Anonymity in voting revisited

Hugo Jonker[1,2] and Wolter Pieters[3,⋆]

[1] Eindhoven University of Technology, The Netherlands
[2] University of Luxembourg, Luxembourg
h.l.jonker@tue.nl,hugo.jonker@uni.lu
[3] wolter.pieters@gmail.com

Abstract. According to international law, anonymity of the voter is a fundamental precondition for democratic elections. In electronic voting, several aspects of voter anonymity have been identified. In this paper, we re-examine anonymity with respect to voting, and generalise existing notions of anonymity in e-voting. First, we identify and categorise the types of attack that can be a threat to anonymity of the voter, including different types of vote buying and coercion. This analysis leads to a categorisation of anonymity in voting in terms of a) the strength of the anonymity achieved and b) the extent of interaction between voter and attacker. Some of the combinations, including weak and strong receipt-freeness, are formalised in epistemic logic.

1 Introduction

In the field of peer-to-peer (P2P) networks, much effort has been put into formalizing the concept of anonymity of messages (e.g. [14]). Intuitively, anonymity means that it is impossible to determine who sent which message to whom. Depending on the context, different formalizations of the notion of anonymity seem to be necessary [7].

The concept of anonymity is also of importance in electronic voting – often, voters should have the ability to vote without anybody else knowing which option they voted for (although in some countries, such as the United Kingdom and New Zealand, this is ultimately not the case). In the electronic voting community, the property expressing precisely that is usually called "privacy" instead of anonymity [6]. In voting, however, enabling privacy is not sufficient, as this does not prevent vote buying. To prevent vote buying, an election needs to require privacy – no voter should be able to convince any other party of how she voted.

The concept of receipt-freeness expresses that a voter cannot convince any other party of how she voted by creating a receipt. The notion has been introduced by [2], after which various receipt-free voting protocols were proposed, such as [9, 17]. Delaune et al. [5] provide a definition of receipt-freeness based on observational equivalence. Independently, Jonker and De Vink [10] provide an alternate definition that allows identification of receipts. Juels et al. note in [11]

⋆ Research was carried out partially at Radboud University, Nijmegen, The Netherlands, and supported by NWO, the Netherlands Organisation for Scientific Research.

that receipt-freeness is not sufficient to prevent coercion in electronic elections, and they introduced the notion of coercion-resistance. This broader notion is again formalized by Delaune et al. in [6].

Given the differences in approaches and in notions, the question arises whether these notions capture the specific needs for anonymity in voting. The three main levels of anonymity that have been identified in voting, capture progressively more strict notions of anonymity. The notion of receipt-freeness was motivated as necessary to provide secret-ballot elections. If receipts can be obtained, using a voting booth makes no difference to the secrecy: Votes can be bought, and voters can be coerced.

To address the question of whether or not the notion of receipt-freeness is sufficient, we reexamine voter influencing, focusing on vote buying. What is vote buying, when can an action be called vote buying and when is it an election promise? As this is, ultimately, a subjective issue, the goal is not to provide a yes-or-no test. Instead, we aim to arrive at a charactarisation of vote buying / election promises, which will enable election officials to decide which practices are allowed and which should be abolished. Based on these findings, we then reexamine the concept of receipt-freeness and adapt it to encompass uncovered issues.

1.1 Related work

Distinctions between vote buying and election promises have been investigated by economists, philosophers and political scientists before.

Van Acker [1] discusses the relation between the notions of coercion, forced abstention, randomisation and simulation. However, he includes vote buying in the concept of coercion.

Kochin and Kochin [12] discuss the issue of giving benefits to individual voters versus giving benefits to identifiable groups. They also consider the difference between benefits offered through the normal processes of government (related to being elected) versus benefits offered through private arrangements. Thirdly, they mention that trading votes for or against proposals between parties or members in parliament is acceptable.

The latter practice is also mentioned by Hasen [8] and called "legislative logrolling". Hasen further differentiates the issues of corporate vote buying, payments to increase turnout, campaign promises and campaign contributions, and vote buying in so-called "special district"[1] elections.

