SoK: Surveying definitions of coercion resistance

Thomas Haines[1] and Ben Smyth[2]

[1]Mathematics Department, Norwegian University of Science and Technology, Trondheim, Norway
[2]Interdisciplinary Centre for Security, Reliability and Trust, University of Luxembourg, Esch-sur-Alzette, Luxembourg

Abstract—**We explore formal definitions of coercion resistance (WPES'05, FC'09, CRYPTO'10, and CSF'10), conceived to capture the strongest privacy notions achievable by voting systems. We find all but one is unsuitable, demonstrating difficulties faced by our community in formalising this property and raising questions over the security of schemes striving to deliver coercion resistance. We find the remaining definition to be reliant on burdensome combinatorial analysis, prohibiting immediate application. We propose a variant that simplifies application. We also patch an unsuitable definition and introduce sufficient conditions to simplify proofs.**

I. Introduction

Coercion resistance is the strongest notion of privacy a voting system can deliver. It asserts that no voter can prove they followed a coercer's instructions; ensuring voters can evade coercion and vote freely. The notion is rather intricate; definitions have proven elusive: Intuitively, no voting system ought to be coercion-resistant without anonymous channels. when all ballot-collection channels are controlled. Indeed, a coerced voter instructed to abstain can prove compliance, when channels are controlled. Accordingly, we should expect surveyed definitions (and systems proven to satisfy them) to make assumptions about ballot-collection channels. Alternatively, definitions may consider weaker forms of coercion resistance, without protection against forced abstention attacks.

Coercion resistance was introduced by Okamoto [1] and first formalised by Juels, Catalano & Jakobsson [2], [3], [4], with a plethora of coercion-resistant voting systems now in existence [5], [6], [7], [8], [9], [10], [11], [12]. Coercion resistance strengthens receipt-freeness by considering an adversary that instructs a voter – possibly with instructions that deviate from the prescribed ballot casting procedure – rather than merely asking the voter for proof [13], [14], [15], [16]. In turn, receipt-freeness strengthens ballot secrecy, wherein the adversary's capabilities are limited to controlling ballot collection [17]. In its strongest form, coercion resistance includes protection against forced abstention attacks, whereby a coercer instructs a voter to abstain, yet the voter is able to evade coercion and vote freely (assuming the coercer does not control all ballot-collection channels).

Following the first definition of coercion resistance by Juels, Catalano & Jakobsson, further definitions have been proposed by Gardner, Garera & Rubin [18], Unruh & Müller-Quade [19], and Küsters, Truderung & Vogt [20], [21]. We will explore each of those definitions in the context of syntax by Juels, Catalano & Jakobsson, which is common to all definitions. Indeed, the definition by Küsters, Truderung & Vogt is stated independently of any particular syntax and the definition by Gardner, Garera & Rubin is largely syntax independent, hence, those definitions can be considered in the narrower context of syntax by Juels, Catalano & Jakobsson. The definition by Unruh & Müller-Quade is stated in terms of a particular syntax (but to a lesser extent than the definition by Juels, Catalano & Jakobsson) and we cast their definition into the context of syntax by Juels, Catalano & Jakobsson. Using a common syntax simplifies our exploration and facilitates comparisons between definitions.

We limit ourselves to voting systems that centralise tallying and output the number of votes for each candidate. The first limitation implies that trust is required, since tallying an individual voter's ballot will reveal their vote, violating privacy [22]. By comparison, such violations are impossible for voting systems that distribute tallying (assuming at least one honest participant) [23], [24], [25], [26], [27], [28]. Albeit, at the cost of complexity, scalability, and understandability. The second limitation is a functional requirement of many nations, which comes at a privacy cost. Indeed, revealing the number of votes for each candidate leaks more information than, for instance, revealing only the winning candidate [29], [30], [31].

We remark that privacy and verifiability are considered independently: Definitions of coercion resistance do not include voters checking whether their ballots are collected (individual verifiability), voters and other stakeholders checking whether votes expressed in ballots are counted (universal verifiability), nor voters checking whether ballots correctly express their votes (cast-as-intended) [32], [33], [34], [35], [14], [36], [37], [38]. These aspects may have implications for ballot secrecy [39], receipt-freeness, and coercion resistance, and extending definitions to include these notions of verifiability may be an interesting arena for future research.

a) Contribution and structure.: We critic definitions of coercion resistance by Juels, Catalano & Jakobsson (§III), Gardner, Garera & Rubin (§IV), Unruh & Müller-Quade (§V), and Küsters, Truderung & Vogt (§VI). Discovering the first to be too strong, the last burdensome to apply, and the others too weak:

Juels, Catalano & Jakobsson's definition does not appear to be satisfiable. At least, not by the voting system due to the definition's authors. We make a small change, which fills a hole in the system's security proof.

Gardner, Garera & Rubin's definition does not consider tallying; privacy violations arising from tallying go unnoticed.

Unruh & Müller-Quade's definition does not consider voters giving up coins; privacy violations go unnoticed.

Küsters, Truderung & Vogt's definition requires analysts to conduct a burdensome combinatorial analysis to establish a measure of coercion resistance that can be expected for a particular context. (Such analysis is only tangentially related to the security of voting systems and, as such, should surely be in the remit of the definition's authors, rather than analysts.) The authors do provide a minimal measure of coercion resistance, but we show that measure is unsuitable for the analysis of the seminal voting system by Juels, Catalano & Jakobsson and its Civitas variant.

Beyond surveying definitions of coercion resistance, Section VII proceeds as follows: We propose a variant of the definition by Küsters, Truderung & Vogt, which eliminates combinatorial analysis to simplify application. Compare that variant to the (patched) definition by Juels, Catalano & Jakobsson, finding the latter to be strictly stronger (when some generalities are ignored). And introduce sufficient conditions for our patched definition, essentially reducing the burden of proof to a checklist. Section VIII discusses the effect of a malicious bulletin board on coercion resistance. The remaining sections introduce syntax (§II) and present a brief conclusion (§IX), Sidebar 1 introduces games and standard notation. To ensure academic honesty, footnotes are scattered throughout, to clarify subtle details and assumptions, highlight and justify minor discrepancies between the original definitions and our presentation, and to add additional thoughts.

Based on our analysis, we suggest that the patched variant of Juels, Catalano & Jakobsson should be used where applicable. In other settings the definition by Küsters, Truderung & Vogt should be used; however, care must be taken over the choice of parameters and particularly the suitability of the δ-bound. More information on our recommendations can be found in the conclusion.

II. Election scheme syntax

We will consider definitions of coercion resistance in the context of syntax by Juels, Catalano & Jakobsson, more precisely, we adopt the variant of their syntax by Smyth, Frink & Clarkson [35], which clarifies several details. The syntax captures a class of voting systems that consist of the following four steps. First, a tallier generates a key pair. Secondly, a registrar generates credentials for voters. Thirdly, each voter constructs and casts a ballot for their vote. These ballots are recorded on a bulletin board. Finally, the tallier tallies the recorded ballots and announces the outcome as a frequency distribution of votes. (The chosen representative is derived from this distribution, which suffices for both first-past-the-post and ranked-choice voting systems.)

Sidebar 1 Preliminaries: Notation and games [17]

We let $A(x_1, \dots, x_n; r)$ denote the output of probabilistic algorithm A on inputs $x_1, \dots, x_n$ and coins r, and we let $A(x_1, \dots, x_n)$ denote $A(x_1, \dots, x_n; r)$, where coins r are chosen uniformly at random from the coin space of algorithm A. Moreover, we let $x \leftarrow T$ denote assignment of T to x, and $x \leftarrow_R S$ denote assignment to x of an element chosen uniformly at random from set S, similarly, $x \leftarrow_R D$ denotes assignment to x of an element chosen according to the distribution D. Furthermore, we let $x[i]$ denote component i of vector x and let $|x|$ denote the length of vector x. Finally, we write $(x_1, \dots, x_{|T|}) \leftarrow T$ for $x \leftarrow T; x_1 \leftarrow x[1]; \dots; x_{|T|} \leftarrow x[|T|]$, when T is a vector, and $x, x' \leftarrow_R S$ for $x \leftarrow_R S; x' \leftarrow_R S$.

A game is a probabilistic algorithm that outputs a boolean. Using our notation, we can formulate the following game, denoted $\mathsf{Exp}(H, S, \mathcal{A})$, which tasks an adversary $\mathcal{A}$ to distinguish between a function H and a simulator S: $m \leftarrow \mathcal{A}(); \beta \leftarrow_R \{0,1\};$ **if** $\beta = 0$ **then** $x \leftarrow H(m)$ **else** $x \leftarrow S(m)$; $g \leftarrow \mathcal{A}(x);$ **return** $g = \beta$. Adversaries are *stateful*, i.e., information persists across invocations of an adversary in a game. In particular, adversaries can access earlier assignments. For instance, the adversary's second instantiation in game Exp has access to any assignments made during its first instantiation. An adversary *wins* a game by causing it to output true ($\top$) and the adversary's *success* in a game $\mathsf{Exp}(\cdot)$, denoted $\mathsf{Succ}(\mathsf{Exp}(\cdot))$, is the probability that the adversary wins, that is, $\mathsf{Succ}(\mathsf{Exp}(\cdot)) = \Pr[\mathsf{Exp}(\cdot) = \top]$. We focus on computational security, rather than information-theoretic security, and tolerate adversary wins in non-polynomial time or with negligible probability, since such wins are infeasible in practice. Game Exp captures a single interaction between the challenger and the adversary. We can extend games with oracles to capture arbitrarily many interactions. For instance, we can formulate a strengthening of Exp as follows: $\beta \leftarrow_R \{0,1\}; g \leftarrow \mathcal{A}^{\mathcal{O}}(x);$ **return** $g = \beta$, where $\mathcal{A}^{\mathcal{O}}$ denotes $\mathcal{A}$'s access to oracle $\mathcal{O}$ and $\mathcal{O}(m)$ computes **if** $\beta = 0$ **then** $x \leftarrow H(m)$ **else** $x \leftarrow S(m)$; **return** x. Oracles may access game parameters such as bit β.

Definition 1 (Election scheme [35])**.** *An* election scheme *is a tuple of probabilistic polynomial-time algorithms* $(\mathsf{Setup}, \mathsf{Register}, \mathsf{Vote}, \mathsf{Tally})$ *such that:*[1]

Setup, *denoted* $(pk, sk, mb, mc) \leftarrow \mathsf{Setup}(\kappa)$*, is run by the tallier. The algorithm takes a security parameter* κ *as input and outputs a key pair* pk, sk*, a maximum number of ballots* mb*, and a maximum number of candidates* mc*.*

$\mathsf{Register}$, *denoted* $(pd, sd) \leftarrow \mathsf{Register}(pk, \kappa)$*, is run by the registrar. The algorithm takes a public key* pk *and security parameter* κ *as input and outputs a public credential* pd *and a private credential* sd*.*

Vote, *denoted* $b \leftarrow \mathsf{Vote}(sd, pk, v, nc, \kappa)$*, is run by voters. The algorithm takes as input a private credential* sd*, a*

[1] The syntax bounds the number of ballots mb, respectively candidates mc, to broaden the correctness definition's scope. The syntax represents votes as integers, rather than alphanumeric strings, for brevity. Finally, the syntax employs sets, rather than multisets or lists, to preclude construction (and consequently modelling) of schemes vulnerable to attacks that arise due to duplicate ballots.

public key pk, a voter's vote v, some number of candidates nc, and a security parameter κ. The vote should be selected from: a sequence $1, \ldots, nc$ of candidates; candidate ϕ, representing abstention; and (optionally) candidate λ, representing casting a vote with an invalid credential. The algorithm outputs a ballot b or error symbol $\perp$.

