Usability Evaluation of Loki as a Coercion-Resistant Internet Voting System

Christina Nissen[0009−0005−7912−9249], Oksana Kulyk[0000−0003−4218−1658], and Rosario Giustolisi[0000−0002−2917−9601]

IT University of Copenhagen, Copenhagen, Denmark
{chfn,okku,rosg}@itu.dk

Abstract. A human-targeted threat to internet voting is voter coercion, where an adversary forces a voter to cast a ballot in a particular manner. Loki is an internet voting system that offers coercion resistance through flexible vote updating, that is, allowing the voter to change the vote cast under coercion to their actual preferred candidate. We evaluate Loki via an online experiment, ($N = 172$), asking participants to perform a variety of tasks, including two use scenarios involving asking the participants to behave as if they were under coercion. We conclude that the usability of Loki varies depending on the use scenario, ranging from a 100% success rate for casting a single vote without coercion to a 59% success rate in more complex scenarios involving coercion, and achieving a medium- to high-level of perceived trust, transparency, understandability, and user experience. We conclude that while Loki is usable for most voters, future work is needed to improve voters' mental models of the system.

Keywords: Coercion-resistance · Usability · Internet voting

1 Introduction

Internet voting systems are an opportunity to increase voter turnout and make elections more accessible voters However, implementing internet voting introduces new threats due to the uncontrolled voting environment. One human-targeted threat is voter coercion, where an adversary force the voter to cast their vote in a particular manner, undermining the integrity of election results [19].

Various coercion-resistant voting systems have been proposed to mitigate the problem of voter coercion [8, 19, 25]. However, the protection in these systems relies on the voter being able to apply certain techniques to deceive the coercer, making it appear as if they are complying with the adversary's instructions. If successful, the voter is able to vote according to their actual intention, while the coercer has no possibility of detecting it. Therefore, the effectiveness of these coercion-resistant internet voting systems relies heavily on the voter being able to successfully apply such techniques. However, this dependency is known to be problematic due to unrealistic assumptions on voters' behaviour, which some of these techniques depend upon [21, 33].

Loki [16] is an internet voting system that protects voters against coercion via *flexible vote updating*, implemented through *vote identifiers* assigned to each cast ballot and known only to the voter. The system allows voters to cast multiple votes in an election, only one of which will count towards the final tally. Thus, a coerced voter can deceive the coercer by first casting the instructed vote under observation and later replacing it with a vote of their choosing by providing the correct vote identifier(s) once the coercer has left. Alternatively, if the coercer observes the voter until the end of the election, the voter can cast their true vote first and later invalidate a coerced vote by entering an incorrect identifier.

A preliminary evaluation of Loki [20] has shown promising results with regard to its usability, with 100% of voters being able to cast an initially true without coercion, and 90% of participants being able to withstand coercion by casting an invalid vote. However, the study has several limitations, making it difficult to generalise the results (e.g., small sample size, short period between vote casting and simulated coercion, and investigating a single voting scenario). Furthermore, the study has a narrow focus on usability without investigating other widely acknowledged factors within internet voting, such as trust and transparency [4, 26, 51]. While one other study has investigated trust in a coercion-resistant internet voting system, it only compares voters' trust with other voting methods, and it does not employ validated measures or focus on trust in the voting system itself [31].

In this work, we address these gaps and conduct an evaluation of Loki ($N = 172$) via a between-subject online experiment. We evaluate voters' ability to use Loki in three scenarios, each consisting of two tasks: (S1) casting one's true vote, then casting a vote as if they were under coercion; (S2) casting a vote as if they were under coercion, then updating it with their true vote; and (S3) casting a vote, then updating it without coercion involved. We furthermore evaluate the effect of risk framing in the user interface by providing two versions of the system: one with *implicit* framing presenting the coercion-resistant mechanism as a general security feature (i.e., without coercion), and one with *explicit* framing that motivates how the required actions work to mitigate coercion. We finally test for their user experience with the system (using the UEQ questionnaire [48]), as well as their perceived transparency, understandability and trust in the system (using the TDIV and TVS questionnaires [2, 3]). Therefore, we evaluate the following research questions:

RQ1 How effectively can voters use Loki in the three evaluated scenarios?
RQ2 What effect do the risk framing have on effectiveness, user experience, understandability, transparency and trust?
RQ3 How do voters perceive Loki with regard to user experience, understandability, transparency and trust?

We conclude that Loki is effective when voters cast a single vote without coercion in an election. However, the system is less effective in scenarios where voters experience coercion after having already cast an initial true vote, as well as in cases where voters need to update their vote – either to correct a coerced vote or simply to change their preference. Furthermore, we show that risk framing do not affect effectiveness, user experience, understandability, transparency, and trust. Nevertheless, voters generally perceive Loki as trustworthy, transparent, and understandable, while the user experience also indicates positive perceptions.

2 Background

Coercion-resistance in Internet Voting Coercion-resistance in internet voting refers to the ability for voters to deceive an adversary into thinking that they have behaved as instructed while casting a vote according to their true preference despite being exposed to threats and intimidation [19]. This ability closely aligns with the United Nations International Human Rights Standards on Elections, requiring voters to be protected from any form of coercion in political elections [36]. In the literature, Juels et al. were the first to formalise the notion of coercion-resistance in electronic voting [19]. They describe how coercion-resistance should defend against coercers with different goals (e.g., forcing them to vote for a specific candidate, abstaining them from voting, or casting an invalid ballot) and capabilities (e.g., observing the voter cast their ballot or requesting proof that their vote followed the coercer's instructions). However, achieving such a strong notion of coercion resistance still requires specific assumptions, and many technical solutions provide weaker guarantees in practice. In this study, we focus on a threat model in which an adversary attempts to coerce a voter to cast their ballot for a particular candidate either through physical presence or remote observation.

To mitigate the problem of voter coercion, several coercion-resistant systems exist that differ in how they protect the voter against coercion [6, 8, 9, 11, 16–19, 24, 25, 45, 52, 53]. However, these systems all require additional actions from the voter to achieve coercion-resistance. One of these systems is Loki, which implements a *flexible vote updating* strategy to achieve coercion-resistance [16]. In this study, we focus on evaluating Loki [16], which provides strong guarantees of ballot privacy and verifiability given the validity of the underlying cryptographic assumptions, and, to the best of our knowledge, is the only internet voting system employing a flexible vote updating strategy for coercion-resistance, implemented through vote identifiers assigned to each cast ballot and known only to the voter. A flexible vote updating-strategy provides, as the name implies, flexibility towards different coercive scenarios. In particular, this strategy protects the voter against last-minute coercion where the adversary instructs the voter minutes before the end of the election, making it impossible for the voter to update their vote according to their true preference. Thus, this type of strategy makes it possible for the voter to cast their true vote before coercion occurs and not just after coercion, emphasising the security advantage of Loki. Moreover, the system enables all voters to update their votes through this security mechanism, including those not subject to coercion but who simply revise their preferences. From the voter's perspective, the process of casting a vote is as follows:

First-time voting: The voter selects their preferred candidate and casts a ballot. Afterwards, the system displays a vote identifier that must be memorised for potential future vote updates. To preserve coercion-resistance, the voter should not store or record this identifier.

