Vote buying revisited: implications for receipt-freeness

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

[1] Institute for Computing and Information Sciences
Radboud University Nijmegen
P.O. Box 9010, 6500 GL Nijmegen, The Netherlands
wolterp@cs.ru.nl

[2] Department of Mathematics and Computer Science
Eindhoven University of Technology
P.O. Box 513, 5600 MB Eindhoven, The Netherlands
h.l.jonker@tue.nl

[3] Faculty of Sciences, Communication and Technology
University of Luxembourg
6, rue Richard Coudenhove-Kalergi, L-1359 Luxembourg
hugo.jonker@uni.lu

Abstract. In this paper, we analyse the concept of vote buying based on examples that try to stretch the meaning of the concept. Which examples can still be called vote buying, and which cannot? We propose several dimensions that are relevant to qualifying an action as vote buying or not. As a means of protection against vote buying and coercion, the concept of receipt-freeness has been proposed. We argue that, in order to protect against a larger set of vote buying activities, the concept of receipt-freeness should be interpreted probabilistically. We propose a general definition of probabilistic receipt-freeness by adapting existing definitions of probabilistic anonymity to voting.

Keywords: vote buying, probabilistic anonymity, probabilistic receipt-freeness

1 Introduction

By law[4], it is prohibited to offer money or other personal benefits in exchange for votes in an election. On the other hand, promising something to a group of people on the condition of being elected, without explicitly requiring their votes, is perfectly legal. In this paper, we investigate the different dimensions of vote buying, and we discuss implications for the notion of receipt-freeness.

Clarifying the distinction between what constitutes vote-buying and what constitutes an election promise is necessary to understand what we consider threats to voting systems. Until now, the absence of a receipt that would prove

[4] E.g. Dutch criminal law (Wetboek van Strafrecht) art. 126.

how a voter voted has been considered sufficient to ensure that vote buying cannot occur. However, this does not capture all threats related to voter persuasion – something already indicated in e.g. [9] (introducing coercion-resistance) and [8] (introducing strong receipt-freeness).

The goal of this paper is to approach the subject not by extending previously found threats, but by determining the possibilities for influencing voters and analysing where the boundary between allowed and illegal influencing is. The paper thus clearly states the considerations an election system has to take into account and the choices an election system can make with respect to this distinction. We use these considerations to propose stronger definitions of voting system properties aimed at preventing vote buying and coercion.

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 [10] 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 [7] 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"[5] elections.

Schaffer [12] 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, 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

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

the elections are being held and the participants of the elections. In the end, it seems that 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 introduce several dimensions on which voter influencing activities can be classified. Furthermore, we distinguish sets of voters that can be targets to vote buyers. We present several examples, and classify these according to our dimensions. We conclude by listing the attributes of a voter influencing activity that make it likely for the action to be considered an instance of vote buying.

In section 3, we discuss implications of our analysis of vote buying for computer science. Most importantly, we argue that existing concepts defined to describe the protection of voting systems against vote buying should be strengthened. In section 4, we present probabilistic definitions of anonymity and receipt-freeness in voting systems, by adapting the definitions from [3, 5]. In section 5, we present our conclusions and suggestions for future work.

2 Coercion, buying, persuasion

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. At the polling station, I give each voter 100 euros together with mentioning my candidacy.

Example 2. The district with the highest percentage of votes for me gets a theme park.

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).

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. Given that, example 2 can be considered the most dubious of this list.

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. 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 play a role in flipping the acceptability, we construct an attack tree (see [13, 11]) of voter influencing in Section 2.2.

This attack tree deviates slightly from the norm. The purpose of the attack tree is to determine the characteristics that distinguish acceptable from unacceptable influence. To elucidate these detailed characteristics, details need to be explicit in the attack tree. Hence, we do not use attributes (which would hide the exact characteristics that are of interest). Instead of attributes, leaves are used.

2.2 Classifying vote buying

Based on the literature, the examples and the analysis above, the attack tree in Figure 1 was constructed and dimensions of vote buying were clarified. The main goal in the attack tree is to get a "Vote for You" (abbreviated V4Y). This means that the attack tree is from the perspective of a vote buyer. As not all aspects of vote buying are of direct interest to the buyer (e.g. how sure is delivery of the reward), there are several dimensions which do not emerge from the attack tree. The dimensions that do emerge are listed below.

