Verifiability of electronic voting

between confidence and trust

Wolter Pieters

University of Twente

Abstract When computing scientists speak about electronic voting, it is often in terms of trust. But there are two contradictory statements. First, they argue that it should not be necessary to trust e-voting systems, which would be the case if they are provably secure. Second, for an e-voting system to be successful, the public must trust it. When we unravel the confusing concept of trust, we find that there are two quite different meanings: relying on something that one does not understand and does not really choose (confidence), or relying on something that one does understand and has consciously chosen (trust). The distinction is due to the German sociologist Niklas Luhmann. In this contribution, we analyse how this distinction can help in analysing the controversies around electronic voting. It is argued that because of the controversy, paper voting and e-voting now tend to be seen as radically different alternatives, which require comparison and a conscious decision. Trustworthiness, as opposed to reliability only, has thereby become a major requirement of electronic voting systems, leading to the implementation of various verification options. This increasingly applies to other systems that handle sensitive data as well. We will discuss the various types of verifiability in electronic voting systems, and how these can contribute to trustworthiness of data processing in general.

1 Introduction

In many countries, controversies exist or have existed on the acceptability of electronic forms of voting in elections. These may include both electronic voting machines at polling stations and Internet voting. In the Netherlands, the use of electronic voting machines in elections was discontinued after a pressure group had raised concerns about the secrecy of the ballot and the verifiability of election results. Internet voting experiments were halted as well. The trajec-

tory from the broadcasting of the first findings until the abandonment of the machines was full of discussion and differences in risk estimation.

In the dynamics of such controversies, trust plays a major role. Do we trust electronic voting machines to accurately count the votes? And do we trust government measurements of possible secrecy problems due to radiation of the machines? In this chapter, we analyse the importance of trust from the perspective of the Dutch electronic voting controversy. We also address the relation with different types of verifiability in electronic voting. A more extensive overview is found in Pieters (2008).

In section 2, we analyse the Dutch electronic voting based on the distinction between confidence and trust, as introduced by the German sociologist Niklas Luhmann. It is argued that the controversy has initiated a transition in the requirements of electronic voting from reliability to trustworthiness. In section 3, we show which mechanisms are available to increase the trustworthiness of electronic voting by adding verifiability features. We distinguish between different types of verifiability. In section 4, we analyse the trust assumptions of the different forms of verifiability in more detail, and discuss which normative choices need to be made when introducing electronic voting systems. We also investigate how the analysis of trust and verifiability can be applied beyond the context of electronic voting. The last section draws conclusions from the presented results.

2 Trust

2.1 Good and bad trust

In the computing science literature, there seem to exist two different conceptions of trust (Pieters, 2006a). On occasion, they even appear in the same article. In a section named "*Increasing* trust" [our italics] in Evans and Paul (2004), the following sentence is found: "One

way to *decrease* the trust voters must place in voting machine software is to let voters physically verify that their intent is recorded correctly." [our italics] But was the intent not to *increase* trust? Do we want to increase and decrease trust at the same time? What is happening here?

A similar paradox is found in Randell and Ryan (2006). The authors state that recent cryptographic voting schemes "require only a minimum amount of public trust in voting devices or voting officials." On the same page, they say that their "ultimate goal is an e-voting system that isn't only completely trustworthy – doesn't lose, add, or alter ballots, for example, or violate ballot secrecy – but is also trusted by voters to have these properties." Again, are we aiming for a minimum or a maximum amount of trust?

Apparently, computing scientists stem from a tradition in which minimising trust is the standard. "In computer security literature in general, the term is used to denote that something must be trusted [...]. That is, something trusted is something that the users are necessarily dependent on." (Nikander and Karvonen, 2001) Because we *must* trust certain parts of the system for the whole system to be verifiably correct according to the computing science models, we wish to minimise the size of the parts we have to trust, thus minimising trust itself. However, from a psychological perspective, or even a marketing perspective, it is desirable that users trust the *whole* system. Maximising trust seems to lead to more fluent interaction between the user and the system, and is therefore desirable. In Nikander (2001), Matt Blaze says: "I've always wanted trust, as a security person, to be a very simple thing: I trust something if it's allowed to violate my security; something that's trusted is something that I don't have to worry about and if it is broken, I am broken. So I want as little trust in the system as possible, and so security people are worried about minimising trust and now suddenly we have this new set of semantics that are concerned with maximising trust, and I'm terribly confused."

Apparently, two different definitions of trust have to be distinguished (cf. Nikander and Karvonen (2001)):

- trust as something that is *bad*, something that people establish because they *have to*, *not* because the system is trustworthy;
- trust as something that is *good*, something that people establish

because they *want to*, because the system *is* trustworthy.

In order to understand the origins of this difference, we need to explain which concepts lie behind the interpretation of trust as bad and good, respectively. We use an analogy here with a discussion in political science, where the situation with respect to the concept of freedom is similar to the situation we have here with respect to security. In political science, there is a well-known distinction between *negative freedom* and *positive freedom*. Negative freedom means the absence of interference by others; positive freedom means the opportunity for people to pursue their own goals in a meaningful way.[1] We see a parallel here with two possible concepts of safety and security, namely a negative and a positive one:

- negative safety/security: absence of everything that is unsafe/insecure;
- positive safety/security: opportunity to engage in meaningful trust relations.

When people use a negative concept of security, trust has to be minimised, since it denotes a dependence on (possibly) insecure systems. By removing everything that is insecure – thus increasing security – trust defined in this way can indeed be minimised, just as constraints imposed by others have to be minimised to increase negative freedom. In a setting where security is defined positively, however, trust suddenly forms an essential precondition for security, because security then requires the possibility to engage in trust relations. This is precisely the approach that comes from psychology, as opposed to the dominantly negative approach of computing science (remove all insecurities).