Vote updating: The voter must provide the correct vote identifier(s) corresponding to all prior ballot(s). Since a new vote identifier is issued for every cast ballot, all previously issued vote identifiers must be supplied when updating a vote. The voter then casts a new ballot reflecting their current preference. To ensure coercion-resistance, the system does not indicate whether the submitted vote identifiers are correct. Thus, even if incorrect vote identifiers are provided, the ballot appears successfully cast to the voter.

Voting under coercion: The voter's actions depend on whether coercion occurs before or after a true vote has already been cast. If coercion occurs after a true vote has been cast, the voter follows the coercer's instructions but provides incorrect vote identifier(s), thereby invalidating the coerced ballot. If coercion occurs before any true vote has been cast, the voter complies with the coercer's instructions and later updates the coercive vote by providing the correct vote identifier(s) once the coercer is no longer present.

Human Factors in Internet Voting While usability has been studied extensively for internet voting, the majority of studies focus on either verifiable internet voting systems or code-voting solutions and do not address the threat of coercion or its associated usability challenges, such as lack of consistent system feedback [10,13,22,27,30,50]. Only a few studies have examined the usability of coercion-resistant voting systems [12,20,31,32]. All of these studies, except one, investigate systems implementing other approaches than flexible vote updating against coercion [12, 31, 32]. A preliminary lab study of Loki shows promising usability results with high success rates for participants casting a single vote without coercion (100%) and a subsequent invalid vote to protect their first vote against coercion (90%) [20]. The study furthermore finds a mean SUS score of 73. However, the evaluation had a small sample size ($N = 30$) of university students and had further limitations concerning: (1) its study design (e.g., a short period between casting the true vote and simulated coercion); (2) provided instructions (e.g., explicitly telling the participants to cast an invalid vote rather than letting them figure out by themselves how to avoid coercion); and (3) a single evaluated use case scenario. Thus, we address these gaps by conducting the first large-scale usability evaluation of a Loki.

In addition to usability, trust is also a critical factor for voters' acceptance of internet voting systems [4, 26, 51]. Transparency has been shown to have a significant positive effect on trust [3], while both policymakers and researchers also have acknowledged transparency as an important aspect in internet voting systems [1,4,14,26,51]. Furthermore, improving users' perceived understandability of system functionality can positively influence perceptions of transparency and trust [3, 54]. This emphasises that perceived trust, transparency, and understandability are important aspects of an internet voting system. However, exposing security mechanisms directly in the interface can have mixed effects in internet voting. Some studies show improved user experience when security features such as verifiability are made visible [13]. Other studies indicate that unfamiliar mechanisms such as vote updating may reduce trust [26]. However, to our knowledge, no prior studies have investigated the trust and transparency of coercion-resistant internet voting systems. This emphasises our motivation to study how trustworthy voters perceive coercion-resistant mechanisms, as these might pose a barrier to user acceptance in the real world.

3 Methodology

Prototype We designed our prototype using the interface of Loki proposed in previous research as a starting point [20]. To refine the prototype, we followed an iterative design process based on (1) feedback sessions with lay voters and usable security experts, and (2) feedback from an online pre-study ($N = 166$) [1]. We describe the resulting prototype below.

General Information and Authentication At the beginning, the voter is provided with general information about the voting system and needs to authenticate themselves. The authentication is implemented as credentials in the form of a voter ID and password that the voter is assumed to receive before the election. In real-world elections, however, authentication can be implemented in various ways depending on the digital infrastructure available in the country where the elections take place. Following recommendations in [33], the voter is advised to use a private browser before logging in, which is essential for coercion resistance, making it impossible for the coercer to detect whether the voter has accessed the voting system in their absence [2].

Voted Before As in [20], the voter is asked upon authentication whether they have cast any vote in the election (see Figure 1a). The voters who answer "no" (i.e., those either casting their initial true vote or lying to the coercer to hide their previously cast vote) will proceed directly to vote casting; those who click "yes" will be asked to identify previously cast ballots as the next step.

Identifying previously cast ballot(s) For vote updating to be successful, the voter must be able to recognise their vote identifier(s) (see Figure 1b). We decided to use pictures of objects with their associated word as vote identifier(s), as prior research demonstrates better recall of pictures compared to words, as they encode both visually and verbally according to the dual-coding theory [37, 38]. However, we acknowledge that other vote identifiers exist (e.g., numbers or abstract symbols), potentially with higher effectiveness. We used the same image database as in [20]. To ensure coercion resistance, this screen does not provide any feedback on whether their selection of vote identifier(s) is correct.

[1] The pre-study followed the same setup as the main study. In addition to providing feedback for the prototype, it has also been used to refine the study instructions to make them more understandable to the participants. For the sake of brevity, we omit a detailed discussion of the pre-study in the paper.

[2] Note, for the sake of our study and the ability to track user interactions with the system, our participants were asked to refrain from doing so in our study instructions.

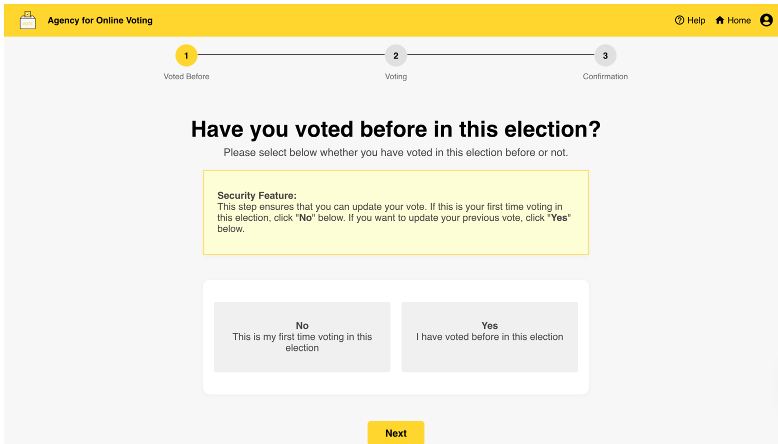

(a) Answer previous voting

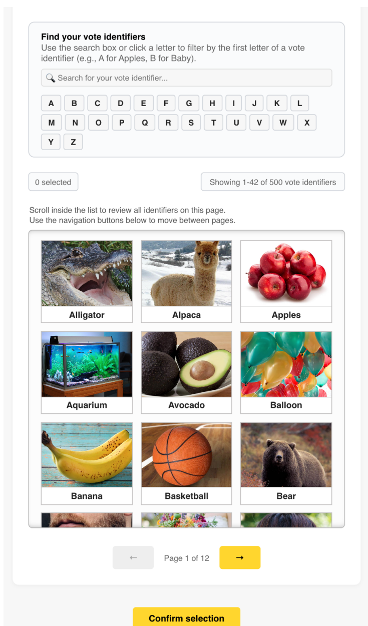

(b) Identify previous ballot(s).