Tally*, denoted* $(\mathfrak{v}, pf) \leftarrow \mathsf{Tally}(sk, \mathfrak{bb}, L, nc, \kappa)$*, is run by the tallier. The algorithm takes as input a private key sk, a bulletin board $\mathfrak{bb}$, an electoral roll L, some number of candidates nc, and a security parameter κ, where $\mathfrak{bb}$ and L are sets. And outputs an election outcome $\mathfrak{v}$ and a non-interactive tallying proof pf demonstrating that the outcome corresponds to votes expressed in ballots on the bulletin board. The election outcome $\mathfrak{v}$ should be a vector of length nc such that $\mathfrak{v}[v]$ indicates the number of votes for candidate v.*

Election schemes must satisfy correctness*: there exists a negligible function* negl*, such that for all security parameters κ, integers nb and nc, and votes $v_1, \ldots, v_{nb} \in \{1, \ldots, nc\}$, it holds that, given a zero-filled vector $\mathfrak{v}$ of length nc, we have:* $\Pr[(pk, sk, mb, mc) \leftarrow \mathsf{Setup}(\kappa);$ **for** $1 \leq i \leq nb$ **do** $\{ (pd_i, sd_i) \leftarrow \mathsf{Register}(pk, \kappa); b_i \leftarrow \mathsf{Vote}(sd_i, pk, v_i, nc, \kappa); \mathfrak{v}[v_i] \leftarrow \mathfrak{v}[v_i] + 1; \}$ $(\mathfrak{v}', pf) \leftarrow \mathsf{Tally}(sk, \{b_1, \ldots, b_{nb}\}, \{pd_1, \ldots, pd_{nb}\}, nc, \kappa) : nb \leq mb \wedge nc \leq mc \Rightarrow \mathfrak{v} = \mathfrak{v}'] > 1 - \mathsf{negl}(\kappa)$.

We omit algorithm Verify above, because we focus on privacy, rather than verifiability, and we want to avoid defining an algorithm which isn't used anywhere in this paper.

a) Syntax scope.: Syntax by Juels, Catalano & Jakobsson captures voting systems with centralised tallying that output the number of votes for each candidate. The syntax trivially generalises to distributed tallying.[2] The syntax similarly generalises to distributed registration.

Without private credentials, algorithm Vote cannot be used to achieve receipt-freeness nor coercion resistance, hence, private credentials are necessary without distributing ballot construction [17, §7]. By comparison, distribution gives way to straightforward constructions for receipt-free election schemes from schemes satisfying ballot secrecy, whereas coercion-resistance requires private input such as credentials [17, §7]. Whether such distribution is an interesting line of enquiry is an open question.

The syntax has proven useful in analysing [35], [40] voting systems by Juels, Catalano & Jakobsson [4] and by Cortier *et al* [34], and in guiding construction of the Athena voting system [12]. Moreover, without private credentials, the syntax has proven useful in analysing [17], [41], [42] the Helios [43] and Helios Mixnet [44], [45], [46] systems. The syntax also proved useful in establishing a generic construction for secure, verifiable auction systems from voting systems [47]. Nonetheless, the syntax is incomplete: Not every voting system can be modelled.

III. Juels, Catalano & Jakobsson [2], [3], [4]

Juels, Catalano & Jakobsson formulate a simulation-based definition of coercion resistance, wherein coercion resistance is derived from indistinguishably of real and fake private credentials. The definition has evolved over time,[3] and we consider the most recent [4], with the following notable patch: We input the bulletin board length to the adversary (game JCJ-\$, Line 27). Without this patch the coercion-resistance proof given by Juels, Catalano & Jakobsson (for their voting system) does not hold – their system does not satisfy the definition. Their proof does not elaborate on how to simulate the bulletin board to the adversary. Although ciphertexts on the bulletin board can be simulated, the number of ciphertexts must be known. The patch is noteworthy, since it is required for the aforementioned proof and future proofs for other voting systems (including Civitas [5]), and since it alters our intuition – without our patch, the definition intuitively asserts the adversary has negligible advantage over what is revealed by the election outcome, whereas, the patched definition intuitively asserts that the adversary has negligible advantage over what is revealed by the tally *and* the bulletin board length.

Definition 2. *Let* $\Gamma = (\mathsf{Setup}, \mathsf{Register}, \mathsf{Vote}, \mathsf{Tally})$ *be an election scheme, na, nc and nv be integers, and D be a distribution over $\{1, \ldots, nc, \phi, \lambda\}^{nv}$. We say Γ satisfies* JCJ *with respect to na, nc, nv, D, if for all probabilistic polynomial-time adversaries $\mathcal{A}$, there exists probabilistic polynomial-time algorithms $\mathcal{B}$, algorithm* $\mathsf{fakecred}$*, and a negligible function* negl*, such that for all security parameters κ, we have* $|\mathsf{Succ}(\mathsf{JCJ}(\Gamma, \mathcal{A}, na, nc, nv, \mathsf{fakecred}, D, \kappa)) - \mathsf{Succ}(\mathsf{JCJ\text{-}\$}(\Gamma, \mathcal{B}, na, nc, nv, D, \kappa))| \leq \mathsf{negl}(\kappa)$*, where games* JCJ *and* $\mathsf{JCJ\text{-}\$}$ *are defined in Figure 1,[4] and algorithm* $\mathsf{fakecred}$ *takes a public key, a public credential and a private credential as input, and outputs a (fake) private credential.*

The definition captures forced abstention attacks, hence, some intuitively coercion-resistant election schemes (without protection against forced abstention attacks) cannot satisfy their definition. Indeed, the voting system by Juels, Catalano & Jakobsson can only satisfy their definition assuming registration proceeds without adversarial interference, since an adversary that can block registration can trivially force abstention. It remains an open question as to whether coercion resistance should be defined in its strongest form (with protection against forced abstention attacks), or whether a weaker form is tolerable.

[2] Generalisation to distributed tallying can, for instance, be achieved as follows: algorithm Setup can be run by each tallier to compute key shares and those shares can be combined to derive a public key $pk = (pk_1, \ldots, pk_{|pk|})$; algorithm Tally can be tweaked to compute partial tallying proofs, e.g., $pf_i \leftarrow \mathsf{Tally}(sk, \mathfrak{bb}, L, nc, \kappa)$; and algorithm Verify can compute outcomes from partial proofs, e.g., $\mathfrak{v} \leftarrow \mathsf{Verify}(pk, \mathfrak{bb}, L, nc, pf, \kappa)$, where $pf = (pf_1, \ldots, pf_{|pf|})$.

[3] Preprints prior to [2] consider a passive adversary in game JCJ-\$ and do not consider ideal tallying. Also, there is variance in whether the adversary gets to choose the target vote, voter, or both. Further variants are proposed by others, e.g., [48], [49].

[4] Distinctions between games JCJ and JCJ-\$, and subsequent games, are highlighted in yellow.

Fig. 1 Games JCJ and JCJ-$

```
$\mathsf{JCJ}(\Gamma, \mathcal{A}, na, nc, nv, \mathsf{fakecred}, D, \kappa) =$
$(pk, sk, mb, mc) \leftarrow \mathsf{Setup}(\kappa)$;
$V \leftarrow \mathcal{A}(pk, \kappa)$;
for $1 \leq i \leq nv$ do
    $(pd_i, sd_i) \leftarrow \mathsf{Register}(pk, \kappa)$;
$L \leftarrow \{pd_1, \ldots, pd_{nv}\}$;
$M \leftarrow \{(i, sd_i) \mid i \in V \wedge 1 \leq i \leq nv\}$;
$(j, v) \leftarrow \mathcal{A}(M, L)$;
if $|V| \neq na \vee j \notin \{1, \ldots, nv\} \setminus V \vee v \notin \{1, \ldots, nc\} \cup \{\phi\}$ then
    return $0$;
$\mathfrak{bb}_1 \leftarrow \emptyset$;
$\beta \leftarrow_R \{0, 1\}$;
if $\beta = 0$ then
    if $v \neq \phi$ then
        $b \leftarrow \mathsf{Vote}(sd_j, pk, v, nc, \kappa)$;
        $\mathfrak{bb}_1 \leftarrow \mathfrak{bb}_1 \cup \{b\}$;
    $sd' \leftarrow \mathsf{fakecred}(pk, pd_j, sd_j)$;
else
    $sd' \leftarrow sd_j$;
for $i \in \{1, \ldots, nv\} \setminus (\{j\} \cup V)$ do
    $v \leftarrow_R D$;
    if $v \neq \phi$ then
        if $v = \lambda$ then
            $v \leftarrow_R \{1, \ldots, nc\}$;
            $sd_i \leftarrow \mathsf{fakecred}(pk, pd_i, sd_i)$;
        $b \leftarrow \mathsf{Vote}(sd_i, pk, v, nc, \kappa)$;
        $\mathfrak{bb}_1 \leftarrow \mathfrak{bb}_1 \cup \{b\}$;
$\mathfrak{bb}_2 \leftarrow \mathcal{A}(sd', \mathfrak{bb}_1)$;
$(\mathfrak{v}, pf) \leftarrow \mathsf{Tally}(sk, \mathfrak{bb}_1 \cup \mathfrak{bb}_2, L, nc, \kappa)$;
$g \leftarrow \mathcal{A}(\mathfrak{v}, pf)$;
return $\beta = g$;
```

```
$\mathsf{JCJ\text{-}\$}(\Gamma, \mathcal{B}, na, nc, nv, D, \kappa) =$
$(pk, sk, mb, mc) \leftarrow \mathsf{Setup}(\kappa)$;
$V \leftarrow \mathcal{B}(pk, \kappa)$;
for $1 \leq i \leq nv$ do
    $(pd_i, sd_i) \leftarrow \mathsf{Register}(pk, \kappa)$;
$L \leftarrow \{pd_1, \ldots, pd_{nv}\}$;
$M \leftarrow \{(i, sd_i) \mid i \in V \wedge 1 \leq i \leq nv\}$;
$(j, v) \leftarrow \mathcal{B}(L)$;
if $|V| \neq na \vee j \notin \{1, \ldots, nv\} \setminus V \vee v \notin \{1, \ldots, nc\} \cup \{\phi\}$ then
    return $0$;
$\mathfrak{bb}_1 \leftarrow \emptyset$;
$\beta \leftarrow_R \{0, 1\}$;
if $\beta = 0$ then
    if $v \neq \phi$ then
        $b \leftarrow \mathsf{Vote}(sd_j, pk, v, nc, \kappa)$;
        $\mathfrak{bb}_1 \leftarrow \mathfrak{bb}_1 \cup \{b\}$;
    $sd' \leftarrow sd_j$;
else
    $sd' \leftarrow sd_j$;
for $i \in \{1, \ldots, nv\} \setminus (\{j\} \cup V)$ do
    $v \leftarrow_R D$;
    if $v \neq \phi$ then
        if $v = \lambda$ then
            $v \leftarrow_R \{1, \ldots, nc\}$;
            $sd_i \leftarrow \mathsf{fakecred}(pk, pd_i, sd_i)$;
        $b \leftarrow \mathsf{Vote}(sd_i, pk, v, nc, \kappa)$;
        $\mathfrak{bb}_1 \leftarrow \mathfrak{bb}_1 \cup \{b\}$;
$\mathfrak{bb}_2 \leftarrow \mathcal{B}(sd', M, |\mathfrak{bb}_1|)$;
$\mathfrak{v} \leftarrow \mathsf{Ideal\text{-}Tally}(sk, \mathfrak{bb}_1 \cup \mathfrak{bb}_2, L, nc, \kappa)$;
$g \leftarrow \mathcal{B}(\mathfrak{v})$;
return $\beta = g$;
```

Function Ideal-Tally tallies $\mathfrak{bb}_1$ in the normal way to derive $\mathfrak{v}$ and tallies $\mathfrak{bb}_2 \setminus \mathfrak{bb}_1$ specially: for each $b \in \mathfrak{bb}_2 \setminus \mathfrak{bb}_1$ constructed using private credential $sd \in M$ and vote $v \in \{1, \ldots, nc\}$, compute $\mathfrak{v}[v] \leftarrow \mathfrak{v}[v] + 1$, disregarding any double votes.[a] Moreover, if $\beta = 1$ and there exists ballot $b \in \mathfrak{bb}_2 \setminus \mathfrak{bb}_1$ constructed using private credential sd' and vote $v \in \{1, \ldots, nc\}$, then compute $\mathfrak{v}[v] \leftarrow \mathfrak{v}[v] + 1$.[b]

[a] Juels, Catalano & Jakobsson explicitly specify that function Ideal-Tally should not count any ballot $b \in \mathfrak{bb}_2 \setminus \mathfrak{bb}_1$ constructed using private credential $sd \notin M \setminus \{sd'\}$. They also specify that no double vote should be counted. They do not specifically specify that ballots constructed using private credentials $sd \in M$ should be counted. (Cf. [4, pp49–50].) We believe this was an oversight and we count such ballots.

