not counted. Additionally, any double vote is not counted, i.e., *ideal-tally* performs the weeding of double votes that normally occurs during the tallying procedure. Finally, *ideal-tally* does the following based on the value of the secret bit b . If $b = 0$, then *ideal-tally* does not count any ballot cast (by the adversary) using private key \tilde{sk} . If $b = 1$, then *ideal-tally* does include in the final tally a ballot cast using \tilde{sk} (excluding double votes).

Our definition of *ideal-tally* here assumes that every ballot has a unique corresponding private key. This is true of most natural ballot structures (and true of our proposed scheme). This definition, of course, also assumes ideal functionality in *ideal-tally*, namely the ability to extract private keys and plaintext votes from ballots. We do not specify in our definition how this “oracle” power is achieved. In our proofs, we construct a simulator capable of performing this functionality required from *ideal-tally*.

Note that although \mathcal{A}' receives the secret key of a coerced voter in our ideal experiment, this is really just a technical step. This secret key in fact provides \mathcal{A}' with no information useful in voting, since the ideal function *ideal-tally* ensures against misuse of keys; also, this secret key can provide no information useful in learning votes, since \mathcal{A}' never sees \mathcal{BB} .

We are now ready to present the experiment *c-resist-ideal* that characterizes the success of \mathcal{A}' .

Experiment $\text{Exp}_{\text{ES}, \mathcal{A}, H}^{c\text{-resist-ideal}}(k_1, k_2, k_3, n_V, n_A, n_C)$

$V \leftarrow \mathcal{A}'$ (voter names, “control voters”);	% \mathcal{A}' corrupts voters
$\{(sk_i, pk_i) \leftarrow \text{register}(SK_{\mathcal{R}}, i, k_2)\}_{i=1}^{n_V}$;	% voters are registered
$(j, \beta) \leftarrow \mathcal{A}'$ (“set target voter and vote”);	% \mathcal{A}' sets coercive target
if $ V \neq n_A$ or $j \notin \{1, 2, \dots, n_V\} - V$ or $\beta \notin \{1, 2, \dots, n_C\} \cup \phi$ then	% outputs of \mathcal{A}' checked for validity
output ‘0’;	
$b \in_U \{0, 1\}$;	% coin is flipped
if $b = 0$ then	% voter evades coercion
$\mathcal{BB} \leftarrow \text{vote}(sk_j, PK_{\mathcal{T}}, n_C, \beta, k_2)$;	
$\tilde{sk} \leftarrow sk_j$;	
$\mathcal{BB} \leftarrow \text{vote}(\{sk_i\}_{i \neq j, i \notin V}, PK_{\mathcal{T}}, n_C, D_{n_U, n_C}, k_2)$;	% ballots posted for honest voters
$\mathcal{BB} \leftarrow \mathcal{A}'(\tilde{sk}, \{sk_i\}_{i \in V}, \text{“cast ballots”})$;	% \mathcal{A}' specifies vote choices
$(\mathbf{X}, P) \leftarrow \text{ideal-tally}(SK_{\mathcal{T}}, \mathcal{BB}, n_C, \{pk_i\}_{i=1}^{n_V}, k_3)$;	% election results are tallied
$b' \leftarrow \mathcal{A}(\mathbf{X}, \text{“guess } b\text{”})$;	% \mathcal{A}' guesses coin flip
if $b' = b$ then	% experimental output determined
output ‘1’;	
else	
output ‘0’;	

Definition 1. We define an election scheme ES as *coercion resistant* if for any polynomially-bounded adversary \mathcal{A} , any parameters n and n_C , and any probability distribution D_{n,n_C} , the quantity

$$\mathbf{Adv}_{\text{ES},\mathcal{A}}^{c\text{-resist}} = \left| \mathbf{Succ}_{\text{ES},\mathcal{A}}^{c\text{-resist}}(\cdot) - \mathbf{Succ}_{\text{ES},\mathcal{A}}^{c\text{-resist-ideal}}(\cdot) \right|$$

is negligible in all security parameters.

Viewed intuitively, this definition means that in a real protocol execution, the adversary effectively learns nothing more than the election tally X . The adversary cannot learn any significant information from the protocol execution itself, even when mounting an active attack.

3.1 The need for uncertain adversarial knowledge

Even if an election scheme ES is coercion resistant by our definition, i.e., $|\mathbf{Succ}_{\text{ES},\mathcal{A}}^{c\text{-resist}}(\cdot) - \mathbf{Succ}_{\text{ES},\mathcal{A}}^{c\text{-resist-ideal}}(\cdot)|$ is negligibly close to 0, some (thankfully small) degree of coercion remains possible for many natural distributions D_{n,n_C} . In particular, in many settings, we may well have $\mathbf{Succ}_{\text{ES},\mathcal{A}}^{c\text{-resist-ideal}}(\cdot) \gg 0$. This seems counterintuitive, but in fact reflects a critical observation: The degree of possible coercion resistance is bounded below by adversarial uncertainty about the behavior of honest voters. This is true for *any* election scheme whatever.

For example, suppose that an adversary knows that a targeted voter aims to vote “Democrat” and that every other voter will vote “Republican.” Then coercion is *unavoidable*. If a “Democrat” vote turns up, then the adversary will necessarily know that the targeted voter has succeeded in registering a vote. Similarly, if the adversary is attempting to coerce one voter in a given election and knows that all hundred of the other eligible voters will cast ballots, then the adversary can mount an abstention attack straightforwardly. The adversary in this case simply threatens the voter in the case that the total tally for the election is one hundred and one.

Viewed another way, coercion resistance depends on “noise” or statistical uncertainty in the adversary’s view of how honest voters will vote. In other words, it hinges on the distribution D_{n,n_C} . The more entropy in D_{n,n_C} , the better the level of attainable coercion resistance. D_{n,n_C} serves the purpose in our experiments of defining the distribution of the “noise” that conceals the behavior of voters targeted by the adversary for coercion.

To our benefit, it is natural to expect that in a real-world election an adversary can obtain only fragmentary knowledge about the likely behavior of voters. An adversary may know, for instance, that most of the voters in a given district will vote “Republican,” but probably won’t know exactly how many. This means that coercion-resistance is a viable possibility. (Additionally, it is possible for voting authorities – or indeed any entity – intentionally to inject “chaff” in the form of blank and invalid ballots into an election system.)

4 A Coercion-Resistant Election Protocol

We are now ready to introduce our protocol proposal. We begin by describing the cryptographic building blocks we employ. Where appropriate, we model these as ideal primitives, as discussed in appendix D.

Threshold cryptosystem with re-encryption: Our first building block is a threshold public-key cryptosystem CS that permits re-encryption of ciphertexts with knowledge only of public parameters and keys. The private key for CS is held by \mathcal{T} in our construction.

To describe our aim in the ideal, we would like any ciphertext E to be perfectly hiding. We would like decryption to be possible only by having a majority of players in \mathcal{T} agree on a ciphertext to be decrypted. We model this latter ideal property as in terms of a special decryption oracle denoted by \tilde{DEC} . We assume further that any decryption performed by \tilde{DEC} is publicly verifiable.

Selected cryptosystem: At first El Gamal [23] may seem a natural choice of cryptosystem for our purposes. However this solution is not completely satisfying in our setting for the following reason.

