

# Preventing Coercion in E-Voting: Be Open and Commit

Wojciech Jamroga<sup>1</sup> and Masoud Tabatabaei<sup>2</sup>

<sup>1</sup> Institute of Computer Science, Polish Academy of Sciences

<sup>2</sup> Interdisciplinary Centre for Security and Trust, University of Luxembourg  
w.jamroga@ipipan.waw.pl, masoud.tabatabaei@uni.lu

**Abstract.** We present a game-theoretic approach to coercion-resistance from the point of view of an honest election authority that chooses between various protection methods with different levels of resistance and different implementation costs. We give a simple game model of the election and propose a preliminary analysis. It turns out that, in the games that we look at, Stackelberg equilibrium for the society does not coincide with maxmin, and it is always more attractive to the society than Nash equilibrium. This suggests that the society is better off if the security policy is publicly announced, and the authorities commit to it.

## 1 Introduction

It was recognised early on in the history of voting that *ballot privacy* is an essential property of voting systems to counter threats of coercion or vote buying. More recently, cryptographers and security experts have been looking at using cryptographic mechanisms to provide *voter-verifiability*, i.e. the ability of voters to confirm that their votes are correctly registered and counted. This strengthens to integrity properties, but, if it is not done carefully, new threats to ballot secrecy can be introduced. The observation lead to the introduction of more sophisticated privacy-style notions: *receipt-freeness* and *coercion-resistance*. The latter is the strongest property and can be defined informally as: a voting system provides coercion-resistance if the voter always has a strategy to vote as they intend while appearing to comply with all the coercer's requirements. The coercer is assumed to be able to interact with the voter throughout the voting process: before, during and after.

Achieving coercion-resistance is extremely challenging, especially in the context of internet and remote voting (e.g. postal). A number of schemes have been proposed that provide it, but typically this comes at a cost, in particular in terms of usability. In this paper, we take a game theoretic approach to analyse the trade-offs between the costs of implementing coercion-resistance mechanisms on the one hand, and on the other hand the cost the to society regarding the threats to the legitimacy of the outcome due to coercion attacks.

Unlike most existing papers, we neither propose a new coercion-resistant voting scheme nor prove that a scheme is secure in that respect. Instead, we

focus on the context of coercion attempts in e-voting, namely costs and benefits of involved parties. The main question is: *Should the society invest in protection against coercion attempts, and if so, in what way?*. We do not aim at devising a secure voting procedure, but rather at exposing conditions under which security of a procedure is relevant at all.

Our game models rely on several simplifying assumptions. We do not represent ballot privacy explicitly, and we do not investigate its relation to coercion. Furthermore, we do not differentiate on different coercion scenarios. Instead, we model the level of coercion attempts and coercion-resistance as simple scalars. The former refers to how many voters the coercer(s) attempt to coerce, i.e., indicates the scale of coercion in the election. The latter indicates how much effort/cost is needed to break the protection measures. Although an actual voting system might consist of a set of authorities with possibly different interests, we assume a single agent that we call the “election authority” whose interests are in line with what we consider “the common good of the society”. This agent’s interests might or might not represent the preferences of the actual authorities of the election. But by modelling it this way we can study the question of what strategy should the authorities collectively choose, *if* they want to benefit the society as a whole. Finally, we assume that all the potential coercers fully cooperate so that they can be represented by a single “coercer” player. Thus, the scenario can be modelled as a two-player game with largely conflicting incentives.

**Related work on preventing coercion in elections.** The related work can be roughly divided into three strands: definitions of coercion-resistance and its relation to privacy, proposals of coercion-resistant voting procedures, and studies of the context of coercion-resistance. The notion of coercion-resistance was first introduced in [11]. In [6], a formalization of coercion-resistance was proposed, and its relation to receipt-freeness and privacy was studied. [8] gave a formal definition of coercion-resistance for the end-to-end voting schemes. In [15], a game-based cryptographic definition of coercion-resistance was proposed.<sup>3</sup> The same authors added a game-based cryptographic definition of privacy in [16], and showed that the relationship between privacy and coercion-resistance can be more subtle than it is normally assumed. [7] provided formal definitions of various privacy notions in applied pi calculus, and showed how they are related to each other. Finally, [10] used CSP to fit a wide range of definitions and properties given in the literature for coercion-resistance.

The second strand overlaps with the first: [11, 8] all propose voting protocols that satisfy their definitions of coercion-resistance while [15] proves coercion-resistance of two previously existing protocols. Another coercion-resistant voting scheme was introduced in [2]. Several other papers proposed voting schemes which provably satisfy privacy as an intuitive argument for coercion-resistance,

---

<sup>3</sup> The definition was game-based in the technical sense, i.e., the security property was defined as the outcome of an abstract game between the “verifier” and the “adversary”. In this paper, we use game models to study the interaction between the actual participants of the protocol.

cf. e.g. [24]. Several works such as [21, 18, 1, 17, 14, 28, 3, 25] have developed weaker, more practical or more efficient ways to realize coercion-resistance.

Putting coercion-resistance in a broader economic or social context has been, to our best knowledge, largely left untouched. The only paper in this strand that we are aware of is [5]. The authors compare two voting systems using game models, more precisely zero-sum two-player games based on attack trees. Two actions are available for the attacker (performing the attack or not); the authority is presumably choosing one of the two voting systems. The utility of the attacker is the expected probability of successful coercion minus the expected probability of being caught. The value is computed for the two systems using empirical data. In contrast, we consider a more general game where coercion – and resistance measures – come at a *cost* (instead of simply assuming probability distributions for the possible events), and we look for the rational choices of the players using game-theoretic solution concepts. We also argue that the coercion game is not zero-sum, with important consequences for the best policy to be chosen.

**Game-based analysis of similar application domains.** Our analysis is based on two game-theoretic solution concepts: Nash equilibrium and Stackelberg equilibrium. Nash equilibrium corresponds to the behaviour of players that should emerge “organically” when they adapt to the behaviour observed from the other players over a period of time. It is often used to analyse how the policies of multiple interacting users are likely to converge in the long run. The typical application is to so called *energy games* where dynamic pricing schemes are proposed and studied in order to balance the supply and demand of electrical energy in a small-scale distributed market, cf. e.g. [26, 23, 31, 20].

Stackelberg equilibrium corresponds to a scenario where a designated “leader” commits openly to a selected strategy and thus forces the response from the other players. Stackelberg games have become very popular in design and analysis anti-terrorist policies [12, 13, 27, 29]. Our study comes close to that line of research, but differs in two important ways. First, anti-terrorist games focus on protection of multiple tangible resources (planes, airport buildings, etc.), while our coercion games address protection of “the good of the society” as a whole. Secondly, because of the inherent differences between the two application domains, we only use Stackelberg equilibria in pure (deterministic) strategies, whereas the main solutions in the research on strategic prevention of terrorism are based on mixed (randomised) Stackelberg strategies.

