# Acceptance of Voting Technology: between Confidence and Trust[*]

Wolter Pieters

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

**Abstract.** Social aspects of security of information systems are often discussed in terms of "actual security" and "perceived security". This may lead to the hypothesis that e-voting is controversial because in paper voting, actual and perceived security coincide, whereas they do not in electronic systems. In this paper, we argue that the distinction between actual and perceived security is problematic from a philosophical perspective, and we develop an alternative approach, based on the notion of trust. We investigate the different meanings of this notion in computer science, and link these to the philosophical work of Luhmann, who distinguishes between familiarity, confidence and trust. This analysis yields several useful distinctions for discussing trust relations with respect to information technology. We apply our framework to electronic voting, and propose some hypotheses that can possibly explain the smooth introduction of electronic voting machines in the Netherlands in the early nineties.

## 1 Introduction

Electronic voting is one of the most interesting examples of the application of security-sensitive information technology in society. Democracy is one of the foundations on which western culture is built, and it is no wonder that the introduction of new technology into this domain has raised a considerable amount of discussion. Controversies are strengthened by the media coverage of security leaks and viruses in many different information technology applications, and by the advent of Internet voting as the future election platform. Many scientific papers have appeared covering the security of e-voting systems [1, 4, 10, 12, 13, 15, 27], but yet, there is no consensus over which aspects to take into account in such an analysis.

[*] This article appeared in K. Stølen et al. (Eds.): iTrust 2006, LNCS 3986, pp. 283–297, 2006. ©Springer-Verlag Berlin Heidelberg 2006. This work is supported by a Pionier grant from NWO, the Netherlands Organisation for Scientific Research. The author wishes to thank (in alphabetical order) Marcel Becker, Iwan Bos, Bart Jacobs, Erik Poll, Sanne van der Ven and Martijn Warnier for useful comments on drafts of this paper, and Gerard Alberts for inspiring brainstorm sessions on the topic.

There is some agreement, however, about the fact that both technical and social aspects of security should be covered [8, 21, 22, 25, 30]. The social aspects are then often labelled "trust", and the implementation of these aspects in concrete systems is delegated to user interface experts and communication departments. They are assigned the task of transforming "trustworthiness", created by the technical experts, into "trust".

Although the distinction between the trustworthiness established by the technical experts and the trust established by the user interface and marketing experts seems a very business-oriented way to think about these matters, the scientific literature – to the best of our knowledge – takes this distinction for granted as well. In fact, all of the papers that we consulted about social aspects of security reflect this view by using a distinction between "actual security" and "perceived security" in their analyses of security-sensitive systems [8, 21, 22, 25, 30].

In this paper, we provide an alternative analysis of the issue of trust in information systems, with emphasis on the case of electronic voting. First of all, we explain why we are not satisfied with the prevailing distinction between actual and perceived security. Then, we develop a new model of trust, based on literature in both computer science and philosophy. In section 4, we describe the role of technology in the trust relation, focused on information technology. In the final section, we combine the results in an analysis of trust in electronic voting. The paper ends with conclusions.

## 2 Actual and perceived security

In the prevailing approach to social aspects of information system security, the notions of "actual security" and "perceived security" are used. "Actual security" can be assessed by technical experts, and "perceived security" is a more or less distorted version of this in the mind of a member of the non-technical community. From this point of view, trust is based on "perceived security", as opposed to "actual security". It can easily be determined to be either justified or unjustified depending on the agreement between the perceived and actual security of the system.

This distinction can also be applied to election systems. Xenakis and Macintosh [30] argue that "[s]ince procedural security is *evident and understandable* to voters, it has a comparative advantage when it comes to developing and supporting the social acceptance for the new e-processes" [our italics]. In the case of procedural security (as opposed to technical security), the actual security of the system can apparently be perceived by the voters, such that trust can easily be established and justified.[1] This yields the hypothesis that resistance to

[1] There is a remarkable resemblance here to Descartes conceiving certain ideas as "clear and distinct". It is supposed, in both cases, that there are certain things that are understandable by just common sense, as opposed to derived or expert knowledge. These things can be directly extracted from experience, such that "perceived" and "actual" coincide. However, many researchers after Descartes' time have confirmed that there is much more "constructed" about our experience of even common

electronic voting is explained by the fact that the paper system is evident and understandable to voters and electronic systems are not.