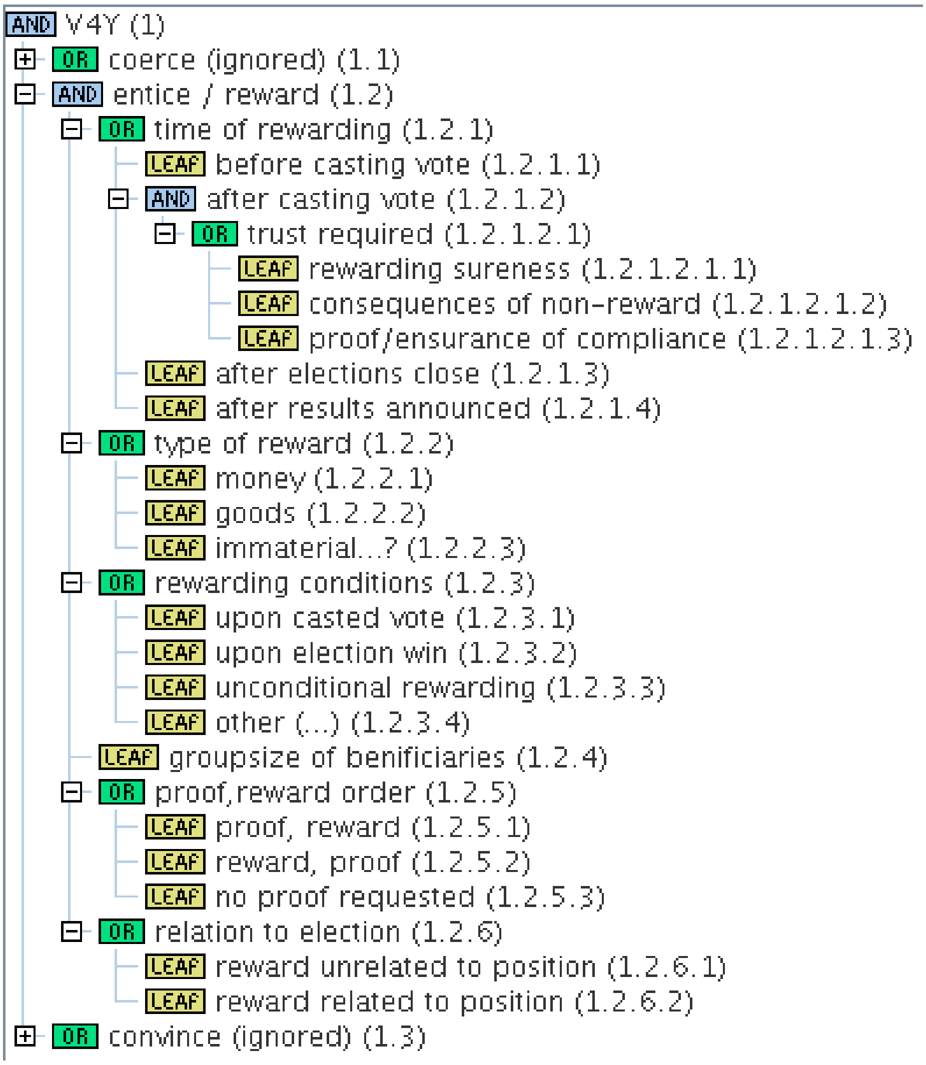

```
AND V4Y (1)
⊞ OR coerce (ignored) (1.1)
⊟ AND entice / reward (1.2)
    ⊟ OR time of rewarding (1.2.1)
        LEAF before casting vote (1.2.1.1)
        ⊟ AND after casting vote (1.2.1.2)
            ⊟ OR trust required (1.2.1.2.1)
                LEAF rewarding sureness (1.2.1.2.1.1)
                LEAF consequences of non-reward (1.2.1.2.1.2)
                LEAF proof/ensurance of compliance (1.2.1.2.1.3)
        LEAF after elections close (1.2.1.3)
        LEAF after results announced (1.2.1.4)
    ⊟ OR type of reward (1.2.2)
        LEAF money (1.2.2.1)
        LEAF goods (1.2.2.2)
        LEAF immaterial...? (1.2.2.3)
    ⊟ OR rewarding conditions (1.2.3)
        LEAF upon casted vote (1.2.3.1)
        LEAF upon election win (1.2.3.2)
        LEAF unconditional rewarding (1.2.3.3)
        LEAF other (...) (1.2.3.4)
    LEAF groupsize of benificiaries (1.2.4)
    ⊟ OR proof,reward order (1.2.5)
        LEAF proof, reward (1.2.5.1)
        LEAF reward, proof (1.2.5.2)
        LEAF no proof requested (1.2.5.3)
    ⊟ OR relation to election (1.2.6)
        LEAF reward unrelated to position (1.2.6.1)
        LEAF reward related to position (1.2.6.2)
⊞ OR convince (ignored) (1.3)
```