Schaffer [18] distinguishes between *instrumental*, *normative* and *coercive* compliance in relation to vote buying. Instrumental compliance covers tangible benefits in exchange for votes, normative compliance means voting based on a feeling of obligation, and coercive compliance denotes voting based on threats. Schaffer also mentions the possibility that money is offered for *not* changing voting behaviour. In order to check compliance, a buyer may monitor the individual vote,

[1] "a special purpose unit of government assigned the performance of functions affecting definable groups of constituents more than other constituents"

monitor the aggregate turnout, prevent people from voting altogether, make the rewards dependent on his election, make voters believe in his goodness or make voters feel personally obliged. The applicability of these strategies is dependent on the mode of compliance the buyer is seeking. From the perspective of voters, benefits can be received in the form of payment, gift or wage, with different explicit and implicit meanings in terms of modes of compliance.

From these papers it is clear that what exactly constitutes acceptable influence and what does not, depends on the type of elections, the society in which the elections are being held and the participants of the elections. In the end, the matter is ultimately a subjective one. However, by determining the various characteristics of vote buying, and their respective ranges, it is possible to establish a pre-election consensus on allowed and disallowed practices. Such a pre-election consensus enables putting precise requirements on voting systems to support the one type of behaviour, while preventing the other type.

1.2 Outline of the paper

In Section 2, we identify and categorise the types of attack that can be a threat to anonymity of the voter, including different types of vote buying and coercion. This analysis leads to a categorisation of anonymity in voting in terms of a) the strength of the anonymity achieved and b) the extent of interaction between voter and attacker, which is presented in Section 3. In Section 4, some of the combinations, including weak and strong receipt-freeness, are formalised in epistemic logic. The last section presents conclusions and future work.

2 Characteristics of voter influencing

In this section, we investigate the characteristics of voter influencing. The examples below are used as supporting guidelines throughout the section. These examples are deliberately without context – in lieu of what was established in Section 1.1. The reason for this is that the aim is to discover the generic characteristics involved, irrespective of social and electoral context. The examples are not meant to capture any precise attempt at influencing voters, but rather they convey a broad idea of a, possibly controversial, attempt at changing the outcome of an election by targeting the voters.

Example 1 (handout). At the polling station, I give each voter 100 euros together with mentioning my candidacy.

Example 2 (theme park). The district with the highest percentage of votes for me gets a theme park as my first act as elected official.

Example 3 (zalmsnip). If I get elected, everyone gets 100 euros tax refund.

Example 4 (election promise). If I get elected, disabled child prodigies get 100 euros (i.e. children with a physical handicap, who are members of Mensa).

Example 5 (rf). I give 100 euro for anyone voting for the Democratic Party.

Example 6 (non-rf). I give 100 euro for anyone not voting for the Democratic Party.

Example 7 (reimburse). I provide a reimbursement for voters for time away from work.

The *zalmsnip* example is based on a tax rebate that occurred in the Netherlands in 1998 and 1999. The *rf* and *non-rf* examples are inspired by the notion of receipt-freeness. The *reimburse* example is inspired by this practice occurring in the 19th century in the Netherlands.

Note that – as long as there is no request to vote in a specific way – example 1 can be considered legitimate. Examples 3 and 4 are fabrications resembling possible election promises. Example 2 is dubious; examples 5–7 are outright illegal.

2.1 Legal and illegal influencing

Influencing voters can be done either legally or illegally. To avoid a legal discussion on what is allowed by which laws, we only focus upon characterising what is desirable. As established in the introduction, this is a subjective notion. The aim here is to outline the range of possibilities available, indicate where the boundary between desirable and undesirable lies and give a supportive reasoning for where we feel this boundary lies.

Note that, in general, there are two methods to influence a voter's vote:

coercion where voters are threatened to ensure compliance;
enticement where voters are seduced into compliance.

Whereas persuasion is allowed, buying and coercion are not. Both buying and coercion require proof of compliance, whereas persuasion does not. Both buying and persuasion are dependent on voluntary cooperation of the voter, coercion is not.

Voter influencing can be considered acceptable or unacceptable. What is considered acceptable depends on culture and the nature of the elections. That there can exist both acceptable and unacceptable variants of the above two methods is illustrated by the following list.

- *acceptable coercion* claiming that all other candidates have significantly worse plans for the voter
- *unacceptable coercion* threatening with physical violence in case of non-compliance
- *acceptable enticement* promising to lower taxes
- *unacceptable enticement* paying a voter to vote for you

The above list clearly indicates, that there is a distinction between acceptable influence and unacceptable influence. To establish the characteristics that

together determine the acceptability, we construct an objective tree of voter influencing in Section 2.2. Objective trees are attack trees (see [19, 13]), but focus upon meeting goals instead of achieving attacks.

