Preventing Coercion in E-Voting: Be Open and Commit

Wojtek Jamroga, Polish Academy of Sciences
(joint work with Masoud Tabatabaei and Peter Y. A. Ryan)

LAMAS Seminar on INteraction Gdansk
24th of September 2015

Outline

Coercion Resistance

- Desirable properties of voting schemes: privacy, anonymity, receipt-freeness, coercion resistance
- In this work, we focus on coercion resistance

Coercion Resistance

- Desirable properties of voting schemes: privacy, anonymity, receipt-freeness, coercion resistance
- In this work, we focus on coercion resistance
- Standard definition:

Coercion resistance: The voter cannot cooperate with a coercer to prove to him that she voted in a certain way.

Coercion Resistance

- We look at a more fundamental property
- CR $\approx$ voter's ability to... well, resist coercion ☺

Coercion Resistance

- We look at a more fundamental property
- CR $\approx$ voter's ability to... well, resist coercion 🙂

Coercion resistance: The system should provide good prerequisites for the voter to cast her vote according to her free intent, despite potential efforts of the coercer.

Coercion as a Game

- We will model coercion as a **game** where different participants have possibly conflicting interests

Coercion as a Game

- We will model coercion as a **game** where different participants have possibly conflicting interests
- In general: very complex
- An exhaustive model should include the incentives of: multiple voters, multiple coercers, possibly also social groups, business conglomerates, government agencies, etc.
- ...Also, we would have to define the interaction between incentives and behaviors of different groups (competition, collusion, etc.)

Coercion as a Game

- In this work, we settle for something much simpler
- We see coercion resistance as a game between:
  1. a single voting authority (approximating the interests of the society as a whole),
  2. and a single coercer (approximating the interests of potential coercers and their groups)

Coercion as a Game

- In this work, we settle for something much simpler
- We see coercion resistance as a game between:

  1. a single voting authority (approximating the interests of the society as a whole),
  2. and a single coercer (approximating the interests of potential coercers and their groups)

$\leadsto$ We look at 2-player games with largely conflicting interests

Coercion as a Game

Note:

We do **not** propose a new coercion resistant voting scheme, but a model of interaction that involves coercion!

Outline

Game Models: Strategic Games

Definition 1 (Strategic game)

A **strategic game** G is a tuple $(N, \{\Sigma_i | i \in N\}, o, W)$ that consists of a nonempty finite set of players N, a nonempty set of strategies Σ_i for each player $i \in N$, a nonempty set of outcomes W, an outcome function $o : \prod_{i \in N} \Sigma_i \to W$ which associates an outcome with every strategy profile, and a utility function $o : N \times W \to \mathbb{R}$ which assigns agent's payoffs (or: utility values) to each possible outcome.

Example: “Twisted” Battle of Sexes

$Bob \backslash Sue$	Bar	Th
Bar	$2, 1$	$0, 0$
Th	$3, 0$	$1, 2$

Solution Concepts

- **Solution concepts** define which collective behaviors are **rational**
- Formally, a solution concept is modelled as a subset of strategy profiles (= cells in the payoff table)

Solution Concepts

- **Solution concepts** define which collective behaviors are **rational**
- Formally, a solution concept is modelled as a subset of strategy profiles (= cells in the payoff table)
- We will use two solution concepts: Nash equilibrium and Stackelberg equilibrium

Nash Equilibrium

We look for strategy profiles which are stable under unilateral deviations

$Bob\backslash Sue$	Bar	Th
Bar	$2,1$	$0,0$
Th	$3,0$	$1,2$

Nash Equilibrium

We look for strategy profiles which are stable under unilateral deviations

$Bob \backslash Sue$	Bar	Th
Bar	$2, 1$	$0, 0$
Th	$3, 0$	$1, 2$

Stackelberg Equilibrium

We look for the **best response to best responses**

$Bob \backslash Sue$	Bar	Th
Bar	$2, 1$	$0, 0$
Th	$3, 0$	$1, 2$

Stackelberg Equilibrium

We look for the **best response to best responses**

$Bob \backslash Sue$	Bar	Th
Bar	$\mathbf{2, 1}$	$0, 0$
Th	$3, 0$	$1, 2$

Stackelberg Equilibrium

We look for the **best response to best responses**

$Bob \backslash Sue$	Bar	Th
Bar	$2,1$	$0,0$
Th	$3,0$	$\mathbf{1,2}$

Stackelberg Equilibrium

We look for the **best response to best responses**

$Bob \backslash Sue$	Bar	Th
Bar	$\mathbf{2, 1}$	$0, 0$
Th	$3, 0$	$\mathbf{1, 2}$

Stackelberg Equilibrium

We look for the **best response to best responses**

$Bob\backslash Sue$	Bar	Th
Bar	$\mathbf{2, 1}$	$0, 0$
Th	$3, 0$	$1, 2$

Stackelberg Equilibrium

We look for the **best response to best responses**

$Bob\backslash Sue$	Bar	Th
Bar	$\mathbf{2, 1}$	$0, 0$
Th	$3, 0$	$1, 2$

Nash vs. Stackelberg

- Nash equilibrium captures the outcome of mutual long-run adaptation of players to each others' strategies