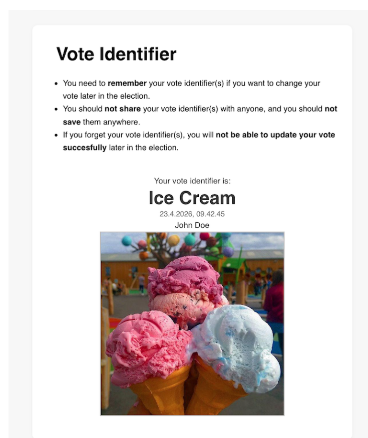

(c) Receival of vote identifier.

Vote casting The voter selects a candidate and cast their ballot.

Confirmation with vote identifier After vote casting, the voter receives a confirmation together with the vote identifier associated with the cast ballot (see Figure 1c). To protect against coercion, the screen displays the selected candidate from the vote-casting step regardless of whether the voter correctly identifies their vote identifier(s). Consequently, system feedback may not reliably reflect whether a vote update was successful. Furthermore, the screen emphasises that the voter must not save or store this identifier, and they must remember it to update their vote.

Risk framing We developed two variations of UI instructions that present different coercion risk framings: *explicit* and *implicit*. The risk framing refers to whether the interface communicates the security threat that motivates a voter action. While voters are exposed to the coercion-resistant mechanisms in both variants, they differ in whether coercion is explicitly presented as the risk that the mechanism is intended to address. Previous studies find that explicitly framed security warnings and unfamiliar security interactions (i.e., vote updating) may reduce trust [7, 26], a critical factor for user acceptance in e-voting systems [3]. At the same time, exposing underlying security mechanisms in the interface may improve the user experience in verifiable e-voting systems [13]. Given these mixed findings, we designed two variants of the interface to investigate how the explicitness of coercion-resistant security mechanism influences the voters' perceptions:

- Explicit coercion risk framing: The user interface explicitly mentions coercion as the motivating threat behind specific required voter actions, such as identifying previously cast ballots (see Figure 2a).
- Implicit coercion risk framing: The user interface presents the coercion-resistant mechanism as a general security feature without referring to coercion as the underlying risk motivating the required voter actions, except where strictly necessary (see Figure 2b).

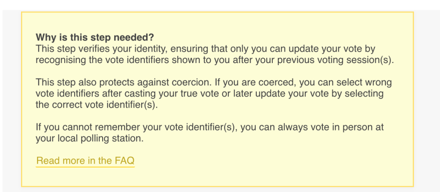

(a) *Explicit framing* that explicitly mentions threat of coercion.

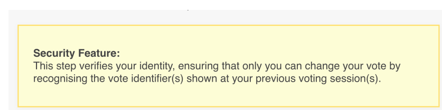

(b) *Implicit variant* that frames the step as necessary for security reasons without mentioning coercion.

Fig. 2: Different variants of risk framing for identifying previously cast ballot(s).

Study Design Our study employed a 3×2 factorial design, crossing three different voting scenarios (described in detail below), with the two risk framings. We conducted an online between-subjects design in which participants engaged in unmoderated usability testing of our prototype and answered a questionnaire. We chose a between-subjects design because it more closely reflects a real-world election where voters only experience a single voting scenario. Furthermore, we chose this design to reduce practice effects.

Choice of Voting Scenarios As we described in Section 2, the voter's action to deceive the coercer depends on whether they have cast a true vote before the coercion occurs. To capture these two distinct user journeys, we decided to evaluate them in their simplest form. While more complex scenarios exist we omitted them because if voters struggle with the simplest scenarios, they are also unlikely to succeed in more complex ones. In addition, we wanted to examine the vote-updating mechanism in a scenario without coercion to understand how voters interact with the mechanism during normal voting. Consequently, we investigated the following three voting scenarios in this study:

– Scenario 1 (S1): The voter has already cast a true vote and must protect it against coercion.
– Scenario 2 (S2): The voter follows the coercer's instructions and casts an initial coercive vote. After the coercer disappears, the voter updates the ballot to reflect their true preference.
– Scenario 3 (S3): The voter updates a previously cast vote to another candidate without any coercion.

We conducted two moderated lab tests and a pilot study with 30 participants on Prolific, five for each condition.

Study Procedure Our study consisted of two phases of unmoderated remote usability testing, conducted seven hours apart to simulate a 15-hour online election (see study procedure in [34]). We chose this length of election and interval between the two phases for several reasons: (1) it somehow resembles real-world polling station opening hours for general elections, which are usually 12-14 hours [15, 39]; (2) considerations about participant return; and (3) study logistics, as the gap allowed us to set up the necessary materials for phase two. We used SurveyXact for the survey [46]. Participants were randomly assigned to one of six conditions, reflecting our 3x2 study design that combined the three voting scenarios with the two risk framings.

Study: Part 1 We welcomed participants and asked for their consent (see consent form in [34]). Afterwards, the participants were instructed to complete *Task 1*. The study instructions included an attention check asking participants to identify their assigned task, as well as an embedded hidden-response bot check. Participants, who were assigned to either *S1* or *S3* were instructed to cast a true vote for the candidate 'John Doe' (see study instructions in [34])[3]. In *S2*, participants were instructed to imagine experiencing coercion and follow the coercer's instruction and cast an initial vote for the candidate 'John Doe'. To succeed in *Task 1*, all participants had to answer "no" when asked if they had cast a vote previously in the election. Subsequently, they had to cast a vote for 'John Doe' and were shown their associated vote identifier, which was an image of an ice-cream (see Figure 1c). Afterwards, the participants were directed to SurveyXact to provide their prolific ID and a random number copied from the final page of the prototype's study instructions. We needed this information to track and identify the participants in the databases. Upon completion, participants were directed back to Prolific.

Study: Part 2 After seven hours, participants were re-contacted to complete the second phase, using their first-phase random number to assign them to the

[3] For the sake of vote privacy and avoiding biasing the participants, the list of candidates consisted of randomly generated names.

same condition. After providing informed consent, we instructed the participants to complete *Task 2*. Identical to the first study part, we included an attention check asking participants to identify their assigned task, as well as an embedded hidden-response bot check. In *S1*, we instructed participants to imagine experiencing coercion and protect their first vote by using the system's security feature. To succeed in this task, they could either (1) answer "no" when asked if they had cast a vote previously in the election, cast a vote, and receive a vote identifier, or (2) answer "yes" to the question and then provide incorrect vote identifier(s) before vote casting and receival of vote identifier. In *S2*, we instructed participants to update their initial coercive vote for the candidate 'Sofia Lee'. This required the participants to answer 'yes" when asked if they had cast a vote previously in the election, select the vote identifier that they saw when casting their initial coerced vote (i.e., ice-cream), cast the vote for 'Sofia Lee', and receive a vote identifier. In *S3*, we instructed participants to update their vote to the candidate 'Sofia Lee', which required the same actions as those in *S2*.