We have two main objections against this approach. First of all, the idea that some people have access to the actual security of a system and others do not is problematic. Science and technology studies have been showing for a long time that scientific research and technological development are full of matters of perception and acceptance themselves,[2] and we think that it is essential for understanding the issues to provide a more integrated perspective.

Moreover, the "actual security" of a system is often verified in terms of a model of security that has its own limitations. It is well known in the field of information security that even small security protocols may contain major flaws that go unnoticed for decades, precisely because verification is done using a limited model of security.[3] Also, security assessment of systems involves looking into the future, trying to guess what attackers will be up to. Thus, the tools available for assessing "actual security" are inherently fallible.

Next to the general objections, we also think that this approach does not reach the core of the matter in the case of electronic elections. Although there is a difference in degree of complexity between the paper system and electronic systems, that does not mean that, in the paper case, everyone just knows what is happening and what the risks are. The paper system is complex enough in itself to reach beyond the understanding of the average citizen, and not to be "evident and understandable". Instead, the security sensitivities of the traditional procedural voting system have been "black-boxed"[4] in our experience of democracy. Only when something goes wrong in an election, the black box of the "evident and understandable" paper election system is opened, and risks are exposed. Meanwhile, electronic election systems have not been black-boxed yet, and their vulnerabilities are out there for everyone to discuss. The whole phenomenon of the traditional system being black-boxed, and therefore being "evident and understandable", is already based on trust. Trust is not a derivative of "actual security" here, but it *defines* "actual security". And this actual security is perceived as well.

These are the main arguments supporting our view that there is no meaningful distinction between "actual security" and "perceived security". Security is always perceived security. Of course, the perception of the security of a system, and the reasons why the system is believed to be or not to be secure, may differ

sense issues than people in the Enlightenment age would have admitted. The apparent clear-and-distinctness of certain things is nothing more than our self-initiated reduction of complexity.

[2] Cf. [14], p. 7: "Modern technology studies have opened up the "black box" of technological development and revealed the intimate intertwinement of technology and society. Both scientific facts and technological artifacts appear to be the outcome of negotiations, in which many diverse actors are involved."

[3] For example, the Needham-Schroeder protocol was thought to be correct for 17 years, until it was eventually broken [16].

[4] Cf. the actor-network theory of Bruno Latour, as explained in [29]: ch. 6.

from person to person based on the tools[5] of analysis that are available, which are different for an expert than for a layman. However, there is no such thing as "actual security" to be considered apart from the tools that were used to determine it. Just as "actual intelligence" is not an objective property measured by an IQ-test, but rather defined in terms of the test, security is defined in terms of the available tools for analysis.

As philosophers, we conclude that the distinction between actual and perceived security gives a too naive realist account of the matter.[6] Instead, we start from the observation that people *experience* an environment as more or less secure. This can be seen as a phenomenological approach to the issue [11, 29]. Based on this perspective, we argue that trust is the primary factor in the relations between humans and systems when it comes to security, and not a derivative of an objective kind of security. We will develop this approach further in the following sections.

## 3 Trust

First of all, we need to make sure that the reader is familiar with the distinction between *safety* and *security.* Safety refers to limited effects of the possible failure of a system under normal circumstances. Security refers to limited effects of an attacker deliberately trying to make the system fail in the worst possible way. In scientific research into technological systems, the first property is estimated by verifying the *correctness* of the design. Security is assessed by verifying the *tamper-resistance* of the design. In computer science, these two branches of research can be distinguished easily.[7]

In society, trust relations with respect to systems are characterised by being concerned with either safety or security. In the first case, the trust relation involves trust in the limited effects of failure; in the second case, it involves trust in the limited effects of attack. For example, trust in a nuclear power plant is composed of trust in the *safety* of the plant (e.g. it does not explode under normal circumstances) and trust in the *security* of the plant (e.g. it does not explode under terrorist bombing).

A second distinction that we wish to draw here is connected to the scope of the effects of failure or attack. These effects may be either *private* or *public.* When I drive a car, I trust in the limited effects of failure of the car for my own health. When I vote, I trust in the limited effects of failure of the election system for the whole country. The same holds for the effects of attack: an attack on a

[5] Tools are meant in a pragmatist way here; these include all that people have at their disposal for problem solving purposes.

[6] Recent attempts to derive a taxonomy for dependability and security of information systems suggest a realist view as well [2]. We do not doubt the value of such approaches as a tool for analysis. We do think, however, that a philosophical approach may shed more light on the origin of the concepts used in such a taxonomy.