As an example, consider the difference between a proprietary and secret computer program and an open-source project. In the former, we need to trust the vendor with respect to the security of the program, and would rather not need to do so, which might be the case if the software can be proven to be secure. In the latter, a community will be formed around the program, in which a consensus may

[1]Cf. Cunningham (2002), pp. 36-39. The notion was originally introduced by Isaiah Berlin (1969 [1958]).

emerge about how secure the program is. Trust can then be based on this social process, and may actually be considered a good thing.

We will label these two conceptions of trust *bad trust* and *good trust*, respectively. We deliberately avoid the terms negative and positive in our distinction of trust, because these are used in the definitions of both freedom and security as indicators of how the concepts are defined (certain things *not* being there vs. certain things *being* there), not of their desirability. Bad and good instead indicate whether we should try to minimise or maximise the associated appearance of trust. Thus, we linked the two different interpretations of trust to two different conceptions of security. Bad trust is linked to a negative conception of safety and security, and good trust to a positive conception.

2.2 Confidence and trust

A similar distinction was made by the German sociologist Niklas Luhmann. Luhmann (1979) provides an extensive model of trust, based on the view of systems theory. According to Luhmann, trust is a mechanism that helps us to reduce social complexity.[2] Without reducing complexity, we cannot properly function in a complex social environment. Luhmann distinguishes several types of trust relations. First of all, he distinguishes between *familiarity* and *trust*. Familiarity reduces complexity by an orientation towards the past. Things that we see as familiar, because "it has always been like that", are accepted – we do engage in relations with those – and things that we see as unfamiliar are rejected – we do not engage in relations with those. For example, especially elderly people often refuse to use ATM's or ticket vending machines, precisely because they are not used to them.[3]

---

[2]The function of trust as a means for reduction of complexity seems to be known in computing science. For example, Nikander and Karvonen (2001) mention this aspect. However, this paper does not refer to the work on trust by Luhmann.

[3]One may argue instead that the reason is not that they are not used to them, but rather the fact that it is harder for them to learn new things. Yet this is precisely one of the conditions that invites relying on familiarity rather than trust.

Trust, on the contrary, has an orientation towards the future: it involves expectations. We trust in something because we expect something. For example, we use ATM's because we expect these machines to provide us with money faster than a bank employee behind the counter.

In later work, Luhmann (1988) also draws a distinction between *trust* and *confidence*. Both confidence and trust involve the formation of expectations with respect to contingent events. But there is a difference. According to Luhmann, trust is always based on assessment of risks, and a decision whether or not to accept those. Confidence differs from trust in the sense that it does not presuppose a situation of risk. Confidence, instead, neglects the possibility of disappointment, not only because this case is rare, but also because there is not really a choice. Examples of confidence that Luhmann gives are expectations about politicians trying to avoid war, and of cars not suddenly breaking down and hitting you. In these cases, you cannot decide for yourself whether or not to take the risk.

When there *is* a choice, trust takes over the function of confidence. Here, the risky situation is evaluated, and a decision is made about whether or not to take the risk: "If you do not consider alternatives [...] you are in a situation of confidence. If you choose one action in preference to others [...], you define the situation as one of trust." (Luhmann, 1988) If you choose to drive a car by evaluating the risks and accepting them, this is a form of trust.

Apparently, Luhmann ascribes the same negative characteristics to confidence that are ascribed to bad trust from a computing science perspective, in the sense that people do not have a choice. People *have to* have confidence in "trusted" parts of the system. Moreover, what Luhmann calls trust has the positive connotation of our good trust, in the sense that people can decide for themselves whether they want to trust something. Trust is then necessary for a system to be successful. We have to note, however, that Luhmann does not regard confidence as a bad thing in general; it is even necessary for society to function. Still, with respect to information systems, confidence means accepting a system without knowing its risks, and computer scientists are generally not willing to do this.

Thus, Luhmann distinguishes between two kinds of relations of self-assurance, based on whether people engage in these relations

because they have to or because they want to. Luhmann calls these two relations confidence and trust, respectively. These observations also cover the situation we described in computing science. This means that the distinction we made is not something that characterises social aspects of security in information systems only, but something that can be considered a general characteristic of trust relations.

Computing scientists generally try to replace confidence with trust, i.e. exchange unconscious dependence on a system for explicit evaluation of the risks, and minimising the parts in which we still have to have confidence.[4] Philosophers (and social scientists), instead, recognise the positive aspects of confidence, and may evaluate positively people having a relation of self-assurance with the system without exactly knowing its risks (i.e. confidence).

Based on the distinction between confidence and trust, we also propose a distinction between reliability and trustworthiness. A system acquires *confidence* if it is *reliable*, and it acquires *trust* if it is *trustworthy*.[5] A reliable system is a system that people can use confidently without having to worry about the details. A trustworthy system is a system that people can assess the risks of and that they still want to use.

2.3 Trust in e-voting

Based on the analysis in the previous section, we will show how the distinction between reliability and trustworthiness influenced the electronic voting debate in the Netherlands.

Electronic voting systems *may* be seen as alternatives to paper voting systems. Whether this is indeed the case depends on the situation. If they are merely seen as improvements upon the paper system and they seem to behave correctly, the confidence in the paper sys-

[4]This general approach is not without exceptions; cf. Nikander (2001).

[5]Reliability is used in the more limited sense of continuity of correct service in Aviz, □ienis, Laprie, Randell, and Landwehr (2004). Our notion of reliability roughly corresponds to the “alternate definition of dependability” in their taxonomy, whereas trustworthiness corresponds to the “original definition of dependability”.

tem may easily be transferred to the electronic systems. If they are seen as alternatives, people suddenly get the option to *choose* a voting system. This invites actively assessing the risks of the different systems, and basing the decision on an analysis of these risks. This means that *trust* now becomes the dominant form of self-assurance, as opposed to confidence. This has as a consequence that voting systems are required to be *trustworthy* rather than reliable only.