After completing *Task 2*, the participants had to answer a questionnaire (see full questionnaire [34]). Identical to the first study phase, participants had to provide their Prolific ID and the random number from the final page of the prototype. To investigate *user experience*, participants then had to answer the standard User Experience Questionnaire (UEQ) [23]. We furthermore asked if they had stored their vote identifier, despite the voting system instructing them not to do so. To investigate *understandability* and *transparency*, participants answered items from their respective subscales of the Transparency Dimensions for Internet Voting (TDIV) measurement [5] using a 7-point Likert scale ranging from "strongly disagree" to "strongly agree"[4]. To investigate *trust*, we similarly had participants answer items on a 7-point Likert scale within the overall trust subscale in the Trust in Voting Systems (TVS) measurement [2]. Within these items, we included an attention check for quality control, following Prolific recommendations [42]. The participants furthermore answered demographic questions, whether they had prior experience with internet voting, and their general attitude towards internet voting, as this can distort user study results [29]. Finally, we asked them about any final remarks before their return to Prolific.

Measurements Our independent variables are the assigned voting scenario (*S1-S3*) and risk framing (*implicit* and *explicit*). Our dependent variables are defined as follows:

1. Effectiveness is measured as the binary variable *overall effectiveness* indicating whether participants successfully completed all their tasks in their assigned voting scenario (1 = succeeded in both tasks, 0 = failed at least one). We only considered participants succeeding if they correctly completed all tasks in their assigned scenario, as this would be needed from the voters to cast a ballot according to their true preferences if such scenarios were to happen in the real world.

[4] We excluded the subscales *information availability, monitoring & verification, remedial measures and testing*, as these are either not included or not relevant for our designed prototype.

2. User experience is measured as the mean of attractiveness, pragmatic quality, and hedonic quality, following the UEQ [23].
3. Understandability and transparency are measured as the mean of the items in their respective dimensions, following the TDIV questionnaire [5].
4. Trust is measured as the mean of the items within the overall trust subscale, following the TVS questionnaire [2].

Recruitment and Ethics We used the platform Prolific to recruit participants [40]. Before the data collection, we conducted a power analysis with G*Power for a two-factor between-subjects ANOVA with interaction (df = 2, α = 0.05, effect size = 0.25). The analysis indicated that 158 participants would be required to detect a medium effect with 80% power. Thus, we decided to recruit 300 in total, respectively 50 for each condition, as we anticipated some participant attrition with our two-phase study procedure. All participants were fluent in English and were residents of the United Kingdom. We recruited participants from one country to eliminate it as a confounding effect on trust. Based on the median times from the pilot and pre-study, we paid them, respectively, £0.75 and £1.95 for their participation in each study phase, which is above the hourly rate recommended by Prolific [41]. We coordinated the process alongside the ethical guidelines and received approval from the ethical committee of our institution. Additionally, we used an institution-approved, GDPR-compliant survey platform, and while we used Back4app to track the participants' interactions, we did not store any personal or sensitive data in the database. Furthermore, we provided participants with consent forms outlining the study's purpose, withdrawal options, and data handling procedures. Contact details for the researchers were also provided for any enquiries (see consent forms in [34]). As some of our study tasks included imagining being coerced while using the voting system, it could make participants feel discomfort. Therefore, they had the option to leave the study without any consequences, as described in the consent forms.

Data Analysis We used the statistical software R 4.3.0 (R Studio) to analyse our data [43, 44]. In addition, we used the Excel data analysis tool for the UEQ provided by [48]. In our analysis, we included participants who completed their *Task 2* before 10 PM and passed all attention checks. From 300, we excluded 21 for failing at least one of the attention checks, 17 for failing to perform *Task 2* before 10 PM, and 88 for not completing both study phases. Additionally, we excluded two participants who did not follow the instructions, making it impossible to identify their data from the first study phase. After applying these exclusion criteria, our final sample consisted of 172 participants.

Effectiveness We used data from Back4App to compute successful and unsuccessful outcomes for each study task, as well as successful completion of both study tasks. *Task 1* was coded as correct if participants' initial vote in the database for 'John Doe' across all conditions. All other cases were coded as incorrect. For *Task 2*, *S1* was coded as correct if participants had a final vote for 'John Doe' cast during the time of Task 1, while *S2* and *S3* required a vote for 'Sofia Lee' with the correct identifier (i.e., ice-cream). All other cases were coded as incorrect. We conducted a logistic regression model with a binomial link, with the interaction of scenario and risk framing as predictors to our dependent variable *overall effectiveness*. Furthermore, we included the time between the two study phases as a main effect in our model, as participants had seven hours to complete *Task 2*, resulting in different return times. We assessed the model assumptions and observed a mild degree of overdispersion (dispersion parameter = 1.27). Otherwise no substantive violations of model assumptions were detected. Furthermore, we analysed why participants failed their voting tasks.

User Experience For each condition, we calculated the mean scores of pragmatic quality, hedonic quality, and attractiveness. Additionally, we assessed internal consistency using Cronbach's alpha. Results indicated acceptable to good reliability across most subscales of the UEQ ($\alpha = 0.71$–0.91), except for dependability, which showed lower reliability ($\alpha = 0.64$), indicating reduced internal consistency for this subscale. We conducted a two-way MANOVA using Pillai's Trace, with scenario and risk framing and their interaction as independent variables. The dependent variables were pragmatic quality, hedonic quality, and attractiveness. We checked the assumptions of MANOVA, and no substantive violations were detected.

Trust, Transparency, and Understandability For each condition, we calculated the mean scores of trust, transparency, and understandability. Additionally, we assessed internal consistency using Cronbach's alpha, which indicated good reliability ($\alpha = 0.85 - 0.96$). We conducted a two-way MANOVA with scenario and risk framing and their interaction as independent variables. The dependent variables were trust, transparency, and understandability. These constructs have previously been shown to be correlated [3]. Consistent with this, Pearson correlations among the dependent variables indicated strong associations ($r = 0.71 - 0.78$), supporting a multivariate analytic approach. We evaluated assumptions of MANOVA with Shapiro-Wilk tests indicating minor deviations from normality in some conditions across. However, these deviations were not consistent across all conditions. Homogeneity of covariance matrices was assessed using Box's M test, which indicated a significant violation ($\chi^2 = 51.20$, $df = 30$, $p = 0.009$). Given these results, Pillai's Trace was used for MANOVA

4 Results

In this section, we report the results from 172 participants. Table 1 provides the distribution of participants across the conditions. The average time it took the participants to return for phase 2 was 8 hours 20 minutes. Among participants, 95 were women, 76 were men, and one was self-described. The biggest age group was 35-44 years, with 48 participants (see [34] for full demographics). Across all conditions, 29 participants had a negative attitude towards internet voting, 28 were not sure, and 80 had a positive attitude. Similar attitudes were present for postal voting, with 33 participants having a negative attitude, 16 being unsure, and 123 having a positive attitude. Across all conditions, 26 participants reported having tried online voting before.