Our objective tree deviates slightly from normal attack trees. The purpose of our tree is to determine characteristics that distinguish acceptable from unacceptable influence. To elucidate these detailed characteristics, details need to be explicit in the tree. Hence, where normally attributes would be used, we promote these characteristics to leaves. This makes these characteristics explicit.

2.2 Classifying vote buying

Based on the literature, the examples and the analysis above, the tree in Figure 1 was constructed and dimensions of vote buying were clarified. The main goal in the tree is to buy a vote, by means of persuasion. The tree is thus from the perspective of a vote buyer. Where necessary, the range of possible values has been indicated in the tree (as leaves).

- **or** reward time
  - **leaf** before vote casting
  - **and** later
    - ∗ **or** trust required
      - · **leaf** rewarding sureness
      - · **leaf** consequences of non-reward
      - · **leaf** ensurance of compliance
    - ∗ **or** hand out reward
      - · **leaf** after vote casting
      - · **leaf** after ballot box closes
      - · **leaf** after vote counting
- **or** type of reward
  - **leaf** money
  - **leaf** goods
  - **leaf** immaterial
- **or** rewarding conditions
  - **leaf** cast vote
  - **leaf** election win
  - **leaf** unconditional
  - **leaf** complex
- **leaf** group size of reward receivers
- **or** proving compliance
  - **leaf** before rewarding
  - **leaf** after rewarding
  - **leaf** not required
- **leaf** reward related to election

Fig. 1. Objective tree for vote buying

The tree is to be read as follows: before each entry, the type of the entry is marked ('or' for or-nodes, 'and' for and-nodes, 'leaf' for leaves). Or-nodes are nodes of which at least one of the branches must be satisfied, for an and-node to be satisfied, all branches must be satisfied.
In the tree, the characteristic *reward related to election* poses the question whether or not the reward can only be handed out by the election winner. An example of such rewards is granting amnesty, which is not possible unless elected to office.

As not all aspects of vote influencing are of direct interest to the buyer (e.g. how sure is delivery of the reward), other trees for different objectives have been constructed. Noting that this objective tree focuses upon vote acquisition only via rewarding leads to the following dimension:

type of compliance how is compliance achieved?
instrumental (by rewarding), normative (by convincing), coercive.

A vote buyer, however, is not interested in acquiring one vote, but in acquiring many votes. Specifically, a vote buyer's goal is to acquire enough votes to make a difference. This goal of a vote buyer encompasses the above objective of acquiring one vote. A vote buyer could use various means to attempt vote acquisition of many votes (as indicated by the examples). An analysis of this, similar to the above one, uncovered the following dimensions of approaches to vote buying:

non-compliance what impact does non-compliance have?
no rewarding, rewarded anyway.
focusability can only the relevant people be the sole targets of the vote buying method?
highly targetable, less targetable.
scalability how easy is it to employ this rewarding on a large scale?
high, medium, low.
costs do costs vary with the number of acquired votes or voters convinced?
variable, more fixed.
openness/publicity is the persuasion attempt general knowledge?
secret (known to buyer and seller(s)), not hidden, general knowledge.

Considering the point of view of the vote seller introduced the following dimensions for vote buying:

rewarding certainty can rewarding be avoided?
unavoidable, avoidable.
consequences of non-reward what impact will not rewarding have?
high impact, low impact.

Note that these two dimensions are closely related; they can be combined as **commitment to reward**.

Considering these various objectives together (i.e. vote seller, vote buyer for one vote and for a group of votes) led to the following dimension:

proof of compliance who should prove compliance?
proof by buyer, proof by seller. in case of vote buying, proof by seller is expected; in case of promises, proof by buyer is expected

One remarkable observation, given these dimensions, is that absence of receipts (receipt-freeness, see e.g. [2]) is not sufficient to prevent vote buying – it only suffices to prevent proving compliance. Forms of vote buying exist that do not require proving compliance by the individual voter, as can be seen in example 2.