Finally, we mention [30] that applies Stackelberg games to prevent manipulation of elections, but its focus is on the computational complexity of preventing Denial of Service type attacks.

## 2 Game-Theoretic Preliminaries

In this paper, we propose a preliminary game-theoretic analysis of coercion prevention in an election. The main idea is to model the election as a simple strategic game between the society and coercer(s). We begin by a gentle introduction to

<i>Bob \ Sue</i>	<i>Bar</i>	<i>Home</i>	<i>Theater</i>
<i>Bar</i>	<b>3, 2</b>	<u>2</u> , 1	<b>1</b> , 0
<i>Theater</i>	<u>4</u> , 0	0, 0	<span style="border: 1px solid black;">2, 3</span>

**Fig. 1.** Example two-player strategic game. The only Nash equilibrium is indicated by the black frame, maxmin for Bob is highlighted in bold, and Stackelberg equilibrium for Bob is set on yellow background. The players’ best responses to the other player’s strategies are underlined.

the basic concepts of noncooperative game theory. A more detailed exposition can be found in numerous textbooks, cf. e.g. [22, 19, 4].

## 2.1 Strategic Games

**Definition 1 (Strategic game).** A strategic game (called also normal form game) is a tuple  $\Gamma = (N, \Sigma, u)$ , consisting of:

1. a finite set of agents or players  $N = \{A_1, \dots, A_{|N|}\}$ ,
2. a set of strategy profiles  $\Sigma = \Sigma_{A_1} \times \dots \times \Sigma_{A_{|N|}}$ , where  $\Sigma_{A_i}$  collects the available strategies of player  $A_i \in N$ ,
3. a utility profile  $u = \{u_1, \dots, u_{|N|}\}$  with  $u_i : \Sigma \rightarrow \mathbb{R}$  being the utility function of player  $A_i$  that assigns the “payoffs” of  $A_i$  to strategy profiles.

When needed, we will refer to  $A_i$ ’s part of strategy profile  $\sigma$  by  $\sigma_i$ , and to the other players’ part of the profile by  $\sigma_{-i}$ .

A strategic game captures a “bird’s-eye view” of interaction, where  $A_i$ ’s strategies represent her possible behaviours in a game. Strategies are treated as atomic: we are not interested in their internal structure, and can as well view them as simple actions. The combined behaviour of all the players is represented by a strategy profile, i.e., a tuple of individual strategies. Given a strategy profile  $\sigma$ ,  $u_i(\sigma)$  defines how much the outcome of the game is “worth” to player  $A_i$ . Thus, the utility profile is meant to represent the incentives (or preferences) of each player. An example strategic game – a slightly modified variant of the “Battle of the Sexes” – is shown in Figure 1. Two players (Bob and Sue) are choosing in parallel whether to go to the local bar, or to the theater. The strategies and utilities of Sue are set in grey font for better readability.

When modelling interaction by a strategic game, we implicitly assume *complete information*, i.e., that the structure of the game is common knowledge among the players. In particular, players know each others’ preferences and the available actions of the opponents. Especially the former assumption is often unrealistic. We will come back to this issue and relax the assumption in Section 4.

## 2.2 Solution Concepts

In game theory, *solution concepts* are used to define which collective behaviours are “rational” and should (or may) be selected by players in the game. Formally,

a solution concept maps each game to a subset of strategy profiles. Different solution concepts encode different assumptions about the deliberation process that leads to selecting one or another strategy. In this paper, we compare the predictions obtained by three solution concepts: Nash equilibrium, maxmin, and Stackelberg equilibrium, presented briefly below.

**Nash equilibrium.** A strategy profile  $\sigma$  is a *Nash equilibrium* if it is stable under unilateral deviations of players, i.e., if player  $A_i$  changed her part of  $\sigma$  (and the other players stuck to their strategies) then the payoff of  $A_i$  would decrease or stay the same. Formally, for every  $A_i \in N$  and  $\sigma'_i \in \Sigma_i$ , it must hold that  $u_i(\sigma) \geq u_i(\sigma'_i, \sigma_{-i})$ . Equivalently,  $\sigma$  is a Nash equilibrium if each  $\sigma_i$  is the best response to  $\sigma_{-i}$ . As an example, consider the game in Figure 1. The strategy profile  $(Theater, Bar)$  is not stable because Sue can improve her payoff from 0 to 3 by changing her strategy to *Theater*. On the other hand,  $(Theater, Theater)$  is stable because both players can only lose when they change their minds: Bob would then decrease his utility from 2 to 1, and Sue analogously from 3 to 0.

Intuitively, Nash equilibrium represents a collective behaviour that can emerge when players play the game multiple times, and adapt their choices to what they expect from the other players. Thus, it captures the “organic” emergence of behaviour through a sequence of strategy adjustments from different players that leads to a point when nobody is tempted to change their strategy anymore.

**Maxmin.** *Maxmin for player  $A_i$*  aims at the largest value that the player can ensure regardless of what the other players do. Formally, it is the strategy profile  $\sigma^*$  such that  $\sigma_i^* = \operatorname{argmax}_{\sigma_i} \min_{\sigma_{-i}} u_i(\sigma_i, \sigma_{-i})$  and  $\sigma_{-i}^* = \operatorname{argmin}_{\sigma_{-i}} u_i(\sigma_i^*, \sigma_{-i})$ . Intuitively, maxmin captures decision making of “paranoid” agents who always look at the worst possible outcome of their choices.

The maxmin for Bob in Figure 1 is  $(Bar, Theater)$ , since playing *Bar* guarantees the payoff of at least 1 to Bob, while playing *Theater* may obtain 0.

**Stackelberg equilibrium.** Finally, *Stackelberg equilibrium for player  $A_i$*  represents rational play in 2-player games where a designated player (the *leader*) makes her choice first. Formally, it is the strategy profile  $\sigma^*$  for which  $\sigma_i^* = \operatorname{argmax}_{\sigma_i} u_i(\sigma_i, \operatorname{argmax}_{\sigma_{-i}} u_{-i}(\sigma_i, \sigma_{-i}))$  and  $\sigma_{-i}^* = \operatorname{argmax}_{\sigma_{-i}} u_{-i}(\sigma_i^*, \sigma_{-i})$ . That is, for every strategy  $\sigma_i$  of the leader we find the response  $\operatorname{resp}(\sigma_i)$  that maximizes the utility of the opponent; then, we select the  $\sigma_i$  which maximizes  $u_i(\sigma_i, \operatorname{resp}(\sigma_i))$ . In our example, *Bar* is Sue’s best response to Bob’s strategy *Bar*, and *Theater* is Sue’s best response to Bob’s *Theater*. Thus, the Stackelberg equilibrium is  $(Bar, Bar)$  because it obtains 3 for Bob, whereas  $(Theater, Theater)$  obtains only 2.