Fig. 1. Attack tree for vote buying

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

time of rewarding before casting a vote — after casting a vote — after elections close — after results announced

type of reward monetary — goods — immaterial

rewarding conditions upon casted vote — upon election win — unconditional rewarding — other (consider, for example, the theme park example)

group size of beneficiaries the number of people targeted in a single action

proving compliance can happen before the reward is received, after the reward is received or need not be considered.

reward related to election i.e. can the reward only be given by the winner of the election?

As mentioned, there are additional characteristics which did not fit this particular attack (a vote buyer buying one vote for himself). Below is a list of these other dimensions found.

non-compliance no rewarding — rewarding for group members irrespective of compliance (proof not necessary)

focusability highly targetable — less targetable: is it possible to target relevant individuals/groups only?

scalability low — high: how easy is it to employ this rewarding on a large scale?

costs only variable — more fixed: are costs related to the number of acquired votes? variable in terms of number of votes targeted or number of voters convinced?

rewarding certainty unavoidable — avoidable (need not reward): vote buying gives more certainty than election promises

consequences of non-reward high impact — low impact: what happens if the promised reward is not given?
(note that this and the previous dimension are closely related; they can be combined as “commitment to reward (high—low)”)

proof of 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

openness/publicity is the persuasion attempt general knowledge? must it remain secret, known only to seller and buyer? or is it not general knowledge, but there are no special restrictions, no special efforts to hide the buying attempt?

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.

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:

1. the entire population;
2. eligible voters;
3. registered voters;
4. voters casting votes;
5. voters casting valid votes;
6. voters casting valid votes for the influencer's choice.

Each class is a subset of the previous classes, and a superset of the following (as depicted in Figure 2 – the set of voters casting valid votes (set 5 is divided into those votes in favour and those not in favour of the influencer's choice). Influencing can be directed towards any level, and need not be confined to the set 5/6.

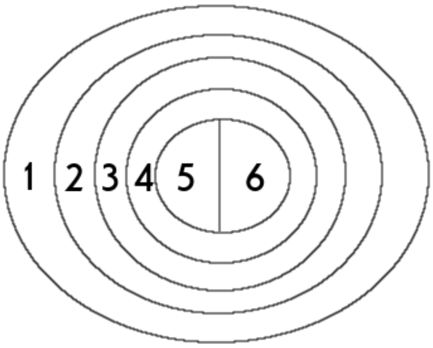

Fig. 2. Voter classes

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. Hence, people can be classified as follows:

will accept reward yes / no
initial preference V4Y / not V4Y
awareness of attempt none / partially aware (heard rumours) / fully aware
is a desired target yes / no
cast vote V4Y / not V4Y

This classification can be applied to each of the sets depicted in Figure 2. 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 'V4Y' to 'not V4Y'. 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 established above.

Example 5. At the polling station, I give each voter 100 euros together with mentioning my candidacy.
Individual, direct, unrelated, unconditional, irrespective, not targetable, low scalability, variable w.r.t. turnout, unavoidable, no impact, no proof, buyer delivers first, targets voters casting votes.

Example 6. The district with the highest percentage of votes for me gets a theme park.
Collective, postponed, related, conditional (but not necessarily upon election), respective (but collective!), not targetable, high scalability, fixed costs, relatively unavoidable (with respect to "normal" promises), high impact (therefore relatively unavoidable), proof by buyer, seller delivers first, targets registered voters.

Example 7 (zalmsnip). If I get elected, everyone gets 100 euros tax refund.
collective, postponed, related (loosely), conditional, irrespective, untargeted, high scalability, fixed costs, avoidable, high impact, proof by seller, seller delivers first, targets registered voters.