[b] Juels, Catalano & Jakobsson do not specify which vote should be counted when multiple ballots are constructed using private credential sd'.

Game JCJ models a real-world instance of an election scheme which captures information leakage and game JCJ-$ captures an ideal-world election with more modest information leakage.[5] According to the definition, a scheme satisfies coercion resistance if for any adversary against a real-world instance of the scheme (JCJ) there is an adversary against an ideal-world election (JCJ-$) that can learn an equivalent amount of information. Given that any adversary against an ideal world learns – by definition – minimal information, we can infer that any adversary against a real-world instance of the scheme also learns minimal information. Hence, the definition captures coercion resistance.

[5] Untypically for simulation-based definitions, the definition by Juels, Catalano & Jakobsson includes some game-based aspects: Ideal worlds should capture minimal information leakage, which game JCJ-$ seemingly violates, in particular, not only is an election outcome and bulletin-board length revealed, but also a public key, public credentials, and even some private credentials. Nonetheless, despite appearances to the contrary, we will see (§VII-A) that game JCJ-$ actually leaks minimal information.

The first six lines of games JCJ and JCJ-\$ are identical: The challenger generates a key pair (Line 1),[6] the adversary selects a set of corrupt voters (Line 2),[7] and the challenger generates credentials for all voters (Lines 3 & 4). Next, the adversary chooses a voter to coerce along with the voter's (preferred) vote, with private voter credentials in game JCJ and without credentials in game JCJ-\$ (Line 6 & 7). (Although it is somewhat unnatural for the adversary to specify the voter's vote, this essentially quantifies over all votes.) Lines 8–15 are also identical in both games: the challenger checks that the adversary corrupted (exactly) na voters, chose to coerce a voter that it did not corrupt, and chose a valid vote (Lines 8 & 9), initialises an empty bulletin board (Line 10), and flips a coin (Line 11). If that coin flip produces a zero (Line 12) , then the challenger constructs and casts a ballot on behalf of the coerced voter, except if the voter wants to abstain (Lines 13 –15). Moreover, the challenger constructs a fake private credential to evade coercion in game JCJ (Line 16), whereas the real private credential is used in game JCJ-\$. Otherwise (the coin flip produces a one), the challenger uses the real private credential in both games (Lines 17 & 18). Lines 19–26 are identical in both games: for each non-corrupt, non-coerced voter (Line 19), a vote is sampled (Line 20), and the challenger constructs and casts a ballot for that vote, except when the vote signifies abstention or casting with an invalid credential (Lines 21–26).[8] Next, the adversary constructs a set of ballots (Line 27), which might include ballots constructed using the coerced voter's private credential, the fake credential (only in JCJ), and corrupt voters' private credentials.[9] The challenger initialises a set of public credentials (Line 5) and tallies the ballots (Line 28), using algorithm Tally in JCJ and Ideal-Tally in JCJ-\$.[10] Finally, the adversary is given the election outcome (and proof of correct computation in JCJ) and attempts to determine whether the coin flip resulted in zero or one (Lines 29 & 30).[11] Intuitively, any scheme satisfying JCJ leaks at most a negligible amount of information on the bulletin board over what is revealed by the tally, even if the adversary controls corrupt voters' private credentials and knows the coerced voter's private credential (or a fake credential). In essence, the coercer learns nothing more then the tally (and the number of items on bulletin board) and hence we consider the scheme coercion resistant.

The definition of coercion resistance by Juels, Catalano & Jakobsson leaves analysts with a conundrum: For which parameters should an election scheme be proven secure? Intuitively, the number of corrupt voters (na) and candidates (nc), along with the total number of voters (nv), should all be upper-bound by a polynomial in the security parameter. (Lines 8 & 9 check whether na voters are corrupt and the coerced voter is amongst the honest voters, implicitly checking na is strictly less than nv.) Distributions (D) should be universally quantified over. These details are not stated by Juels, Catalano & Jakobsson, but are necessary to analysts applying their definition. To make details explicit, we say an election scheme satisfies JCJ, if for all integers na, nc, and nv, and all efficiently sampleable distributions D over $\{1, \ldots, nc, \phi, \lambda\}^{nv}$, the election scheme satisfies JCJ with respect to na, nc, nv, D, when those integers are upper-bound by a polynomial in the security parameter.

IV. Gardner, Garera & Rubin [18]

Gardner, Garera & Rubin formulate a game-based definition of coercion resistance that challenges an adversary to distinguish a ballot for the adversary's preferred vote v_0 from a ballot for the voter's preferred vote v_1, wherein ballots are constructed using coins provided by the adversary and the latter ballot is constructed using inputs that may have been modified for the purposes of evading coercion.

Definition 3. *We say an election scheme* $\Gamma = (\mathsf{Setup}, \mathsf{Register}, \mathsf{Vote}, \mathsf{Tally})$ *satisfies* GGR*, if there exists a probabilistic polynomial-time algorithm* evade *such that for all probabilistic polynomial-time adversaries* $\mathcal{A}$*, there exists a negligible function* negl *and for all security parameters* κ*, we have* $\mathsf{Succ}(\mathsf{GGR}(\Gamma, \mathcal{A}, \mathsf{evade}, \kappa)) < \frac{1}{2} + \mathsf{negl}(\kappa)$*, where game* GGR *is defined in Figure 2, and algorithm* evade *takes the inputs to algorithm* Vote *and some coins as input, and outputs a private credential, a public key, an integer, and some coins.*

Game GGR proceeds as follows: The challenger generates a key pair and a credential (Lines 1 & 2), and initialises an empty set of coins (Line 3). Next, the adversary chooses their preferred vote, the voter's preferred vote, some number of candidates, and some coins (Line 4).[12,13] The challenger adds

[6] Key generation is implicit, rather than explicit, in the original presentation by Juels, Catalano & Jakobsson.

[7] The adversary inputs a set of corrupt voters' private credentials in the original presentation, whereas we consider a set of pairs to maintain the relationship between indexes in set V and the corresponding private credentials.

[8] Juels, Catalano & Jakobsson do not define casting with an invalid credential. By comparison, we formalise this aspect as casting a ballot using a vote selected uniformly at random (Line 23) and using a fake private credential (Line 24).

[9] Juels, Catalano & Jakobsson seemingly suggest that the adversary in JCJ-\$ learns all of the private credentials [4, §3, penultimate paragraph], but only the coerced voter's private credential and the corrupt voters' private credentials are provided as input to the adversary in their presentation of game JCJ-\$.

[10] The presentation by Juels, Catalano & Jakobsson computes a tallying proof in both JCJ and JCJ-\$, yet their presentation of function Ideal-Tally does not appear to output a tallying proof. Moreover, the tallying proof is not used. We believe this is a typo and we omit the tallying proof from JCJ-\$.

[11] Beyond the aforementioned differences between our presentation of the coercion resistance definition by Juels, Catalano & Jakobsson and the original, we also make some minor changes in notation.

[12] Gardner, Garera & Rubin do not specify the adversary's precise inputs used (by the challenger and oracle) to construct ballots. We presume that the coerced voter controls some of the inputs to algorithm Vote, in particular, we presume that voter controls their private credential and public information, including the public key and the security parameter.

[13] Gardner, Garera & Rubin consider voting systems in which ballots can contain some information that is only available to the verifier, whereas such ballots are excluded by our election scheme syntax. Hence, we omit the verifier from our formalisation.

Fig. 2 Game GGR

```
$\mathsf{GGR}(\Gamma, \mathcal{A}, \mathsf{evade}, \kappa) =$
1  $(pk, sk, mb, mc) \leftarrow \mathsf{Setup}(\kappa)$;
2  $(pd, sd) \leftarrow \mathsf{Register}(pk, \kappa)$;
3  $\mathfrak{r} \leftarrow \emptyset$;
4  $(v_0, v_1, nc, s) \leftarrow \mathcal{A}^{\mathcal{O}}(pk, pd)$;
5  $\mathfrak{r} \leftarrow \mathfrak{r} \cup \{s\}$;
6  $\beta \leftarrow_R \{0, 1\}$;
7  **if** $\beta = 0$ **then**
8      $(sd', pk', nc', \kappa') \leftarrow \mathsf{evade}(sd, pk, v_0, nc, \kappa, s, v_1)$;
9      $b \leftarrow \mathsf{Vote}(sd', pk', v_1, nc', \kappa'; s)$;
10 **else**
11     $b \leftarrow \mathsf{Vote}(sd, pk, v_0, nc, \kappa; s)$;
12 $g \leftarrow \mathcal{A}^{\mathcal{O}}(b)$;
13 **return** $g = \beta \wedge s \notin \mathfrak{r}$;
```

Oracle $\mathcal{O}$ is defined such that $\mathcal{O}(v, nc, r)$ computes $b \leftarrow \mathsf{Vote}(sd, pk, v, nc, \kappa; r); \mathfrak{r} \leftarrow \mathfrak{r} \cup \{r\}$ and outputs b.

those coins to the set of coins (Line 5).[14] The challenger flips a coin (Line 6), constructs a ballot on behalf of the coerced voter using inputs that may have been modified if the coin flip produces zero (Lines 7–9) and the adversary's input otherwise (Lines 10 & 11). Finally, the adversary is given the constructed ballot and attempts to determine whether the coin flip resulted in zero or one (Lines 12 & 13).

The definition is too weak. In particular, tallying may leak information that can violate privacy. For instance, as an extreme example, suppose tallying leaks the tallier's private key, thereby enabling tallying of individual ballots to reveal each voter's vote. Perhaps Gardner, Garera & Rubin intended to capture coercion resistance of ballot casting, rather than coercion resistance of the entire voting system. Indeed, they write, "We introduce a new definition of a coercion resistant *vote casting* protocol" (emphasis added). Yet, they go on to write, "we are able to address coercion enabled by examination of the protocol's final output," which suggests that the final output – surely including the election outcome and tallying proof – should have been considered. Moreover, they are critical of the definition by Moran & Naor [13] because "[it] focuses on the adversary's view of a voter's interactions with a machine and allows privacy leaks in the final output in the protocol," which seemingly suggests their definition should detect such leaks. We shared our findings with Garera & Rubin (email, 6 Jul 2018), but have not received a response.

V. Unruh & Müller-Quade [19]

Unruh & Müller-Quade formulate a game-based definition of coercion resistance that challenges an adversary to distinguish between a voter following a coercer's instructions to cast a vote v^* preferred by the adversary and the voter deviating from those instructions to cast a vote v preferred by the voter (whilst producing evidence that the instructions were followed), with probability greater than the adversary's ability to distinguish the election outcomes produced in each setting. The definition requires a counter-strategy that deviates from the adversary's instructions, which is captured using an algorithm C, which must produce a ballot for v.