Table 1: Participant distribution across scenarios and risk framing conditions.

Scenario	**Implicit**	**Explicit**	**Total**
S1	24	27	51
S2	26	35	61
S3	32	28	60
Total	82	90	172

Table 2: Number and percentage of participants successfully completing both tasks in their respective scenario/risk framing

	Implicit	Explicit	**Overall**
S1	14/24 (58%)	16/27 (59%)	**30/51 (59%)**
S2	20/26 (77%)	21/35 (60%)	**41/61 (67%)**
S3	24/32 (75%)	21/28 (75%)	**45/60 (75%)**
Overall	**58/82 (71%)**	**58/90 (64%)**	**116/172 (67%)**

Effectiveness Across all conditions, 67% of our participants completed their voting scenario, and 33% failed. Table 2 provides the rates of participants succeeding in both tasks of their scenario for each condition, while Figure 3 details the completion rates for the individual tasks.

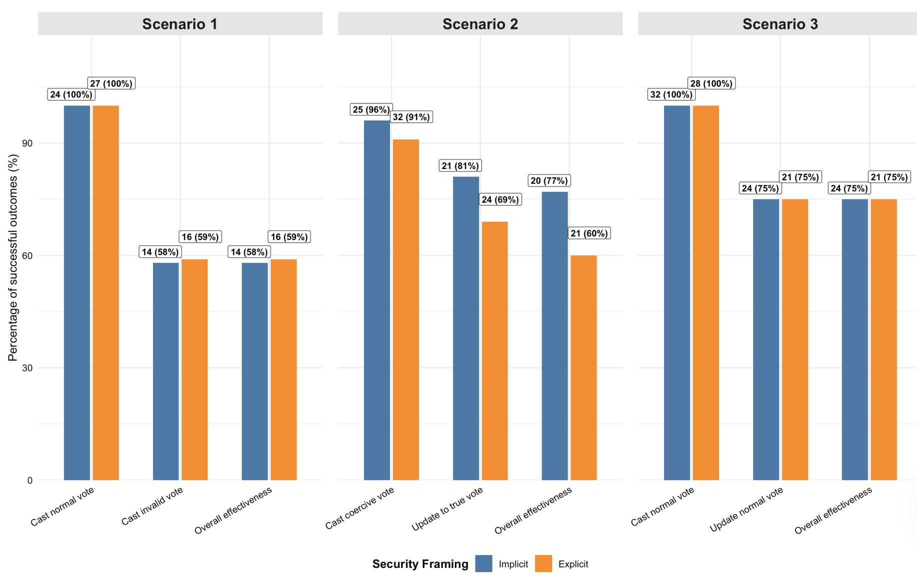

Fig. 3: Overall effectiveness in percentages across the six conditions.

The logistic regression model indicates no significant main effect of scenario on overall effectiveness ($\chi^2(2) = 3.30, p = 0.192$), nor a significant effect of risk framing in the UI ($\chi^2(1) = 0.63, p = 0.426$). There is also no evidence of an interaction between scenario and risk framing ($\chi^2(2) = 0.99, p = 0.610$). Finally, time between the two voting phases do not significantly predict overall effectiveness ($\chi^2(1) = 2.43, p = 0.119$) Additionally, the logistic regression including only main effects shows no significant effect of scenario or risk framing in the UI (both $p > 0.05$), consistent with the full model.

Reasons for Failure: Across all conditions, four participants failed *Task 1* (2%), while 52 participants failed *Task 2* (30%). All participants who failed *Task 1* were assigned to *S2*. One participant failed because they did not cast their vote for John Doe. Two participants failed because they answered "yes" to having voted. One participant failed due to both answering "yes" to having voted before and failing to cast a vote for John Doe. For *Task 2* in *S1*, all 21 participants failed because they updated their vote instead of retaining their original true vote for John Doe. All except one participant subsequently re-selected John Doe. In *S2*, 11 participants failed to update their vote because they did not provide the correct vote identifier from their initial coercive vote. Additionally, three participants failed because they answered "no" to having voted before. One participant failed because they saved their vote identifier. In *S3*, 12 participants failed to provide the correct vote identifier from their initially cast vote. Additionally, four participants failed because they had saved their vote identifier.

User Experience Table 3 provides the mean scores for the pragmatic and hedonic quality together with attractiveness. Across all conditions, the pragmatic quality mean scores are higher than the hedonic quality ones; furthermore, conditions with an *explicit* risk framing also show higher means than conditions with implicit risk framing. The two-way MANOVA reveals no significant multivariate effects on the combined dependent variables (pragmatic quality, hedonic quality, and attractiveness) of the scenario (Pillai's Trace = 0.024, $F(6, 330) = 0.666$, $p = 0.677$) or the risk framing in the UI (Pillai's Trace = 0.038, $F(3, 164) = 1.966$, $p = 0.121$). Furthermore, the interaction between scenario and risk framing in the UI also shows a non-significant multivariate effect (Pillai's Trace = 0.024, $F(6, 330) = 0.664$, $p = 0.679$).

Table 3: Pragmatic quality, hedonic quality, and attractiveness mean scores across the conditions (scale from -3 (low) to 3 (high)).

Dimension	**S1-Impl.**	**S1-Expl.**	**S2-Impl.**	**S2-Expl.**	**S3-Impl.**	**S3-Expl.**
Pragmatic Quality	1.23	1.46	1.11	1.33	1.10	1.73
Hedonic Quality	0.807	1	0.688	1.15	0.668	1
Attractiveness	1.31	1.51	1.13	1.39	1.02	1.63

Trust, Transparency, and Understandability Table 4 summarises the mean scores for trust, transparency, and understandability across all conditions. All mean scores range from 5.03 to 5.48 on a 7-point Likert scale (1 = low, 7 = high). In addition, *explicit* risk framing was associated with slightly higher mean scores for trust, transparency, and understandability in *S1* and *S2*, while *S3* showed a marginally higher trust score in the *implicit* condition. The two-way MANOVA reveals no significant multivariate effects on the combined dependent variables (trust, transparency, and understandability) of the scenario (Pillai's Trace = 0.027, $F(6,330) = 0.741$, $p = 0.617$) or the risk framing in the UI (Pillai's Trace = 0.019, $F(3,164) = 1.056$, $p = 0.370$). Furthermore, the interaction between scenario and risk framing in the UI also shows a non-significant multivariate effect (Pillai's Trace = 0.021, $F(6,330) = 0.597$, $p = 0.733$).

Table 4: Mean scores for trust, transparency, and understandability across the conditions (scale from 1 (low) to 7 (high)).