Intuitively, analysis based on Stackelberg equilibrium assumes that the leader can either execute her strategy before the other player, or irrevocably commit to her choice. Moreover, the choice of  $\sigma_i$  becomes common knowledge before the opponent chooses his strategy. Such commitment is typically possible in case of public institutions and agencies that can commit to a chosen policy through suitable legislation. Note that, when Stackelberg equilibrium coincides with maxmin, it is actually irrelevant for the leader whether her choice will be known to the opponent or not. Conversely, when Stackelberg equilibrium is

different from maxmin, the leader is better off *publicly committing to her policy*, because this way she forces the other player to respond in a desirable way.

### 2.3 Pure vs. Mixed Strategies

So far, we have mentioned only *pure strategies* of players, i.e., the choices explicitly given in sets  $\Sigma_{A_i}$  of the game model. More sophisticated behaviour of players can be represented by so called *mixed strategies* that model randomized play. Formally, a mixed strategy for player  $A_i$  is a probability distribution over  $\Sigma_{A_i}$ , with the idea that the player will randomize her choice according to that distribution. A *mixed strategy profile* is a combination of mixed strategies, one per player. Note that such a strategy profile uniquely determines a joint probability distribution over  $\Sigma$  (assuming that individual probability distributions are independent), and hence also the expected utility of each player. Thus, each normal form game induces an infinite payoff table where the rows and columns are given by the mixed strategies, and the cells contain vectors of the expected utility values. This way, solution concepts like Nash equilibrium, maxmin, and Stackelberg equilibrium are easily extended to analysis of randomized play.

Randomization makes it harder for the opponents to predict the player’s next action, and to exploit the prediction. Moreover, the importance of randomized strategies in game theory stems from the fact that Nash equilibrium is guaranteed to exist in mixed strategies, whereas no such guarantee applies to pure strategies. We notice that Stackelberg equilibrium in mixed strategies, while theoretically elegant, is often questionable in practice. This is because the leader’s commitment to her strategy must be believable to the opponent. However, commitment to a randomized strategy is hard to verify unless the game is played very frequently. This condition is satisfied, e.g., in case of anti-terrorist policies for deployment of air marshals on domestic flights [12], with multiple flights every day. On the other hand, elections are run way too infrequently to achieve the same effect. Thus, we will limit our analysis of Stackelberg equilibrium to pure strategies of the leader.

We also note that all but one of our coercion models have Nash equilibria and maxmins in pure strategies.

## 3 A Simple Game Model of Coercion

Consider an election with a set of candidates  $\Omega = \{\omega_1, \dots, \omega_g\}$  and a set of  $n$  voters. We model the election as a strategic game  $\langle \{A, C\}, \Sigma, (u_A, u_C) \rangle$ , where  $\Sigma = \Sigma_A \times \Sigma_C$  with the ingredients defined below.

### 3.1 Players, Strategies, Utilities

**Players.**  $A$  and  $C$  are the players. Player  $A$  is an honest election authority who acts on behalf of the society. We assume that the goal of  $A$  is in line with “the good of the society” as a whole.  $A$  has no preference for any of the candidates,

and tries to make the result of the election as close as possible to the result when no coercion occurs, i.e., when the voters vote according to their own preferences.

Player  $C$  represents the coercer. The coercer tries to change the result of the election by threatening or bribing voters in order to make them vote according to his plan, rather than to the voters' own preferences over the candidates. In general, several coercers can try to change the result of the election simultaneously. We adopt the worst case assumption that they all collude, and hence may be represented by actions and preferences of a single player  $C$ .

Note that we do not consider candidates and voters as players in the game, but rather as parameters of the model.

**Strategies.**  $\Sigma_A = \{\alpha_0, \dots, \alpha_{\text{Max}}\}$  is the set of protection methods that can be implemented by the election authority  $A$ . These represent the protection measures that can prevent, or make it harder for the coercer to discover the actual values of votes. It is assumed that  $\alpha_0$  represents the case of no protection.

$\Sigma_C = \{0, \dots, n^*, \dots, n\}$  is the set of strategies for  $C$ , indicating the number of voters that the coercer attempts to bribe or threaten to bribe according to his wish. The minimal number of voters that the coercer needs to coerce in order to change the result of the election in his favor is  $n^*$ . We assume that the value of  $n^*$  is common knowledge; we will relax the assumption in Section 4.

**Preferences.** Preferences are represented by utility functions over possible combinations of strategies. We define the utility of the election authority  $A$  as  $u_A(\alpha_j, k) = v_A(\text{out}(\alpha_j, k)) - \text{imp}(\alpha_j)$  where:

- $\text{imp}(\alpha_j)$  is the cost of implementing the protection method  $\alpha_j$ . We assume that  $\text{imp}(\alpha_0) = 0$ , and  $t < t'$  implies  $\text{imp}(\alpha_t) \leq \text{imp}(\alpha_{t'})$ .
- $\text{out}(\alpha_j, k)$  is the outcome of the election when  $A$  implements  $\alpha_j$  and  $C$  attempts to coerce  $k$  voters.
- $v_A(\omega)$  is the social value of the election outcome  $\omega$ . We assume that  $v_A(\omega) = v_A^*$  if the outcome of the election is the same as it would be without coercion, and  $v_A(\omega) = v_A^* - \epsilon_A$  otherwise. Moreover,  $\epsilon_A > \text{imp}(\alpha_i)$  for all  $\alpha_i \in \Sigma_A$ .

The utility of the coercer is  $u_C(\alpha_j, k) = v_C(\text{out}(\alpha_j, k)) - k \cdot \text{cost}_C(\alpha_j)$  where:

- $v_C(\omega)$  is the value of the election outcome  $\omega$  from the coercer's point of view. We assume that  $v_C(\omega) = v_C^*$  if the outcome of the election is in favor of the coercer, and  $v_C(\omega) = v_C^* - \epsilon_C$  otherwise, for some  $\epsilon_C > 0$ .
- $\text{cost}_C(\alpha_j) = d_C(\alpha_j) + \beta_C$  is the total cost that the coercer must bear when coercing one voter, where  $d_C(\alpha_j)$  is the cost of overcoming the protection method, and  $\beta_C$  is the bribing cost. We assume that  $d_C(\alpha_0) = 0$  and  $d_C(\alpha_j)$  increases with  $j$ . Moreover,  $\beta_C$  is constant.

We also assume that at least the strongest protection method  $\alpha_{\text{Max}}$  induces so high costs of coercion that effective coercing becomes unprofitable, formally:  $\text{cost}_C(\alpha_{\text{Max}}) \cdot n^* > \epsilon_C$ .