Definition 4. *Let* $\Gamma = (\mathsf{Setup}, \mathsf{Register}, \mathsf{Vote}, \mathsf{Tally})$ *be an election scheme,*[15] *and* nc *and* nv *be integers. We define games* UM *and* $\mathsf{UM\text{-}\$}$ *in Figure 3, and say* Γ *satisfies* UM *with respect to* nc, nv*, if there exists a probabilistic polynomial-time algorithm* C *such that for all probabilistic polynomial-time adversaries* $\mathcal{A}$*, efficiently sampleable distributions* D *over* $\{1, \ldots, nc, \phi\}^{nv-1}$*, votes* $v, v^* \in \{1, \ldots, nc, \phi\}$*, and integers* $j \in \{1, \ldots, nv\}$*, there exists a negligible function* negl *and the following conditions hold for all security parameters* κ*: First, the election outcome* $\mathfrak{v}$ *computed in game* $\mathsf{UM}(\Gamma, \mathcal{A}, nc, nv, v, v^*, D, j, \kappa)$ *is computationally indistinguishable from the election outcome computed from the votes sampled by the oracle in game* $\mathsf{UM\text{-}\$}(\Gamma, C, \mathcal{A}, nc, nv, v, v^*, D, j, \kappa)$ *and the vote* v*. Secondly,* $|\mathsf{UM}(\Gamma, \mathcal{A}, nc, nv, v, v^*, D, j, \kappa) - \mathsf{UM\text{-}\$}(\Gamma, C, \mathcal{A}, nc, nv, v, v^*, D, j, \kappa)| \leq \max_{v^* \in \{1,\ldots,nc,\phi\}} \Delta(D_v, D_{v^*}) + \mathsf{negl}(\kappa)$*, where* Δ *denotes statistical distance and* D_v*, respectively* D_{v^*}*, is the distribution over* $\{1, \ldots, nc, \phi\}^{nv}$ *that chooses* $nv - 1$ *votes according to* D *and uses* v*, respectively* v^**, for the* nvth *vote.*

The first condition captures the counter-strategy, represented as algorithm C, casting a ballot for v, and the second captures the adversary's inability to distinguish the jth voter following instructions (UM) and deviating from them using algorithm C ($\mathsf{UM\text{-}\$}$).

The games are identical except for Line 7: The challenger generates a key pair and credentials (Lines 1–4),[16] and initialises an empty bulletin board (Line 5). Moreover, if adversarial preferred vote v^* does not represent abstention, then the challenger constructs a ballot for that vote in game UM and constructs a ballot using the counter-strategy in game $\mathsf{UM\text{-}\$}$,[17] the constructed ballot is added to the bulletin board in both games (Lines 6–8). Next, the adversary instructs the oracle to construct ballots on behalf of non-coerced voters and add those ballots to the bulletin board (Line 9).[18] Finally, the challenger tallies the bulletin board and the adversary is given

[14] The presentation by Gardner, Garera & Rubin does not seem to include the adversary's coins in the set of coins, yet this is necessary to ensure the definition is satisfiable. We believe this is an omission and we include the coins in the set.

[15] Unruh & Müller-Quade consider a tallying algorithm that only outputs an election outcome, hence, we consider a variant of algorithm Tally, denoted $\mathfrak{v} \leftarrow \mathsf{Tally}(sk, \mathfrak{bb}, L, nc, \kappa)$, in this section.

[16] Key generation is (seemingly) implicit, rather than explicit, in the original presentation by Unruh & Müller-Quade. (Alternatively, it is possible that Unruh & Müller-Quade only intended their results to be applicable to decentralised voting systems, rather than centralised ones.)

[17] Unruh & Müller-Quade suggest that the adversary *instructs* the coerced voter, but do not specify further details. Here, we assume the adversary's ability to instruct is limited to choosing the adversarial preferred vote, which we model by universal quantification over all votes.

[18] The oracle is not explicitly defined by Unruh & Müller-Quade, but it is useful to capture their requirements that the adversary controls the start of tallying and that the voters other than the coerced voter all following the protocol.

Fig. 3 Games UM and UM-\$

```
$\mathsf{UM}(\Gamma, \mathcal{A}, nc, nv, v, v^*, D, j, \kappa) =$
1  $(pk, sk, mb, mc) \leftarrow \mathsf{Setup}(\kappa)$;
2  for $1 \leq i \leq nv$ do
3      $(pd_i, sd_i) \leftarrow \mathsf{Register}(pk, \kappa)$;
4  $L \leftarrow \{pd_1, \ldots, pd_{nv}\}$;
5  $\mathfrak{bb} \leftarrow \emptyset$;
6  if $v^* \neq \phi$ then
7      $b \leftarrow \mathsf{Vote}(sd_j, pk, v^*, nc, \kappa)$;
8      $\mathfrak{bb} \leftarrow \mathfrak{bb} \cup \{b\}$;
9  $x \leftarrow \mathcal{A}^{\mathcal{O}}(pk, v^*, L, \kappa)$;
10 $\mathfrak{v} \leftarrow \mathsf{Tally}(sk, \mathfrak{bb}, L, nc, \kappa)$;
11 $g \leftarrow \mathcal{A}(\mathfrak{v})$;
12 return $g$;
```

```
$\mathsf{UM\text{-}\$}(\Gamma, C, \mathcal{A}, nc, nv, v, v^*, D, j, \kappa) =$
1  $(pk, sk, mb, mc) \leftarrow \mathsf{Setup}(\kappa)$;
2  for $1 \leq i \leq nv$ do
3      $(pd_i, sd_i) \leftarrow \mathsf{Register}(pk, \kappa)$;
4  $L \leftarrow \{pd_1, \ldots, pd_{nv}\}$;
5  $\mathfrak{bb} \leftarrow \emptyset$;
6  if $v^* \neq \phi$ then
7      $b \leftarrow C(sd_j, pk, v, v^*, nc, \kappa)$;
8      $\mathfrak{bb} \leftarrow \mathfrak{bb} \cup \{b\}$;
9  $x \leftarrow \mathcal{A}^{\mathcal{O}}(pk, v^*, L, \kappa)$;
10 $\mathfrak{v} \leftarrow \mathsf{Tally}(sk, \mathfrak{bb}, L, nc, \kappa)$;
11 $g \leftarrow \mathcal{A}(\mathfrak{v})$;
12 return $g$;
```

Oracle $\mathcal{O}$ is defined such that $\mathcal{O}(i)$ computes $v \leftarrow_R D$; **if** $v \neq \phi$ **then** $b \leftarrow \mathsf{Vote}(sd_i, pk, v, nc, \kappa)$; $\mathfrak{bb} \leftarrow \mathfrak{bb} \cup \{b\}$, where $i \in \{1, \ldots, nv\} \setminus \{j\}$. Moreover, we require that $1 \leq j \leq nv$ and that oracle $\mathcal{O}$ is called with integer i at most once.

the election outcome and challenged to determine whether the ballot constructed by the adversary was for v or v^* (Lines 10–12).

Definition UM is similar to the receipt-freeness definition by Delaune, Kremer & Ryan [50], which Smyth casts from the symbolic model to the computational model of cryptography, resulting in a pair of games Receipt-Freeness-A and Receipt-Freeness-B [17, §7].[19] The former captures the adversary's inability to distinguish between voters following instructions and deviating from them using a counter-strategy (which is similar to the second condition of definition UM), and the latter over-approximates the counter-strategy casting a ballot for vote preferred candidates (which is similar to the first condition of definition UM). Hence, definition UM seemingly captures receipt-freeness rather than coercion resistance. But, upon closer inspection, aspects of Smyth's definition are omitted from the definition by Unruh & Müller-Quade. Indeed, Smyth permits the adversary to learn coins used to construct ballots, whereas Unruh & Müller-Quade do not. Consequently, definition UM does not capture receipt-freeness, thus, coercion resistance is not captured either.[20] We shared our findings with Unruh (email, 29 Jun 2019), who acknowledged some problems.

VI. Küsters, Truderung & Vogt [20], [21]

Küsters, Truderung & Vogt formulate a game-based definition of coercion resistance that challenges an adversary to distinguish between a voter following a coercer's instructions and the voter deviating from those instructions to cast their preferred vote using a counter-strategy. We stress that our presentation casts their definition into the syntax of election schemes (§II), whereas the original definition considers a broader class of voting systems. Moreover, we consider the strongest form of their definition, which is intended to provide protection against forced abstention attacks.

The original definition considers concurrent interaction between talliers, registrars, and voters, whereas we sequentialise actions: First, a tallier computes a key pair; secondly, a registrar computes voter credentials; thirdly, voters vote; and, finally, the tallier computes an election outcome. In the context of election schemes, sequentialising the first, second, and last actions does not lose generality, because the second and last actions are dependent on previous actions. Ordering voters does not lose generality either, given that ballots are recorded in an (unordered) set and the adversary is assumed not to observe individual ballots being cast.

To broaden applicability of the definition in the context of our syntax, we introduce an algorithm O (parameterised by public information) to append items to the bulletin board, thereby capturing additional tallier, registrar, and voter behavior, such as adding additional ciphertexts to the bulletin board, for instance. (The original definition remains more general.)

Definition 5. *Let* $\Gamma = (\mathsf{Setup}, \mathsf{Register}, \mathsf{Vote}, \mathsf{Tally})$ *be an election scheme,* C *be a (stateful) algorithm,* O *be an algorithm,* na*,* nc *and* nv *be integers,* D *be an distribution over* $\{1, \ldots, nc, \phi\}$*, and* v *be a vote in* $\{1, \ldots, nc, \phi\}$*. We say* Γ *satisfies* δ*-*KTV *with respect to* na, nc, nv, v, C, O, D*, if for all probabilistic polynomial-time algorithms* $\mathcal{A}$ *there exists a negligible function* negl *such that for all security parameters* κ*, we have* $\mathsf{Succ}(\mathsf{KTV}(\Gamma, \mathcal{A}, na, nc, nv, v, C, O, D, \kappa)) - \mathsf{Succ}(\mathsf{KTV\text{-}\$}(\Gamma, \mathcal{A}, na, nc, nv, v, C, O, D, \kappa)) \leq \delta + \mathsf{negl}(\kappa)$*, where games* KTV *and* $\mathsf{KTV\text{-}\$}$ *are defined in Figure 4.*

Beyond the requirements specified in Definition 5, Küsters, Truderung & Vogt remark that "[algorithm C must be] defined in such a way that...the coerced voter achieves [their] own goal, e.g., votes for [their] favorite candidate, despite what the coercer tells [them] to do. The concrete definition of [C]

[19] Smyth considers election schemes without registration; it is straightforward to adapt his games to include registration.

[20] We leave consideration of whether the definition by Moran & Naor [13] – upon which the definition by Unruh & Müller-Quade is based – captures coercion resistance as a possible direction for future work.