Dimension	**S1-Impl.**	**S1-Expl.**	**S2-Impl.**	**S2-Expl.**	**S3-Impl.**	**S3-Expl.**
Trust	5.03	5.32	5.04	5.06	5.41	5.21
Transparency	5.32	5.43	5.04	5.36	5.28	5.35
Understandability	5.24	5.48	5.12	5.26	5.43	5.48

5 Discussion and Conclusion

RQ1: Effectiveness in various scenarios Our results show that 67% of our participants can complete their respective voting scenario. Still, the failure rate remains relatively high at 33%, with reasons depending on the specific voting scenario. We discuss these different voting scenarios and their implications below. In the use case where the voter *casts a single vote without coercion*, Loki is highly usable, as all participants successfully do so. Therefore, in online elections where this scenario is the most common, Loki can provide an effective voting experience for the voters. In the use case where the voter needs to *update their vote*, either due to prior coercion or a change of mind, most of our participants complete it (67% and 75%, respectively). However, the reason of failure is mainly because the participants don't provide the correct identifier from their initial vote. This finding indicates that memorability is a key issue in such scenarios, preventing voters from success. This finding further aligns with previous research that identifies memorability as a usability challenge for such a coercion-resistant system as Loki [21, 33]. Further supporting the memorability issue, a few participants saved their vote identifier despite being explicitly instructed not to do so. This behaviour is well known in usable security research (e.g., password storage) and typically emerges when systems impose excessive cognitive burdens on users. In the context of coercion-resistant voting, storing secrets can undermine security guarantees if coercers gain access to the stored vote identifier. Therefore, similar risks apply to other coercion-resistant systems whenever voters store required secrets. Although our study investigated only one type of secret (i.e., vote identifiers consisting of an image and an associated word), other types of secrets may offer improved memorability. Nevertheless, our findings suggest that memorability remains an inherent challenge affecting the effectiveness of coercion-resistant systems relying on secrets. Future work should therefore investigate alternative designs for such secrets to improve their effectiveness. In the use case where a voter needs to *cast an invalid vote under coercion* to protect their initially true vote, our participants struggle the most, with only 59% completing it. This finding undermines the main motivation for Loki, which is protecting voters against last-minute coercion. However, except for one, all participants fail because they update their vote, indicating issues with understanding the scenario and possible actions. The failure can therefore potentially be explained by misunderstanding of the study instructions rather than lack of system usability. Future work should explore how to improve voters' understanding and why they fail (e.g., whether their behaviour is contingent on specific study instructions or has reasons other than their lack of understanding). Our finding, furthermore, contrasts with a preliminary study, showing a 90% success rate of casting an invalid vote under coercion [20]. However, this study had a different design, including other study instructions and the researcher's physical presence in the study lab. Therefore, these differences highlight the importance of evaluating coercion-resistant systems under varying conditions rather than relying solely on results from controlled lab studies.

RQ2: Effect of risk framing We did not find any significant effects of the risk framing on either of the evaluated metrics, showing that an explicit framing neither help voters to use the system more effectively nor have a significant effect on voters' perceptions of trust, transparency, understandability or user experience. We note, however, that lack of statistical significance finding in our study does not necessarily imply that risk framing does not have any effect. Future work would need to take a closer look at different framings, e.g. evaluating whether changes in the wording such as providing further details about the coercion threat would affect the findings, or conducting qualitative research to better understand the users' reaction to the framing. At the same time, for most of the scenarios, explicit framing performed better (albeit not reaching the level of statistically significant difference) in terms of trust, transparency and understandability, as well as in terms of user experience. The finding confirms the results of previous research, showing a potential positive effect of a transparency-based approach [3] that consists of informing voters about the security guarantees of voting systems.

RQ3: Trust, transparency, understandability, user experience Our participants gave positive scores of trust, transparency, and understandability across all conditions. Similarly, Merino et al. found that voters perceived a coercion-resistant voting system using fake credentials as trustworthy [31]. However, their lack of utilising a validated trust measure limits comparability. Table 5 compares our results with other evaluated voting systems using the same measurement instruments [3, 35]. These systems provide different security properties (e.g., verifiability) and do not support coercion resistance; nevertheless, the comparison suggests that introducing coercion-resistant features does not appear to negatively affect perceived trust, transparency, or understandability relative to systems with alternative security guarantees. As this study is among the first to evaluate trust, transparency, and understandability in a coercion-resistant system using validated measurement instruments, the observed positive scores suggest that coercion-resistant mechanisms may be compatible with user acceptance. While the present work focuses on online voting, this may indicate

Table 5: Comparison of trust, transparency, and understandability scores across evaluated voting systems (7-point Likert scales).

System	Trust	Transp.	Underst.
Loki (variants)	5.03–5.41	5.04–5.43	5.12–5.48
Helios (variants) [3]	4.38–5.11	4.77–5.22	4.76–5.21
Tracking-code system [35]	4.98	5.13	5.37

potential applicability in other coercion-sensitive domains. However, further empirical work is required to assess generalisability across domains with differing threat models and user expectations. For the user experience, our participants gave positive scores of pragmatic quality and attractiveness across all conditions. However, hedonic quality scores were negative in the two of the three *implicitly* conditions. Similarly, other internet voting studies using the UEQ report negative scores for novelty [28, 30] and stimulation [28], belonging under hedonic quality. Although these systems focus on verifiable voting rather than coercion resistance, the pattern suggests that voting-related systems tend to perform better on pragmatic than hedonic aspects of user experience. This pattern may be explained by the domain context: internet voting systems support high-stakes civic processes where functionality, reliability, and institutional legitimacy are prioritised over hedonic qualities such as stimulation or novelty. Consequently, lower hedonic quality scores may be an expected characteristic of such systems rather than a design limitation. Extending this perspective, if coercion-resistant mechanisms are deployed in high-threat environments such as governmental services, similar patterns of comparatively lower hedonic quality may be observed. However, future work is needed to evaluate the generalisability of this effect across application domains.

Limitations Our study was conducted with participants from Prolific, and the sample skewed towards a younger and more educated population. While Prolific studies can be representative of attitudes related to security and privacy [47], and a younger population in particular is known to be first adopters of new election technologies [49], a more representative sample might have resulted in different findings. Our sample furthermore only consisted of UK citizens who are unfamiliar with online voting. Results might therefore be different for voters who are familiar with internet voting (e.g., Estonia). Moreover, participants' behaviour in a study simulating an election might vary from their behaviour during an actual election, both under coercion and in its absence. In a real election, voters might be more likely to read the voting instructions and make an effort to memorise crucial information. However, they also might be less likely to attempt to deceive the coercer if they are not confident in their abilities to do so and are concerned about retaliation. Furthermore, we had a fixed seven-hour interval between the two study phases. However, the time between these phases is likely to vary across voting scenarios in practice, which may limit the generalisability of our findings. Similarly, we only considered two simulated coercion scenarios, without investigating more complex use cases. We furthermore only investigated one possibility of designing the voter identifier, following previous research on

Loki [20]. Other designs (e.g. a different set of visuals) might therefore have a different effect on memorability.