Our construction will be proved to achieve coercion resistance under the *Decisional Diffie-Hellman assumption* (details of this proof can be found in appendix D). The basic idea of our proof is very simple: One assumes that

there is an adversary \mathcal{A} that has some advantage μ , in distinguish experiment *c-resist-real* from experiment *c-resist-ideal*, and constructs another algorithm \mathcal{S} which, receives on input a challenge for the decisional Diffie-Hellman problem [7, 44] (g^a, g^b, g^c) , and “using” \mathcal{A} , should be able to “solve” the challenge with a related advantage μ' . Now, as it will become apparent in section D, in order for our proof to go through correctly we require the simulator to be able to correctly decrypt a given ciphertext even when this ciphertext is constructed from the received challenge. Unfortunately basic El Gamal does not allow this. Indeed, if one encrypts a message m as $(g^a, g^c m)$ the simulator cannot decrypt simply because it does not know the discrete logarithm of g^c in base g^a .

To overcome this problem we adopt a modified version of the basic El Gamal scheme which can be seen as a simplified version of the well known Cramer-Shoup [17] cryptosystem (but only providing semantic security with respect to a passive adversary). We let \mathcal{G} denote the algebraic group over which we employ this modified El Gamal (which we'll simply call M-El Gamal), and q denote the group order. For semantic security, we require that the Decision Diffie-Hellman assumption hold over \mathcal{G} . The public key is (g_1, g_2, h) where g_1, g_2, h are elements in \mathcal{G} . The secret key is $x_1, x_2 \in Z_q$ such that $h = g_1^{x_1} g_2^{x_2}$.

To encrypt m one simply computes $(A, B, C) = (g_1^r, g_2^r, h^r m)$ for random r . Decryption is similar to plain El Gamal: one computes $m = C/(A^{x_1} B^{x_2})$. One can easily verify (see [17]) that this variant allows correct decryption even when ciphertexts are constructed from the received challenge.

We let \in_U here and elsewhere denote uniform, random selection from a set. A ciphertext in M-El Gamal on message $m \in \mathcal{G}$ takes the form $(\alpha, \beta, \gamma) = (m h^r, g_1^r, g_2^r)$ for $r \in_U Z_q$. For succinctness of notation, we sometimes let $E_h[m]$ denote a ciphertext on message m under public key h (assuming that g_1 and g_2 are considered public parameters).

Further details on the security of the scheme may be found in appendix E. An important feature of the M-El Gamal cryptosystem is that, exactly as the original version, it may be easily implemented in a threshold setting. In other words, the private keys x_1, x_2 may be distributed such that decryption can be performed by any quorum of share holders, without leakage of additional information. We exploit this distributed form of M-El Gamal in our proposed election scheme. As explained above, rather than focusing on a particular embodiment, we model the process by a decryption oracle denoted by \tilde{DEC} . We refer the reader to appendix E and to [12] for further discussion of threshold decryption in (plain) El Gamal.

Plaintext Equivalence Test (PET): A *plaintext equivalence test* (PET) [26, 31] is cryptographic primitive that operates on ciphertexts in a threshold cryptosystem. The input to PET is a pair of ciphertexts; the output is a single bit indicating whether the corresponding plaintexts are equal or not. PET may be realized as an efficient distributed protocol that reveals no additional, non-negligible information about plaintexts. For a detailed description of efficient methods to perform this verification, along with proofs of the properties of the construction, see [31]. Rather than focusing on a specific embodiment of PET, we model the ideal properties of the primitive by means of an oracle denoted by \tilde{PET} , and with the property of public verifiability.

Mix network: A (re-encryption) mix network (MN) is a distributed protocol that takes as input an ordered set $\mathbf{E} = \{E_1, E_2, \dots, E_d\}$ of ciphertexts generated in a cryptosystem like El Gamal that permits re-encryption. The output of MN is an ordered set $\mathbf{E}' = \{E'_{\pi(1)}, E'_{\pi(2)}, \dots, E'_{\pi(d)}\}$. Here, $E'_{\pi(i)}$ is a re-encryption of E_i , while π is a uniformly random, secret permutation. This is to say that MN randomly and secretly permutes and re-encrypts inputs. Thus, the special privacy property of a mix network is this: An adversary cannot determine which output ciphertext corresponds to which input ciphertext, i.e., which inputs and outputs have common plaintexts. Stated another way, an adversary cannot determine $\pi(j)$ for any j with probability non-negligibly better than a random guess. A number of good mix network constructions have been proposed that offer privacy and robustness against a static, active adversary capable of corrupting any minority of the n players (servers) performing the mix network operation, e.g., [22] and Neff [34]. These constructions can offer the additional property of *verifiability*. In other words, a proof is output that is checkable by any party and demonstrates, relative to \mathbf{E} and the public key of the ciphertexts that \mathbf{E} is correctly constructed. It is convenient to conceptualize MN as an ideal primitive in terms of an oracle \tilde{MN} for MN with the property of public verifiability.

Proofs of knowledge: As sketched above, we use NIZK (non-interactive zero-knowledge) proofs of knowledge [6] in a number of places. We do not describe these tools in detail, as they are standard in the literature. Instead, we refer the reader to, e.g. [15], for discussion of construction and logical composition of such protocols, and [11] for notation and discussion of efficient realization. As usual, our use of NIZK proofs requires the random oracle model in security proofs for our scheme [4].

4.1 Our proposed protocol

Setup: The key pairs $(SK_{\mathcal{R}}, PK_{\mathcal{R}})$ and $(SK_{\mathcal{T}}, PK_{\mathcal{T}})$ are generated (in an appropriately trustworthy manner, as described above), and $PK_{\mathcal{T}}$ and $PK_{\mathcal{R}}$ are published along with all system parameters.

Registration: Upon proof of eligibility from V_i , the registrar \mathcal{R} generates and transmits to V_i a random string $\sigma_i \in_U \mathcal{G}$ that serves as the credential of the voter. \mathcal{R} then adds $S_i = E_{PK_{\mathcal{T}}}[\sigma_i]$ to the voter roll L .⁵ The voter roll L is maintained on the bulletin board \mathcal{BB} and digitally signed as appropriate by \mathcal{R} .

We assume the trustworthiness of \mathcal{R} as explained above. Still, if desired, \mathcal{R} can furnish V_i with a proof that S_i is a valid ciphertext on σ_i . Where erasure of secrets by voters is not automatic, a *designated verifier proof* [28] is needed for coercion resistance. We note that credentials may be used for multiple elections.

Candidate-slate publication: \mathcal{R} or some other appropriate authority publishes a candidate slate C containing the names and unique identifiers in \mathcal{G} for n_C candidates, with appropriate integrity protection. This authority also publishes a unique, random election identifier ϵ .

Voting: Voter V_i casts a ballot for candidate c_j comprising M-El Gamal ciphertexts $(E_1^{(i)}, E_2^{(i)})$ respectively on choice c_j and credential σ_i . In particular, for $a_1, a_2 \in_U Z_q$:

$$E_1^{(i)} = (\alpha_1, \alpha'_1, \beta_1) = (g_1^{a_1}, g_2^{a_1}, c_j h^{a_1}), E_2^{(i)} = (\alpha_2, \alpha'_2, \beta_2) = (g_1^{a_2}, g_2^{a_2}, \sigma_i h^{a_2}).$$

The first is a ciphertext on the candidate choice of the voter, the second a ciphertext on the credential of the voter.