Fig. 4 Games KTV and KTV-$

```
$\mathsf{KTV}(\Gamma, \mathcal{A}, na, nc, nv, v, C, O, D, \kappa) =$
1  $(pk, sk, mb, mc) \leftarrow \mathsf{Setup}(\kappa)$;
2  for $1 \leq i \leq nv$ do
3      $(pd_i, sd_i) \leftarrow \mathsf{Register}(pk, \kappa)$;
4  $\mathfrak{bb} \leftarrow \emptyset$;
5
6
7
8  for $na < i < nv$ do
9      $v \leftarrow_R D$;
10     if $v \neq \phi$ then
11         $b \leftarrow \mathsf{Vote}(sd_i, pk, v, nc, \kappa)$;
12         $\mathfrak{bb} \leftarrow \mathfrak{bb} \cup \{b\}$;
13 $L \leftarrow \{pd_1, \ldots, pd_{nv}\}$;
14 $M \leftarrow (sd_1, \ldots, sd_{na})$;
15
16 $\mathfrak{bb} \leftarrow \mathfrak{bb} \cup O(pk, L, \mathfrak{bb}, \kappa)$;
17 $\mathfrak{bb} \leftarrow \mathfrak{bb} \cup \mathcal{A}(pk, L, M, sd_{nv}, \mathfrak{bb}, \kappa)$;
18 $(\mathfrak{v}, pf) \leftarrow \mathsf{Tally}(sk, \mathfrak{bb}, L, nc, \kappa)$;
19 $g \leftarrow \mathcal{A}(\mathfrak{v}, pf)$;
20 return $g$;
```

```
$\mathsf{KTV\text{-}\$}(\Gamma, \mathcal{A}, na, nc, nv, v, C, O, D, \kappa) =$
1  $(pk, sk, mb, mc) \leftarrow \mathsf{Setup}(\kappa)$;
2  for $1 \leq i \leq nv$ do
3      $(pd_i, sd_i) \leftarrow \mathsf{Register}(pk, \kappa)$;
4  $\mathfrak{bb} \leftarrow \emptyset$;
5  if $v \neq \phi$ then
6      $b \leftarrow C(sd_{nv}, pk, v, nc, \kappa)$;
7      $\mathfrak{bb} \leftarrow \mathfrak{bb} \cup \{b\}$;
8  for $na < i < nv$ do
9      $v \leftarrow_R D$;
10     if $v \neq \phi$ then
11         $b \leftarrow \mathsf{Vote}(sd_i, pk, v, nc, \kappa)$;
12         $\mathfrak{bb} \leftarrow \mathfrak{bb} \cup \{b\}$;
13 $L \leftarrow \{pd_1, \ldots, pd_{nv}\}$;
14 $M \leftarrow (sd_1, \ldots, sd_{na})$;
15 $sd \leftarrow C(pk, pd_{nv}, sd_{nv})$;
16 $\mathfrak{bb} \leftarrow \mathfrak{bb} \cup O(pk, L, \mathfrak{bb}, \kappa)$;
17 $\mathfrak{bb} \leftarrow \mathfrak{bb} \cup \mathcal{A}(pk, L, M, sd, \mathfrak{bb}, \kappa)$;
18 $(\mathfrak{v}, pf) \leftarrow \mathsf{Tally}(sk, \mathfrak{bb}, L, nc, \kappa)$;
19 $g \leftarrow \mathcal{A}(\mathfrak{v}, pf)$;
20 return $g$;
```

depends on the specific goals one wants the coerced voter to be able to achieve... We therefore do not fix this...up front."

The first four lines of games KTV and KTV-$ are identical: The challenger generates a key pair and credentials (Lines 1–3), and initialises an empty bulletin board (Line 4). In game KTV-$, the challenger constructs a ballot using the counter-strategy and adds that ballot to the bulletin board, except if the coerced voter wants to abstain (Lines 5–7). The next seven lines are identical in both games: for each non-corrupt, non-coerced voter (Line 8), a vote is sampled (Line 9), and the challenger constructs and casts a ballot for that vote, except when the vote signifies abstention (Lines 10–12). Moreover, the challenger initialises a set of public credentials (Line 13) and a set of corrupt voters' private credentials (Line 14). On the next line, algorithm O is allowed to add additional entries to the bulletin board (Line 16). Next, the adversary constructs a set of ballots (Lines 15–17), which might include ballots constructed using corrupt voters' private credentials or the coerced voter's private credential in KTV, respectively a fake credential in KTV-$. Finally, the challenger tallies the ballots and the adversary is given the election outcome, along with a proof of correct computation, and is challenged to determine whether the coerced voter gave-up their private credential or followed a strategy to evade coercion (Lines 18–20).

Care should be taken when applying the definition by Küsters, Truderung & Vogt to derive theorems: Quantitative definitions give measures of security. Measure suitability is left to analysts, which is a little dangerous – an analyst bearing ill-will could prove a system satisfies a quantitative definition, without achieving a suitable measure of security. Küsters, Truderung & Vogt defend against this. They define a minimal measure of coercion resistance (δ_{min}) [21, §4], parameterised by some number of corrupt voters (na), candidates (nc), and honest voters (nv) along with a distribution (D) over $\{1, \ldots, nc, \phi\}$, and universally quantify over those parameters in theorems. (Henceforth, we say an election scheme satisfies δ_{min}-KTV, if there exists algorithms C, O such that for all integers na, nc and nv, vote $v \in \{1, \ldots, nc, \phi\}$ and distributions D over $\{1, \ldots, nc, \phi\}$, the election scheme satisfies $\delta_{min}(na, nc, nv, D)$-KTV with respect to na, nc, nv, C, O, D, where C enables a coerced voter to achieve their own goal.) Küsters, Truderung & Vogt prove that the Bingo voting system [51] achieves this measure, i.e., satisfies δ_{min}-KTV. They were unable to do so for the ThreeBallot voting system [52], which they prove satisfies another, weaker measure of coercion resistance. Which measures suffice to deliver coercion resistance in elections is unknown and is likely heavily situational.

No election scheme can achieve 0-KTV for all na, nc, nv, v, D and some C and O, because parameters influence the degree of coercion resistance that can be achieved. Indeed, suppose, for instance, no honest voter votes for the candidate the coerced voter is instructed to vote for. If the coerced voter does not follow the coercer's instruction, then evasion will be detected. The best an election scheme can achieve is δ_{min}-KTV, which is why the definition by Küsters, Truderung & Vogt includes a δ-bound – to capture inevitable leaks.

Küsters, Truderung & Vogt designed their definition for a broad class of voting systems. At present, three schemes (Bingo, ThreeBallot, and Scantegrity II) have been proven to satisfy the weaker form of their definition (which does not consider forced abstention). These schemes share several commonalities, in particularly, they are schemes for in-person, poll-station voting, rather than remote voting. It is unknown how broadly applicable their definition is to remote voting systems, systems satisfying coercion resistance in its strongest form (with protection against forced abstention), or both.

We consider the remote voting system by Juels, Catalano & Jakobsson (and similar systems such as Civitas) and show that it cannot satisfy δ_{min}-KTV, because that measure does not take into account leaks arising from the bulletin board length, suggesting that more complicated δ bounds (and voter goals) are needed in general; in particular, in the proofs by Küsters, Truderung & Vogt the δ bound depends only on the number of candidates, honest voters, and the distribution. A more complicated δ bound would depend on additional parameters including perhaps the behavior of some parties in the system.

a) $\delta_{\min}$-KTV cannot be satisfied by Juels, Catalano & Jakobsson's voting system.: Application of Küsters, Truderung & Vogt's definition to Civitas (and the voting system by Juels, Catalano & Jakobsson, upon which Civitas is based) is rather complicated: the definition cannot be satisfied if the δ-bound comes solely from the information leaked by the election outcome (δ_{min}).

For the voting system by Juels, Catalano & Jakobsson, an adversary can trivially determine whether a coerced voter gave-up a valid credential, when no other honest voter casts a vote with an invalid credential. Indeed, while the voting system hides which ballots were cast with invalid credentials, it does not hide how many. The standard solution is to add noise to the bulletin board to hide the voter's actions, which can be achieved using algorithm O. However, in schemes like that by Juels, Catalano & Jakobsson the cost of tallying is quadratic in the number of items on the bulletin board which puts a hard limit on the amount of noise that can be efficiently added. This limit is too small to make the difference in the coerced voter's actions negligible. Hence, the tally reveals less information than the security games, which suffices to show that the voting system by Juels, Catalano & Jakobsson, and similar systems, cannot satisfy δ_{min}-KTV. Thus, to analyse such systems a new definition of δ is required, which seemingly requires a rather burdensome combinatorial analysis.

The next section presents a variant of the definition by Küsters, Truderung & Vogt which eliminates the need for combinatorial analysis, and requires only that for any adversary against the real world there is an equally successful adversary against the ideal world. We then compare this new definition to the one by Juels, Catalano & Jakobsson.

VII. Simplifying proofs of coercion resistance

A. Variant of δ-KTV

We simplify the definition by Küsters, Truderung & Vogt so that application does not require combinatorial analysis:[21] Our simplified definition folds games KTV and KTV-\$ into a single *combined* game (KTV-C), which samples β, behaves as per KTV-\$ when $\beta = 0$, and KTV otherwise ($\beta = 1$). We require game KTV-C to be indistinguishable from game Ideal, which captures our ideal world. Our simplification allows elimination of the δ-bound, since information leaked from election outcomes is captured implicitly by indistinguishability. (Re-introducing the δ-bound – for the purposes of proving security of schemes allowing forced randomisation attacks, for instance – is trivial, but hinders our comparison with the definition by Juels, Catalano & Jakobsson.)

Definition 6. *Let $\Gamma = (\mathsf{Setup}, \mathsf{Register}, \mathsf{Vote}, \mathsf{Tally})$ be an election scheme, C be a (stateful) algorithm, O be an algorithm, na, nc and nv be integers, $v \in \{1, \dots, nc, \phi\}$ be a vote, and D be an efficiently sampleable distribution over $\{1, \dots, nc, \phi\}$. We say Γ satisfies* KTV-S *with respect to na, nc, nv, v, C, O, D, if for all probabilistic polynomial-time algorithms $\mathcal{A}$, there exists a probabilistic polynomial-time algorithm $\mathcal{B}$ and a negligible function* negl *such that for all security parameters κ and votes $v' \in \{1, \dots, nc, \phi\}$, we have* $\mathsf{Succ}(\mathsf{KTV\text{-}C}(\Gamma, \mathcal{A}, na, nc, nv, v, C, O, D, \kappa)) - \mathsf{Succ}(\mathsf{Ideal}(\Gamma, \mathcal{B}, na, nc, nv, v, v', D)) \leq \mathsf{negl}(\kappa)$*, where games* KTV-C *and* Ideal *are defined in Figure 5,[22] and C enables a coerced voter to achieve their own goal.*

We say Γ satisfies KTV-S*, if there exists (stateful) algorithms C and O, and for all integers na, nc and nv upper-bound by a polynomial in the security parameter, votes $v \in \{1, \dots, nc, \phi\}$, and efficiently sampleable distributions D over $\{1, \dots, nc, \phi\}$, we have Γ satisfies* KTV-S *with respect to na, nc, nv, v, C, O, D.*[23]

Our simplified definition is equivalent to the original when the δ-bound is taken from the ideal world, i.e., when δ is the minimal measure of coercion resistance δ_{min}.

Theorem 1. *An election scheme satisfies* KTV-S *with respect to na, nc, nv, v, C, O, D iff the scheme satisfies $\delta_{\min}$-*KTV *with respect to na, nc, nv, v, C, O, D, when vote $v \in \{1, \dots, nc\}$, D is an efficiently sampleable distribution over $\{1, \dots, nc\}$, and na, nc and nv are integers upper-bound by a polynomial in the security parameter.*

Proof sketch. It suffices to prove

$$\begin{aligned}\mathsf{Succ}(\mathsf{KTV\text{-}C}(\Gamma, \mathcal{A}, na, nc, nv, v, C, O, D, \kappa)) = \\ \mathsf{Succ}(\mathsf{KTV}(\Gamma, \mathcal{A}, na, nc, nv, v, C, O, D, \kappa)) \\ - \mathsf{Succ}(\mathsf{KTV\text{-}\$}(\Gamma, \mathcal{A}, na, nc, nv, v, C, O, D, \kappa))\end{aligned}$$

and $\delta_{min} = \max_{v' \in \{1, \dots, nc\}}(\mathsf{Succ}(\mathsf{Ideal}(\Gamma, \mathcal{B}, na, nc, nv, v, v', D)))$. The former is trivial, because KTV-C is equal to KTV when $\beta = 0$ and KTV-\$ when $\beta = 1$. Let us consider the latter: Game Ideal is an ideal world and δ_{min} is a minimal measure of coercion resistance. They both capture the adversary's advantage in detecting compliance in the context of a particular

[21] Our simplification is made in the context of election-scheme syntax, it can also be applied independently of syntax, i.e., in the original context conceived by Küsters, Truderung & Vogt.

[22] Similarities between games KTV-C and Ideal, and subsequent games, are highlighted in grey.

[23] The S in KTV-S stands for simplified.