Conclusion We conducted an online study ($N = 172$) to examine the usability of the coercion-resistant internet voting system Loki across three voting scenarios and two risk framings in the user interface. We conclude that Loki is usable for voters who cast a single true vote in an election. However, it is less usable for voters facing more complex situations, such as coercion or the need to update their vote for other reasons. We further find that voters' perceptions of trust, transparency, and understandability in Loki are generally positive, indicating that these factors are unlikely to constitute barriers to its acceptance. Overall, the user experience is also positive. Future research should focus on improving voters' mental models of Loki, especially for voters exposed to coercion after casting their initial preferred vote.

Acknowledgments. The work reported in this paper was supported by Independent Research Fund Denmark, grant number 10.46540/2101-00006B

References

1. Federal Constitutional Court of Germany. Decisions: Or- der of 03 March 2009 - 2 BvC 3/07. `http://www.bundesverfassungsgericht.de/SharedDocs/Entscheidungen/EN/2009/03/cs20090303_2bvc000307en.html` (2009)
2. Acemyan, C.Z., Kortum, P., Oswald, F.L.: The trust in voting systems (tvs) measure. International Journal of Technology and Human Interaction **18**(1) (2022). `https://doi.org/10.4018/IJTHI.293196`
3. Agbesi, S., Budurushi, J., Dalela, A., Nissen, C., Kulyk, O.: How to increase transparency and trust in internet voting systems: An experimental study. In: Proceedings of the 13th Nordic Conference on Human-Computer Interaction (2024)
4. Agbesi, S., Dalela, A., Budurushi, J., Kulyk, O.: "What Will Make Me Trust or Not Trust Will Depend Upon How Secure the Technology Is": Factors Influencing Trust Perceptions of the Use of Election Technologies. E-Vote-ID 2022 p. 1 (2022)
5. Agbesi, S., Dalela, A., Budurushi, J., Kulyk, O.: Investigating transparency dimensions for internet voting. In: International Joint Conference on Electronic Voting. vol. 14230, pp. 1–17. Springer, Germany (2023), `https://link.springer.com/book/10.1007/978-3-031-43756-4`
6. Aranha, D.F., Battagliola, M., Roy, L.: Faster coercion-resistant e-voting by encrypted sorting. Cryptology ePrint Archive, Paper 2023/837 (2023)
7. Bartsch, S., Volkamer, M., Theuerling, H., Karayumak, F.: Contextualized web warnings, and how they cause distrust. In: Huth, M., Asokan, N., Čapkun, S., Flechais, I., Coles-Kemp, L. (eds.) Trust and Trustworthy Computing. pp. 205–222. Springer Berlin Heidelberg, Berlin, Heidelberg (2013)
8. Chaieb, M., Yousfi, S.: Loki vote: A blockchain-based coercion resistant e-voting protocol. In: Themistocleous, M., Papadaki, M., Kamal, M.M. (eds.) Information Systems. EMCIS 2020. vol. 402, pp. 151–168. Springer, Cham (2020)
9. Clark, J., Hengartner, U.: Selections: Internet voting with over-the-shoulder coercion-resistance. In: Danezis, G. (ed.) Financial Cryptography and Data Security. pp. 47–61. Springer Berlin Heidelberg, Berlin, Heidelberg (2012)
10. Cortier, V., Gaudry, P., Goetschmann, A., Lemonnier, S.: Belenios with cast-as-intended: Towards a usable interface. In: Electronic Voting: 9th International Joint Conference, E-Vote-ID 2024. p. 1–19. Springer-Verlag, Berlin, Heidelberg (2024)

11. Cortier, V., Gaudry, P., Yang, Q.: Is the jcj voting system really coercion-resistant? Cryptology ePrint Archive, Paper 2022/430 (2022)
12. Cristiano, L., Longo, R., Spadafora, C.: Click and cast - assessing the usability of vote app. In: E-Vote-ID 2024, pp. 67–84. Gesellschaft für Informatik, Bonn (2024). https://doi.org/10.18420/e-vote-id2024_05
13. Distler, V., Zollinger, M.L., Lallemand, C., Roenne, P.B., Ryan, P.Y.A., Koenig, V.: Security - visible, yet unseen? how displaying security mechanisms impacts user experience and perceived security. In: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems. p. 1–13. CHI '19, Association for Computing Machinery, New York, NY, USA (2019)
14. of Europe, C.: Recommendation cmrec(2017)5[1] of the committee of ministers to member states on standards for e-voting. https://search.coe.int/cm/Pages/result_details.aspx?ObjectId=0900001680726f6f. Last accessed February 5th, 2024 (2017)
15. European Commission: Report on the 2024 elections to the european parliament. Tech. rep., European Commission (2025), https://www.wahlbeobachtung.org/wp-content/uploads/2025/09/ec-swd-report-on-the-2024-elections-to-the-european-parliament.pdf
16. Giustolisi, R., Garjan, M.S., Schuermann, C.: Thwarting last-minute voter coercion. Cryptology ePrint Archive, Paper 2023/1876 (2023)
17. Haines, T., Mueller, J.: How not to VoteAgain: Pitfalls of scalable coercion-resistant e-voting. Cryptology ePrint Archive, Paper 2020/1406 (2020)
18. Haines, T., Müller, J., Querejeta-Azurmendi, I.n.: Scalable coercion-resistant e-voting under weaker trust assumptions. In: Proceedings of the 38th ACM/SIGAPP Symposium on Applied Computing. p. 1576–1584. SAC '23, Association for Computing Machinery, New York, NY, USA (2023)
19. Juels, A., Catalano, D., Jakobsson, M.: Coercion-Resistant Electronic Elections, pp. 37–63. Springer Berlin Heidelberg, Berlin, Heidelberg (2010)
20. Juul Harder, A.X.X., Poulsen, L.Q., Nissen, C.F., Kulyk, O., Giustolisi, R.: Flexible vote updating: Lessons learned from a lab-based study. In: Workshop on Advanced Security in Electronic Voting (2026)
21. Kulyk, O., Neumann, S.: Human factors in coercion resistant internet voting – a review of existing solutions and open challenges (2020)
22. Kulyk, O., Neumann, S., Budurushi, J., Volkamer, M.: Nothing comes for free: How much usability can you sacrifice for security? IEEE Security & Privacy **15**(3), 24–29 (2017)
23. Laugwitz, B., Held, T., Schrepp, M.: Construction and evaluation of a user experience questionnaire. In: Holzinger, A. (ed.) HCI and Usability for Education and Work. pp. 63–76. Springer Berlin Heidelberg, Berlin, Heidelberg (2008)
24. Longo, R., Spadafora, C.: Amun: Securing e-voting against over-the-shoulder coercion. Cryptology ePrint Archive, Paper 2021/851 (2021)
25. Lueks, W., Querejeta-Azurmendi, I., Troncoso, C.: VoteAgain: A scalable coercion-resistant voting system. In: 29th USENIX Security Symposium (USENIX Security 20). pp. 1553–1570. USENIX Association (Aug 2020)
26. Marky, K., Gerber, P., Günther, S., Khamis, M., Fries, M., Mühlhäuser, M.: Investigating State-of-the-Art practices for fostering subjective trust in online voting through interviews. In: 31st USENIX Security Symposium (USENIX Security 22). pp. 4059–4076 (Aug 2022)
27. Marky, K., Kulyk, O., Renaud, K., Volkamer, M.: What did i really vote for? on the usability of verifiable e-voting schemes. In: Proceedings of the 2018 CHI Con-

ference on Human Factors in Computing Systems. p. 1–13. CHI '18, Association for Computing Machinery, New York, NY, USA (2018)