Additionally, V_i includes NIZK proofs of knowledge of σ_i and c_j , and also a NIZK proof that $c_j \in C$, i.e., that c_j represents a valid candidate choice. The latter can be accomplished, for example, using a disjunctive proof that the ciphertext constitutes a valid encryption of a candidate choice in C . It is needed because an invalid candidate choice is like a write-in: It can effectively serve as a receipt. These NIZK proofs, which we denote collectively by Pf , may be accomplished efficiently using standard techniques. As is standard practice, the challenge values for Pf are constructed using a call to a cryptographic hash function, modeled in our security analysis by a random oracle $\tilde{O}W$. Input to $\tilde{O}W$ for these challenge values includes ϵ, E_1, E_2 and commitment values required for realization of the NIZK proofs. V_i posts $B_i = (E_1, E_2, Pf)$ to \mathcal{BB} via an anonymous channel.

Note: This is not a receipt! By our definitions, the adversary may not know whether or not a given voter has even posted a vote. The adversary therefore can't simply coerce by requesting decryption information.

Tallying: To tally the ballots posted to \mathcal{BB} , the authority \mathcal{T} performs the following steps:

1. **Checking proofs:** \mathcal{T} verifies the correctness of all proofs on \mathcal{BB} . Any ballots with invalid proofs are discarded. For the valid, remaining ballots, let A_1 denote the list of ciphertexts on candidate choices (i.e., the E_1 ciphertexts), and let B_1 denote the list of ciphertexts on credentials (i.e., the E_2 ciphertexts).
2. **Eliminating duplicates:** The tallying authority \mathcal{T} performs pairwise PETs on all ciphertexts in B_1 , and removes duplicates according to some pre-determined policy, using e.g., order of postings to \mathcal{BB} . When an element is removed from B_1 , the corresponding element (i.e., that with the same index) is removed from A_1 . We let B'_1 and A'_1 be the resulting “weeded” vectors. This is equivalent to retaining at most one ballot per given credential.

⁵ In our definitions above, we use the common terminology of private and public keys – with corresponding notation sk_i and pk_i – to describe the credentials associated with voters. Shifting from a general exposition to our specific protocol, we now use σ_i instead of sk_i to denote a voter credential, and S_i instead of pk_i to denote a public representation thereof. This change of notation aims to reflect the fact that voters do not employ a conventional form of public-key authentication in our scheme.

3. **Mixing:** \mathcal{T} applies MN to A'_1 and B'_1 (using the same, secret permutation for both). Let A_2 and B_2 be the resulting lists of ciphertexts.
4. **Checking credentials:** \mathcal{T} applies mix network MN to the encrypted list L of credentials from the voter roll. \mathcal{T} then compares each ciphertext of B_2 to the ciphertexts of L using PET. \mathcal{T} retains a vector A_3 of all ciphertexts of A_2 for which the corresponding elements of B_2 match an element of L according to PET. This step achieves the weeding of ballots based on invalid voter credentials.
5. **Tallying:** \mathcal{T} decrypts all ciphertexts in A_3 and tallies the final result.

How to cheat a coercer: One possible implementation of the function `fakekey` is simply for the coerced voter V_i to select and reveal a random group element $\tilde{\sigma}_i$, claiming that this is the credential σ_i . (If coerced multiple times, the voter V_i would release the same value $\tilde{\sigma}_i$.) We discuss the process of faking voting keys in more detail in appendix C.

Due to space constraints, we offer further discussion of security and a formal security proof in appendix D.

References

1. Proxyvote.com: Shareholder election website, 2005. URL: www.proxyvote.com.
2. Vote-auction, 2005. URL: www.vote-auction.net.
3. O. Baudron, P.-A. Fouque, D. Pointcheval, J. Stern, and G. Poupard. Practical multi-candidate election system. In *PODC 2001*, pages 274–283. ACM Press, 2001.
4. M. Bellare and P. Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *1st ACM Conference on Computer and Communications Security*, pages 62–73. ACM, 1993.
5. J.C. Benaloh and D. Tuinstra. Receipt-free secret-ballot elections (extended abstract). In *26th ACM STOC*, pages 544–553, 1994.
6. Manuel Blum, Alfredo De Santis, Silvio Micali, and Giuseppe Persiano. Noninteractive zero-knowledge. *SIAM J. Comput.*, 20(6):1084–1118, 1991.
7. D. Boneh. The Decision Diffie-Hellman problem. In *ANTS '98*, pages 48–63. Springer-Verlag, 1998. LNCS no. 1423.
8. D. Boneh and P. Golle. Almost entirely correct mixing with applications to voting. In V. Atluri, editor, *ACM CCS '02*, pages 68–77. ACM Press, 2002.
9. S. Brands. *Rethinking Public Key Infrastructures and Digital Certificates: Building in Privacy*. MIT Press, 2000.
10. J. Camenisch and A. Lysyanskaya. An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In B. Pfitzmann, editor, *EUROCRYPT '01*, pages 93–118. Springer-Verlag, 2001. LNCS no. 2045.
11. J. Camenisch and M. Stadler. Efficient group signature schemes for large groups. In B. Kaliski, editor, *CRYPTO '97*, pages 410–424. Springer-Verlag, 1997. LNCS no. 1294.
12. R. Canetti, R. Gennaro, S. Jarecki and H. Krawczyk, and T. Rabin. Adaptive security for threshold cryptosystems. In M. Wiener, editor, *CRYPTO '99*, pages 98–115. Springer-Verlag, 1999. LNCS no. 1666.
13. Ran Canetti, Cynthia Dwork, Moni Naor, and Rafail Ostrovsky. Deniable encryption. In B. Kaliski, editor, *CRYPTO '97*, pages 90–104, 1997. LNCS no. 1294.
14. D. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2):84–88, 1981.
15. R. Cramer, I. Damgard, and B. Schoenmakers. Proofs of partial knowledge and simplified design of witness hiding protocols. In Y. Desmedt, editor, *CRYPTO '94*, pages 174–187. Springer-Verlag, 1994. LNCS no. 839.
16. R. Cramer, R. Gennaro, and B. Schoenmakers. A secure and optimally efficient multi-authority election scheme. In W. Fumy, editor, *EUROCRYPT '97*, pages 103–118. Springer-Verlag, 1997. LNCS no. 1233.
17. R. Cramer and V. Shoup. A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In H. Krawczyk, editor, *CRYPTO '98*, pages 13–25. Springer-Verlag, 1998. LNCS no. 1462.
18. M. Cross. Public domain. *The Guardian: Guardian Unlimited Online*, 10 June 2004. Available at <http://www.guardian.co.uk/online/insideit/story/0,13270,1234942,00.html>.
19. J. Fowler. Switzerland tests virtual democracy in national referendum. *Technology Review*, 26 September 2004. AP Newswire. Available at http://www.technologyreview.com/articles/04/09/ap_092604.asp.
20. A. Fujioka, T. Okamoto, and K. Ohta. A practical secret voting scheme for large scale elections. In J. Seberry and Y. Zheng, editors, *ASIACRYPT '92*, pages 244–251. Springer-Verlag, 1992. LNCS no. 718.
21. J. Furukawa. Efficient, verifiable shuffle decryption and its requirement of unlinkability. In Bao et al., editor, *PKC 04*, pages 319–332. Springer-Verlag, 2004. LNCS no. 2947.
22. J. Furukawa and K. Sako. An efficient scheme for proving a shuffle. In J. Kilian, editor, *CRYPTO '01*, volume 2139 of *Lecture Notes in Computer Science*, pages 368–387. Springer-Verlag, 2001.
23. T. El Gamal. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE Transactions on Information Theory*, 31:469–472, 1985.
24. R. Gennaro, S. Jarecki, H. Krawczyk, and T. Rabin. The (in)security of distributed key generation in dlog-based cryptosystems. In J. Stern, editor, *EUROCRYPT '99*, pages 295–310. Springer-Verlag, 1999. LNCS no. 1592.