2.3 Classifying the targets

The set of possible targets for voter influencing can be characterised on various levels. From large to small, we distinguish the following:

$$\begin{aligned} & \textit{population} && (1)\\ \supseteq\ & \textit{eligible voters} && (2)\\ \supseteq\ & \textit{registered voters} && (3)\\ \supseteq\ & \textit{voters casting votes} && (4)\\ \supseteq\ & \textit{voters casting valid votes} && (5)\\ \supseteq\ & \textit{voters casting valid votes for vote buyer} && (6) \end{aligned}$$

Additionally, preferences with respect to elections and vote buying differ from person to person. Perhaps some individuals do not mind selling their votes, while others may find the practice so repugnant they will not vote for anyone involved with the practice – even if it is their preferred candidate. We distinguish the following dimensions for voters:

will accept reward	yes / no
initial preference	buyer's choice / other
awareness of attempt	none / heard rumours / fully aware
is a desired target	yes / no
cast vote	buyer's choice / other

This classification can be applied to each of the sets 1–6. This classification extends the work of Acker [1], who classified voters targeted by election promises as follows:

A. Already compliant voters — these would have voted for the coercer without the election promise
B. Voters who change their votes — these vote for the coercer due to the election promise
C. Non-compliant voters — these do not vote for the coercer, despite the election promise

One category missing in that classification is explicitly included by our new classification: the set of voters who, as a result of the vote buying attempt,

change their vote from "buyer's choice" to "other". Intuitively, these voters can be characterised as the voters who find vote buying so repugnant, that they will not vote for anyone involved with the practice.

2.4 Classification of the examples

Below we classify our examples of voter influencing, according to the dimensions[2] established above. In addition, for each example we note which subset of voters is targeted.

Example 1 (handout). **conditions:** unconditional; **group size:** individual; **related:** no; **type:** instrumental; **non-compliance:** rewarded; **focus:** not targetable; **scalable:** low; **costs:** number of attempts; **publicity:** not hidden; **commitment:** unavoidable, low; **proof:** none.
Targetted at voters casting votes (class 4).

Example 2 (theme park). **conditions:** conditional; **group size:** collective; **related:** yes; **type:** instrumental; **non-compliance:** rewarded; **focus:** not targetable; **scalable:** high; **costs:** fixed; **publicity:** public; **commitment:** avoidable, high; **proof:** none.
Targetted at registered voters (class 3).

Example 3 (zalmsnip). **conditions:** conditional; **group size:** collective; **related:** yes; **type:** instrumental; **non-compliance:** rewarded; **focus:** not targetable; **scalable:** high; **costs:** fixed; **publicity:** public; **commitment:** avoidable, high; **proof:** none.
Targetted at registered voters (class 3).

Example 4 (election promise). **conditions:** conditional; **group size:** collective; **related:** yes; **type:** instrumental; **non-compliance:** rewarded; **focus:** highly targetable; **scalable:** high; **costs:** fixed; **publicity:** public; **commitment:** avoidable, high; **proof:** none.
Targetted at a subgroup of registered voters (class 3).

Example 5 (rf). **conditions:** conditional; **group size:** individual; **related:** no; **type:** instrumental; **non-compliance:** unrewarded; **focus:** highly targetable; **scalable:** low; **costs:** number of acquired vote; **publicity:** not hidden; **commitment:** avoidable, low; **proof:** by voter.
Targetted at voters casting valid votes (class 5).

Example 6 (non-rf). **conditions:** conditional; **group size:** individual; **related:** no; **type:** instrumental; **non-compliance:** unrewarded; **focus:** highly targetable; **scalable:** low; **costs:** number of compliant voters; **publicity:** not hidden; **commitment:** avoidable, low; **proof:** by seller.
Targetted at voters casting valid votes (class 5).

[2] Both *time of reward* and *type of reward* have been left out, as these are already explicit in the examples.

Example 7 (reimburse). **conditions:** unconditional; **group size:** individual; **related:** no; **type:** instrumental; **non-compliance:** rewarded; **focus:** not targetable; **scalable:** medium; **costs:** number of requesting voters; **publicity:** public; **commitment:** unavoidable, –; **proof:** not required.
Targetted at registered voters (class 3).

We find that, based on our distinctions, we can easily classify these examples. The question remains which attributes indicate acceptable and unacceptable forms, respectively. From our examples and their intuitive acceptability, we propose that benefits that are related to the contested position, are unconditional and openly announced, are most likely to be found legitimate.

2.5 Conclusions on vote influencing

We conclude that vote buying involves much more than offering money in exchange for a proof of compliance. Attributes that make it likely for an action to be considered vote buying include:

- unrelated to contested position;
- reward independent of being elected;
- reward conditional on compliance (therefore proof by seller);
- highly targetable;
- variable costs (related to individual payment);
- secrecy of the attempt.

Individuality does *not* make things worse; buying a whole district is in itself no better than buying votes one by one. The publication of any election result on a level lower than strictly necessary (e.g. per polling station) facilitates collective vote buying, and would best be eliminated from this perspective. Electronic voting can facilitate such a transition, by storing votes independently from the place where they were cast.