In the Netherlands, electronic voting machines were seen as an instrumental technology to do the same thing in a better way. The confidence that people had in the paper voting system was smoothly transferred to the electronic machines. Only after the offence of the pressure group Wij Vertrouwen Stemcomputers Niet (We Don't Trust Voting Computers) in 2006, the two methods of voting came to be perceived as really different. Drawing this distinction was a major achievement of the pressure group, and it prominently featured in their publications. One important achievement was the replacement in public discussions of the term "voting machine" by "voting computer". The main difference is that a computer is programmed to perform a certain task, and can therefore also be programmed to do different things. This contributed to the perception of e-voting being really different from paper voting.

Thus, the pressure group initiated a shift in perception from e-voting as an improvement to e-voting as an alternative. Due to this shift, electronic voting systems were now conceived as subject to a decision, and needed to be trustworthy (suitable for trust) rather than reliable (suitable for confidence) only. This meant that properties such as verifiability and secrecy had to be expressed in measurable terms, and needed to be compared for different voting systems.

In the Netherlands, the ensuing discussion was decided by the compromising radiation emitted by electronic devices (tempest). Because this might endanger the secrecy of the ballot, the Netherlands abolished electronic voting. This is a risk that seems to be specific to the Dutch risk perception. In other countries, verifiability is more on the foreground of electronic voting politics.

Paper voting seems to have a natural advantage when it comes to trustworthiness, because it is easier to understand how the security mechanisms work. Still, electronic voting may be more *reliable*, because it eliminates human error in for example the counting of the

ballots, and may thus contribute to confidence in the election procedures and the political results. For electronic voting, the *trust* relations are more complex, and necessarily involve expert judgement about the properties of the voting system. People will need to have confidence in the electronic devices based on their confidence in the experts, who in turn should have trust in the procedure based on their analysis and comparisons.

If such more complex trust relations are judged to be unacceptable for elections, for example because they might give too much power to the experts (they can misuse the confidence they get from the public), paper voting will be the only option. When they are judged to be acceptable, various improvements in the trustworthiness of electronic voting are possible. In terms of the software used in elections, trustworthiness with respect to verifiability may be increased by introducing the option to verify one's vote or even the total count. This makes at least sure that the voters can have a role in verifying the result of the election, which may give more confidence than expert judgement only, even though the voters may not understand the details of the procedure. Several such mechanisms have been proposed in the literature. In the following, we investigate the primary differences between such mechanisms, and relate those to the problems they may solve with respect to the trustworthiness of electronic voting.

3 Verifiability

Traditionally, verification of electronic voting in the Netherlands meant that each type of machine was tested by the testing agency TNO. Additionally, a small number of machines were tested before each election. Interestingly, TNO had also been involved in the design of the machines, and even in the drafting of the legislation, making the independence of the testing questionable (Hermans and van Twist, 2007).

This type of testing is mainly directed towards finding unintentional errors in the design of the machines. If a malicious programmer wants to insert manipulated program chips, this does not help much. Of course, requirements concerning the security of the ma-

chines against such attacks might have been included in the legislation, but security was hardly addressed there.

Moreover, verifiability of the *machines* is of a completely different category than verifiability of the *results*. When the demands changed from reliability to trustworthiness, it became apparent that manufacturers could no longer get away with verification of the machines only. Each election result would need to be verifiable. It is this type of verifiability that is the focus of the following analysis. We investigate the concept of verifiability vis-a-vis the scientific literature and the concrete developments in the Netherlands. We propose a distinction between various concepts of verifiability.

3.1 Voter-Verifiable Elections

Verifiability of electronic voting systems has achieved a great deal of attention in computing science literature. In the context of electronic voting machines (DRE's), much discussion has taken place – especially in the US – around the solution of a voter-verified paper audit trail (VVPAT) (Mercuri, 2002). Typically, this includes a paper copy of each vote being kept as a backup trail for recovery or recount. This should increase trust in the proper operation of the black-box DRE machines. More than half of the states in the US have now passed legislation making a paper trail mandatory.

Some people argue that a VVPAT does not help much in improving security, because people will have a hard time checking their vote, due to the large number of races on which they have to vote in a single election in the US. It has been suggested to use an audio trail instead (Selker and Goler, 2004). Also, an important question is what to do if the electronic trail and the paper trail differ. Which one has to be preferred? It could be argued that for small differences, the electronic trail will probably be the more reliable one, whereas for larger differences, the paper trail may be more trustworthy.

A VVPAT anchors the verifiability of electronic voting in organisational features, which should make sure that the paper copies are indeed (statistically) checked for correspondence to the electronic result. Such a procedure aims at re-establishing the trust voters had in the paper counting system. In Internet voting, a paper trail is not

feasible, because "[t]he voter is not at the point of vote summarization to examine a receipt" (Saltman, 2006, p. 211). For such applications, one needs to look into software solutions. For these purposes, various cryptographic receipts have been proposed, e.g. in Chaum (2004). In the following, we will focus on these software-oriented approaches.

In the Netherlands, several experiments with online voting have been conducted during the last couple of years. In the European Elections 2004, Dutch citizens staying abroad were allowed to vote online. The system used, called KOA (Kiezen Op Afstand), was designed by Logica CMG for the Dutch Ministry of Domestic Affairs. Meanwhile, a second system was being developed by the "waterschap" (public water management authority) of Rijnland, in cooperation with the company Mullpon. This system was labelled RIES (Rijnland Internet Election System), and has been used in the elections of the "waterschappen" Rijnland and Dommel in fall 2004 (Hubbers et al., 2005).

There are several interesting features offered by the systems experimented with in the Netherlands. For example, the KOA system uses personalised (randomised) ballots, in order to prevent attacks by e.g. viruses residing on the voter's computer. Moreover, the counting software, written at the Radboud University Nijmegen, was specified and verified using formal methods (Hubbers et al., 2004). Unfortunately, the KOA system does not offer verifiability to the voters.