Fig. 5 Games KTV-C and Ideal

```
KTV-C$(\Gamma, \mathcal{A}, na, nc, nv, v, C, O, D, \kappa)$ =
$(pk, sk, mb, mc) \leftarrow \mathsf{Setup}(\kappa)$;
for $1 \leq i \leq nv$ do
    $(pd_i, sd_i) \leftarrow \mathsf{Register}(pk, \kappa)$;
$L \leftarrow \{pd_1, \ldots, pd_{nv}\}$;
$M \leftarrow (sd_1, \ldots, sd_{na})$;
$\mathfrak{bb} \leftarrow \emptyset$;
$\beta \leftarrow_R \{0, 1\}$;
if $\beta = 0$ then
    if $v \neq \phi$ then
        $b \leftarrow C(sd_{nv}, pk, v, nc, \kappa)$;
        $\mathfrak{bb} \leftarrow \mathfrak{bb} \cup \{b\}$;
    $sd \leftarrow C(pk, pd_{nv}, sd_{nv})$;
else
    $sd \leftarrow sd_{nv}$;
for $na < i < nv$ do
    $v \leftarrow_R D$;
    if $v \neq \phi$ then
        $b \leftarrow \mathsf{Vote}(sd_i, pk, v, nc, \kappa)$;
        $\mathfrak{bb} \leftarrow \mathfrak{bb} \cup \{b\}$;
$\mathfrak{bb} \leftarrow \mathfrak{bb} \cup O(pk, L, \mathfrak{bb}, \kappa)$;
$\mathfrak{bb} \leftarrow \mathfrak{bb} \cup \mathcal{A}(pk, L, M, sd, \mathfrak{bb}, \kappa)$;
$(\mathfrak{v}, pf) \leftarrow \mathsf{Tally}(sk, \mathfrak{bb}, L, nc, \kappa)$;
$g \leftarrow \mathcal{A}(\mathfrak{v}, pf)$;
return $\beta = g$;
```

```
Ideal$(\Gamma, \mathcal{B}, na, nc, nv, v, v', D)$ =
$\mathfrak{v} \leftarrow (0, \ldots, 0)$;
$\beta \leftarrow_R \{0, 1\}$;
if $\beta = 0$ then
    if $v \neq \phi$ then
        $\mathfrak{v}[v] \leftarrow \mathfrak{v}[v] + 1$;
else
    if $v' \neq \phi$ then
        $\mathfrak{v}[v'] \leftarrow \mathfrak{v}[v'] + 1$ ;
for $na < i < nv$ do
    $v \leftarrow_R D$;
    if $v \neq \phi$ then
        $\mathfrak{v}[v] \leftarrow \mathfrak{v}[v] + 1$;
$g \leftarrow \mathcal{B}(\mathfrak{v})$;
return $\beta = g$;
```

tally. There is nonetheless a subtle distinction, namely, game Ideal captures both forced abstention and a voter wishing to abstain being coerced to vote (forced participation), whereas δ_{min} captures neither. Nevertheless, by restricting votes v to set $\{1, \ldots, nc\}$, game Ideal disregards forced participation and abstention. Our result follows. □

Our variant simplifies analysis by eliminating combinatorial analysis and is closer to the definition by Juels, Catalano & Jakobsson, as can be observed from the following comparison.

Figure 6 provides a side-by-side presentation of games modelling the real world for each definition. Differences are as follows. First, our game considers the first na voters to be corrupt, whereas the other lets the adversary pick corrupt voters; since voters are identical up to the order they vote, the games are equivalent in this respect. Secondly, their game considers the adversary picking vote v, whereas our definition quantifies over all votes v, which is equivalent. Finally and most significantly, our game uses a stateful algorithm C to manage the coerced voter's behaviour and O to add noise to the bulletin board which is more general than the definition by Juels, Catalano & Jakobsson.

Lemma 2. *For all election schemes Γ, adversaries $\mathcal{A}$, number of corrupted voters na, candidates nc, and voters nv, ways of faking credentials* fakecred*, distributions D, and security parameters κ, we have* $\mathsf{Succ}(\mathsf{KTV\text{-}C}(\Gamma, \mathcal{B}, na, nc, nv, v, C, O, D', \kappa) \geq \mathsf{Succ}(\mathsf{JCJ}(\Gamma, \mathcal{A}, na, nc, nv, \mathsf{fakecred}, D, \kappa))$*, where $C(sd_{nv}, pk, v, nc, \kappa)$ computes $b \leftarrow$* $\mathsf{Vote}(sd_{nv}, pk, v, nc, \kappa)$ *and outputs b, $C(pk, pd_{nv}, sd_{nv})$ computes $b \leftarrow$* $\mathsf{fakecred}(pk, pd_{nv}, sd_{nv})$ *and outputs b, D' is D except for the probability of sampling λ and ϕ, which are zero and the collective probability of sampling ϕ and λ in D, respectively, and O simulates the behaviour induced by voters drawing λ.*

Proof sketch. Having eliminated the generality afforded by algorithms C and O, the games are equivalent. □

Proving the inverse is straightforward.

Figure 7 provides a similar presentation for games modelling the ideal world. In this instance, games appear quiet different. In our game, the bulletin board is a set of (plaintext) votes when $\beta = 0$ (i.e., the voter resists coercion) and when $\beta = 1$ they vote as the adversary wishes. In the other game, the bulletin board contains ballots, rather than plaintext votes. The adversary may insert ballots for each of the dishonest voters

Fig. 6 Games modeling real worlds for our definition and that by Juels, Catalano & Jakobsson

```
$\mathsf{KTV\text{-}C}(\Gamma, \mathcal{A}, na, nc, nv, v, C, O, D, \kappa) =$
1  $(pk, sk, mb, mc) \leftarrow \mathsf{Setup}(\kappa)$;
2
3  for $1 \le i \le nv$ do
4      $(pd_i, sd_i) \leftarrow \mathsf{Register}(pk, \kappa)$;
5  $L \leftarrow \{pd_1, \ldots, pd_{nv}\}$;
6  $M \leftarrow (sd_1, \ldots, sd_{na})$;
7
8
9
10 $\mathfrak{bb} \leftarrow \emptyset$;
11 $\beta \leftarrow_R \{0,1\}$;
12 if $\beta = 0$ then
13     if $v \neq \phi$ then
14         $b \leftarrow C(sd_{nv}, pk, v, nc, \kappa)$;
15         $\mathfrak{bb} \leftarrow \mathfrak{bb} \cup \{b\}$;
16     $sd \leftarrow C(pk, pd_{nv}, sd_{nv})$;
17 else
18     $sd \leftarrow sd_{nv}$;
19
20 for $na < i < nv$ do
21     $v \leftarrow_R D$;
22     if $v \neq \phi$ then
23         $b \leftarrow \mathsf{Vote}(sd_i, pk, v, nc, \kappa)$;
24         $\mathfrak{bb} \leftarrow \mathfrak{bb} \cup \{b\}$;
25
26
27 $\mathfrak{bb} \leftarrow \mathfrak{bb} \cup O(pk, L, \mathfrak{bb}, \kappa)$;
28 $\mathfrak{bb} \leftarrow \mathfrak{bb} \cup \mathcal{A}(pk, L, M, sd, \mathfrak{bb}, \kappa)$;
29 $(\mathfrak{v}, pf) \leftarrow \mathsf{Tally}(sk, \mathfrak{bb}, L, nc, \kappa)$;
30 $g \leftarrow \mathcal{A}(\mathfrak{v}, pf)$;
31 return $\beta = g$;
```

```
$\mathsf{JCJ}(\Gamma, \mathcal{A}, na, nc, nv, \mathsf{fakecred}, D, \kappa) =$
1  $(pk, sk, mb, mc) \leftarrow \mathsf{Setup}(\kappa)$;
2  $V \leftarrow \mathcal{A}(pk, \kappa)$;
3  for $1 \le i \le nv$ do
4      $(pd_i, sd_i) \leftarrow \mathsf{Register}(pk, \kappa)$;
5  $L \leftarrow \{pd_1, \ldots, pd_{nv}\}$;
6  $M \leftarrow \{(i, sd_i) \mid i \in V \wedge 1 \le i \le nv\}$;
7  $(j, v) \leftarrow \mathcal{A}(M, L)$;
8  if $|V| \neq na \vee j \notin \{1, \ldots, nv\} \setminus V \vee v \notin \{1, \ldots, nc\} \cup \{\phi\}$
   then
9      return $0$;
10 $\mathfrak{bb}_1 \leftarrow \emptyset$;
11 $\beta \leftarrow_R \{0,1\}$;
12 if $\beta = 0$ then
13     if $v \neq \phi$ then
14         $b \leftarrow \mathsf{Vote}(sd_j, pk, v, nc, \kappa)$;
15         $\mathfrak{bb}_1 \leftarrow \mathfrak{bb}_1 \cup \{b\}$;
16     $sd' \leftarrow \mathsf{fakecred}(pk, pd_j, sd_j)$;
17 else
18     $sd' \leftarrow sd_j$;
19
20 for $i \in \{1, \ldots, nv\} \setminus (\{j\} \cup V)$ do
21     $v \leftarrow_R D$;
22     if $v \neq \phi$ then
23         if $v = \lambda$ then
24             $v \leftarrow_R \{1, \ldots, nc\}$;
25             $sd_i \leftarrow \mathsf{fakecred}(pk, pd_i, sd_i)$;
26         $b \leftarrow \mathsf{Vote}(sd_i, pk, v, nc, \kappa)$;
27         $\mathfrak{bb}_1 \leftarrow \mathfrak{bb}_1 \cup \{b\}$;
28 $\mathfrak{bb}_2 \leftarrow \mathcal{A}(sd', \mathfrak{bb}_1)$;
29 $(\mathfrak{v}, pf) \leftarrow \mathsf{Tally}(sk, \mathfrak{bb}_1 \cup \mathfrak{bb}_2, L, nc, \kappa)$;
30 $g \leftarrow \mathcal{A}(\mathfrak{v}, pf)$;
31 return $\beta = g$;
```

on line 29. However, the ideal tally functionality reveals only the tally, and the bulletin board is never shown to the adversary though its length is shown.

Lemma 3. *For all election schemes Γ, adversaries $\mathcal{B}$, number of corrupted voters na, candidates nc, and voters nv, votes v and v', and distribution D, there exists an adversary $\mathcal{A}$ such that for all algorithms* $\mathsf{fakecred}$ *and security parameters κ, we have* $\mathsf{Succ}(\mathsf{JCJ\text{-}\$}(\Gamma, \mathcal{A}, na, nc, nv, \mathsf{fakecred}, D, \kappa)) \geq \mathsf{Succ}(\mathsf{Ideal}(\Gamma, \mathcal{B}, na, nc, nv, v, v', D))$.

Proving the inverse is not possible, because the bulletin board length may leak information. Interestingly, the public and secret keys seen by the adversary in $\mathsf{JCJ\text{-}\$}$ do not afford it any advantage since they are generated by registration which runs only public information. A variant of the definition which allows registration and fakecred to accept more inputs could allow very insecure schemes to be proven secure.