However, this is by itself not sufficient to stop collective vote buying. If a buyer wants to buy a set of votes, and knows with 70% certainty that an individual voter complies, she can be fairly sure that if she buys a large amount of votes, 70% of the votes will be hers, due to the law of large numbers.

Conversely, if in a particular set of votes 70% is for the buyer, she can derive that a voter whose vote is in this set has voted for her with a 70% probability. This particular observation gives rise to the notion of *probabilistic vote buying* – where a buyer requires not exact votes, but is satisfied by a (significant) change in the distribution of votes.

Furthermore, the *non-rf* example showed that it is possible to *discourage* voting for a certain party or candidate in an action of vote buyinng. This is an example of voter influencing that cannot be directly described in terms of the intuitive notion of receipt-freeness, but is close to it.

In the next section, we categorise anonymity in voting based on these observations.

3 Dimensions of anonymity in voting

Traditionally, the concept of vote buying has been related to the possibility of providing a proof of one's choice. The notion of receipt-freeness was proposed to prevent such a proof. However, our framework developed in the previous sections shows that a proof is not always necessary. The following actions would be possible without a proof of the voter's choice in a strict sense:

- rewarding the voter if she does *not* vote for a specific party or candidate (related to negative proof);
- rewarding the voter if it is *likely* that she made a certain choice.

It can be enough for a buyer to hand out the reward if a voter can show that she did *not* vote for two of the buyer's opponents (example 6). The buyer could also pay a voter if after observing the outcome, it is *more* likely that this voter voted for him than that another voter voted for him. If this can be observed from messages sent in the voting protocol, this should be addressed by computer science verification methods. One could also derive this from voter behaviour [4], but that is hard to prevent using computer science tools.

Each of the notions of privacy, receipt-freeness and coercion-resistance can be investigated with respect to these scenarios:

	weak	strong	probabilistic
privacy	this work	this work	Barghava et al.
receipt-freeness	this work	this work	–
coercion-resistance	Delaune et al.	–	–

The distinction between the three notions of anonymity depends on the relation between voter and attacker. In case of privacy, no cooperation of the voter is assumed; the attacker tries to find out about the voter's choice without cooperation the voter. In case of receipt-freeness, the voter cooperates by sending information to the attacker. In case of coercion-resistance, the voter even accepts instructions from the attacker.

The *weak* variant of the notion then concerns the situation that an attacker cannot be sure of the voter's choice. The *strong* variant of the notion concerns the situation that an attacker cannot be sure of what the voter did *not* choose either. The *probabilistic* variant of the notion concerns the situation what an attacker cannot deduce anything from the *probability distribution* of a voter's choices.

In the next section, we formalise the notions of strong and weak privacy and receipt-freeness in an epistemic framework. Probabilistic privacy has been investigated in [3] (where it is called probabilistic anonymity). The notion of coercion-resistance is defined in [6]. Filling the remaining fields in the table based on these definitions is future work.

4 Formalising anonymity properties

Garcia, Hasuo, Pieters and Van Rossum defined a framework for describing anonymity properties of protocols in epistemic logic [7]. We use this framework as the basis for our definitions of privacy and receipt-freeness. Their definitions are based on a formalisation of runs in protocols. For the formal definitions, we refer to the original paper.

Based on observational equivalence of runs, the following notions are defined formally in the original paper. We include the informal definitions here. The formula A Sends m to B means: at some stage in the run, A sends a message to B which *contains* m *as a subterm*. A Sends m means that A sends the message m to someone. The formula A Possesses m means: after the run has finished, A is capable of constructing the message m. The formula A Originates m means that A Sends m, but A is not relaying. More precisely, m does not appear as a subterm of a message which A has received before.

The formula $\Box A\varphi$ is read as "after the run is completed, the agent A *knows* that φ is true". The formula $\Diamond A\varphi$ is short for $\neg\Box A\neg\varphi$ and read as "after the run is completed, the agent A *suspects* that φ is true".

Garcia et al. define the information hiding properties of sender anonymity, unlinkability and plausible deniability using the notion of an *anonymity set* (a collection of agents among which a given agent is not identifiable) in epistemic logic as follows:

Definition 1. *(Sender anonymity) Suppose that r is a run of a protocol in which an agent B receives a message m. We say that r provides* sender anonymity *with anonymity set* AS *if it satisfies*

$$r \models \bigwedge_{X \in \mathsf{AS}} \Diamond B(X \text{ Originates } m).$$

This means that, as far as B is concerned, every agent in the anonymity set could have sent the message.