The RIES system does offer verifiability, and people seem to appreciate this.[6] However, the kind of verifiability that is offered by RIES seems to be quite different from the verifiability that is offered in more advanced cryptographic systems in the literature. In some sense, RIES seems to be *too* verifiable to provide resistance against coercion or vote buying.

Traditionally, two types of verifiability have been distinguished in research on electronic elections. When a system establishes *individual verifiability*, every voter can check if her vote has been prop-

[6]Much depends on the interface though. Before RIES was actually used in an election, a trial session revealed that a too difficult verification procedure *decreases* trust in the system among voters. The user-friendliness of the verification procedure was improved after the trial.

erly counted. In *universal verifiability*, anyone can check that the calculated result is correct (Kim and Oh, 2004; Malkhi et al., 2002). Typically, a bulletin board or some other electronic means is used to publish a document that represents the received votes. Voters can look up their own vote there, and people interested in the results can do correctness checks on the tally.

However, these types of verifiability have been implemented in very different ways. We think that at least one more conceptual distinction is necessary to categorise the different systems appropriately (Pieters, 2006b). We will introduce this distinction via an analysis of the relation between verifiability and receipt-freeness.

3.2 Verifiability and Receipt-Freeness

One of the basic requirements of election systems is the resistance against coercion and vote buying. Therefore, people should not be able to prove how they voted, even if they want to. This makes it impossible for someone who forces them to vote in a certain way, or someone who buys their vote, to check if they actually complied. This requirement is hard, if not impossible, to realise in an environment without public control, as opposed to the classical polling booth. People can watch over your shoulder if you are not guaranteed a private environment for voting, and thereby obtain proof of your vote (Pieters and Becker, 2005).[7] Some scientists hold the view that this and other security problems make it advisable not to implement Internet voting at all (Jefferson et al., 2004).

There is empirical evidence, however, that vote buying may "survive the secret ballot", despite isolating the voter in a polling booth (Brusco et al., 2004). This means that buying *does* happen, even if individual votes are secret. Brusco, Nazareno, and Stokes (2004) mention three possible explanations for the fact that voters

[7]Some systems introduce "practice ballots" or similar measures to prevent such attacks. However, these measures severely limit verifiability, because the tallier still needs to be able to distinguish real ballots from practice ballots, whereas the attacker should not be able to detect this via the means of verification offered to the voter. See e.g. http://zoo.cs.yale.edu/classes/cs490/03-04b/adam.wolf/Paper.pdf, consulted December 9, 2005.

comply to the buyer's wishes in spite of the secret ballot. These include the expectation of future benefits if enough people in a district vote for the desired party, feelings of moral obligation of the voters, and the preference of immediate benefits over vague political promises. Similar effects may exist for coercion.

Thus, some may argue that the fact that people vote in a non-controlled environment does not need to be a fundamental problem compared to the current situation. In any case, the risks of vote buying and coercion are the same as those involved in postal ballots. Organisational and legal measures may be put in place to minimise the risks.

If we accept this argument, there is still a second problem involved. For it is one thing that people physically present at the act of voting can influence the voter, the possibility to prove *remotely* that you voted for a certain party is worse. This means that people could provide proof to a coercer or get money for their votes *after* they voted themselves. This is more convenient for an attacker than buying or stealing access codes and casting all votes herself. There is a trade-off between verifiability and resistance against coercion here. If every voter can check if her vote has been counted correctly, i.e. if the vote in the results corresponding to her own vote maps to the right party or candidate, then she can also show this check to a coercer or buyer as a proof. Thus, we generally do not want a voter to be able to show a proof of her vote *after* the election is over. In the literature, this restricted property is often called receipt-freeness (Benaloh and Tuinstra, 1994; Hirt and Sako, 2000).[8]

Some systems, among which the RIES system, do indeed allow a voter to check after the elections for which party or candidate her vote has been counted (Baiardi et al., 2004, 2005; Hubbers et al., 2005; Malkhi et al., 2002; Storer and Duncan, 2004). These systems are therefore not receipt-free in the technical sense. Although the fact that people can see what they voted for after the elections may increase trust in the system, the lack of resistance against coercion and vote buying makes these systems debatable candidates in elec-

[8]If a system is resistant against coercion even if the coercer can interact with the voter during voting, the term coercion-resistance is sometimes used instead of receipt-freeness (Juels et al., 2005). In order to avoid confusion, we consequently use the term receipt-freeness here.

tions for which we cannot be sure that the chances of buying and coercion are low.

In many systems (Chaum, 2004; Joaquim et al., 2003; Kim and Oh, 2004), this is remedied by allowing a voter to check *that* her vote has been counted, but not *how*. The idea is that it is impossible, or at least computationally infeasible, for an attacker to make the system count a different vote for this voter in case the check turns out to be OK. Receipt-freeness can thus be provided by limiting the information that a voter can retrieve about her vote after the election, while still assuring cryptographically that this is indeed a proof that the vote has been counted for the party or candidate that was chosen during the election.

Thus, the relation between individual verifiability and receipt-freeness gives rise to a distinction between two different types of individual verifiability. In the following section, we discuss the different options for verifiability in remote electronic elections based on this observation.

3.3 Variants of Verifiability

Following the analysis of the relation between individual verifiability and receipt-freeness, we observed a distinction between two kinds of individual verifiability. We will label these two types based on an analogy with the distinction between classical logic and constructive logic. In classical logic, one can prove an existential formula without actually showing an instance in the domain that satisfies this formula.[9] In constructive logic, one has to produce a witness in order to prove the existential formula. We argue that there is a similarity with verifiability in electronic voting here.[10]

When a voter can only verify *that* her vote has been counted, this amounts to showing that a certain vote exists in the results that can be attributed to this voter. However, the actual witness (i.e. the

[9]Equivalently, one shows that the negation of the formula does not hold for all instances.