We will consider two possible settings for the coercion game. In Section 3.2, we assume that a perfect protection method is available to  $A$ , and if it is used then any coercion attempt will inevitably fail. In Section 3.3, we analyze the

$A \setminus C$	0	$n^*$
$\alpha_0$	$v_A^*, v_C^* - \epsilon_C$	$v_A^* - \epsilon_A, v_C^* - \beta_C \cdot n^*$
$\alpha_1$	$v_A^* - \mathbf{imp}(\alpha_1), v_C^* - \epsilon_C$	$v_A^* - \mathbf{imp}(\alpha_1), v_C^* - \epsilon_C - \beta_C \cdot n^*$

**Fig. 2.** Game model for perfect protection. The maxmin profiles and the Stackelberg equilibrium for A are shown. The game has no Nash equilibrium in pure strategies.

other variant where any protection method can be broken if the coercer invests enough money and effort.

### 3.2 Coercion against Perfect Protection

We first study the case where the election authority has a choice between no protection (strategy  $\alpha_0$ ) and perfect protection against coercion ( $\alpha_1$ ). When A plays  $\alpha_1$  then the coercer cannot change the result of the election no matter how many voters he attempts to bribe, as there is no way for him to verify the values of the votes. Therefore the utility of the coercer in this case is  $v_C^* - \epsilon_C - k \cdot \beta_C$ , where  $k$  is the number of voters he attempts to bribe. We assume that a coercion attempt is successful only if the coercer can verify the votes.

Note that, for player C, the strategies 1 to  $n^* - 1$  are all dominated by strategy 0. That is, C gets a higher payoff playing 0 no matter what the other player chooses. In consequence, they never belong to any rational solution, and can be omitted from the game table. Similarly, the coercer's strategies from  $n^* + 1$  to  $n$  are dominated by strategy  $n^*$ . Thus, it suffices to consider only choices 0 and  $n^*$ . The resulting game table is shown in Figure 2. In all the strategic games from now on, we will underline the best response strategies of both players', indicate Nash equilibria by putting them in black frames, highlight the maxmin for A by bold font, and point out the Stackelberg equilibrium for A by the yellow background.

The game has no Nash equilibrium in pure strategies. The unique Nash equilibrium in randomized strategies is as follows: the authority chooses "no protection" with probability  $p = \frac{\beta_C \cdot n^*}{\epsilon_C}$  and "perfect protection" with probability  $1 - p$ , whereas the coercer attempts to coerce  $n^*$  voters with probability  $q = \frac{\mathbf{imp}(\alpha_1)}{\epsilon_A}$  and 0 voters with probability  $1 - q$ . This yields the expected utility of  $v_A^* - \mathbf{imp}(\alpha_1)$  for the society. The maxmin for A is strategy  $\alpha_1$  which provides exactly the same payoff for the society, and the same holds for the Stackelberg equilibrium. Thus, it does not matter whether A adapts to the coercer's strategy (i.e., plays the Nash equilibrium), publicly commits to strategy  $\alpha_1$  of maximal protection method (i.e., plays the Stackelberg equilibrium), or simply chooses  $\alpha_1$  and sticks to it (i.e., follows the maxmin).

### 3.3 Coercion Game for Breakable Protection

The analysis in Section 3.2 did not bring very interesting conclusions, but the assumption of a perfect protection method was not very realistic either. From now

$A \setminus C$	0	$n^*$
$\alpha_0$	$v_A^*, v_C^* - \epsilon_C$	$\underline{v_A^* - \epsilon_A, v_C^* - \text{cost}_C(\alpha_0) \cdot n^*}$
$\alpha_{m^*-1}$	$v_A^* - \text{imp}(\alpha_{m^*-1}), v_C^* - \epsilon_C$	$v_A^* - \text{imp}(\alpha_{m^*-1}) - \epsilon_A, \underline{v_C^* - \text{cost}_C(\alpha_{m^*-1}) \cdot n^*}$
$\alpha_{m^*}$	$\underline{v_A^* - \text{imp}(\alpha_{m^*}), v_C^* - \epsilon_C}$	$v_A^* - \text{imp}(\alpha_{m^*}) - \epsilon_A, v_C^* - \text{cost}_C(\alpha_{m^*}) \cdot n^*$
$\alpha_{\text{Max}}$	$v_A^* - \text{imp}(\alpha_{\text{Max}}), \underline{v_C^* - \epsilon_C}$	$v_A^* - \text{imp}(\alpha_{\text{Max}}) - \epsilon_A, v_C^* - \text{cost}_C(\alpha_{\text{Max}}) \cdot n^*$

**Fig. 3.** Game model for breakable protection

on, we will assume that the election authority can implement several alternative protection methods, none of them fully coercion-proof. In other words, the coercer can successfully coerce against any protection method. The costs of both  $A$  and  $C$  increase with implementation of (resp. coercion against) more advanced methods. As before, we assume that the structure of the game is common knowledge, in particular, the value of  $n^*$  (the amount of voters needed to be coerced in order to change the result of the election in favour of the coercer) is known to both players. The resulting strategic game is depicted in Figure 3. Similarly to Figure 2, we omit dominated strategies from the table for better readability. Best responses, maxmin, Nash equilibrium, and Stackelberg equilibrium are indicated in the same way as before.

Like in the previous game model, the only undominated strategies for  $C$  are 0 and  $n^*$ , i.e., it makes only sense to coerce either 0 or  $n^*$  voters. Moreover, as the authority changes the protection method from  $\alpha_0$  to  $\alpha_{\text{Max}}$ , the difficulty of coercing for the coercer increases. For a given  $\alpha$ , if  $v_C^* - \text{cost}_C(\alpha) \cdot n^*$  is larger than  $v_C^* - \epsilon_C$  then  $C$  prefers coercing over not coercing. Note that, from some protection method  $\alpha_m$  on, the cost of coercing for the coercer is more than  $\epsilon_C$ . In that case the coercer, although being able to coerce successfully, prefers not to tamper with the election. It is easy to observe the following.

**Theorem 1.** *For the coercion game with breakable protection, the Nash equilibrium and the maxmin for  $A$  is  $(\alpha_0, n^*)$ , whereas the Stackelberg equilibrium for  $A$  is  $(\alpha_m, 0)$ . Moreover,  $u_A(\alpha_0, n^*) < u_A(\alpha_m, 0)$ .*

The unique Nash equilibrium in pure strategies is  $(\alpha_0, n^*)$ : the coercer attempts to coerce sufficiently many voters, and the authority chooses the cheapest protection method, leaving the election open to manipulation. Thus, when the players mutually adapt to each other's play, the outcome is clearly undesirable for the society. The Stackelberg equilibrium  $(\alpha_m, 0)$  is much better in this respect: the authority invests in the minimal sufficient protection that makes coercion unprofitable, and the coercer gives up coercion. Thus,  $A$  should choose its strategy in advance and stick to it, without adapting to  $C$ 's play. Moreover, the maxmin for  $A$  in the game coincides with the Nash equilibrium and not the Stackelberg equilibrium, so in order to end up in the latter, the authority must publicly and believably commit to strategy  $\alpha_m$ .