[7] Note that we do not consider correctness and tamper-resistance as objective properties apart from the (scientific) tools that were used to determine them.

nuclear power plant may have private effects (if I or some of my friends live near it) and public effects (changes in politics due to the attack).

Having set this general background, we now investigate the concept of trust itself a bit further. One of the most confusing things that emerges from the literature is the existence of two different conceptions of trust. On occasion, they even appear in the same article [8]. Although the analysis of trust in voting systems that is presented there covers many concrete risks involved in using these systems, the conception of trust that is used is apparently not completely coherent. In a section named "*Increasing* trust" [our italics], 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?

Apparently, computer 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." [21] Because we *must* trust certain parts of the system for the whole system to be verifiably correct according to the computer science models, we want 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 [20], 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."

In the following, we try to alleviate this confusion by explicating the assumptions found in both approaches to trust, and placing them within a larger (philosophical) context. Apparently, two different definitions of trust have to be distinguished (cf. [21]):

- 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.

How can we conceptualise this difference? 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.[8] We see a parallel here with two possible concepts of safety and security, namely a negative and a positive one:

[8] Cf. [6], pp. 36-39. The notion was originally introduced by Isaiah Berlin [3].

– 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, trust defined in this way can indeed be minimised. 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 computer science (remove all insecurities).

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. In philosophy, distinctions between different modes of trust have been drawn before. We will use such a distinction to further clarify the differences.

Luhmann [17] 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.[9] 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.[10]

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. Luhmann distinguishes personal trust, i.e. trust in interpersonal relations, from system trust, i.e. trust in the general functioning of a non-personal system. We may expect something from a person, or we may expect something from society as a whole or from a machine. Since we

[9] The function of trust as a means for reduction of complexity seems to be known in computer science. For example, Nikander and Karvonen [21] mention this aspect. However, this paper does not refer to the work on trust by Luhmann.

[10] 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.

are interested in technological systems here, trust in this paper is always system trust.

In later work [18], Luhmann 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." [18] 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 computer 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 computer 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.

From now on, we will use *relations of self-assurance* as a general notion. Confidence and trust will only be used in Luhmann's sense. We describe relations of self-assurance based on three distinctions:

- self-assurance with respect to safety vs. self-assurance with respect to security;
- self-assurance with respect to private effects vs. self-assurance with respect to public effects;
- confidence vs. trust.

Computer 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.[11] 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). Our point of view in this discussion is that, because society is too complex for everyone to understand all the risks, there should be a balance between the trust experts have in the system, based on their analysis of the risks, and the confidence the users have in the system. This ensures that there *is* knowledge of the detailed workings and risks of the system within the *social* system in which it is embedded, but there is no need for everyone in the social system to know exactly what these risks are, precisely because there is a relation between expert trust and public confidence. How to establish such a relation is a question that we do not discuss further here.

Based on the distinctions we discussed in this section, we will now turn our attention to trust in technology.

## 4 Trust in technology

When discussing security aspects of technology, reliability and trustworthiness are often mentioned. First of all, we propose a distinction between reliability and trustworthiness. A system acquires *confidence* if it is *reliable*, and it acquires *trust* if it is *trustworthy*.[12] 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.

There is a fairly subtle relation between reliability and trustworthiness. On the one hand, trustworthiness is a stronger notion than reliability. Before they give their trust to a system, people will perform a risk analysis. People who establish confidence in a system do not do this. In this sense, it is harder for a system to *acquire* trust than to *acquire* confidence. However, *maintaining* trust is easier than *maintaining* confidence. When people trust a certain system, they are already conscious of the risks and decide to use it anyway. This means that trust is not necessarily broken if something fails. In the case of reliability, however, people have put their confidence in a system because they do not see alternatives, and they will probably not accept any failures. Trustworthiness is therefore the stronger notion for the short term, and reliability is the stronger notion for the long term.

How are reliability and trustworthiness established? As we have made clear in the introduction, we argue that they are not objective properties of a sys-

[11] This general approach is not without exceptions; cf. [20].

[12] Reliability is used in the more limited sense of continuity of correct service in [2]. 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 that are reflected in subjective confidence and trust. Instead, we take a phenomenological point of view, i.e. we conceive the relation between persons and a system as primary to the objective and subjective aspects [11, 29]. The objective aspects of reliability and trustworthiness and the subjective aspects of confidence and trust emerge from the relation between people and the system. The way in which they are established depends on the analytic tools that are available to the person. If a person is just using the system, the outcome will probably be different than in case an expert performs a full security audit based on her expertise.