[10]The analogy does not hold for computational issues around finding a witness. Still, we think that it is useful for understanding what the difference is between the two types of verifiability.

choice this voter made) cannot be recovered from the verification procedure. Here, the voter will believe that her vote was recorded correctly if the election authority can show something that proves the existence of a vote by this voter in the results, without re-examining the original vote.[11] Proving the existence of something without showing a witness can be done in classical logic. We will label this type of verifiability *classical individual verifiability*.

On the other hand, some systems allow a voter to check afterwards *for which candidate* her vote has been counted. This means that the actual instance of a vote is shown as a proof to the voter. Here, the voter does not believe the election authority unless she can reproduce the original vote from the results. This corresponds to the proof of an existential formula in constructive logic. Therefore, we will label this type of verifiability *constructive individual verifiability*.

Definition 1 *Classical individual verifiability is the property of an election system that a voter can verify that her vote has been counted correctly based on a document representing the received votes,* without *being able to reconstruct her choice from that document.*[12]

Definition 2 *Constructive individual verifiability is the property of an election system that a voter can verify that her vote has been counted correctly by reconstructing her choice from a document representing the received votes.*

The first type of individual verifiability has become fairly standard in computing science discussions on voting systems. However, the second type has been used in practice as well, and we think these developments deserve some consideration from both a scientific and a political perspective.

For universal verifiability we can make a similar distinction. We take universal verifiability, to prevent confusion, to mean that any observer can verify that the *final tally* is correct, *given a document representing the received votes*. Thus, universal verifiability does

[11]Equivalently, one shows that it is not the case that one's vote has not been counted.

[12]All types of proof discussed in this section may be relative to cryptographic assumptions.

not necessarily mean that anyone can check that all cast votes have been included in this document.

Definition 3 *Classical universal verifiability is the property of an election system that it can be shown that the tally is correct given a document representing the received votes, without all the data necessary to perform the calculation being publicly accessible.*

Definition 4 *Constructive universal verifiability is the property of an election system that all data necessary for calculating the result from a document representing the received votes are publicly accessible, and that a verifier can compute the tally from this set independently of the election authorities.*

Systems in which votes are encrypted with public keys of election authorities typically establish classical universal verifiability, e.g. via so-called zero-knowledge proofs by these authorities that show that they did their job correctly, or via homomorphic encryption schemes (Chaum, 2004; Kim and Oh, 2004; Neff, 2001). This proves that there is a set of votes corresponding to the published document and to the tally, but the calculation of the tally from the document is not public. Constructive universal verifiability is not possible in this case, unless the private keys are made public after the elections. However, this typically violates secrecy requirements; the encryption is usually *intended* to maintain secrecy of the individual votes.

Systems which only use public functions to calculate the result from the set of received votes typically do establish constructive universal verifiability (Hubbers et al., 2005; Malkhi et al., 2002; Storer and Duncan, 2004). However, these systems need special measures to prevent the votes from being linked to individual voters. Because the received votes are used in public calculations of results, without any intermediate trusted computations that scramble them, the link between voter and vote should be destroyed in a non-trusted environment beforehand. In the UK, the situation is even more complicated due to the requirement that this link can be recovered in special cases (Storer and Duncan, 2004).

Moreover, all the systems we included in our research that offered constructive universal verifiability, *also* offered constructive individual verifiability, and are therefore not receipt-free. For exam-

ple, the RIES system used in the Netherlands (Hubbers et al., 2005) establishes both constructive individual verifiability and constructive universal verifiability. In technical terms, hash functions are used to publish the links between all possible votes and the corresponding candidates before the elections. The original votes are only derivable from a secret handed to the voter. The confidentiality of these secrets is achieved via organisational security measures, in the same way that identification codes for bank cards are handed out. After the elections a table of received votes is published. By computing hashes, individual voters can check for which party or candidate their vote has been registered, and any observer can calculate the result from the list of received votes.

Thus, systems that allow constructive individual verifiability and constructive universal verifiability are beginning to be used in practice, in small-scale or low-risk elections. Meanwhile, many advanced cryptographic systems that establish classical individual verifiability and classical universal verifiability are being developed. We also saw that when the latter type of systems is adapted in order to offer constructive universal verifiability, constructive individual verifiability seems to appear as a side-effect, and receipt-freeness is thereby sacrificed. But which combination of individual and universal verifiability is most desirable? And why do we care?

4 Verifiability and Trust

The dynamics of trust, discussed in the previous section, have shifted the expectations of electronic voting from reliability to trustworthiness. Verifiability may be implemented to provide such trustworthiness. However, technical measures are not by themselves sufficient to establish a relation of trust. The social and political framework has to be taken into account as well, both in terms of the context in which trust can be established and in terms of the social and political effects of technical choices.

4.1 The politics of voting technology

In his famous study “Do artifacts have politics?”, Langdon Winner (1980) showed that technological designs may have political implications. These may occur either intentionally or unintentionally. Winner’s famous example of intentional political effects concerns the building of bridges in New York between 1920 and 1970 that were too low for the buses of public transport, and therefore the lower income classes, to pass underneath. One can easily imagine similar things happening unintentionally as well. Since then, many cases of such influences have been investigated, and many theories about how they come about have been developed in philosophy of technology and science and technology studies (STS).

We may assume similar effects, be they unintentional, occurring in Internet voting technology. Internet voting will undoubtedly, depending on the way in which it is implemented, make certain things possible and others impossible, just as the New York bridges did. One can easily imagine that an Internet voting system will, depending on the types of verifiability that are offered, include different voters in different ways in the election procedure, and thereby change the image of and trust in democracy.