25. M. Hirt and K. Sako. Efficient receipt-free voting based on homomorphic encryption. In B. Preneel, editor, *EUROCRYPT '00*, pages 539–556, 2000. LNCS no. 1807.
26. M. Jakobsson and A. Juels. Mix and match: Secure function evaluation via ciphertexts. In T. Okamoto, editor, *Advances in Cryptology - Asiacrypt '00*, pages 162–177. Springer-Verlag, 2000. LNCS No. 1976.
27. M. Jakobsson, A. Juels, and R. Rivest. Making mix nets robust for electronic voting by randomized partial checking. In D. Boneh, editor, *USENIX '02*, pages 339–353, 2002.
28. M. Jakobsson, K. Sako, and R. Impagliazzo. Designated verifier proofs and their applications. In U. Maurer, editor, *EUROCRYPT '96*, pages 143–154. Springer-Verlag, 1996. LNCS no. 1070.
29. A. Juels and J. Brainard. Client puzzles: A cryptographic countermeasure against connection depletion attacks. In *Proceedings of the 1999 ISOC Network and Distributed System Security Symposium*, pages 151–165, 1999.
30. A. Kiayias and M. Yung. Self-tallying elections and perfect ballot secrecy. In D. Naccache and P. Paillier, editors, *PKC '02*, pages 141–158. Springer-Verlag, 2000. LNCS no. 2274.
31. P. MacKenzie, T. Shrimpton, and M. Jakobsson. Threshold password-authenticated key exchange. In M. Yung, editor, *CRYPTO '02*, pages 385–400, 2002. LNCS no. 2442.
32. E. Magkos, M. Burmester, and V. Chrissikopoulos. Receipt-freeness in large-scale elections without untappable channels. In B. Schmid *et al.*, editor, *First IFIP Conference on E-Commerce, E-Business, E-Government (I3E)*, pages 683–694, 2001.
33. M. Michels and P. Horster. Some remarks on a receipt-free and universally verifiable mix-type voting scheme. In K. Kim and T. Matsumoto, editors, *ASIACRYPT '96*. Springer-Verlag, 1996. LNCS no. 1163.
34. A. Neff. A verifiable secret shuffle and its application to e-voting. In P. Samarati, editor, *ACM CCS '01*, pages 116–125. ACM Press, 2001.
35. V. Niemi and A. Renvall. How to prevent buying of votes in computer elections. In J. Pieprzyk and R. Safavi-Naini, editors, *ASIACRYPT '94*, pages 164–170. Springer-Verlag, 1994. LNCS no. 917.
36. T. Okamoto. An electronic voting scheme. In N. Terashima *et al.*, editor, *IFIP World Congress*, pages 21–30, 1996.
37. T. Okamoto. Receipt-free electronic voting schemes for large scale elections. In B. Christianson *et al.*, editor, *Security Protocols Workshop*, pages 25–35. Springer-Verlag, 1997. LNCS no. 1361.
38. S. Parker. Shaking voter apathy up with IT. *The Guardian*, 11 Dec. 2001.
39. K. Sako and J. Kilian. Receipt-free mix-type voting scheme - a practical solution to the implementation of a voting booth. In L. Guillou and J.-J. Quisquater, editors, *EUROCRYPT '95*, pages 393–403. Springer-Verlag, 1995. LNCS no. 921.
40. C.-P. Schnorr. Efficient signature generation by smart cards. *Journal of Cryptology*, 4(3):161–174, 1991.
41. B. Schoenmakers. A simple publicly verifiable secret sharing scheme and its application to electronic voting. In M. Wiener, editor, *CRYPTO '99*, pages 148–164. Springer-Verlag, 1999. LNCS no. 1666.
42. B. Schoenmakers, 2000. Personal communication.
43. A. Shamir. How to share a secret. *Communications of the Association for Computing Machinery*, 22(11):612–613, November 1979.
44. Y. Tsiounis and M. Yung. On the security of ElGamal-based encryption. In *Workshop on Practice and Theory in Public Key Cryptography (PKC '98)*. Springer, 1998.

A Definitions of Correctness and Verifiability

Correctness: We first consider the property of correctness. This property is in fact twofold: First, it stipulates that an adversary \mathcal{A} cannot pre-empt, alter, or cancel the votes of honest, i.e., voters that are not *controlled*; Second, it stipulates that \mathcal{A} cannot cause voters to cast ballots in such a way as to achieve double voting, i.e., use of one credential to vote multiple times, where more than one vote per credential is counted in the tally.

In our experiment characterizing correctness, we give the adversary powers she does not normally have. Namely, apart from getting to select a set V of voters she will control, we also allow her to choose the candidate-slate size n_C , and to choose what votes will be cast by voters she does not control. The latter voters will indeed vote according to the adversary’s wish – but only for the purposes of our thought experiment defining correctness, of course. If the adversary still cannot cause an incorrect tally to be computed (i.e., one not corresponding to the votes cast), then the scheme has the correctness property even in the real-world scenario in which the adversary has less power. The aim of the adversary is to cause more than $|V|$ ballots to be counted in the final tally on behalf of the controlled voters, or to alter or delete the vote of at least one honest voter. (This corresponds to the condition that: (1) The verification of the tally succeeds, and (2) That either a vote is “dropped” or “added”.) Our definition assumes implicitly that tally is computed correctly by the authority \mathcal{T} . (The next property we consider, namely verifiability, addresses the possibility that this is not so.) In what follows, we let $\langle \mathbf{Y} \rangle$ denote the multiset corresponding to entries in the vector \mathbf{Y} , and $|Y|$ denote the cardinality of set Y .

Experiment $\text{Exp}_{\text{ES},\mathcal{A}}^{\text{corr}}(k_1, k_2, k_3, n_C, n_V)$

$\{(sk_i, pk_i) \leftarrow \text{register}(SK_{\mathcal{R}}, i, k_2)\}_{i=1}^{n_V};$	% voters are registered
$V \leftarrow \mathcal{A}(\{pk_i\}_{i=1}^{n_V}, \text{“choose controlled voter set”});$	% \mathcal{A} corrupts voters
$\{\beta_i\}_{i \notin V} \leftarrow \mathcal{A}(\text{“choose votes for uncontrolled voters”});$	% \mathcal{A} chooses votes for honest voters
$\mathcal{BB} \leftarrow \{\text{vote}(sk_i, PK_{\mathcal{T}}, n_C, \beta_i, k_2)\}_{i \notin V};$	% honest voters cast ballots
$(\mathbf{X}, P) \leftarrow \text{tally}(SK_{\mathcal{T}}, \mathcal{BB}, n_C, \{pk_i\}_{i=1}^{n_V}, k_3);$	% honest ballots are tallied
$\mathcal{BB} \leftarrow \mathcal{A}(\text{“cast ballots”}, \mathcal{BB});$	% \mathcal{A} posts ballots to \mathcal{BB}
$(\mathbf{X}', P') \leftarrow \text{tally}(SK_{\mathcal{T}}, \mathcal{BB}, n_C, \{pk_i\}_{i=1}^{n_V}, k_3);$	% all ballots are tallied
if $\text{verify}(PK_{\mathcal{T}}, \mathcal{BB}, n_C, \mathbf{X}', P') = \text{‘1’}$ and	% does function verify accept?
$(\{\beta_i\} \notin \langle \mathbf{X}' \rangle \text{ or } \langle \mathbf{X}' \rangle - \langle \mathbf{X} \rangle > V)$ then	% did \mathcal{A} successfully tamper?
output ‘1’;	
else	
output ‘0’;	

We say that ES possesses the property of correctness if for all polynomial-time adversaries \mathcal{A} , it is the case that $\text{Succ}_{\text{ES},\mathcal{A}}^{\text{corr}}(k_1, k_2, k_3, n_V)$ is negligible.