The relations that different people have with the system make the objective aspects of reliability and trustworthiness converge into different images of the system. These images then become "objective" properties of the system. The relations that experts have with the system determine what is often called the "actual" security of the system, but this "actual" security is still based on perception and relations of self-assurance, and therefore we rather avoid the term "actual".

How does this analysis of trust in technology apply to computer systems? Computer systems can be characterised as congealed procedures. Such procedures are typically more rigid than human-managed procedures. They are less easy to circumvent, but also less robust. Humans are easy to persuade to abandon the rules, computers are not. Humans can easily find solutions for problems that do not exactly match the rules, computers cannot. Because computers are not flexible, congealed procedures must be specified in more detail than human procedures. Every possible situation should be covered. This, and the fact that most people do not have expert knowledge about computers, makes congealed procedures hard to understand.

As we have seen before, trust in a system requires understanding of the risks involved in using a system. This is usually relatively easy to achieve in human procedures, not necessarily because the systems are less complex, but because we have a good understanding (or at least we think we have a good understanding) of how humans function. Understanding the risks of using a computer system is typically much harder. On the other hand, precisely because congealed procedures are more rigid, the associated systems are generally more reliable, in the sense that they produce fewer errors. This makes them more suitable for being reliable and acquiring confidence, while less suitable for being trustworthy and acquiring trust. Thus, automation implies a transition from public trust to public confidence. This makes it all the more important that a balance is established between expert trust and public confidence, in order to have public confidence still reflect a risk analysis in some way.

Luhmann observes the same tendency of replacing trust by confidence in functional differentiation of society. Because people in a functionally differentiated environment have knowledge of only a very small part of the complex structures that surround them, establishing trust is hard, and confidence is extremely important. This also requires procedures to be increasingly rigid, precisely because they need to maintain confidence. This may be seen as a first step in the freezing of procedures; automation is then a continuation of this process, by entering these semi-frozen procedures into machines, and thereby fixing all the details even further.[13]

The concepts of reliability and trustworthiness extend the conceptual framework we introduced in the previous section. We will now investigate whether this framework yields new results when we apply it to voting technology.

## 5 Trust in voting systems

Voting is a way to surrender oneself to a representational body in a democracy. It is at the same time a reconfirmation of the social contract between the rulers and the ruled, and a reconfirmation of the autonomous individual that engages in this contract. In this act, the Enlightenment ideals are established over and over again. The reconfirmation of the social contract and the autonomous individual has the character of a ritual. The associated relation of self-assurance is primarily based on familiarity, for a ritual always has an orientation towards the past.

But this ritual dimension is not the only relation of self-assurance in democratic politics. There are also expectations involved about the functionality of the political system, for example the expectation that the desires of the public are accurately represented in policy by this system. Engaging in political activities such as voting requires confidence or trust that these expectations will be fulfilled. Finally, there is also a need for trust and confidence in the present government, which represents the people based on expectations not about the political system in general, but about the current policy.

Thus, elections involve at least three different relations of self-assurance: the familiarity with democracy that is established by means of a ritual, the confidence or trust that people have in the government and confidence or trust in the political system.[14]

However, trust and confidence in the election procedures themselves are also necessary. These in turn co-evolve with the relation of self-assurance that people have with the government and the political system. This means that a lack of trust or confidence in election procedures may reduce trust or confidence in the

[13] Interestingly, this transformation of trust in human procedures into confidence in congealed procedures goes against the tendency that Luhmann observes in liberalism. According to Luhmann, “liberalism focuses on the individual responsibility for deciding between trust and distrust [...]. And it neglects the problems of attribution and the large amount of confidence required for participation in the system” [18]. From this point of view, either information technology is a threat to liberalism, or liberalism should revise its goals.

[14] Generally, people have confidence with regard to politics rather than trust, in Luhmann’s sense. It is precisely the phenomenon of elections that may turn political confidence into trust: “A relation of confidence may turn into one of trust if it becomes possible (or is seen to be possible) to avoid that relation. Thus elections may to some extent convert political confidence into political trust, at least if your party wins. Conversely, trust can revert to mere confidence when the opinion spreads that you cannot really influence political behaviour through the ballot.” [18].

government or the political system, but also the other way around. The specific characteristics of this relation are a topic for further research. In this section, we will focus on trust and confidence in the election system. We will primarily discuss the differences that can be observed from this point of view between the traditional paper systems and electronic variants.