## 4 Coercion with Incomplete Information

In the previous section, we assumed that the players have complete information about the structure of the game. In many scenarios the assumption is not realistic, as players are not certain about some aspects of the game they are playing. For example, they may be uncertain about the available strategies of other players, their preferences, etc. We have deliberately defined the utility functions  $u_A, u_C$  based on several basic parameters instead of fixing concrete utility values, and specified as few constraints as possible about the relationships between the parameters. Since our results hold for *all* the admissible values of the parameters, our conclusions are valid even if the players do not know the exact numerical values.<sup>4</sup> By and large, this seems a justifiable level of abstraction except for one point: typically, neither the election authority nor the coercer will know the *precise* number of voters that need to be coerced in order to swing the outcome of the election. The coercer is also unlikely to know exactly which voters are the right targets of coercion (for instance, it makes little sense to coerce voters that plan to vote for the coercer’s favourite candidate). What the players know instead is some probabilistic information, obtained e.g. from pre-election polls. We incorporate the observation in this section and extend our game model to include probabilistic uncertainty of the players about the  $n^*$  parameter.

Formally, we will model the uncertainty by assuming that the players take into account not one, but a set  $T$  of strategic games for different possible values of  $n^*$ . The current belief of each player is represented by a probability distribution over  $T$ , and possibly also over the probability distributions held as beliefs by the other players. Such models are known as *Bayesian games*. Again, we refer the interested reader to [19, 4] for details.

In what follows, we assume that the coercer and the election authority hold the same beliefs about  $n^*$  (represented by the same probability distribution). In general, this may not be true, but in the case of an election the players’ beliefs are usually based on public opinion polls which are equally accessible to everyone. Thus, the assumption seems acceptable in our application domain. At the same time, it greatly simplifies the analysis, as we will only need to take into account the players’ factual beliefs, and not their beliefs about each others’ beliefs, beliefs about beliefs about beliefs, and so on.

### 4.1 Bayesian Game for Coercion

We consider the Bayesian game  $\langle \{A, C\}, \Omega, \Sigma, T, \tau, p, (\hat{u}_A, \hat{u}_C) \rangle$  with the following elements:

**Players and strategies.** The sets of players and their available strategies are defined as before (cf. Section 3).

**States of the world.**  $\Omega = \{1, \dots, n\}$  is the set of possible *states of the world* (sometimes also called *states of nature*). In our scenario, each state of the world

---

<sup>4</sup> It suffices that the constraints are common knowledge among the players.

corresponds to one possible value of  $n^*$ , i.e., the number of voters needed to be coerced to swing the outcome of the election. Note that the same strategies are available to players in all states of the world.

**Preferences.**  $\hat{u}_A, \hat{u}_C : \Omega \times \Sigma \rightarrow \mathbb{R}$  are utility functions of the players. The only difference to the complete information setting is that  $\hat{u}_{A_i}(n^*, \alpha_j, k)$  depends not only on the strategy profile  $(\alpha_j, k)$ , but also on the actual value of  $n^*$ .

**Player types and signaling.** In Bayesian games, the set of *type profiles*  $T = T_A \times T_C$  is used to construct higher-order beliefs of players, i.e., beliefs about beliefs etc. We define  $T_A = \{t_A\}$  and  $T_C = \{t_C\}$ . That is, players' uncertainty about each others' beliefs is irrelevant. The *signaling functions*  $\tau_A : \Omega \rightarrow T_A$  and  $\tau_C : \Omega \rightarrow T_C$  are trivial and can be also omitted from our analysis.

**Players' beliefs.** The probabilistic beliefs of  $A$  and  $C$  are represented by a single probability distribution  $p \in \Delta(\Omega)$  over the states of nature. In this work, we consider two cases of such probabilistic beliefs, based on the uniform distribution (Section 4.2) and the normal distribution (Section 4.3). Although the values of  $n^*$  are discrete, when the number of voters is large we can use continuous probability distributions to estimate the probability of different intervals of  $n^*$ .

**Solution concepts.** In order to apply solution concepts to Bayesian games, we use the standard transformation into strategic games [9]. That is, we transform the Bayesian game  $\langle \{A, C\}, \Omega, \Sigma, T, \tau, p, (\hat{u}_A, \hat{u}_C) \rangle$  into a strategic game  $\langle \{A, C\}, \Sigma, (u_A, u_C) \rangle$  such that, for every strategy profile  $s \in \Sigma$ ,

$$u_A(s) = \mathbf{E}_{\omega \in \Omega}[\hat{u}_A(\omega, s)] \quad \text{and} \quad u_C(s) = \mathbf{E}_{\omega \in \Omega}[\hat{u}_C(\omega, s)].$$

## 4.2 Uniform Probabilistic Beliefs

Our first approach is to assume the players' beliefs in the form of a uniform probability distribution in range  $[n_a, n_b]$ , where  $0 \leq a \leq b \leq n$ . Thus, we assume that  $A$  and  $C$  can rule out some values of  $n^*$ , but apart from that they consider all the possible states of nature equally likely. In order to transform the model to a strategic game, we need to compute  $u_A(\alpha, n)$  and  $u_C(\alpha, n)$  for a protection method  $\alpha$  and the number of voters to coerce  $k$ .

**Utility of the coercer.** We consider three ranges for  $k$  and compute  $u_C(\alpha, k)$  in each range separately:

- If  $k < a$  then in all states of the nature  $k < n^*$ , therefore:

$$u_C(\alpha, k) = v_C^* - \epsilon_C - k \cdot \text{cost}_C(\alpha).$$

In this range the strategy 0 is the best response of player  $C$ . By choosing this strategy the utility of the coercer is  $v_C^* - \epsilon_C$ .

- If  $k \geq b$  then in all states of the nature  $k \geq n^*$ , therefore:

$$u_C(\alpha, k) = v_C^* - k \cdot \text{cost}_C(\alpha).$$

In this range the strategy  $b$  is the best response of the player  $C$ , which corresponds to the utility  $v_C^* - b \cdot \text{cost}_C(\alpha)$  for the coercer.