Verifiability: As explained above, an election system has the property of correctness if computation of tally always yields a valid tabulation of ballots. Given the ability of an adversary \mathcal{A} , however, to corrupt some number of authorities among \mathcal{T} , we cannot be assured that tally is always computed correctly. The property of verifiability is the ability for any player to check whether the tally \mathbf{X} has been correctly computed, that is, to detect any misbehavior by \mathcal{T} in applying the function tally.

A strong security definition for verifiability is appropriate given the high level of auditability required for trustworthy elections. Such a definition considers an attacker \mathcal{A} capable of controlling *all* of the voters and tallying authorities in \mathcal{T} . This attacker seeks to construct a set of ballots on \mathcal{BB} and a corresponding tally \mathbf{X} and proof P of correct tabulation such that the proof is accepted by `verify`, but the tally is in fact incorrect. By an incorrect tally, we mean one in which all of the valid ballots of a particular voter (i.e., corresponding to a particular credential) are discounted, or else where multiple votes are tallied that could have been generated by the same voting credential. Our experiment characterizing verifiability is as follows.

Experiment $\text{Exp}_{\text{ES},\mathcal{A}}^{\text{ver}}(k_1, k_2, k_3, n_C, n_V)$

$\{(sk_i, pk_i) \leftarrow \text{register}(SK_{\mathcal{R}}, i, k_2)\}_{i=1}^{n_V};$	% voters are registered
$(\mathcal{BB}, \mathbf{X}, P) \leftarrow \mathcal{A}(SK_{\mathcal{T}}, \{(sk_i, pk_i)\}_{i=1}^{n_V}, \text{“forge election”});$	% \mathcal{A} concocts full election
$(\mathbf{X}', P') \leftarrow \text{tally}(SK_{\mathcal{T}}, \mathcal{BB}, n_C, \{pk_i\}_{i=1}^{n_V}, k_3);$	% tally is taken on \mathcal{BB}
if $\mathbf{X} \neq \mathbf{X}'$	% does \mathcal{A} 's tally differ from correct \mathcal{BB} tally?
and $\text{verify}(PK_{\mathcal{T}}, \mathcal{BB}, n_C, \mathbf{X}, P) = \text{‘1’}$ then	% does function verify accept?
output ‘1’;	
else	
output ‘0’;	

We say that ES possesses the property of verifiability if for all positive integers n_V and all adversaries \mathcal{A} with polynomial running time, the quantity $\text{Succ}_{\text{ES},\mathcal{A}}^{\text{ver}}(k_1, k_2, k_3, n_V)$ is negligible. A technical strengthening of this definition and that for correctness is possible, and discussed in the next section, appendix B, of this paper.

Another aspect of verifiability that we do not formally define, but do mention here and incorporate into our proposed protocol is that of verification against voter rolls. In particular, it may be desirable for any election observer to check that credentials were assigned only to voters whose names are on a published roll. This is not technically a requirement if we rule out corruption of players \mathcal{R} , but may still be desirable for high assurance of election integrity. Our definitions can be modified accordingly.

B Remark on strong verifiability

We set forth our definitions of correctness and verifiability in appendix A to meet the minimal requirements for a fair election and to achieve some measure of conceptual simplicity. These definitions are adequate for most election scenarios, but have a technical deficiency that may be of concern in some cases. In particular, our definitions allow for the possibility that a voter controlled by \mathcal{A} casts a ballot corresponding to vote β , but that the ballot gets counted as a vote for β' . Since \mathcal{A} can choose the vote cast by a controlled voter in any case, this technical deficiency only means that \mathcal{A} can potentially cause the votes of *controlled voters only* to change in the midst of the election process. It does not provide \mathcal{A} with control of a larger number of votes. Most importantly, we note that this definitional weakness does not apply to our proposed protocol, which meets the stronger definition we now set forth.

Nonetheless, one can envisage some (somewhat artificial) scenarios in which stronger guarantees may be desirable. For example, \mathcal{A} might have the aim of causing the victor in an election to win by the slimmest possible margin. In this case, if \mathcal{A} controls a majority of \mathcal{T} , then \mathcal{A} might seek to decrypt all of the ballots cast in an election and alter the votes of controlled voters so as to favor the losing candidate.

We discuss now how our definition of verifiability may be modified to discount the possibility of this type of attack. (Analogous modifications may be made to the definition of correctness.) In particular, we can require that P be a proof that every tallied vote corresponds uniquely to a credential for which a valid ballot has been cast. For this, we require a natural technical restriction on `vote`. Let $\langle \text{vote}(\cdot) \rangle$ denote the set of possible outputs for the randomized function `vote` on a particular input. We require that an output ballot be wholly unambiguous with respect to both the vote β and the credential sk . In other words, we require $\langle \text{vote}(sk_0, PK_{\mathcal{T}}, n_C, \beta_0, k_2) \rangle \cap \langle \text{vote}(sk_1, PK_{\mathcal{T}}, n_C, \beta_1, k_2) \rangle = \emptyset$ if $\beta_0 \neq \beta_1$ or $sk_0 \neq sk_1$.

To achieve our strengthened definition of verifiability, we alter experiment $\text{Exp}_{\text{ES}, \mathcal{A}}^{\text{ver}}(k_1, k_2, k_3, n_V)$ such that if the following conditions 1 and 2 are met, then the output of the experiment is '1'. Otherwise it is '0'.

1. $\text{verify}(PK_{\mathcal{T}}, \mathcal{BB}, n_C, \mathbf{X}, P) = '1'$
2. For every injective mapping $f : \langle \mathbf{X} \rangle \rightarrow Z_{n_V}$ one of two conditions holds:
 - (a) $\exists B : B \in \mathcal{BB}, B \in \langle \text{vote}(sk_i, PK_{\mathcal{T}}, n_C, \beta, k_2) \rangle, \forall j f(j) \neq i$
 - (b) $\exists \beta \in \mathbf{X} : f(\beta) = i, \forall B \in \mathcal{BB}, B \notin \langle \text{vote}(sk_i, PK_{\mathcal{T}}, n_C, \beta, k_2) \rangle$

Conditions 2(a) and 2(b) here respectively specify that the adversary has successfully defeated the verifiability of the system either by causing all of the valid ballots associated with a particular credential not to be counted or else enabling multiple votes to be tallied for a single credential.

Given use of a verifiable mix network, our proposed protocol meets this stronger security definition for verifiability.

C The Faking of Voting Keys

We provide some more detail here on the process whereby a voter fakes a voting credential in our proposed protocol. Upon receiving a claimed credential $\tilde{\sigma}_i$, the adversary would like to verify if it is correct. Let us consider the possibility of doing so under each of our three possible assumptions on the registration phase discussed in the body of the paper; in doing so, recall that we always assume that the adversary can corrupt only a minority of servers in \mathcal{T} , and so, will not be able to decrypt any of the semantically secure encryptions of credentials.