Why do we want electronic voting? The rigidness of technology is often an argument. Errors with paper ballots, as in the Florida presidential elections in 2000, may be a reason to switch to the supposedly more reliable Direct Recording Electronic (DRE) machines. Indeed, electronic machines may be more reliable and trustworthy with respect to *safety* than paper systems are, because possibilities for error, both in casting and in counting, are reduced.

However, reliability and trustworthiness with respect to *security* are not as straightforward, especially when there is little transparency in the design, e.g. when the source code is kept secret. Acquiring trust in security, as opposed to trust in safety, is hard when things are secret. Insider attacks against security, e.g. an employee of the manufacturer changing something in the software, are indeed pretty easy in such a case, and experts evaluating the risks will at some point notice this. This not only includes possibilities for altering the results, but also the possibility to deduce a relation between a voter and a vote.[15] This lack of transparency may make it hard as well to maintain public confidence in the long run, since this confidence is often influenced by expert trust.

Besides the distinction between security and safety, we also proposed a distinction between private effects and public effects. Self-assurance with respect to private effects in voting amounts to trust or confidence that one's own vote is handled correctly, e.g. kept confidential. Self-assurance with respect to public effects means trust or confidence that the results are calculated correctly. *Both* kinds need to be acquired by an election system. People may have confidence in electronic systems in the sense that they calculate the results correctly in general, but if they are not sure what happens to their own vote – e.g. doubt the secrecy of their vote – the whole system may not acquire confidence anyway.

In the previous section, we argued that congealed procedures are more suitable for confidence, whereas human procedures are more suitable for trust. Still, because the paper system has been the only option for a long time, the relation of self-assurance people had with the paper system was largely based on confidence. Confidence in the election system, confidence in the government and confidence in the political system supported each other. Now, what happens when electronic voting comes into play?

First of all, electronic voting systems *may* be seen as alternatives to the existing system. Whether this is indeed the case depends on the situation. 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

[15] We do not discuss vote buying and vote coercion in this paper; see e.g. [23] for discussions on this issue for the case of Internet elections.

as a consequence that voting systems are required to be *trustworthy* rather than reliable only. This, again, leads to the traditional paper system becoming *more* attractive, because it is based on human procedures, and human procedures more easily acquire trust than congealed procedures. On the other hand, if the new technologies are not seen as an alternative, but as an improvement of existing procedures, electronic devices are more attractive, because they are more reliable and thus more easily acquire confidence.

If various alternatives are available, and citizens cannot assess the risks themselves, it can be desirable to establish a balance between expert trust and public confidence, in order to establish a relation of self-assurance between citizens and the election system again. This is important for maintaining people's confidence in the government and the political system. However, if people do not see these options as alternatives, risk analysis may instead break their confidence in the existing system by exposing the risks, and thereby destroy confidence. Thus, the role of the expert in these matters is extremely important. This role will be a topic of ongoing research.

As an example of the value of our approach for the analysis of concrete developments, we propose some hypotheses as explanations for the fact that in the Netherlands, electronic voting machines have been introduced in the early nineties without much discussion about their security. It was not regarded a serious problem that the design was secret, and only the independent voting system licenser TNO knew the details. Most of the concern was about whether all citizens would be able to operate the machines. Possible hypotheses for the smooth and uncontroversial introduction are:

- the ritual of going to the polling station, identifying oneself and casting a vote remained fairly stable (as opposed to online voting), maintaining familiarity;[16] also, the Dutch machines have a layout that is very similar to the paper ballots used before;[17]
- confidence in the government was relatively high, which led to confidence in the election systems proposed by the government as well;
- trust and confidence in information systems were more related to safety than to security at the time; people knew that calculators were reliable, and probably no one had ever attacked a calculator;
- voters paid more attention to the election outcome (public effects) than to what happened to their own vote (private effects); they knew that computers were able to calculate reliably, and therefore had confidence in the computers with respect to the public effects; focusing instead on the private

[16] In relatively new democracies, such as Estonia, tradition (and thus familiarity) are less important. This may explain why Estonia already implemented Internet voting. See e.g. http://www.euractiv.com/Article?tcmuri=tcm:29-145735-16&type=News, consulted November 17, 2005.

[17] This means, among other things, that all candidates have their own button on the machine, as opposed to PC software in which one first chooses a party and then a candidate.

effects of a machine "stealing" or revealing one's vote will expose the lack of transparency and probably undermine confidence;
- the electronic systems were not seen as *alternatives* to the existing procedures, but rather as automated versions of existing procedures; this made it easy to transfer confidence to the new systems; nowadays, trust *is* an issue in other countries: e-voting is really seen as an *alternative*, instead of just automating a known process;
- risk evaluation of computer systems was not as mature at the time as it is now; this made it harder for computer scientists to transform confidence into trust by making explicit the risks involved.