$A \setminus C$	0	$b$
$\alpha_0$	$v_A^*, v_C^* - \epsilon_C$	$v_A^* - \epsilon_A, v_C^* - \beta_C \cdot b$
$\alpha_{m^*-1}$	$v_A^* - \text{imp}(\alpha_{m^*-1}), v_C^* - \epsilon_C$	$v_A^* - \text{imp}(\alpha_{m^*-1}) - \epsilon_A, v_C^* - b \cdot \text{cost}_C(\alpha_{m^*-1})$
$\alpha_{m^*}$	$v_A^* - \text{imp}(\alpha_{m^*}), v_C^* - \epsilon_C$	$v_A^* - \text{imp}(\alpha_{m^*}) - \epsilon_A, v_C^* - b \cdot \text{cost}_C(\alpha_{m^*})$
$\alpha_{\text{Max}}$	$v_A^* - \text{imp}(\alpha_{\text{Max}}), v_C^* - \epsilon_C$	$v_A^* - \text{imp}(\alpha_{\text{Max}}) - \epsilon_A, v_C^* - b \cdot \text{cost}_C(\alpha_{\text{Max}})$

**Fig. 4.** Coercion game with incomplete information, where the number of voters needed to coerce is estimated by a uniform probability distribution

– If  $a \leq k < b$  then

$$\begin{aligned}
u_C(\alpha, k) &= \mathbf{E}_{\omega \in \Omega}[\hat{u}_C(\omega, (\alpha, k))] = v_C^* - k \cdot \text{cost}_C(\alpha) - \frac{b-k}{b-a} \cdot \epsilon_C \\
&= v_C^* - \frac{b}{b-a} \cdot \epsilon_C + k \cdot \left( \frac{\epsilon_C}{b-a} - \text{cost}_C(\alpha) \right).
\end{aligned}$$

If  $\frac{\epsilon_C}{b-a} - \text{cost}_C(\alpha)$  is positive then  $u_C(\alpha, k)$  is increasing in  $k$  and otherwise it is decreasing in  $k$ .

**Utility of A.** Again, we consider three possible ranges of  $k$ :

– If  $k < a$  then in all states of the nature  $k < n^*$ , therefore:

$$u_A(\alpha, k) = v_A^* - \text{imp}(\alpha).$$

– If  $k \geq b$  then in all states of the nature  $k \geq n^*$ , therefore:

$$u_A(\alpha, k) = v_A^* - \text{imp}(\alpha) - \epsilon_A.$$

– If  $a \leq k < b$  then

$$u_A(\alpha, k) = \mathbf{E}_{\omega \in \Omega}[\hat{u}_A(\omega, (\alpha, k))] = v_A^* - \text{imp}(\alpha) - \frac{k-a}{b-a} \cdot \epsilon_A.$$

**Best responses and equilibria.** By observing the values of  $u_C$ , we can see that in the range  $[0, a]$ , and also when  $k > b$ ,  $u_C(\alpha, k)$  is decreasing in  $k$ . In the range  $[a, b]$ , based on the sign of  $(\frac{\epsilon_C}{b-a} - \text{cost}_C(\alpha))$  it can be increasing or decreasing in  $k$ . So the best response of the coercer is always one of the strategies 0 or  $b$  (strategy  $a$  is always dominated by 0). Therefore we need only to consider these two strategies for player  $C$ . We have that  $u_C(\alpha, 0) = v_C^* - \epsilon_C$  and  $u_C(\alpha, b) = v_C^* - b \cdot \text{cost}_C(\alpha)$ . The coercer profits more by coercing  $b$  voters when  $\text{cost}_C(\alpha) < \frac{\epsilon_C}{b}$ , and otherwise would prefer to not to coerce. We assume that from  $\alpha_0$  to  $\alpha_{m^*-1}$ , it holds that  $\text{cost}_C(\alpha) < \frac{\epsilon_C}{b}$  and from  $\alpha_{m^*}$  on, it holds that  $\text{cost}_C(\alpha) > \frac{\epsilon_C}{b}$ . Figure 4.2 shows the resulting strategic game for the uniform distribution of  $n^*$ .

**Theorem 2.** *For the coercion game with uniform beliefs, the Nash equilibrium and the maxmin for A is  $(\alpha_0, b)$ , while the Stackelberg equilibrium for A is  $(\alpha_{m^*}, 0)$ . Moreover,  $u_A(\alpha_0, b) < u_A(\alpha_{m^*}, 0)$ .*

Thus, similar to the game in Section 3.3, this game has a unique pure Nash equilibrium  $(\alpha_0, b)$ . Again, the equilibrium is undesirable, and the authority should instead prefer the Stackelberg equilibrium which is at  $(\alpha_{m^*}, 0)$ . As the Stackelberg equilibrium is different from the maxmin for A, player A needs to commit to strategy  $\alpha_m$ , and to make this commitment public.

### 4.3 Normal Probabilistic Beliefs

In our second approach, we assume that the players' beliefs about the value of  $n^*$  are represented by a normal probability distribution with mean  $\mu$  and standard deviation  $\sigma$ .

**Utility of the coercer.** When  $n^*$  has a normal distribution with mean  $\mu$  and standard deviation  $\sigma$ , the probability of a chosen  $k$  being more than  $n^*$  is:

$$Pr[n^* \leq k] = \frac{1}{2} [1 + \operatorname{erf}(\frac{k - \mu}{\sigma\sqrt{2}})]$$

where:

$$\operatorname{erf}(x) = \frac{1}{\sqrt{\pi}} \int_{-x}^x e^{-t^2} \cdot dt.$$

Therefore  $u_C(\alpha, k)$  can be calculated as:

$$\begin{aligned} u_C(\alpha, k) &= \mathbf{E}_{\omega \in \Omega} [\hat{u}_C(\omega, (\alpha, k))] = v_C^* - k \cdot \operatorname{cost}_C(\alpha) - \frac{1}{2} [1 - \operatorname{erf}(\frac{k - \mu}{\sigma\sqrt{2}})] \cdot \epsilon_C \\ &= v_C^* - \mu \cdot \operatorname{cost}_C(\alpha) - \frac{\epsilon_C}{2} + \gamma(k). \end{aligned}$$

where:

$$\gamma(k) = \frac{\epsilon_C}{2} \cdot \operatorname{erf}(\frac{k - \mu}{\sigma\sqrt{2}}) - (k - \mu) \cdot \operatorname{cost}_C(\alpha).$$

Analysing the changes of function  $\gamma(k)$  shows that if  $\operatorname{cost}_C(\alpha) > \frac{\epsilon_C}{\sigma\sqrt{2}}$  then  $\gamma(k)$  is decreasing in  $k$ . In this case  $u_C(\alpha, k)$  has its maximum at  $k = 0$ . If  $\operatorname{cost}_C(\alpha) < \frac{\epsilon_C}{\sigma\sqrt{2}}$  then  $\gamma(k)$ , and hence  $u_C(\alpha, k)$ , has a maximum at

$$k_\alpha^{max} = \mu + \sqrt{2\sigma^2 \ln(\frac{\epsilon_C}{\sqrt{2\pi} \cdot \operatorname{cost}_C(\alpha) \cdot \sigma})}.$$

Notice that this number is decreasing in  $\operatorname{cost}_C(\alpha)$ . We denote the value of  $u_C(\alpha, k)$  at this point by  $u_C^{max, \alpha}$ , where:

$$u_C^{max, \alpha} = v_C^* - k_\alpha^{max} \cdot \operatorname{cost}_C(\alpha) - \frac{1}{2} [1 - \operatorname{erf}(\frac{k_\alpha^{max} - \mu}{\sigma\sqrt{2}})] \cdot \epsilon_C$$

$u_C^{max, \alpha}$  is positive and is increasing in  $\sigma$  and decreasing in  $\operatorname{cost}_C(\alpha)$ .