Nash vs. Stackelberg

- Nash equilibrium captures the outcome of mutual long-run adaptation of players to each others' strategies
- Stackelberg equilibrium captures the outcome in games where one player (the *leader*) exposes her strategy first

Nash vs. Stackelberg

- Nash equilibrium captures the outcome of mutual long-run adaptation of players to each others' strategies
- Stackelberg equilibrium captures the outcome in games where one player (the *leader*) exposes her strategy first
- Applicability of Stackelberg: the leader must be able to
  1. either complete her strategy before the other players start,
  2. or irrevocably commit to her strategy in advance.

Are Leaders Always at Advantage?

$Bob \backslash Sue$	H	T
H	$1,0$	$0,1$
T	$0,1$	$1,0$

Are Leaders Always at Advantage?

$Bob \backslash Sue$	H	T
H	$1, 0$	$0, 1$
T	$0, 1$	$1, 0$

- No pure Nash equilibrium

Are Leaders Always at Advantage?

$Bob \backslash Sue$	H	T
H	$1, 0$	$0, 1$
T	$0, 1$	$1, 0$

- No pure Nash equilibrium
- Unique mixed Nash equilibrium (everybody plays at random, with equal probabilities), promising each player the expected payoff of 0.5

Are Leaders Always at Advantage?

$Bob \backslash Sue$	H	T
H	$1, 0$	$\mathbf{0, 1}$
T	$\mathbf{0, 1}$	$1, 0$

- No pure Nash equilibrium
- Unique mixed Nash equilibrium (everybody plays at random, with equal probabilities), promising each player the expected payoff of 0.5
- Two Stackelberg equilibria, each promising Bob the payoff of $\mathbf{0}$

Outline

Coercion as a Game

Main idea:

- Coercion resistance comes at a cost

Coercion as a Game

Main idea:

- Coercion resistance comes at a cost
- The society should balance the cost of anti-coercion measures vs. damage from successful coercion attacks

Coercion as a Game

Main idea:

- Coercion resistance comes at a cost
- The society should balance the cost of anti-coercion measures vs. damage from successful coercion attacks
- Coercer: costs vs. benefits of coercion

Coercion as a Game

Main idea:

- Coercion resistance comes at a cost
- The society should balance the cost of anti-coercion measures vs. damage from successful coercion attacks
- Coercer: costs vs. benefits of coercion

Question:

Should society invest in anti-coercion measures?

Coercion as a Game

Main idea:

- Coercion resistance comes at a cost
- The society should balance the cost of anti-coercion measures vs. damage from successful coercion attacks
- Coercer: costs vs. benefits of coercion

Question:

Should society invest in anti-coercion measures? If so, how much?

Coercion as a Game

Main idea:

- Coercion resistance comes at a cost
- The society should balance the cost of anti-coercion measures vs. damage from successful coercion attacks
- Coercer: costs vs. benefits of coercion

Question:

Should society invest in anti-coercion measures? If so, how much? ...And, in what way?

Game Model for Coercion Resistance

2 players:

- A: the honest election authority

Game Model for Coercion Resistance

2 players:

- A: the honest election authority
- C: the coercer

Game Model for Coercion Resistance

2 players:

- A: the honest election authority
- C: the coercer

Strategies:

- A: choose one of anti-coercion measures $a_0, \ldots, a_m$

Game Model for Coercion Resistance

2 players:

- A: the honest election authority
- C: the coercer

Strategies:

- A: choose one of anti-coercion measures $a_0, \ldots, a_m$
- C: choose how many voters to coerce $c_0, \ldots, c_n$

Utility of the Society

$$u_A(a_i, c_i) \;=\; v_A(c_i) - imp(a_i) - \delta \cdot c_i,$$ where:

- $v_A(c_i)$: “quality” of the election outcome (v_A^* if undisturbed, $v_A^* - \epsilon_A$ if disturbed)

Utility of the Society

$$u_A(a_i, c_i) = v_A(c_i) - imp(a_i) - \delta \cdot c_i,$$ where:

- $v_A(c_i)$: "quality" of the election outcome (v_A^* if undisturbed, $v_A^* - \epsilon_A$ if disturbed)
- $imp(a_i)$: cost of implementing the anti-coercion measure

Utility of the Society

$$u_A(a_i, c_i) \;=\; v_A(c_i) - imp(a_i) - \delta \cdot c_i,$$ where:

- $v_A(c_i)$: "quality" of the election outcome (v_A^* if undisturbed, $v_A^* - \epsilon_A$ if disturbed)
- $imp(a_i)$: cost of implementing the anti-coercion measure
- δ: corruption damage per coerced voter