1. Assume that there is a mechanism forcing erasure of voter information no longer needed at the end of the registration phase, and that only a minority of servers in \mathcal{R} may be corrupted. At the end of the registration process, each voter will erase information specifying what part of the transcript leading to the credential σ_i he got from what registration server. Without proofs or transcripts from individual servers of \mathcal{R} , it is not possible for the adversary to verify the correctness of $\tilde{\sigma}_i$.
2. Assume that the adversary cannot corrupt *any* server in \mathcal{R} . As mentioned, the registration servers may if desired use designated verifier proofs to prove to each voter that the share they send is authentic (i.e., will be part of the recorded transcript S_i). While the voter will be convinced of these proofs, the adversary will not; in fact, he

cannot distinguish between real such proofs and proofs simulated by V_i . Therefore, V_i can convincingly release full *simulated* transcripts from the registration phase, corresponding to a credential $\tilde{\sigma}_i$.

3. Assuming that the user knows what (minority of) servers in \mathcal{R} are corrupted, but is not necessarily able to erase data, he can present the adversary with registration transcripts that are consistent with the view of the servers he knows to be corrupted, but inconsistent (in terms of the real share of σ_i) with the view of the servers that are not. The latter transcripts will be accompanied by simulated designated verifier proofs. Since the adversary may only corrupt a minority of servers in \mathcal{R} , and a majority is required to compute the credential σ_i , there will be at least one share of σ_i that V_i can change to obtain a fake credential $\tilde{\sigma}_i \neq \sigma_i$, without the detection of the adversary.

D Proving Coercion-Freeness

In this section, we provide a detailed outline for proof of the property of coercion-freeness in our proposed election protocol. (We do not consider correctness or verifiability here, as these are more standard properties, and the associated proofs are more straightforward.) For the purposes of this proof, we assume the use of the M-El Gamal cryptosystem over a preselected group \mathcal{G} of order q . The coercion-freeness of our scheme is dependent on the Decision-Diffie Hellman (DDH) assumption on \mathcal{G} . Briefly stated, this assumption states that no algorithm with running-time polynomial in the security parameters for \mathcal{G} can distinguish between the two distributions D and D' with non-negligible probability: Here, D is the distribution of tuples of the form (y_1, g_1, y_2, g_2) , where $g_1, g_2 \in_U \mathcal{G}$, $y_1 = g_1^x$, and $y_2 = g_2^x$ for $x \in_U \mathbb{Z}_q$; i.e., the pair (y_1, g_1) and (y_2, g_2) are related by a common exponent. D' is the distribution of random tuples, i.e., tuples of the form (y_1, g_1, y_2, g_2) , where $y_1, g_1, y_2, g_2 \in_U \mathcal{G}$. For detailed treatment of this assumption (expressed in an alternative, equivalent form), see, e.g., [7].

D.1 Assumptions

As explained above, we simplify our analysis by assuming ideal constructions for a number of components in our election protocol. Our aim in doing so is twofold: (1) Our protocol is flexible enough to accommodate a range of cryptographic building blocks from the literature and (2) We wish to retain a focus on the conceptual and definition elements of our paper, and not on protocol details. Hence, we assume the availability of oracles for the four following cryptographic operations in our protocol: mixing, plaintext equivalence testing (PET), threshold ciphertext decryption, and calls to the one-way or hash function required for NIZK proofs. As in the main body of the paper, denote these oracles respectively by $\tilde{M}N$, $\tilde{P}ET$, $\tilde{D}EC$ and $\tilde{O}W$. Although the functioning of these oracles should be clear from our protocol description, we present it again here:

- The oracle $\tilde{M}N$ performs exactly the same function as a mix network. It accepts as input an ordered list $\mathbf{E} = \{E_1, E_2, \dots, E_d\}$ of ciphertexts under the public key $PK_{\mathcal{T}}$ of the tallying authorities. Its output on \mathbf{E} is an ordered set $\mathbf{E}' = \{E'_{\pi(1)}, E'_{\pi(2)}, \dots, E'_{\pi(d)}\}$ for a secret, random permutation π , where $E'_{\pi(i)}$ represents a re-encryption of ciphertext E_i .
- The oracle $\tilde{P}ET$ takes as input a pair of ciphertexts (E, E') under $PK_{\mathcal{T}}$. It outputs a ‘1’ if E and E' have identical corresponding plaintexts, and outputs ‘0’ otherwise.
- The oracle $\tilde{D}EC$ takes as input a ciphertext E under $PK_{\mathcal{T}}$. It outputs the corresponding plaintext.
- The oracle $\tilde{O}W$ takes as input a query value in $\{0, 1\}^*$, and outputs a random value $\{0, 1\}^{k_4}$, where k_4 is a security parameter (that may depend on k_1, k_2 and k_3). The output of $\tilde{O}W$ is consistent, in the sense that a given input value always yields the same output value. This oracle may be viewed as the ideal embodiment of a cryptographic hash function.

Each of these oracles accepts publicly viewable input from all participating authorities (talliers). Each tallier may be thought of as having a publicly readable tape to which it may write input values for a given oracle; each tape contains a write portion for each time-step of the protocol, which we assume to be synchronous. At the end of a given timestep, an oracle produces output according to the following procedure. If a majority of talliers have

furnished identical non-null values Z on their tapes, then the oracle processes input Z and yields the corresponding output. If there is no non-null majority input, then the oracle simply outputs the special symbol \perp . The requirement for majority input ensures that the protocol execution is determined by honest players, i.e., effectively reduces \mathcal{A} to an honest-but-curious adversary once the ballot-posting phase for the election is complete.

We additionally assume for simplicity that key setup and registration are performed by a trusted entity. Our proofs may be extended to accommodate more general assumptions in which these two processes are performed in a distributed manner.

D.2 Proof overview

Recall that our definition of coercion-freeness revolves around a game played between an adversary \mathcal{A} and a voter targeted for coercion. The aim of \mathcal{A} is to guess which of the following two behaviors the voter has adopted during the execution of an election system ES: (1) The voter has divulged valid voting credentials and abstained from voting or (2) The voter has divulged fake credentials and cast a ballot. In order to demonstrate that ES possesses coercion-freeness, we must show that \mathcal{A} can guess successfully with probability only negligibly better than a weaker adversary \mathcal{A}' interacting with an ideal election system. This adversary \mathcal{A}' is passive, and its only input is the final tally \mathcal{X} of votes cast by honest voters in the completed election plus T , the number of ballots eliminated for invalid associated credentials.

Our proof strategy is to construct a polynomial-time algorithm \mathcal{S} that takes a set of ballots W of honest voters and simulates the election system ES in the experiment *c-resist*. If the simulation is indistinguishable to \mathcal{A} from use of the true functional components of ES, and \mathcal{A} cannot cause the simulation to deviate from correct execution, then we see that \mathcal{A} learns nothing more than the correct election tally \mathcal{X} and the number of bad ballots T . This means in turn that \mathcal{A} is no more powerful than the ideal adversary \mathcal{A}' characterized in our experiment *c-resist-ideal*. Thus ES is coercion-free.