**Utility of A.** The utility of the society in the transformed game is:

$$u_A(\alpha, k) = v_A^* - \operatorname{imp}(\alpha) - \frac{1}{2} [1 + \operatorname{erf}(\frac{k - \mu}{\sigma\sqrt{2}})] \cdot \epsilon_A.$$

Notice that if we fix  $k$ , this function is decreasing in  $\operatorname{imp}(\alpha)$ .

**Best responses and equilibria.** We can consider two cases: If  $\operatorname{cost}_C(\alpha) > \frac{\epsilon_C}{\sigma\sqrt{2}}$  then the best response for the coercer is 0, and otherwise his best response is

$A \setminus C$	0	$k_{m^*-1}^{max}$	$k_0^{max}$
$\alpha_0$	$\underline{v}_A^*, \underline{v}_C^* - \epsilon_C$	$u_A(\alpha_0, k_{m^*-1}^{max}), u_C(\alpha_0, k_{m^*-1}^{max})$	$\boxed{u_A(\alpha_0, k_0^{max}), u_C^{max, \alpha_0}}$
$\alpha_{m^*-1}$	$v_A^* - imp(\alpha_{m^*-1}), v_C^* - \epsilon_C$	$u_A(\alpha_{m^*-1}, k_{m^*-1}^{max}), u_C^{max, \alpha_{m^*-1}}$	$u_A(\alpha_{m^*-1}, k_0^{max}), u_C(\alpha_{m^*-1}, k_0^{max})$
$\alpha_{m^*}$	$\underline{v}_A^* - imp(\alpha_{m^*}), \underline{v}_C^* - \epsilon_C$	$u_A(\alpha_{m^*}, k_{m^*-1}^{max}), u_C(\alpha_{m^*}, k_{m^*-1}^{max})$	$u_A(\alpha_{m^*}, k_0^{max}), u_C(\alpha_{m^*}, k_0^{max})$
$\alpha_{Max}$	$v_A^* - imp(\alpha_{Max}), v_C^* - \epsilon_C$	$u_A(\alpha_{Max}, k_{m^*-1}^{max}), u_C(\alpha_{Max}, k_{m^*-1}^{max})$	$u_A(\alpha_{Max}, k_0^{max}), u_C(\alpha_{Max}, k_0^{max})$

**Fig. 5.** Coercion game with incomplete information, where the number of voters needed to coerce is estimated by a normal probability distribution

$k_\alpha^{max}$ . We assume that from  $\alpha_0$  to  $\alpha_{m^*-1}$ , it holds that  $cost_C(\alpha) < \frac{\epsilon_C}{\sigma\sqrt{2}}$  and from  $\alpha_{m^*}$  on, it holds that  $cost_C(\alpha) > \frac{\epsilon_C}{\sigma\sqrt{2}}$ .

Figure 5 shows the strategic game for the normal distribution of  $n^*$ . For the choices of the authority, we have only shown four protection measures:  $\alpha_0$ ,  $\alpha_{m^*-1}$ ,  $\alpha_{m^*}$  and  $\alpha_{Max}$ . For the choices of the coercer, we only included the ones that are the best responses to one of the depicted choices of the authority. The choice 0 is the best response for the coercer when authority chooses any protection method from  $\alpha_{m^*}$  on. The choice  $k_0^{max}$  is the best response when authority chooses  $\alpha_0$ , and the choice  $k_{m^*-1}^{max}$  is the best response when authority's choice is  $\alpha_{m^*-1}$ .

The game has a unique pure Nash equilibrium at  $(\alpha_0, k_0^{max})$ , which is clearly a bad outcome for the society. However, if the implementation cost of the protection method  $\alpha_{m^*}$  is less than the expected damage that player  $A$  gets from the coercion at the Nash equilibrium, i.e., if  $imp(\alpha_{m^*}) < \frac{1}{2}[1 + \text{erf}(\frac{k_0^{max} - \mu}{\sigma\sqrt{2}})] \cdot \epsilon_A$ , then the authority can use the Stackelberg equilibrium at  $(\alpha_{m^*}, 0)$  by committing itself to choose the method  $\alpha_{m^*}$  and to make this commitment public.

Now consider that the authority cannot, or does not prefer to implement  $\alpha_{m^*}$  or more secure protection methods (for example because of the high cost of it) and the strongest protection method that can be implemented is a suboptimal protection method  $\alpha_{m^*-1}$ . By announcing its choice and committing to it, the authority can achieve an equilibrium at  $(\alpha_{m^*-1}, b_{k_{m^*-1}^{max}})$ . In this equilibrium the estimated cost of a successfully coerced election for the authority  $(\frac{1}{2}[1 + \text{erf}(\frac{k_0^{max} - \mu}{\sigma\sqrt{2}})] \cdot \epsilon_A)$  is lower than ones in the pure Nash equilibrium of the game. If this reduction of cost is worthwhile for the authority (in comparison to the extra implementation cost of  $\alpha_{m^*-1}$  comparing to  $\alpha_0$ ), the authority can benefit from announcing and committing to its strategy even in a suboptimal protection method.

Notice that by increasing the uncertainty about the number of needed votes to buy, i.e by increasing  $\sigma$ , the value of  $m^*$  decreases. It means that the Stackelberg equilibrium can be moved to a one with lower implementation cost for the authority. This may suggest that the authority can in fact benefit from making very accurate polls *unavailable* to the public before the election.

## 5 Conclusions

In this work, we look at simple game models of protection against coercion in voting procedures. The models are two-person nonzero-sum noncooperative games, where one player represents the society and the other a potential coercer in the election. Our modelling relies on a number of abstractions and simplifying assumptions. Still, even at this level of abstraction some interesting patterns can be observed. Most importantly, we show that in all games that we consider, Stackelberg equilibrium is different from Nash equilibrium. In other words, it is in the interest of the society *not* to adapt to the expected strategy of the coercer. Instead of that, the society should decide on its coercion-resistance policy in advance.

Moreover, for almost all of our models, the Stackelberg equilibrium does not coincide with maxmin. Translating the formal result to intuitive terms, the society will benefit from *announcing its anti-coercion policy openly*. This way, the rational coercer is forced to refraining from coercion altogether. Paraphrasing the well-known slogan, the advice is *not to seek coercion-resistance through obscurity*.

**Acknowledgements.** Wojciech Jamroga acknowledges the support of the National Research Fund (FNR), Luxembourg, under the project GALOT (INTER/DFG/12/06), the support of the 7th Framework Programme of the European Union under the Marie Curie IEF project ReVINK (PIEF-GA-2012-626398), and the support of the National Centre for Research and Development (NCBR), Poland, under the PolLux project VoteVerif (POLLUX-IV/1/2016).