Utility of the Coercer

$$u_C(a_i, c_i) = v_C(c_i) - \beta(a_i) \cdot c_i,$$ where:

- $v_C(c_i)$: "quality" of the election outcome (v_C^* if disturbed, $v_C^* - \epsilon_C$ if undisturbed)

Utility of the Coercer

$$u_C(a_i, c_i) \;=\; v_C(c_i) - \beta(a_i) \cdot c_i, \quad \text{where:}$$

- $v_C(c_i)$: "quality" of the election outcome (v_C^* if disturbed, $v_C^* - \epsilon_C$ if undisturbed)
- $\beta(a_i)$: Cost of coercion per voter (bribery, disclosure of votes, etc.)

Coercion Game

$A \backslash C$	c_0	c^*
a_0	$v_A^*,\ v_C^* - \epsilon_C$	$v_A^* - \epsilon_A - \delta \cdot c^*,\ v_C^* - \beta_C \cdot c^*$
$\vdots$		
a_m	$v_A^* - imp(a_m),\ v_C^* - \epsilon_C$	$v_A^* - \epsilon_A - imp(a_m) - \delta \cdot c^*,$ $v_C^* - \beta(a_1) \cdot c^*$

Note: from the coercer's point of view, it suffices to consider only the actions of no coercion (c_0) and bribing the minimal amount of voters that would swing the result of the election (c^*)

Coercion Game: Example

For

- $m = 1$, $v_A^* = 5$, $\epsilon_A = 3$, $imp(a_0) = 0$, $imp(a_1) = 1$, $\delta = 1$
- $c^* = 1$, $v_C^* = 5$, $\epsilon_C = 2$, $\beta = 3$

we get

$A \backslash C$	c_0	c^*
a_0	$5, 3$	$1, 4$
a_m	$4, 3$	$0, 2$

Coercion Game: Example

For

- $m = 1$, $v_A^* = 5$, $\epsilon_A = 3$, $imp(a_0) = 0$, $imp(a_1) = 1$, $\delta = 1$
- $c^* = 1$, $v_C^* = 5$, $\epsilon_C = 2$, $\beta = 3$

we get

$A \backslash C$	c_0	c^*
a_0	5, 3	1, 4
a_m	4, 3	0, 2

Coercion Game: Example

For

- $m = 1$, $v_A^* = 5$, $\epsilon_A = 3$, $imp(a_0) = 0$, $imp(a_1) = 1$, $\delta = 1$
- $c^* = 1$, $v_C^* = 5$, $\epsilon_C = 2$, $\beta = 3$

we get

$A \backslash C$	c_0	c^*
a_0	$5, 3$	$1, 4$
a_m	$\mathbf{4, 3}$	$0, 2$

Coercion Game: Example

For

- $m = 1$, $v_A^* = 5$, $\epsilon_A = 3$, $imp(a_0) = 0$, $imp(a_1) = 1$, $\delta = 1$
- $c^* = 1$, $v_C^* = 5$, $\epsilon_C = 2$, $\beta = 3$

we get

$A \backslash C$	c_0	c^*
a_0	$5, 3$	$1, 4$
a_m	**$4, 3$**	$0, 2$

Playing Stackelberg is much more profitable than Nash!

Coercion Game: General Result

Theorem 2

Under some mild assumptions, we get the following:

1. *The coercion game has a unique Nash equilibrium in* (a_0, c^*),
2. *The Stackelberg equilibrium is* (a_m, c_0), *and*
3. *Stackelberg equilibrium is preferred to Nash equilibrium, i.e.,* $u_A(a_0, c^*) < u_A(a_m, c_0)$.

Coercion Game: General Result

Theorem 2

Under some mild assumptions, we get the following:

1. *The coercion game has a unique Nash equilibrium in* (a_0, c^*),
2. *The Stackelberg equilibrium is* (a_m, c_0), *and*
3. *Stackelberg equilibrium is preferred to Nash equilibrium, i.e.,* $u_A(a_0, c^*) < u_A(a_m, c_0)$.

Note: the society enforces the coercer not to coerce (c_0) by publicly committing to high-security policy (a_m)

Outline

Conclusions

- The work is very preliminary, but...

Conclusions

- The work is very preliminary, but...
- ...our analysis suggests that the society should not adapt to what it expects from the bad guys

Conclusions

- The work is very preliminary, but...
- ...our analysis suggests that the society should not adapt to what it expects from the bad guys
- Committing publicly to an anti-coercion policy prevents coercing attempts

Conclusions

- The work is very preliminary, but...
- ...our analysis suggests that the society should not adapt to what it expects from the bad guys
- Committing publicly to an anti-coercion policy prevents coercing attempts

No coercion resistance through obscurity!

Thank you for your attention