The inability of the adversary to cause deviation in the experiment from correct execution hinges on our oracle definitions, which require majority agreement on input values. Given this, we show that the simulation produced by \mathcal{S} is indistinguishable by \mathcal{A} from a real experimental execution of *c-resist* under the DDH assumption on \mathcal{G} . Our proof relies on the semantic security of M-El Gamal (see appendix E). In particular, we make use of the following, useful fact implied by the DDH assumption: A poly-time adversary that selects a plaintext m cannot distinguish between the distribution of M-El Gamal ciphertexts on m (A_1, A_2, B) and the distribution of random triplets in \mathcal{G} with non-negligible probability (in the security parameters for \mathcal{G}). In consequence of this observation, it is possible for \mathcal{S} to simulate the election process by substituting *random ciphertexts*, i.e., random triplets of group elements, for the real ciphertexts that would be processed in a true execution of the experiment *c-resist*. In particular, \mathcal{S} can simulate the ballots of voters not controlled by \mathcal{A} with a list of random ciphertexts. Additionally, \mathcal{S} can simulate the oracle \overline{MN} by setting its simulated output to a list of random ciphertexts. Under the DDH assumption, \mathcal{A} cannot distinguish between the random ciphertexts furnished by \mathcal{S} and the ciphertexts that would be processed in a true execution of ES.

D.3 The simulation

We now outline the steps of the simulation of *c-resist* executed by \mathcal{S} . Throughout the simulation, according to the usual technique in the literature, \mathcal{S} maintains state for the simulated oracle \overline{OW} so as to ensure consistency of output values. Let $W \in D_{n_U, n_C}$ represent a set of ballots input into the simulation as representing the posting of honest voters. At the very beginning the simulator receives a quadruple (g_1, g_2, h_1, h_2) which is either a Diffie-Hellman quadruple or a random one, according to some hidden bit d . More formally, $d = 1$ if the quadruple is a DH one and $d = 0$ otherwise. The goal of the simulator is to guess which situation is dealing with.

1. **Setup:** \mathcal{S} chooses uniformly and at random two elements $x_1, x_2 \in_U Z_q$ and sets $h = g_1^{x_1} g_2^{x_2} \bmod p$. \mathcal{S} publishes the public key (g_1, g_2, h) and also a randomized candidate slate $\mathcal{C} = \{c_i\}_{i=1}^{n_C}$ such that $c_i = g_1^{r_i}$ for $r_i \in_U Z_q$. (For technical reasons in our proof, we require that candidate identifiers here be random, rather than comprising the set $\{1, 2, \dots, n_C\}$.)

2. **Registration:** \mathcal{S} simulates the registrar \mathcal{R} , generating a set of credentials $\{\sigma_i = g_1^{s_i}\}$ for $s_i \in_U Z_q$. For the encrypted credential list L_0 , the simulator \mathcal{S} publishes a list of n_V ciphertexts. These ciphertexts are generated using the received Diffie-Hellman challenge as follows. For each σ_i , \mathcal{S} chooses a random $\omega_i \in Z_q$ and sets $(L_{i,1} = h_1^{\omega_i}, L_{i,2} = h_2^{\omega_i}, L_{i,3} = h_1^{\omega_i x_1} h_2^{\omega_i x_2} \sigma_i)$.
3. **Adversarial corruption:** The adversary \mathcal{A} selects a set V of n_A voters to corrupt, as well as a voter j for coercion and a target vote β . If any of these selections are invalid, i.e., if $V \neq n_A$ or $j \notin V - V$ or $\beta \notin \mathcal{C} \cup \phi$, then the simulation is terminated.
4. **Coin flip:** A coin $b \in_U \{0, 1\}$ is flipped.
5. **Credential release:** \mathcal{S} gives \mathcal{A} the set of credentials $\{\sigma_i\}_{i \in V}$ as well as a credential σ for the targeted voter j . If $b = 1$, then \mathcal{S} gives $\sigma = \sigma_j$; otherwise σ is a random string.
6. **Honest voter simulation:** For each ballot element in W , the simulator posts a ballot consisting of two ciphertexts $(\alpha_{i,1}, \alpha'_{i,1}, \beta_{i,1}), (\alpha_{i,2}, \alpha'_{i,2}, \beta_{i,2})$. \mathcal{S} also furnishes the associated NIZK proofs of the form specified above. Since the associated challenges value comes from $\tilde{O}W$, and may therefore be predetermined by \mathcal{S} , the NIZK proof may be simulated using standard techniques. Let A_0 be the list of these ballots. Let A^* be the associated set of plaintext ballot choices in W for which the associated credential is correct, i.e., excluding λ elements. The simulator creates the ciphertexts above as follows. For each ballot element in W , \mathcal{S} chooses two elements r_i, k_i at random in Z_q and sets $(\alpha_{i,1} = h_1^{r_i}, \alpha'_{i,1} = h_2^{r_i}, \beta_{i,1} = h_1^{r_i x_1} h_2^{r_i x_2} c_j), (\alpha_{i,2} = h_1^{k_i}, \alpha'_{i,2} = h_2^{k_i}, \beta_{i,2} = h_1^{k_i x_1} h_2^{k_i x_2} \sigma_i)$.
7. **Adversarial ballot posting:** The adversary \mathcal{A} posts a set of ballots B_0 and associated NIZK proofs.
8. **Decryption of ballots posted by the adversary** \mathcal{S} checks the NIZK proofs in B_0 . Let B_1 be the list of ballots with correct proofs. For each ballot in B_1 and each credential in $\{\sigma_i\}_{i \in V} \cup \sigma_j$, the simulator decrypts using his own private key (see above).
9. **Tallying simulation:** \mathcal{S} simulates the behavior of honest tallying authorities. Since these are a majority, any deviating behavior by tallying authorities in the control of \mathcal{A} may be ignored. This part of the simulation proceeds as follows:
 - (a) **Proof checking:** Let E_0 denote the combined list of input ballots A_0 and B_0 . \mathcal{S} simulates the behavior of honest tallying authorities in rejecting all ballots with invalid associated NIZK proofs. Let E_1 be the resulting ballot list.
 - (b) **Eliminating duplicates:** Since no mixing has yet occurred, \mathcal{S} may simulate the elimination of duplicate ballots using its own decryption key. Let E_2 be the resulting ballot list.
 - (c) **Mixing:** \mathcal{S} simulates the oracle $\tilde{M}N$ as applied to E_2 by outputting an equal-length list E_3 of random ciphertext triples. Likewise, \mathcal{S} simulates the mixing of L_0 by outputting an equal-lengthed list L_1 of random ciphertexts.
 - (d) **Checking credentials:** \mathcal{S} simulates the process of credential checking. In a true protocol execution, this would involve sequential comparison using $P\tilde{E}T$ between each ballot in E_3 (more precisely, the credential ciphertext therein) and the ciphertexts in L_1 . Either a match is found, in which case a ballot is deemed to be based on a valid credential, or else the list L_1 is exhausted, and the ballot is rejected. \mathcal{S} simulates the output of $P\tilde{E}T$ for this phase of the protocol using its own decryption key as before. Let E_4 be the resulting ballot list.
 - (e) **Decryption:** This is done straightforwardly.

Now if the adversary outputs a guess bit b' the simulator returns b' as his own guess for the decisional Diffie-Hellman challenge.

Observe that if the simulator's input is a Diffie-Hellman triplet (that is $d = 1$) then the simulation above is perfectly indistinguishable from the experiment $\text{Exp}_{\mathcal{ES}, \mathcal{A}, H}^{c-resist}$.