Definition 2. *(Unlinkability) A run r provides* unlinkability *for users A and B with anonymity set* AS *iff*

$$r \models \big(\neg\Box\mathsf{spy}\varphi(A,B)\big) \wedge \bigwedge_{X \in \mathsf{AS}} \Diamond\mathsf{spy}\varphi(X,B)\ ,$$

where $\varphi(X,Y) = \exists n.\ \big(X \text{ Sends } n \wedge Y \text{ Possesses } n\big)$.

Intuitively, the left side of the conjunction means that the spy (the adversary) is not certain that A sent something to B. The right conjunct means that, according to the spy, every other user could have sent a message to B. Similarly, unlinkability between a user A and a message m could be defined as $\models \neg\Box\mathsf{spy}(A \text{ Sends } m) \wedge \bigwedge_{X \in \mathsf{AS}} \Diamond\mathsf{spy}(X \text{ Sends } m)$.

In certain circumstances (e.g. relays), agents might be interested in showing that they did not know that they had some sensitive information m. This might be modeled by the following epistemic formula:

Definition 3. *(Plausible deniability) Agent A can* plausibly deny *message m in run r iff*

$$r \models \Box \mathsf{spy} \neg (\Box A(A \ \mathsf{Possesses} \ m)) \ .$$

This formula is read as: the spy knows that A does not know that she possesses m.

We extend this set of definitions by providing the additional property of *receipt-freeness*. Receipt-freeness of an agent A with respect to a message m (e.g. a vote) intuitively means that A cannot send a message m' to the spy that proves that she sent m in the past. For this purpose, the definition of plausible deniability is too strong, since A *does* know that she possesses m. Sender anonymity is particularly useful for providing anonymity of the voter with respect to the election authorities, but in receipt-freeness, A herself tries to communicate with the spy. Instead, it should not be possible to link A to her vote. Thus, unlinkability seems the most natural property to base our definition of receipt-freeness upon.

In the anonymity framework, the concept of anonymity set is used to define the set of entities between which an observer should not be able to distinguish. To apply the framework to votes, we need to adapt the concept of anonymity set. In voting, we are sure that each (actual) voter submits a vote. Therefore, the point is not whether any other user in an anonymity set could have sent the message, but *whether the voter could have submitted any other vote*. Therefore, we define an anonymity set of *messages*, AMS, instead of an anonymity set of agents. This set typically consists of all possible votes.

First of all, we define the notion of (weak) privacy. This can be achieved without referring to the anonymity set.

Definition 4. *(Weak privacy) A run of a protocol is* weakly private *for agent A with respect to message m iff*

$$r \models \neg \Box \mathsf{spy}(A \ \mathsf{Sends} \ m)$$

To be able to define receipt-freeness, we need to have a way to extend a given run with one message: the receipt. We write this as $r.(A \rightarrow B : m))$ for a given run r, message m (the receipt), sender A and receiver B. For A to be able to send the receipt, she needs to have the message in her possessions at the end of the original run. The new run does *not* need to be a run of the protocol. It *does* need to be legitimate with respect to the initial possession function. $\mathsf{Poss}_{\mathsf{IPo}}(r, A, |r| - 1)$ denotes the possessions of agent A at the end of the original run (see [7]).

Definition 5. *(Weak receipt-freeness) A run of a protocol is* weakly receipt-free *for agent A with respect to message m iff for all $m' \in \mathsf{Poss}_{\mathsf{IPo}}(r, A, |r| - 1)$,*

$$r.(A \rightarrow \mathsf{spy} : m') \models \neg \Box \mathsf{spy}(A \ \mathsf{Sends} \ m)$$

Weak receipt-freeness implies that the voter cannot prove to the spy that she sent message m during the protocol, where m is the (part of a) message representing

the vote. However, this notion is still fairly limited. For example, suppose that the spy wants the voter to vote for party X. Suppose, furthermore, that the voter instead chooses to vote Y, which is represented by message m in the above definition. Now, if the voter cannot show that she voted Y, this protocol is receipt-free with respect to the definition above. However, if the spy can acquire information which proves that the voter did *not* vote X, the spy will not be satisfied. Therefore, we introduce stronger notions of privacy and receipt-freeness as well.