In this sense, choosing a particular kind of verifiability in a particular experiment is not a choice that only influences this particular system. Instead, the type of verifiability offered and the surrounding practices in the elections may mediate the idea that people have of elections. For example, if the RIES system is successful in an experiment with elections for the local water management authorities, people may start to think that constructive individual verifiability is a good thing in general. People may also wonder why they cannot verify their choice in the same way in a later election that uses a different system.

Thus, we would like to stress that choosing a particular kind of verifiability in an experiment may have political consequences, not only for the elections that the system is being used in, but also in terms of expectations that are raised about future elections. This may lead to changes in public trust not only in the elections, but also in the democratic system as a whole.

4.2 What Proof Do We Prefer?

How, then, can we decide which kind of verifiability we wish to implement or use? Because of the role of voting systems in people's experience of democracy, basing a decision on technical requirements only is not the way to go. We argue that the choice between different types of verifiability should be the outcome of a political discussion, rather than the unconscious influence of techno-social developments.

Technology, and especially a politically sensitive one such as electronic voting, occupies a place in people's lifeworlds, i.e. their daily experiences and acts (Ihde, 1990). Based on such a phenomenological approach to technological innovation (Ihde, 1990; Verbeek, 2005) and the work on trust by Luhmann (1979); Luhmann (1988), it appears that there are two basic ways of acquiring trust in large-scale technology such as electronic voting:

- connecting to experiences that people are already familiar with (focusing on familiarity of experience);
- connecting to a clear vision of a future good to be achieved, for which democratic support exists (focusing on expectations of action).

In the case of voting, a good example of the former strategy is the introduction of the Nedap voting machines in the Netherlands in the mid-nineties. Because the layout of the interface of the voting machines was very similar to the previously used paper ballots, one of the reasons that the system was so easily accepted may have been the familiarity of the interface. Now that people are already familiar with voting machines, the introduction of a Voter Verified Audit Trail can be considered an example of the latter strategy, since there is a strong public agreement on the beneficial properties of audit trails.

If we choose to implement verifiability features, we have to face the fact that people are generally *not* familiar with vote and result verification, and people will probably not be happy with their verifiability if the complete election system is turned upside down. So how can we maintain familiarity in Internet elections if people are not familiar with verification, but at the same time demand the pos-

sibility of verification of the results? The best we can do is preserve as many of the things that people are familiar with in current elections, while offering verification to make Internet elections acceptable. Two main demands, which are not only functional requirements, but also part of a ritual that establishes familiarity with elections, can be mentioned here:

- the demand of the secret ballot;[13]
- the demand of the public character of vote counting.[14]

How do these requirements relate to the various types of verifiability? In the case of individual verifiability, the demand of the secret ballot implies that constructive individual verifiability is not desirable. Thus, from the perspective of connecting to existing experiences, we should choose classical individual verifiability. This does *not* mean that we argue for this type because of functional requirements, but rather from an "if it ain't broke, don't fix it" perspective. *Unless* there is democratic consensus about the desirability of constructive individual verifiability, either from the point of view of enhancing trust or from the point of view that democracy functions better without the secret ballot (which is held for many representational bodies such as parliament and meetings such as party congresses), we had better stick to the demand of the secret ballot, and implement classical individual verifiability.

However, the existing schemes that offer classical individual verifiability, to the best of our knowledge, also offer classical universal verifiability. The limitation of the ability of result computation to dedicated parts of the system, with accompanied proofs of correctness, goes against the demand of the public character of vote counting. Typically, *any* encryption with a public key implies that the public character of vote counting is being set aside, unless the corresponding private key is made public afterwards, which is generally not the case. As much as the secret ballot is an important part of the ritual of voting, so is the public character of vote counting. Therefore, we think that *constructive* universal verifiability, in which any party can do an independent calculation of the result, is

[13]Cf. Dutch constitution art. 53.2 and Dutch election law ("Kieswet") art. J 15.
[14]Cf. Dutch election law ("Kieswet") art. N 1, N 8 and N 9.

preferable, *unless* there is democratic consensus about arguments for the opposite point of view.

4.3 Beyond electronic voting

Verifiability is a way to increase trustworthiness, as opposed to reliability, of e-voting systems. Implementing verifiability features not only influences trust in the elections themselves; different variants of verifiability also have different trust assumptions on a smaller scale.

Constructive individual verifiability leads to high trustworthiness with respect to the validity of the results, but low trustworthiness with respect to the secrecy of the ballot. In this case, assuring that the individual choices will be kept secret requires high confidence in the voter: she should not succumb to the persuasion of selling her vote, or to coercion. In classical individual verifiability, there is no need to have such confidence in the voter, since she will not be able to prove the contents of here vote. However, in this case, the voter cannot "see" if her vote has been registered correctly; she can only see *that* it has been registered. Accepting the assurance that it must then also be correct requires confidence of the voter in the verification procedure as designed by mathematicians (cryptographers) and computing scientists. These scientists, in turn, must have analysed the procedure and have found it trustworthy. Similar trust assumptions are present in the two types of universal verifiability.

These relations between verifiability and trust can be generalised to other systems and properties (e.g. privacy). In general, verifiability can be necessary when some data processing is done by a third party on behalf of the party who is responsible for the results, or when there are other parties who have interest in correctness and security properties of the results. Constructive verifiability then means repeatability of the data processing by other parties, whereas classical verifiability means a (mathematical) proof obligation on the part of the calculator, without revealing the details of the calculation. In case of privacy, we have the data processor, the processing of which may need to be verified by either the data controller or the individual the data concerns. These correspond to the e-voting provider, the Electoral Commission and the voter in the election process.