## References

1. Riza Aditya, Byoungcheon Lee, Colin Boyd, and Ed Dawson. An efficient mixnet-based voting scheme providing receipt-freeness. In *Trust and Privacy in Digital Business*, pages 152–161. Springer, 2004.
2. Roberto Arajo, Narjes Rajeb, Riadh Robbana, Jacques Traor, and Souheib Youssfi. Towards practical and secure coercion-resistant electronic elections. In Swee-Huay Heng, RebeccaN. Wright, and Bok-Min Goi, editors, *Cryptology and Network Security*, volume 6467 of *Lecture Notes in Computer Science*, pages 278–297. Springer Berlin Heidelberg, 2010.
3. Roberto Araújo, Narjes Ben Rajeb, Riadh Robbana, Jacques Traoré, and Souheib Youssfi. Towards practical and secure coercion-resistant electronic elections. In *Cryptology and Network Security*, pages 278–297. Springer, 2010.
4. G. Bonanno. Game theory (open access textbook with 165 solved exercises). *CoRR*, abs/1512.06808, 2015.
5. Ahto Buldas and Triinu Mägi. Practical security analysis of e-voting systems. In *Proceedings of IWSEC*, volume 4752 of *Lecture Notes in Computer Science*, pages 320–335. Springer, 2007.
6. S. Delaune, S. Kremer, and M. Ryan. Coercion-resistance and receipt-freeness in electronic voting. In *Computer Security Foundations Workshop, 2006. 19th IEEE*, pages 12–pp. IEEE, 2006.

7. Jannik Dreier, Pascal Lafourcade, and Yassine Lakhnech. A formal taxonomy of privacy in voting protocols. In *Communications (ICC), 2012 IEEE International Conference on*, pages 6710–6715. IEEE, 2012.
8. Ryan W. Gardner, Sujata Garera, and Aviel D. Rubin. Coercion resistant end-to-end voting. In Roger Dingledine and Philippe Golle, editors, *Financial Cryptography and Data Security*, volume 5628 of *Lecture Notes in Computer Science*, pages 344–361. Springer Berlin Heidelberg, 2009.
9. John C Harsanyi and Reinhard Selten. A generalized Nash solution for two-person bargaining games with incomplete information. *Management Science*, 18(5-part-2):80–106, 1972.
10. J. Heather and S. Schneider. A formal framework for modelling coercion resistance and receipt freeness. *FM 2012: Formal Methods*, pages 217–231, 2012.
11. A. Juels, D. Catalano, and M. Jakobsson. Coercion-resistant electronic elections. In *Proceedings of the 2005 ACM workshop on Privacy in the electronic society*, pages 61–70. ACM, 2005.
12. C. Kiekintveld, M. Jain, J. Tsai, J. Pita, M. Tambe, and F. Ordonez. Computing optimal randomized resource allocations for massive security games. In *Proceedings of AAMAS*, pages 689–696. IFAAMAS, 2009.
13. D. Korzhuk, Z. Yin, C. Kiekintveld, V. Conitzer, and M. Tambe. Stackelberg vs. Nash in security games: An extended investigation of interchangeability, equivalence, and uniqueness. *Journal of Artificial Intelligence Research*, 41:297–327, 2011.
14. Wei-Chi Ku and Chun-Ming Ho. An e-voting scheme against bribe and coercion. In *e-Technology, e-Commerce and e-Service, 2004. EEE'04. 2004 IEEE International Conference on*, pages 113–116. IEEE, 2004.
15. R. Küsters, T. Truderung, and A. Vogt. A game-based definition of coercion-resistance and its applications. In *Proceedings of the 2010 23rd IEEE Computer Security Foundations Symposium*, pages 122–136. IEEE Computer Society, 2010.
16. R. Küsters, T. Truderung, and A. Vogt. Verifiability, privacy, and coercion-resistance: New insights from a case study. In *Security and Privacy (SP), 2011 IEEE Symposium on*, pages 538–553. IEEE, 2011.
17. Byoungcheon Lee, Colin Boyd, Ed Dawson, Kwangjo Kim, Jeongmo Yang, and Seungjae Yoo. Providing receipt-freeness in mixnet-based voting protocols. In *Information Security and Cryptology-ICISC 2003*, pages 245–258. Springer, 2004.
18. Byoungcheon Lee and Kwangjo Kim. Receipt-free electronic voting scheme with a tamper-resistant randomizer. In *Information Security and CryptologyICISC 2002*, pages 389–406. Springer, 2003.
19. K. Leyton-Brown and Y. Shoham. *Essentials of Game Theory: A Concise, Multi-disciplinary Introduction*. Morgan & Claypool, 2008.
20. Na Li, Lijun Chen, and Steven H. Low. Optimal demand response based on utility maximization in power networks. IEEE, 2011.
21. Emmanouil Magkos, Mike Burmester, and Vassilis Chrissikopoulos. Receipt-freeness in large-scale elections without untappable channels. In *Towards The E-Society*, pages 683–693. Springer, 2001.
22. M. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.
23. M. Pedrasa, T. Spooner, and I. MacGill. Coordinated scheduling of residential distributed energy resources to optimize smart home energy services. *IEEE Transactions on Smart Grid*, 1(2):134–143, 2010.
24. Peter Y. A. Ryan. The computer ate my vote. In *Formal Methods: State of the Art and New Directions*, pages 147–184. Springer, 2010.

25. Michael Schlapfer, Rolf Haenni, Reto Koenig, and Oliver Spycher. Efficient vote authorization in coercion-resistant internet voting. In *E-Voting and Identity: Third International Conference, VoteID 2011, Tallinn, Estonia, September 28-20, 2011, Revised Selected Papers*, volume 7187, page 71. Springer, 2012.
26. C.L. Su and D. Kirschen. Quantifying the effect of demand response on electricity markets. *IEEE Transactions on Power Systems*, 24(3):2009, 1199-1207.
27. M. Tambe. *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press, 2011.
28. Stefan G Weber, Roberto Araujo, and Johannes Buchmann. On coercion-resistant electronic elections with linear work. In *Availability, Reliability and Security, 2007. ARES 2007. The Second International Conference on*, pages 908–916. IEEE, 2007.
29. H. Xu. The mysteries of security games: Equilibrium computation becomes combinatorial algorithm design. In *Proceedings of SECMAS*. IFAAMAS, 2016.
30. Y. Yin, Y. Vorobeychik, B. An, and N. Hazon. Optimally protecting elections. In *Proceedings of SECMAS*. IFAAMAS, 2016.
31. J. Zhang, J.D. Fuller, and S. Elhedhli. A stochastic programming model for a day-ahead electricity market with real-time reserve shortage pricing. *IEEE Transactions on Power Systems*, 25(2):703–713, 2010.