Definition 6. *(Strong privacy) A run of a protocol is* strongly private *for agent A with respect to a message m in anonymity set* AMS *iff*

$$r \models (\neg\Box\mathsf{spy}(A \text{ Sends } m)) \wedge \bigwedge_{m'' \in \mathsf{AMS}} \Diamond\mathsf{spy}(A \text{ Sends } m'')$$

Definition 7. *(Strong receipt-freeness) A run of a protocol is* strongly receipt-free *for agent A with respect to a message m in anonymity set* AMS *iff for all $m' \in \mathsf{Poss}_{\mathsf{IPo}}(r, A, |r| - 1)$,*

$$r.(A \rightarrow \mathsf{spy} : m') \models (\neg\Box\mathsf{spy}(A \text{ Sends } m)) \wedge \bigwedge_{m'' \in \mathsf{AMS}} \Diamond\mathsf{spy}(A \text{ Sends } m'')$$

Here, no matter what information the voter supplies to the spy, *any* vote in the anonymity set is still possible. This is represented by the "suspects" symbol $\Diamond\mathsf{spy}$. In other words, for all possible votes, the spy still suspects that the voter cast this particular vote; or: the spy is not certain she did *not* cast this vote. This requires that at least one message has been received (i.e. at least one vote has been cast) for every message (vote) $m'' \in \mathsf{AMS}$. Otherwise, the spy could observe from the results that no-one, in particular not voter A, cast a certain vote. Thus, for votes, AMS is a subset of the set of candidates who received votes.

Notice that this definition is analogous to the definition of unlinkability of Garcia et al.

Theorem 1. *If a run of a protocol is strongly receipt-free for agent A with respect to message m in anonymity set* AMS*, then it is also weakly receipt-free for agent A with respect to message m.*

Proof. This follows directly from the definitions.

In the framework, it can be defined which part of the messages the spy can observe. If the spy could read all the messages, the voter only needs to supply the secret keys in order to provide a receipt. This is not what is commonly understood by analyzing receipt-freeness. Instead, there are certain messages in the protocol that the spy is not assumed to have access to (when the voter is in a "voting booth").

In our definition, we deviate from the approach by Delaune et al. [6]. Intuitively, receipt-freeness is achieved if a voter does not possess convincing, exclusive evidence of how she voted. The approach by Delaune et al. defines receipt-freeness using two voters (to preserve indistinguishability of the result). By focusing on the actual receipt, our definition only relies on one voter, and thus remains closer to the intuition. The indistinguishability is made explicit in our definition by AMS, which does not need to encompass all candidates but can be confined e.g. to all candidates for whom at least one vote was cast.

5 Conclusions and future work

In this paper, we examined various dimensions of the notion of anonymity in electronic voting. Based on an analysis of the acceptability of vote buying and coercion, we found that many different dimensions contribute to whether an action is classified as vote buying or coercion, or not. Having established these dimensions, we distinguished between *weak*, *strong* and *probabilistic* notions of anonymity in voting. This distinction applies to privacy, receipt-freeness and coercion-resistance.

Following the definitions of anonymity in [7], we introduced an approach to formally verify weak and strong receipt-freeness in epistemic logic. To the best of our knowledge, we are the first to do so. The approach has not been tested on real protocols thus far.

One of the main benefits of our approach is the intuitive definition that it provides for receipt-freeness. As opposed to other approaches, especially [6], the "receipt" can easily be distinguished in our model as a separate message that the voter sends to the spy. Instead of investigating whether the spy can recover the vote from forwarded messages, we judge whether the spy really *knows* what the voter's choice was, based on any possible receipt. This notion of *knows* is characteristic for the epistemic logic approach, and this justifies our choice for the anonymity framework of [7] as a basis.

While our approach captures receipt-freeness, the investigation of the notion of vote influencing clearly indicates that compliance with a technical criterion such as receipt-freeness is insufficient to prevent voter influencing – if unaware, a voter could still be susceptible to voter influencing. Further study in this area (see also [16]) is needed to ensure that security requirements not only enable secure voting, but that voters are sufficiently aware of the security they provide.

In future work, we aim at providing an alternative definition of receipt-freeness in our model, based on the knowledge of the spy instead of on extension of a run. We hope to prove that the two definitions are equivalent. Moreover, we wish to apply the definitions to existing voting protocols, in order to prove (or disprove) receipt-freeness. It may also be interesting to investigate the relation between verifiability [15] and receipt-freeness in epistemic logic, since both receipt-freeness and verifiability are based on an agent's knowledge instead of its possessions.