As an example, we take the profiling of an individual by an organisation, on the basis of which a decision about the individual is made by the organisation. Profiling can be defined as “the process of ‘discovering’ correlations between data in databases that can be used to identify and represent a human or nonhuman subject (individual or group) and/or the application of profiles (sets of correlated data) to individuate and represent a subject or to identify a subject as a member of a group or category” (Hildebrandt, 2008, p. 19). Hildebrandt distinguishes between *group profiling* and *personalised profiling*. In the former, correlations between people are derived from a dataset by constituting a group of similar people and identifying the (probable) characteristics of people belonging to this group. In the latter, individual preferences are recorded and analysed to personalise services.

Both individual and universal verifiability can be useful concepts in such a setting. The processing here concerns the calculation of a decision based on the information derived about an individual. For personalised profiling, this may for example the decision whether or not to issue a creditcard. Laws and regulations may apply that limit what the organisation is allowed to consider in this decision. The individual and other parties such as the government may demand the opportunity to verify that these rules were indeed adhered to. In case of classical individual verifiability, the organisation would have to prove that the calculations that were made concerning an individual – both in terms of creating the profile and in terms of applying it to make a decision – were made according to the rules. In case of constructive individual verifiability, all the data that was used as input as well as the calculation procedure would need to be made available to the individual, such that she can repeat the calculation or ask an independent institution to do so. Which of the types of verifiability is most desirable should be discussed based on the required level of transparency as well as the sensitivity of the information involved.

Similar considerations apply to universal verifiability. Here, the calculations concern groups of persons rather than a single individual (group profiling instead of personalised profiling). Typically, group profiles will affect decisions about individuals that are members of the established groups, whether or not the attributes associated with the group apply to them. For example, if someone lives in

a poor neighbourhood, this may affect decisions about issuing her a creditcard. The acceptability of such a decision depends on a) the correctness of the calculations concerning the group and b) the acceptability of applying the group profile in this particular decision. In classical universal verifiability, the organisation would have to prove that statistical data concerning groups was calculated and applied according to the rules. In constructive universal verifiability, all data that was used to calculate the statistics would have to be made available (possibly in an anonymised form), such that stakeholders can validate the calculations by performing them themselves.

These concepts may be useful in discussions on the alignment of technical and legal aspects of profiling. Depending on the application area, different forms of verifiability may be prescribed. One can imagine that the bigger the consequences of a decision, the stronger the verifiability requirements. This requires classifications of both decisions and verifiability properties. Such verifiability requirements may also make it possible to assure that forms of profiling that are considered illegal – e.g. in relation to race or religion – are indeed absent from the decision-making process. In such cases, data processors would either need to prove that their processing of the data has certain properties, or allow independent parties to redo the calculations and confirm the fairness and correctness.

Similar arguments apply to any form of externalised data processing, such as cloud computing or outsourcing. The data controller inevitably has an interest in assessing whether the externalised calculations conform to certain desirable properties, i.e. establishing verifiability of the calculations. Investigating the implications of the present analysis for these fields is future work.

5 Conclusions

In controversies on electronic voting, like in 2006 in the Netherlands, different types of trust play a major role. In the Netherlands, e-voting was represented as a possibly disastrous alternative to paper voting, and the “blind” trust (confidence) which existed had to give way to rational trust (trust). This is a process that computing scien-

tists can contribute to: they aim at minimising blind trust in information systems, and replacing this by a provably correct behaviour. Although some computing scientists claim that there are no good solutions except for a print of each vote, many people work on advanced methods to provide verifiability in electronic elections.

We distinguished between two types of individual verifiability and two types of universal verifiability in electronic elections, based on scientific literature and concrete developments. We made this distinction based on an analogy with proofs in classical and constructive logic, and labelled the corresponding types of verifiability classical and constructive verifiability, respectively. This distinction is meaningful both for individual and universal verifiability, and we think that it is a useful tool for explicating the hidden assumptions of the way in which verifiability is realised in concrete systems.

We argued that choices for particular kinds of verifiability in experiments may have political implications, not only for the specific election that a system is used in, but also in terms of expectations of future elections. Therefore, it is wise to attempt to arrive at political consensus about the kinds of verifiability that are desirable. We argued that even if verifiability is widely accepted as a good thing, we have to maintain familiarity with elections in order to make the whole system acceptable. The best we can do here is maintain the existing properties of vote secrecy and public counting. This can be done with a system that establishes classical individual verifiability and constructive universal verifiability, which, as far as we are aware, has not been invented yet.

We showed that the types of verifiability we distinguished can also be useful beyond the application area of electronic voting. In that case, they denote which information individuals or organisations may get about data that is stored or processed either concerning them or on behalf of them. This broader notion may for example be used when discussing the acceptability of different forms of profiling. In this sense, we will have learnt something from the controversy, even if the time of electronic voting would be over.

Acknowledgments

The ideas described here were partly developed while the author was employed by Radboud

University Nijmegen, and are based on earlier publications (Pieters, 2006a, 2006b). This work was supported by a Pionier grant from NWO, the Netherlands Organisation for Scientific Research.

References

A. Avižienis, J.C. Laprie, B. Randell, and C. Landwehr. Basic concepts and taxonomy of dependable and secure computing. *IEEE transactions on dependable and secure computing*, 1(1):11–33, 2004.

F. Baiardi, A. Falleni, R. Granchi, F. Martinelli, M. Petrocchi, and A. Vaccarelli. SEAS: a secure e-voting applet system. In K. Futatsugi, F. Mizoguchi, and N. Yonezaki, editors, *Software security — theories and systems*, LNCS 3233, pages 318–329. Springer, Berlin, 2004.

F. Baiardi, A. Falleni, R. Granchi, F. Martinelli, M. Petrocchi, and A. Vaccarelli. SEAS, a secure e-voting protocol: design and implementation. *Computers & Security*, 24:642–652, 2005.

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.