Example 8 (election promise-style). If I get elected, disabled child prodigies get 100 euros (i.e. children with a physical handicap, who are members of Mensa).
collective (but approaching individual?), postponed, related, conditional, irrespective, very targeted, fixed costs, avoidable, impact equal to importance in campaign (important promise implies high impact), proof by seller, seller delivers first, targets (a specific group of) registered voters.

Example 9. collective (but approaching individual?), postponed, related, conditional, irrespective, very targeted, fixed costs, avoidable, impact equal to importance in campaign (important promise implies high impact), proof by seller, buyer delivers first.

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 buying

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.

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 bunch 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.

3 Implications for computer science

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. The issue of negative proof has been addressed before by the notion of *strong receipt-freeness* [8]. Here, we focus on the second case.

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.

As a reward may depend on the probability that a voter voted for the specified candidate, we conclude that vote buying can be probabilistic. Because of this extended concept of vote buying, the traditional notion of receipt-freeness needs to be reinvestigated.

4 Probabilistic anonymity and receipt-freeness

Recently, Bhargava and Palamidessi [3] and Deng, Palamidessi and Pang [6] introduced the notion of probabilistic anonymity. Given anonymous users U, sets of anonymous actions $a(u)$, for $u \in U$, observation sequences O, and a probability measure p, they define anonymity of the system as follows (for the precise definitions, see [3]):

Definition 1. *A fully probabilistic system (M, U, a, B, p) is anonymous if*

$$\forall i, j \in U, \forall o \in O :$$

$$p(a(i)) > 0 \wedge p(a(j)) > 0 \implies p(o \mid a(i)) = p(o \mid a(j))$$

Intuitively, this means that the probability of a certain observation by the attacker does not depend on the anonymous actions.

It is claimed that this notion can be applied to various kinds of systems, such as voting protocols. However, this only captures one aspect of privacy. what if it can be derived that a user performed action b rather than action a? Let's say, a voter voted liberals instead of social democrats. This still meets the indistinguishability requirement in the definition above, since this definition only speaks about indistinguishability of the *same* action by *different* users. In voting, we also need indistinguishability of *different* actions of the *same* user. This second notion is similar to the notion of uncertainty (see e.g. [14] for a description of these two notions in the area of statistical databases).

Instead of the set of actions $a(u)$ in the original definition, we now define a set of actions $a(v, c)$, where v is a voter and c is a candidate (in C). Uncertainty is ensured by:

$$\forall v, w \in V, \forall c, d \in C, \forall o \in O :$$

$$p(a(v, c)) > 0 \wedge p(a(w, c)) > 0 \implies p(o \mid v(v, c)) = p(o \mid a(w, c))) \wedge$$

$$p(a(v, c)) > 0 \wedge p(a(v, d)) > 0 \implies p(o \mid a(v, c)) = p(o \mid a(v, d)))$$

However, the second condition of probabilistic anonymity (we cannot distinguish between different choices of the same voter) can only be assured to a certain degree, since the probability of a certain election outcome is not independent from the choice of an individual voter. If the election is unanimous, we know *for sure* what each voter voted. In this sense, only *weak* probabilistic anonymity could be achieved.

Deng, Palamidessi and Pang [6] describe weak probabilistic anonymity by means of a maximal difference between observation probabilities depending on anonymous actions, termed α. However, when describing anonymity of voting protocols in such a way, the α-value will typically depend on the number of voters and candidates. For example, if there are only two voters and two candidates, a unanimous result for candidate c will be impossible, given that a particular voter voted for candidate d. On the other hand, given that this voter made the "right" choice (i.e. candidate c), the unanimous and the divided outcome will both have probability .5. In this case, the maximum difference between the probabilities is .5. If there are more than two voters, or more than two candidates, the probabilities decrease and hence this difference will decrease.

We argue that this dependency is an undesired property, and we propose an alternative definition of probabilistic anonymity that establishes more desirable properties for voting systems.

Our definition is inspired by the work of Delaune, Kremer and Ryan [5], who formalised the different notions of anonymity in voting (privacy, receipt-freeness and coercion-resistance). They consider anonymity as the property that it is undetectable if two voters *swap* votes. This prevents the problem where a unanimous election is not anonymous according to the definition. In the definition below, c and d are votes, and a is a function mapping voters and votes to actions. For a voter v and vote c, the event $a(v,c)$ means that in the execution of the protocol the action $a(v,c)$ takes place. Probability measures are extended to encompass more than one event, i.e. $p(a;b)$ is the measure that both event a and b occur.