As we have seen, voter influencing can take many forms, and only some can be hindered by technological means. However, being more precise in what we can and cannot achieve in technology can lead to better decisions on which protocols are acceptable and which are not. Again, the latter will depend on what is seen as acceptable in a particular culture, which has been determined during a long history of success and fraud.

References

1. B. van Acker. Remote e-voting and coercion: a risk-assessment model and solutions. In *Electronic Voting in Europe 2004*, pages 53–62, 2004.
2. J.C. Benaloh and D. Tuinstra. Receipt-free secret ballot elections (extended abstract). In *Proc. 26th ACM Symposium on the Theory of Computing (STOC)*, pages 544–553. ACM, 1994.
3. M. Bhargava and C. Palamidessi. Probabilistic anonymity. In Martín Abadi and Luca de Alfaro, editors, *Proceedings of CONCUR 2005*, number 3653 in Lecture Notes in Computer Science. Springer, 2005.
4. V. Brusco, M. Nazareno, and S.C. Stokes. Vote buying in Argentina. *Latin American Research Review*, 39(2):66–88, 2004.
5. S. Delaune, S. Kremer, and M.D. Ryan. Receipt-freeness: Formal definition and fault attacks (extended abstract). In *Proceedings of the Workshop Frontiers in Electronic Elections (FEE 2005)*, Milan, Italy, September 2005.
6. S. Delaune, S. Kremer, and M.D. Ryan. Coercion-resistance and receipt-freeness in electronic voting. In *Proceedings of the 19th IEEE Computer Security Foundations Workshop (CSFW'06)*, Venice, Italy, July 2006. IEEE Computer Society Press.
7. F. D. Garcia, I. Hasuo, W. Pieters, and P. van Rossum. Provable anonymity. In Ralf Küsters and John Mitchell, editors, *3rd ACM Workshop on Formal Methods in Security Engineering (FMSE 2005)*, pages 63–72. ACM Press, 2005.
8. R.L. Hasen. Vote buying. *California Law Review*, 88(5):1323–1371, October 2000.
9. M. Hirt and K. Sako. Efficient receipt-free voting based on homomorphic encryption. In B Preneel, editor, *Proc. EUROCRYPT 2000*, volume 1807 of *LNCS*, pages 539–556, 2000.
10. H.L. Jonker and E.P. de Vink. Formalising receipt-freeness. In Michael Backes, Stefanos Gritzalis, and Bart Preneel, editors, *ISC'06*. To appear, 2006.
11. A. Juels, D. Catalano, and M. Jakobsson. Coercion-resistant electronic elections. In *Proc. WPES'05*. ACM, 2005.
12. M.S. Kochin and L.A. Kochin. When is buying votes wrong? *Public Choice*, 97:645–662, 1998.
13. S. Mauw and M. Oostdijk. Foundations of attack trees. In D. Won and S. Kim, editors, *Proc. 8th Annual International Conference on Information Security and Cryptology, ICISC'05*, number 3935 in LNCS, pages 186–198. Springer, 2006.
14. S. Mauw, J. Verschuren, and E.P. de Vink. A formalization of anonymity and onion routing. In P. Samarati, P. Ryan, D. Gollmann, and R. Molva, editors, *Proc. esorics 2004*, pages 109–124, Sophia Antipolis. LNCS 3193.
15. W. Pieters. What proof do we prefer? variants of verifiability in voting. In P. Ryan, S. Anderson, T. Storer, I. Duncan, and J. Bryans, editors, *Workshop on e-Voting and e-Government in the UK*, pages 33–39, Edinburgh, February 27-28 2006. e-Science Institute, University of St. Andrews.

16. W. Pieters. *La Volonté Machinale – understanding the electronic voting controversy*. PhD thesis, Radboud University, Nijmegen, The Netherlands, 2007.
17. K. Sako and J. Kilian. Receipt-free mix-type voting scheme – a practical solution to the implementation of a voting booth. In L.C. Guillou and J.-J. Quisquater, editors, *EUROCRYPT'95*, volume 921 of *LNCS*, pages 393–403. Springer, 1995.
18. F.C. Schaffer and A. Schedler. What is vote buying? In F.C. Schaffer, editor, *Elections for Sale: The Causes and Consequences of Vote Buying*. Lynne Rienner, Boulder CO, 2007.
19. B. Schneier. Attack trees: Modeling security threats. *Dr. Dobb's journal*, December 1999.