I. Berlin. *Four concepts of liberty*. Oxford University Press, Oxford, 1969 [1958].

V. Brusco, M. Nazareno, and S.C. Stokes. Vote buying in Argentina. *Latin American Research Review*, 39(2): 66–88, 2004.

D. Chaum. Secret-ballot receipts: true voter-verifiable elections. *IEEE Security & Privacy*, 2(1):38–47, 2004.

F. Cunningham. *Theories of democracy: a critical introduction*. Routledge, London, 2002.

D. Evans and N. Paul. Election security: perception and reality. *IEEE Security & Privacy*, 2(1):24–31, January/February 2004.

L.M.L.H.A. Hermans and M.J.W. van Twist. Stemmachines: een verweesd dossier. Rapport van de Commissie Besluitvorming Stemmachines, April 2007. Available online: `http://www.minbzk.nl/contents/pages/86914/rapportstemmachineseenverweesddossier.pdf`, consulted April 19, 2007.

M. Hildebrandt. Defining Profiling: A New Type of Knowledge? In M. Hildebrandt and S. Gutwirth, editors, *Profiling the European Citizen: Cross-disciplinary Perspectives*, pages 17–45. Springer, 2008.

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.

E. Hubbers, B. Jacobs, J. Kiniry, and M. Oostdijk. Counting votes with formal methods. In C. Rattray, S. Maharaj, and C. Shankland, editors, *Algebraic Methodology and Software Technology (AMAST'04)*, number 3116 in LNCS, pages 241–257. Springer, 2004.

E. Hubbers, B. Jacobs, and W. Pieters. RIES – Internet voting in action. In R. Bilof, editor, *Proc. 29th Annual International Computer Software and Applications Conference, COMPSAC'05*, pages 417–424. IEEE Computer Society, July 2005. ISBN 0-7695-2413-3.

D. Ihde. *Technology and the lifeworld*. Indiana University Press, Bloomington, 1990.

D. Jefferson, A.D. Rubin, B. Simons, and D. Wagner. Analyzing internet voting security. *Communications of the ACM*, 47(10):59–64, 2004.

R. Joaquim, A. Zúquete, and P. Ferreira. REVS – a robust electronic voting system. *IADIS International Journal of WWW/Internet*, 1(2), 2003.

A. Juels, D. Catalano, and M. Jakobsson. Coercion-resistant electronic elections. In *Proc. WPES'05*. ACM, 2005.

S. Kim and H. Oh. A new universally verifiable and receipt-free electronic voting scheme using one-way unwappable channels. In C.-H. Chi and K.-Y. Lam, editors, *AWCC 2002*, volume 3309 of *LNCS*, pages 337–345. Springer, 2004.

N. Luhmann. *Trust and power: two works by Niklas Luhmann*. Wiley, Chichester, 1979.

N. Luhmann. Familiarity, confidence, trust: problems and alternatives. In D. Gambetta, editor, *Trust: Making and breaking of cooperative relations*. Basil Blackwell, Oxford, 1988.

D. Malkhi, O. Margo, and E. Pavlov. E-voting without 'cryptography'. In *Financial Cryptography '02*, 2002.

R.T. Mercuri. A better ballot box? *IEEE Spectrum*, 39(10):26–50, 2002.

C.A. Neff. A verifiable secret shuffle and its application to e-voting. In *Proceedings of the 8th ACM Conference on Computer and Communications Security*, pages 116–125. ACM, 2001.

P. Nikander. Users and trust in cyberspace (transcript of discussion). In B. Christianson, B. Crispo, J.A. Malcolm, and M. Roe, editors, *Security Protocols: 8th International Workshop, Cambridge, UK, April 3-5, 2000, Revised Papers*, number 2133 in Lecture Notes in Computer Science, pages 36–42. Springer, 2001.

P. Nikander and K. Karvonen. Users and trust in cyberspace. In B. Christianson, B. Crispo, J.A. Malcolm, and M. Roe, editors, *Security Protocols: 8th International Workshop, Cambridge, UK, April 3-5, 2000, Revised Papers*, number 2133 in Lecture Notes in Computer Science, pages 24–35. Springer, 2001.

W. Pieters. Acceptance of voting technology: between confidence and trust. In K. Stølen, W.H. Winsborough, F. Martinelli, and F. Massacci, editors, *Trust Management: 4th International Conference (iTrust 2006), Proceedings*, volume 3986 of *LNCS*, pages 283–297. Springer, 2006a.

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 2006b. e-Science Institute, University of St. Andrews.

W. Pieters. *La volonté machinale: understanding the electronic voting controversy*. PhD thesis, Radboud University Nijmegen, January 2008.

W. Pieters and M. Becker. Ethics of e-voting: An essay on requirements and values in Internet elections. In P. Brey, F. Grodzinsky, and L. Introna, editors, *Ethics of New Information Technology: Proc. Sixth International Conference on Computer Ethics: Philosophical Enquiry (CEPE'05)*, pages 307–318, Enschede, 2005. Center for Telematics and Information Technology.

B. Randell and P.Y.A. Ryan. Voting technologies and trust. *IEEE Security & Privacy*, 4(5):50–56, 2006.

R.G. Saltman. *The History and Politics of Voting Technology*. Palgrave Macmillan, New York, 2006.

T. Selker and J. Goler. Security vulnerabilities and problems with VVPT. Caltech / MIT Voting Technology Project, Working Paper #16, 2004. Available online: `http://www.vote.caltech.edu/media/documents/wps/vtp_wp16.pdf`, consulted February 10, 2006.

T. Storer and I. Duncan. Practical remote electronic elections for the UK. In S. Marsh, editor, *Proceedings of the Second Annual Conference on Privacy, Security and Trust*, pages 41–45. National Research Council Canada, 2004.

P.P.C.C. Verbeek. *What things do: Philosophical Reflections on Technology, Agency, and Design*. Pennsylvania State University Press, 2005.

L. Winner. Do artifacts have politics? *Daedalus*, 109(1):121–136, 1980.