Definition 2. *A voting protocol with anonymous voters V, candidates (i.e. possible votes) C and possible observation sequences O is* probabilistically anonymous *iff*

$$\forall v, w \in V, \forall c, d \in C, \forall o \in O :$$

$$p(\ a(v,c)\ ;\ a(w,d)\ ) > 0\ \wedge\ p(\ a(v,d)\ ;\ a(w,c)\ ) > 0 \implies$$

$$p(\ o \mid (a(v,c)\ ;\ a(w,d))\ )\ =\ p(\ o \mid (a(v,d)\ ;\ a(w,c))\ )$$

We formalise receipt-freeness as well. Here, the information to be protected consists of two votes and a receipt that the voter sends to the coercer. This receipt may contain private information that the voter received during the protocol. r is a function mapping a voter and a vote to a receipt event.

Definition 3. *A voting protocol with anonymous voters V, candidates C and possible observation sequences O is* probabilistically receipt-free *if*

$$\forall v, w \in V, \forall c, d \in C, \forall o \in O :$$

$$p(a(v,c)\ ;\ a(w,d)\ ;\ r(v,c)) > 0\ \wedge\ p(a(v,d)\ ;\ a(w,c)\ ;\ r(v,c)) > 0 \implies$$

$$p(o \mid (a(v,c)\ ;\ a(w,d)\ ;\ r(i,x)))\ =\ p(o \mid (a(v,d)\ ;\ a(w,c)\ ;\ r(i,x)))$$

Note that the receipt is in the *conditional* part of the probability here. Intuitively, the definition means that if the receipt changes the probability of certain observations depending on who votes what, then the protocol is not receipt-free. Also, if the likelihood of being able to construct the receipt depends on whether an actual vote for that candidate has been cast, then the two probabilities will be different.[6]

Typically, the function r generating the receipt events will be dependent on what the voter knows at the end of the protocol run. If the voter could not have sent a particular receipt in any run of the protocol leading to a particular observation, then the probability of the corresponding observation given the receipt will be zero, and therefore not dependent on the votes. Such constraints should be taken care of in the protocol specification.

The definitions are intentionally general, and do not refer to a particular protocol semantics. In future work, we will try to incorporate these definitions in frameworks for automated verification.

5 Conclusions

Vote buying involves much more than offering money in exchange for a proof of vote. We distinguished many different dimensions. In particular, vote buying need not target individuals. Information about a set of votes may provide probabilistic information about the individuals whose votes are in this set. We argue that based on our analysis, the notion of receipt-freeness needs to be strengthened, and should be defined probabilistically. However, existing definitions of probabilistic anonymity are unsuitable for voting. We adapted the definition by requiring that it should be indistinguishable if two voters swap votes. We applied this definition to receipt-freeness as well. These definitions can be implemented in tools for automated reasoning about voting protocols, which is future work.

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. 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.

[6] There may be better ways to express this independence in case we assume the users to be non-deterministic rather than probabilistic. This is future work.

6. Y. Deng, C. Palamidessi, and J. Pang. Weak probabilistic anonymity. In *Proceedings of the 3rd International Workshop on Security Issues in Concurrency (SecCo)*, Electronic Notes in Theoretical Computer Science. Elsevier Science Publishers, 2005.
7. R.L. Hasen. Vote buying. *California Law Review*, 88(5):1323–1371, October 2000.
8. H.L. Jonker and W. Pieters. Receipt-freeness as a special case of anonymity in epistemic logic. In *Proceedings of the IAVoSS Workshop On Trustworthy Elections (WOTE 2006)*, Robinson College, Cambridge, June 28 – June 30 2006.
9. A. Juels, D. Catalano, and M. Jakobsson. Coercion-resistant electronic elections. In *Proc. WPES'05*. ACM, 2005.
10. M.S. Kochin and L.A. Kochin. When is buying votes wrong? *Public Choice*, 97:645–662, 1998.
11. 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.
12. 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.
13. B. Schneier. Attack trees: Modeling security threats. *Dr. Dobb's journal*, December 1999.
14. Chao Yao, Lingyu Wang, Xiaoyang Sean Wang, and Sushil Jajodia. Indistinguishability: The other aspect of privacy. In *Secure Data Management*, pages 1–17, 2006.