28. Marky, K., Zimmermann, V., Funk, M., Daubert, J., Bleck, K., Mühlhäuser, M.: Improving the usability and ux of the swiss internet voting interface. p. 1–13. CHI '20, Association for Computing Machinery, New York, NY, USA (2020)
29. Marky, K., Zollinger, M.L., Funk, M., Ryan, P.Y.A., Mühlhäuser, M.: How to assess the usability metrics of e-voting schemes. In: Bracciali, A., Clark, J., Pintore, F., Rønne, P.B., Sala, M. (eds.) Financial Cryptography and Data Security. pp. 257–271. Springer International Publishing, Cham (2020)
30. Marky, K., Zollinger, M.L., Roenne, P., Ryan, P.Y.A., Grube, T., Kunze, K.: Investigating usability and user experience of individually verifiable internet voting schemes **28**(5) (Sep 2021)
31. Merino, L.H., Azhir, A., Zhang, H., Colombo, S., Tellenbach, B., Estrada-Galiñanes, V., Ford, B.: E-vote your conscience: Perceptions of coercion and vote buying, and the usability of fake credentials in online voting. In: 2024 IEEE Symposium on Security and Privacy (SP). pp. 3478–3496 (2024)
32. Neto, A.S., Leite, M., Araújo, R., Mota, M.P., Neto, N.C.S., Traoré, J.: Usability considerations for coercion-resistant election systems. In: Proceedings of the 17th Brazilian Symposium on Human Factors in Computing Systems. IHC '18, Association for Computing Machinery, New York, NY, USA (2018). https://doi.org/10.1145/3274192.3274232
33. Nissen, C., Hilt, T., Budurushi, J., Volkamer, M., Kulyk, O.: Voting under pressure: Perceptions of counter-strategies in internet voting. In: International Joint Conference on Electronic Voting. Springer Nature Switzerland (2025), https://e-vote-id.org/
34. Nissen, C.F., Kulyk, O., Giustolisi, R.: Loki supplementary materials (2026), https://anonymous.4open.science/r/Supplementary_Materials_Flexible_Vote_Updating-3E32/
35. Nissen, C., Kulyk, O., Volkamer, M., Fredrich, L., Hermansen, H.: Tracking code-based verification-design and evaluation. In: Proceedings of the Ninth International Joint Conference on Electronic Voting (E-Vote-ID 2024). Gesellschaft für Informatik (2024)
36. Office of the United Nations High Commissioner for Human Rights: Human Rights and Elections: A Handbook on International Human Rights Standards on Elections. United Nations (2021), https://www.un-ilibrary.org/content/books/9789214030522
37. Paivio, A.: Dual coding theory: Retrospect and current status. Canadian Journal of Psychology / Revue canadienne de psychologie **45**(3), 255–287 (1991). https://doi.org/10.1037/h0084295
38. Paivio, A., Rogers, T.B., Smythe, P.C.: Why are pictures easier to recall than words? Psychonomic Science **11**, 137–138 (1968)
39. Potrafke, N., Roesel, F.: Opening hours of polling stations and voter turnout: Evidence from a natural experiment. The Review of International Organizations **15**(1), 133–163 (January 2020). https://doi.org/10.1007/s11558-018-9305-8, https://ideas.repec.org/a/spr/revint/v15y2020i1d10.1007_s11558-018-9305-8.html
40. Prolific: Prolific participant recruitment platform. https://www.prolific.com (2025), accessed: 2025-04-03
41. Prolific: Prolific's payment principles. https://researcher-help.prolific.com/en/article/2273bd (2025), https://researcher-help.prolific.com/en/article/2273bd, accessed: 2026-04-03

42. Prolific: Prolific's attention and comprehension check policy (April 2026), https://researcher-help.prolific.com/en/article/fb63bb, accessed: 2026-04-28
43. R Core Team: R: A language and environment for statistical computing (2025), https://www.r-project.org/, accessed: 2025-08-18
44. RStudio Team: Rstudio: Integrated development environment for r (2025), https://www.rstudio.com/, accessed: 2025-08-18
45. Spadafora, C., Longo, R., Sala, M.: Coercion-resistant blockchain-based e-voting protocol. Cryptology ePrint Archive, Paper 2020/674 (2020)
46. SurveyXact (Rambøll / Xact by Rambøll): Surveyxact (2025), https://www.survey-xact.dk/, accessed: 2025-08-18
47. Tang, J., Birrell, E., Lerner, A.: Replication: How well do my results generalize now? the external validity of online privacy and security surveys. In: Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022). pp. 367–385. USENIX Association (Aug 2022), https://www.usenix.org/conference/soups2022/presentation/tang
48. UEQ Team: User experience questionnaire (ueq) (2024), https://www.ueq-online.org/, accessed: 2026-03-17
49. Vassil, K., Solvak, M., Vinkel, P., Trechsel, A.H., Alvarez, R.M.: The diffusion of internet voting. usage patterns of internet voting in estonia between 2005 and 2015. Government Information Quarterly **33**(3), 453–459 (2016)
50. Volkamer, M., Kulyk, O., Ludwig, J., Fuhrberg, N.: Increasing security without decreasing usability: A comparison of various verifiable voting systems. In: Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022). pp. 233–252 (Aug 2022)
51. Volkamer, M., Spycher, O., Dubuis, E.: Measures to establish trust in internet voting. In: Proceedings of the 5th International Conference on Theory and Practice of Electronic Governance. pp. 1–10 (2011)
52. Yin, Z., Zhang, B., Nastenko, A., Oliynykov, R., Ren, K.: A scalable coercion-resistant blockchain decision-making scheme. Cryptology ePrint Archive, Paper 2023/1578 (2023)
53. Zaghloul, E., Li, T., Ren, J.: Anonymous and coercion-resistant distributed electronic voting. In: 2020 International Conference on Computing, Networking and Communications (ICNC). pp. 389–393 (2020)
54. Zollinger, M.L., Rønne, P.B., Schneider, S., Ryan, P.Y.A., Jamroga, W.: Intelligo ut confido: Understanding, trust and user experience in verifiable receipt-free e-voting. In: Electronic Voting: 9th International Joint Conference, E-Vote-ID 2024. pp. 158–174. Springer (2024). https://doi.org/10.1007/978-3-031-72244-8_10