Proof sketch. We construct adversary $\mathcal{A}$ from $\mathcal{B}$ as follows. First, output vote v on line 7 of $\mathsf{JCJ\text{-}\$}$ and append to the bulletin board a vote for v' using credential in sd' on line 29: this suffices to ensure the tally will be the same. On line 31 pass $\mathfrak{v}$ to $\mathcal{B}$ and return its guess g since the tallies are identical in the two games it follows that the advantage of $\mathcal{A}$ is the same $\mathcal{B}$. □

In summary, the only real difference between definitions is that the adversary learns the bulletin board length in the ideal game by Juels, Catalano & Jakobsson but not in our variant, and that our variant considers a more general counter-coercion strategy and the election system (through the modelling provided by O); the definitions are complementary, some election schemes will satisfy one definition but not the

Fig. 7 Games modelling ideal worlds for our definition and that by Juels, Catalano & Jakobsson

```
$\mathsf{Ideal}(\Gamma, \mathcal{B}, na, nc, nv, v, v', D) =$

$\mathfrak{v} \leftarrow (0, \ldots, 0)$;
$\beta \leftarrow_R \{0, 1\}$;
if $\beta = 0$ then
    if $v \neq \phi$ then
        $\mathfrak{v}[v] \leftarrow \mathfrak{v}[v] + 1$;
else
    if $v' \neq \phi$ then
        $\mathfrak{v}[v'] \leftarrow \mathfrak{v}[v'] + 1$ ;
for $na < i < nv$ do
    $v \leftarrow_R D$;
    if $v \neq \phi$ then
        $\mathfrak{v}[v] \leftarrow \mathfrak{v}[v] + 1$;
$g \leftarrow \mathcal{B}(\mathfrak{v})$;
return $\beta = g$;
```

```
$\mathsf{JCJ}\text{-}\$(\Gamma, \mathcal{B}, na, nc, nv, D, \kappa) =$

$(pk, sk, mb, mc) \leftarrow \mathsf{Setup}(\kappa)$;
$V \leftarrow \mathcal{B}(pk, \kappa)$;
for $1 \leq i \leq nv$ do
    $(pd_i, sd_i) \leftarrow \mathsf{Register}(pk, \kappa)$;
$L \leftarrow \{pd_1, \ldots, pd_{nv}\}$;
$M \leftarrow \{(i, sd_i) \mid i \in V \wedge 1 \leq i \leq nv\}$;
$(j, v) \leftarrow \mathcal{B}(L)$;
if $|V| \neq na \vee j \notin \{1, \ldots, nv\} \setminus V \vee v \notin \{1, \ldots, nc\} \cup \{\phi\}$
  then
    return $0$;
$\mathfrak{bb}_1 \leftarrow \emptyset$;
$\beta \leftarrow_R \{0, 1\}$;
if $\beta = 0$ then
    if $v \neq \phi$ then
        $b \leftarrow \mathsf{Vote}(sd_j, pk, v, nc, \kappa)$;
        $\mathfrak{bb}_1 \leftarrow \mathfrak{bb}_1 \cup \{b\}$;
    $sd' \leftarrow sd_j$;
else
    $sd' \leftarrow sd_j$;
for $i \in \{1, \ldots, nv\} \setminus (\{j\} \cup V)$ do
    $v \leftarrow_R D$;
    if $v \neq \phi$ then
        if $v = \lambda$ then
            $v \leftarrow_R \{1, \ldots, nc\}$;
            $sd_i \leftarrow \mathsf{fakecred}(pk, pd_i, sd_i)$;
        $b \leftarrow \mathsf{Vote}(sd_i, pk, v, nc, \kappa)$;
        $\mathfrak{bb}_1 \leftarrow \mathfrak{bb}_1 \cup \{b\}$;
$\mathfrak{bb}_2 \leftarrow \mathcal{B}(sd', M, |\mathfrak{bb}_1|)$;
$\mathfrak{v} \leftarrow \mathsf{Ideal\text{-}Tally}(sk, \mathfrak{bb}_1 \cup \mathfrak{bb}_2, L, nc, \kappa)$;
$g \leftarrow \mathcal{B}(\mathfrak{v})$;
return $\beta = g$;
```

other. In particular, the voting system by Juels, Catalano & Jakobsson and Civitas can only satisfy the definition by the aforementioned authors. When we disregard the generality afforded by our counter-coercion strategy, we achieve the following result:

Theorem 4. *An election scheme* Γ *satisfying* KTV-S *with respect to* na, nc, nv, v, C, O, D *for all* $v \in \{1, \ldots, nc, \phi\}$*, also satisfies* JCJ *with respect to* na, nc, nv, D*, when algorithm* C *is defined as per Lemma 2.*

Proof sketch. By definition of KTV-S, any adversarial advantage is at most negligibly better in the real world then a corresponding adversary in the ideal world. Lemma 2 considers real worlds and shows that maximal adversarial advantage against KTV-S is greater than or equal the maximal advantage against JCJ, and Lemma 3 considers ideal worlds and shows the maximal adversarial advantage against JCJ is greater than or equal the maximal advantage against KTV-S. Since success in the real world is always greater than success in the ideal, we conclude our proof. □

A hybrid (between definitions) could offer a general counter-coercion strategy whilst revealing the bulletin board length, offering the best of both. Such a hybrid would be satisfied by the voting system by Juels, Catalano & Jakobsson and Civitas. (We reiterate that the definition by Küsters, Truderung & Vogt is more general then our simplified variant, with all the advantages and drawbacks that entails.)

B. *Sufficient conditions for* JCJ

Proving coercion resistance is expensive: It requires a significant devotion of time by experts. Sufficient conditions for coercion resistance are highly desirable and we identify (Theorem 5) such conditions for our variant of the definition by Juels, Catalano & Jakobsson, thereby enabling simpler, less-expensive proofs. Our theorem is reliant on the definition of *universal verifiability* by Smyth, Frink & Clarkson [35], a straightforward adaptation of the definition of *zero-knowledge tallying proofs* by Smyth [17, §5] such that it covers election schemes with registration, and three new properties that we introduce below.

Theorem 5. *Any election scheme satisfying universal verifiability, zero-knowledge tallying proofs, bulletin board indistinguishability, credential indistinguishability, and credential independence, also satisfies* JCJ.

The following proof sketch shows how the (highlighted) distinctions (Figure 1) between games JCJ and JCJ-$ can be eliminated in the presence of our five properties.

Proof sketch. By universal verifiability, election outcomes computed in game JCJ can be replaced by outcomes computed using algorithm Ideal-Tally, i.e., the algorithm used in game JCJ-$. Moreover, since tallying proofs are zero-knowledge (i.e., computed using a non-interactive zero-knowledge proof system), they do not offer an adversarial advantage and need not be input to the adversary in game JCJ. Thus, distinctions between games JCJ and JCJ-$ on Lines 28 & 29 can be eliminated. We proceed by informally introducing each of our three remaining properties and showing how they eliminate the remaining distinctions between our games.

Bulletin board indistinguishability is satisfied if there exists an efficient algorithm fakebb that can simulate the bulletin board to the adversary, i.e., output $\mathsf{fake}\mathfrak{bb}(L, |\mathfrak{bb}_1|)$ and bulletin board $\mathfrak{bb}_1$ are indistinguishable, where L is the electoral roll.

By bulletin board indistinguishability, we can replace $\mathcal{A}(sd', \mathfrak{bb}_1)$ with $\mathcal{A}(sd', \mathsf{fake}\mathfrak{bb}(L, |\mathfrak{bb}_1|))$ on Line 27 of JCJ, which allows us to simulate the real world on Line 27 based solely on information available in the ideal world.

Credential indistinguishably is satisfied if fake and private credentials (computed using algorithms fakecred and Register, respectively) are indistinguishable.

By credential indistinguishably, differences between games on Line 16 can be eliminated. The remaining distinction between games JCJ and JCJ-$ is that the former inputs private credentials of dishonest voter to the adversary on Line 7, whereas the latter does so on Line 27. The only adversary action between those lines is choosing the coerced voter and vote.

Credential independence is satisfied if algorithm fakecred does not depend on the input of private credentials, i.e., $\mathsf{fakecred}(pk, pd, sd) = \mathsf{fakecred}(pk, pd, sd')$.

Taken together, credential independence and credential indistinguishably imply that the adversary gains only negligible information from the private credentials of dishonest voter (compared to some independent and random keys); this follows since the simulator can use credential indistinguishably and independence to create fake credentials for the dishonest voters when it provides these credentials in Line 7 of game JCJ. □

Universal verifiability is a de facto standard property of voting systems and zero-knowledge tallying proofs are widely used, hence, these properties should not limit our theorem's applicability. Bulletin board indistinguishability is satisfied by systems wherein ballots comprise of ciphertexts and zero-knowledge proofs. Credential indistinguishability is seemingly necessary to protect against *simulation attacks* [4, §1.1], whereby a coercer instructs a voter to reveal their private credential and determines whether the credential is valid. Finally, credential independence holds for most systems that do not leak information about private credentials. Thus, our sufficient conditions are not arduous, and Theorem 5 should be applicable to many election schemes. In particular, they hold for Athena [12] and the voting system by Juels, Catalano & Jakobsson (after applying a patch to ensure universal verifiability [35, §6]).

VIII. Discussion: Malicious bulletin boards

A malicious bulletin board (controlled by a coercer), can harm privacy [53], [54], [17]. Yet, the presented definitions consider only honest bulletin boards: A coercer is implicitly prohibited from controlling ballot collection. Indeed, no coercer can remove nor modify honest voters' ballots. (We acknowledge that the original definition by Küsters, Truderung & Vogt can capture a coercer removing and modifying ballots, but it would be unsatisfiable, for the following reason.) Upon reflection it becomes apparent that coercion resistance cannot be achieved for malicious bulletin boards, since protection against coercion resistance cannot be offered when the coercer controls all ballot-collection channels (i.e., when the bulletin board is malicious).

We foresee two approaches to protect against forced abstention attacks when a coercer has some control over the bulletin board. First, we limit the coercer's control to a subset of ballot-collection channels. This seems a little dangerous, since it is unclear what subsets should be considered. Perhaps the subset of channels with honest voters is the most reasonable. Secondly, we introduce a reliance on voters checking whether their ballots are collected and instruct talliers not to proceed when checks fail. This essentially eliminates the possibility that a coercer removes or modifies honest voters' ballots, since individual verifiability enables detection of such actions. Albeit, denial of service attacks become trivial.

Further consideration of these two approaches and, more generally, malicious bulletin boards is a possible direction for future research.

IX. Conclusion

This work was initiated by a desire to establish formal relations between definitions of coercion resistance, which would have been useful to establish suitability and relative strength. As work progressed, we discovered that the definition by Juels, Catalano & Jakobsson does not appear to be satisfiable, and that definitions by Gardner, Garera & Rubin and Unruh & Müller-Quade are satisfiable by voting systems that are not coercion resistant: Our initial desire serves no purpose; formal relations between unsuitable definitions are worthless. Yet, that discovery is more interesting and will bring an end to the use of unsuitable definitions.

With regards the definition by Küsters, Truderung & Vogt, which has been used to analyse coercion resistance (without protection against forced abstention) of three in-person, poll-station voting systems, we found that the definition cannot be immediately applied to the seminal remote voting system by Juels, Catalano & Jakobsson nor the Civitas variant, which achieve the strongest form of coercion resistance (with protection against forced abstention). Or, at least, the definition cannot be applied without performing a rather burdensome combinatorial analysis to establish a measure of coercion resistance, and not without assuming that any such measure is suitable. To overcome that limitation, we propose a variant of the definition which eliminates combinatorial analysis.

We patched the Juels, Catalano & Jakobsson definition and compared that definition to our variant of the one by Küsters, Truderung & Vogt, observing that the definitions are rather similar, in the context of our election scheme syntax. A direct comparison between the patched definition and the original definition by Küsters, Truderung & Vogt is hindered by differences in syntax. Nonetheless, it is immediately clear that applicability of the former is limited to voting systems expressible in our election scheme syntax, whereas the latter is not. We also observed that the definition by Juels, Catalano & Jakobsson can be immediately applied, whereas (in many cases) the definition by Küsters, Truderung & Vogt requires establishing a suitable measure of coercion resistance. To simplify application of the former, we introduce sufficient conditions, which can be trivially checked, essentially eliminating the expense of proving coercion resistance.

Overall, we recommend that analysts make use of the patched Juels, Catalano & Jakobsson definition when considering voting systems that can be expressed in our election scheme syntax and that offer protection against forced abstention attacks, since analysts can then make use of our sufficient conditions which trivialise proofs of coercion resistance. For systems beyond the scope of our syntax or systems without protection against forced abstention attacks, we recommend the original definition by Küsters, Truderung & Vogt; noting that, when adopting a δ-bound other than δ_{min}, analysis will be more burdensome and that care should be taken to ensure the suitability of the δ-bound.