As a matter of fact, assuming $g_1 = g, g_2 = g^a, h_1 = g^b, h_2 = g^{ab}$ for some g , any ciphertext of the form $(\alpha_{i,1} = h_1^{r_i}, \alpha'_{i,1} = h_2^{r_i}, \beta_{i,1} = h_1^{r_i x_1} h_2^{r_i x_2} m)$ is actually a valid one. Indeed $h_1^{r_i} = g^{br_i} = g_1^{br_i}, h_2^{r_i} = g^{abr_i} = g_2^{br_i}$ and $h_1^{r_i x_1} h_2^{r_i x_2} m = g^{br_i x_1} g^{abr_i x_2} m = g_1^{br_i x_1} g_2^{br_i x_2} m = h^{br_i} m$.

This means that

$$\Pr[S = 1 | d = 1] = \Pr[\text{Exp}_{\mathcal{ES}, \mathcal{A}, H}^{c-resist}(\mathcal{V}) = 1] = \text{Succ}_{\mathcal{ES}, \mathcal{A}}^{c-resist}(\mathcal{V})$$

where we denoted with \mathcal{V} the view of the adversary.

On the other hand if the simulator's input is not a Diffie-Hellman triplet (that is $d = 0$) then the view produced by the simulation above does not give any information (in a strong information theoretic sense) about the votes posted by the honest parties. This is because, assuming $g_1 = g, g_2 = g^a, h_1 = g^b, h_2 = g^c$ for some $c \in_U Z_q$, one has that a ciphertext of the form $(\alpha_{i,1} = h_1^{r_i}, \alpha'_{i,1} = h_2^{r_i}, \beta_{i,1} = h_1^{r_i x_1} h_2^{r_i x_2} m)$ actually "masks" the message m perfectly. Indeed $h_1^{r_i} = g^{br_i} = g_1^{br_i}, h_2^{r_i} = g^{cr_i} = g_2^{c'r_i}$ and $h_1^{r_i x_1} h_2^{r_i x_2} m = g^{br_i x_1} g^{cr_i x_2} m = g_1^{br_i x_1} g_2^{c'r_i x_2} m = g_1^{br_i x_1} g_2^{c''r_i x_2} m = h^{br_i} g_2^{c''r_i x_2} m$.

This means that, in this case, the probability that the simulator outputs one is equal to the probability that the adversary outputs one in experiment $\mathbf{Exp}^{c\text{-resist-ideal}}$.

More formally

$$\Pr[\mathcal{S} = 1 | d = 0] = \Pr[\mathbf{Exp}_{\mathcal{ES}, \mathcal{A}, H}^{c\text{-resist-ideal}}(\mathcal{V}) = 1] = \mathbf{Succ}_{\mathcal{ES}, \mathcal{A}}^{c\text{-resist-ideal}}(\mathcal{V})$$

This means that

$$\mathbf{Adv}_{\mathcal{S}}^{\text{ddh}} = |\Pr[\mathcal{S} = 1 | d = 1] - \Pr[\mathcal{S} = 1 | d = 0]| = \mathbf{Adv}_{\mathcal{ES}, \mathcal{A}}^{c\text{-resist}}$$

under the Decisional Diffie-Hellman Assumption this quantity is negligible.

E Some details on primitives

El Gamal: Recall that we let \mathcal{G} denote the algebraic group over which we employ El Gamal, and q denote the group order. For semantic security, we require that the Decision Diffie-Hellman assumption hold over \mathcal{G} [7, 44]. A public/private key pair in El Gamal takes the form $(y(= g^x), x)$, where $x \in_U Z_q$. We let \in_U here and elsewhere denote uniform, random selection from a set. The private key x may be distributed among the n_T players in \mathcal{T} using (t, n_T) -Shamir secret sharing [43] over $GF[q]$, for $t > n_T/2$. This private key may be generated by a trusted third party or via a computationally secure simulation of this process [24]. Each player then holds a public/private key pair $(y_i(= g^{x_i}), x_i)$, where x_i is a point on the polynomial used for the secret sharing. A ciphertext in El Gamal on message $m \in \mathcal{G}$ takes the form $(\alpha, \beta) = (m y^r, g^r)$ for $r \in_U Z_q$. For succinctness of notation in the body of the paper, we sometimes let $E_y[m]$ denote a ciphertext on message m under public key y . To re-encrypt a ciphertext (α, β) , it suffices to multiply it pairwise by a ciphertext on $m = 1$, i.e., to compute a new ciphertext $(\alpha', \beta') = (y^{r'} \alpha, g^{r'} \beta)$ for $r' \in_U Z_q$.

To decrypt a ciphertext (α, β) , the plaintext $m = \alpha/\beta^x$ is computed. To achieve a threshold decryption of ciphertext (α, β) , each active player i publishes a decryption share $\beta_i = \beta^{x_i}$. The value β^x , and thus m , may be computed using standard LaGrange interpolation. Player i may prove the correctness of its share using an NIZK proof of the form $PK\{s : \beta_i = \beta^s \wedge u_i = g^s\}$ – essentially two Schnorr identification proofs [40] with conjunction achieved using techniques described in, e.g., [15]. We omit many details in this description regarding the scheduling of these operations and the use of commitments to avoid adversarial bias. (The reader is referred to, e.g., [12, 24] for some discussion of these issues in relation to key generation.)

Modified El Gamal: As mentioned before our modified version of the El Gamal cryptosystem can be seen as a simplified version of the Cramer-Shoup [17], method. It is rather straightforward to prove that the scheme is actually semantically secure under the decisional Diffie-Hellman assumption. The argument closely follows the one presented in [17]. Here we provided a sketched version of such an argument. Imagine there exists a probabilistic polynomial time algorithm A which can break the semantic security of the proposed scheme. Then our goal is to describe a different algorithm S (a simulator) which uses A to break the decisional DH problem. So assume S receives on input a quadruple (g_1, g_2, h_1, h_2) and has to determine if this is a DDH quadruple or not. S constructs the public key (for the M-El Gamal scheme) as prescribed by the key generation algorithm, i.e. it chooses x_1 and x_2 at random and sets $h = g_1^{x_1} g_2^{x_2}$.

Encryption and decryption remain unchanged

Next when the adversary comes up with the two messages m_0, m_1 he wants to be challenged on S proceeds as follows. It flips a random (private) bit b , and encrypts m_b as follows

$$(h_1^{kx_1} h_2^{kx_2} m, h_1^k, h_2^k)$$

(where k is a random value)

Note that if the given quadruple is a DH one the ciphertext has the right distribution. This is because $h_1^k = g_1^{k'}$ and $h_2^k = g_2^{k'}$ for some k' and $(h_1^{x_1} h_2^{x_2})^k = h^{k'}$ (for the same k')

If, on the other hand, the given quadruple is not a DH one then it is easy to check that the A gains no information at all about the encrypted message (this is because this time to decrypt adv has to know the secret exponents x_1 and x_2 which remains information theoretically hidden by h).

Mix networks: As explained above, there are many good choices of mix networks for our scheme. The examples with the strongest security properties are the constructions of Furukawa and Sako [22] and Neff [34]. Both of these employ El Gamal as the underlying cryptosystem, i.e., an input ciphertext $E_i = (\alpha, \beta) = (my^k, g^k)$ for some public key y and published generator g . Security in these constructions is reducible to the Decision Diffie-Hellman assumption and a random-oracle assumption on a hash function. We also note that the security of these and most other mix network constructions relies on a second input $\mathcal{P} = \{P_1, P_2, \dots, P_d\}$, where P_i is an NIZK proof of knowledge of the plaintext for E_i . This serves the purpose of rendering the cryptosystem chosen-ciphertext-attack secure while still permitting re-encryption.