References

[1] T. Okamoto, "Receipt-Free Electronic Voting Schemes for Large Scale Elections," in *SP'97: 5th International Workshop on Security Protocols*, ser. LNCS, vol. 1361. Springer, 1998, pp. 25–35.

[2] A. Juels, D. Catalano, and M. Jakobsson, "Coercion-Resistant Electronic Elections," Cryptology ePrint Archive, Report 2002/165, 2002.

[3] ——, "Coercion-Resistant Electronic Elections," in *WPES'05: 4th Workshop on Privacy in the Electronic Society*. ACM Press, 2005, pp. 61–70.

[4] ——, "Coercion-Resistant Electronic Elections," in *Towards Trustworthy Elections: New Directions in Electronic Voting*, ser. LNCS, D. Chaum, M. Jakobsson, R. L. Rivest, and P. Y. Ryan, Eds. Springer, 2010, vol. 6000, pp. 37–63.

[5] M. R. Clarkson, S. Chong, and A. C. Myers, "Civitas: Toward a Secure Voting System," in *S&P'08: 29th Security and Privacy Symposium*. IEEE Computer Society, 2008, pp. 354–368.

[6] M. Schläpfer, R. Haenni, R. Koenig, and O. Spycher, "Efficient Vote Authorization in Coercion-Resistant Internet Voting," in *VoteID'11: International Conference on E-Voting and Identity*, ser. LNCS, vol. 7187. Springer, 2011, pp. 71–88.

[7] O. Spycher, R. Koenig, R. Haenni, and M. Schläpfer, "A New Approach Towards Coercion-Resistant Remote E-Voting in Linear Time," in *FC'11: 15th International Conference on Financial Cryptography*, ser. LNCS, vol. 7035. Springer, 2011, pp. 182–189.

[8] J. Clark and U. Hengartner, "Selections: Internet voting with over-the-shoulder coercion-resistance," in *FC'11: 15th International Conference on Financial Cryptography*, ser. LNCS, vol. 7035. Springer, 2011, pp. 47–61.

[9] A. Essex, J. Clark, and U. Hengartner, "Cobra: Toward Concurrent Ballot Authorization for Internet Voting," in *EVT/WOTE'12: Electronic Voting Technology Workshop/Workshop on Trustworthy Elections*. USENIX Association, 2012.

[10] A. T. Haghighat, M. S. Dousti, and R. Jalili, "An Efficient and Provably-Secure Coercion-Resistant E-Voting Protocol," in *PST'13: 11th International Conference on Privacy, Security and Trust*. IEEE Computer Society, 2013, pp. 161–168.

[11] R. Araújo, A. Barki, S. Brunet, and J. Traoré, "Remote electronic voting can be efficient, verifiable and coercion-resistant," in *FC'16: 20th International Conference on Financial Cryptography and Data Security*, ser. LNCS, vol. 9604. Springer, 2016, pp. 224–232.

[12] B. Smyth, "Athena: A verifiable, coercion-resistant voting system with linear complexity," Cryptology ePrint Archive, Report 2019/761, 2019.

[13] T. Moran and M. Naor, "Receipt-Free Universally-Verifiable Voting with Everlasting Privacy," in *CRYPTO'06: 26th International Cryptology Conference*, ser. LNCS, vol. 4117. Springer, 2006, pp. 373–392.

[14] A. Kiayias, T. Zacharias, and B. Zhang, "End-to-end verifiable elections in the standard model," in *EUROCRYPT'15: 34th International Conference on the Theory and Applications of Cryptographic Techniques*, ser. LNCS, vol. 9057. Springer, 2015, pp. 468–498.

[15] P. Chaidos, V. Cortier, G. Fuschbauer, and D. Galindo, "BeleniosRF: A Non-interactive Receipt-Free Electronic Voting Scheme," in *CCS'16: 23rd ACM Conference on Computer and Communications Security*. ACM Press, 2016, pp. 1614–1625.

[16] A. Fraser, E. A. Quaglia, and B. Smyth, "A critique of game-based definitions of receipt-freeness for voting," in *ProveSec'19: 13th International Conference on Provable and Practical Security*, ser. LNCS. Springer, 2019.

[17] B. Smyth, “Ballot secrecy: Security definition, sufficient conditions, and analysis of Helios,” Cryptology ePrint Archive, Report 2015/942, 2019.

[18] R. W. Gardner, S. Garera, and A. D. Rubin, “Coercion Resistant End-to-end Voting,” in *FC’09: 13th International Conference on Financial Cryptography and Data Security*, ser. LNCS, vol. 5628. Springer, 2009, pp. 344–361.

[19] D. Unruh and J. Müller-Quade, “Universally Composable Incoercibility,” in *CRYPTO’10: 30th International Cryptology Conference*, ser. LNCS, vol. 6223. Springer, 2010, pp. 411–428.

[20] R. Küsters, T. Truderung, and A. Vogt, “A Game-Based Definition of Coercion-Resistance and its Applications,” in *CSF’10: 23rd IEEE Computer Security Foundations Symposium*. IEEE Computer Society, 2010, pp. 122–136.

[21] ——, “A Game-Based Definition of Coercion-Resistance and its Applications,” *Journal of Computer Security*, vol. 20, no. 6, pp. 709–764, 2012.

[22] B. Smyth, “A foundation for secret, verifiable elections,” Cryptology ePrint Archive, Report 2018/225, 2018.

[23] B. Schoenmakers, “A simple publicly verifiable secret sharing scheme and its application to electronic voting,” in *CRYPTO’99: 19th International Cryptology Conference*, ser. LNCS, vol. 1666. Springer, 1999, pp. 148–164.

[24] A. Kiayias and M. Yung, “Self-tallying elections and perfect ballot secrecy,” in *PKC’01: 3rd International Workshop on Practice and Theory in Public Key Cryptography*, ser. LNCS, vol. 2274. Springer, 2002, pp. 141–158.

[25] J. Groth, “Efficient maximal privacy in boardroom voting and anonymous broadcast,” in *FC’04: 8th International Conference on Financial Cryptography*, ser. LNCS, vol. 3110. Springer, 2004, pp. 90–104.

[26] F. Hao, P. Y. A. Ryan, and P. Zieliński, “Anonymous voting by two-round public discussion,” *Journal of Information Security*, vol. 4, no. 2, pp. 62 – 67, 2010.

[27] D. Khader, B. Smyth, P. Y. A. Ryan, and F. Hao, “A Fair and Robust Voting System by Broadcast,” in *EVOTE’12: 5th International Conference on Electronic Voting*, ser. Lecture Notes in Informatics, vol. 205. Gesellschaft für Informatik, 2012, pp. 285–299.

[28] S. Khazaei and M. Rezaei-Aliabadi, “A rigorous security analysis of a decentralized electronic voting protocol in the universal composability framework,” *Journal of Information Security and Applications*, vol. 43, pp. 99–109, 2018.

[29] J. Benaloh and M. Yung, “Distributing the Power of a Government to Enhance the Privacy of Voters,” in *PODC’86: 5th Principles of Distributed Computing Symposium*. ACM Press, 1986, pp. 52–62.

[30] A. Hevia and M. A. Kiwi, “Electronic jury voting protocols,” *Theoretical Computer Science*, vol. 321, no. 1, pp. 73–94, 2004.

[31] Y. Desmedt and K. Kurosawa, “Electronic Voting: Starting Over?” in *ISC’05: International Conference on Information Security*, ser. LNCS, vol. 3650. Springer, 2005, pp. 329–343.

[32] S. Kremer, M. D. Ryan, and B. Smyth, “Election verifiability in electronic voting protocols,” in *ESORICS’10: 15th European Symposium on Research in Computer Security*, ser. LNCS, vol. 6345. Springer, 2010, pp. 389–404.

[33] R. Küsters, T. Truderung, and A. Vogt, “Accountability: Definition and relationship to verifiability,” in *CCS’10: 17th ACM Conference on Computer and Communications Security*. ACM Press, 2010, pp. 526–535.

[34] V. Cortier, D. Galindo, S. Glondu, and M. Izabachène, “Election Verifiability for Helios under Weaker Trust Assumptions,” in *ESORICS’14: 19th European Symposium on Research in Computer Security*, ser. LNCS, vol. 8713. Springer, 2014, pp. 327–344.

[35] B. Smyth, S. Frink, and M. R. Clarkson, “Election Verifiability: Cryptographic Definitions and an Analysis of Helios and JCJ,” Cryptology ePrint Archive, Report 2015/233, 2018.

[36] B. Smyth, “Surveying global verifiability,” *Information Processing Letters*, 2020.

[37] ——, “Surveying definitions of election verifiability,” 2020, draft.

[38] P. Roenne, P. Y. A. Ryan, and B. Smyth, “Cast-as-intended: A formal definition and case studies,” 2020, draft.

[39] V. Cortier, J. Lallemand, and B. Warinschi, “Fifty shades of ballot privacy: Privacy against a malicious board,” *IACR Cryptol. ePrint Arch.*, vol. 2020, p. 127, 2020.

[40] E. A. Quaglia and B. Smyth, “Authentication with weaker trust assumptions for voting systems,” in *AFRICACRYPT’18: 10th International Conference on Cryptology in Africa*, ser. LNCS. Springer, 2018.

[41] B. Smyth, “Verifiability of Helios Mixnet,” in *Voting’18: 3rd Workshop on Advances in Secure Electronic Voting*, ser. LNCS. Springer, 2018.

[42] B. Smyth and Y. Hanatani, “Non-malleable encryption with proofs of plaintext knowledge and applications to voting,” *International Journal of Security and Networks*, vol. 14, no. 4, pp. 191–204, 2019.

[43] B. Adida, O. Marneffe, O. Pereira, and J. Quisquater, “Electing a University President Using Open-Audit Voting: Analysis of Real-World Use of Helios,” in *EVT/WOTE’09: Electronic Voting Technology Workshop/Workshop on Trustworthy Elections*. USENIX Association, 2009.

[44] B. Adida, “Helios: Web-based Open-Audit Voting,” in *USENIX Security’08: 17th USENIX Security Symposium*. USENIX Association, 2008, pp. 335–348.

[45] P. Bulens, D. Giry, and O. Pereira, “Running Mixnet-Based Elections with Helios,” in *EVT/WOTE’11: Electronic Voting Technology Workshop/Workshop on Trustworthy Elections*. USENIX Association, 2011.

[46] G. Tsoukalas, K. Papadimitriou, P. Louridas, and P. Tsanakas, “From Helios to Zeus,” *Journal of Election Technology and Systems*, vol. 1, no. 1, 2013.

[47] E. A. Quaglia and B. Smyth, “Secret, verifiable auctions from elections,” *Theoretical Computer Science*, vol. 730, pp. 44–92, 2018.

[48] J. Clark, “Democracy Enhancing Technologies: Toward deployable and incoercible E2E elections,” Ph.D. dissertation, University of Waterloo, 2011.

[49] X. Yia and E. Okamoto, “Practical Internet voting system,” *Journal of Network and Computer Applications*, vol. 36, no. 1, pp. 378–387, 2013.

[50] S. Delaune, S. Kremer, and M. D. Ryan, “Verifying privacy-type properties of electronic voting protocols,” *Journal of Computer Security*, vol. 17, no. 4, pp. 435–487, Jul. 2009.

[51] J.-M. Bohli, J. Müller-Quade, and S. Röhrich, “Bingo Voting: Secure and Coercion-Free Voting Using a Trusted Random Number Generator,” in *VoteID’07: First international conference on e-voting and identity*. Springer, 2007, pp. 111–124.

[52] R. L. Rivest, “The ThreeBallot Voting System,” 2006.

[53] B. Smyth, “Ballot secrecy with malicious bulletin boards,” 2014, cryptology ePrint Archive, Report 2014/822 (version 20141012:004943).

[54] D. Bernhard and B. Smyth, “Ballot secrecy with malicious bulletin boards,” 2015, cryptology ePrint Archive, Report 2014/822 (version 20150413:170300).