Each of these possible causes, which are based on the philosophical analysis in this paper, can serve as a hypothesis for empirical research. Also, the fact that voting machines are now under discussion in the Netherlands as well may be explained by a change in situation with respect to these hypotheses. For example, international developments may have changed the image of voting machines from a simple automation of existing procedures to a real alternative. These hypotheses show the relevance of our conceptual framework for voting system sciences in general. Of course, some of them are related, and further research, both theoretical and empirical, would be useful to determine these interdependencies.

## 6 Conclusions

In this paper, we described a framework for discussing trust in relation to voting procedures. Instead of distinguishing between actual and perceived security, we took a more phenomenological approach, in which subjective and objective aspects of security are seen as constituted from the relation between the people and systems involved. The main concepts were discussed both from a computer science point of view and from a philosophical perspective. Luhmann was the primary source for the latter.

Based on the theory of Luhmann, we distinguished between familiarity, confidence and trust. Luhmann understands these concepts as means for the reduction of social complexity. Familiarity has an orientation towards the past, whereas confidence and trust are based on expectations and thus oriented towards the future. People trust because they want to, based on risk evaluation. People have confidence because they have to, not because they understand the risks. The concepts of confidence and trust are related to the different views on trust that can be found in the computer science literature, namely bad and good trust. These are again related to negative and positive conceptions of security, respectively. Computer scientists generally try to replace confidence with trust by making explicit the risks involved in using the system. This, again, allows the public to base their confidence on expert trust.

The "objective" aspects related to the "subjective" aspects of confidence and trust were labelled reliability and trustworthiness. Human procedures are typically good at being trustworthy (and thus at acquiring trust), whereas the congealed procedures of computers are good at being reliable (and thus at acquiring confidence).

In elections, the traditional election system, whatever it may be, always invites confidence, precisely because it is the established system, and people are not conscious of alternatives. When new technologies for elections are presented, these may be seen as alternatives. Then, election systems suddenly have to be trustworthy instead of reliable only. This is one of the reasons why the demands posed on new election technologies are often more severe than those posed on existing systems. However, the fact that alternatives are now available may also undermine confidence in the existing system, and require this system to earn trust as well.

In this situation, an interdisciplinary approach to matters of trust in election systems is indispensable. The hypotheses we offered for the smooth introduction of voting machines in the Netherlands serve as a modest attempt at illustrating possible results. We hope to have justified trust in the benefits of such an approach here.

## References

1. R.M. Alvarez and T.E. Hall. *Point, click & vote: the future of Internet voting.* Brookings Institution Press, Washington D.C., 2004.
2. 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.
3. I. Berlin. *Four concepts of liberty.* Oxford University Press, Oxford, 1969 [1958].
4. D. Chaum. Secret-ballot receipts: true voter-verifiable elections. *IEEE Security & Privacy*, 2(1):38–47, 2004.
5. K. Chopra and W.A. Wallace. Trust in electronic environments. In *Proceedings of the 36th Hawaii International Conference on System Sciences (HICSS'03)*, 2002.
6. F. Cunningham. *Theories of democracy: a critical introduction.* Routledge, London, 2002.
7. J. Dewey. *The public and its problems.* Swallow Press / Ohio University Press, Athens, 1991 [1927].
8. D. Evans and N. Paul. Election security: perception and reality. *IEEE Security & Privacy*, 2(1):24–31, January/February 2004.
9. D. Fahrenholtz and A. Bartelt. Towards a sociological view of trust in computer science. In M. Schoop and R. Walczuch, editors, *Proceedings of the eighth research symposium on emerging electronic markets (RSEEM 01)*, 2001.
10. E.-M.G.M. Hubbers, B.P.F. 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.
11. D. Ihde. *Technology and the lifeworld.* Indiana University Press, Bloomington, 1990.
12. D. Jefferson, A.D. Rubin, B. Simons, and D. Wagner. Analyzing internet voting security. *Communications of the ACM*, 47(10):59–64, 2004.
13. C. Karlof, N. Sastry, and D. Wagner. Cryptographic voting protocols: a systems perspective. In *Proceedings of the 14th USENIX Security Symposium*, pages 33–50, 2005.

14. J. Keulartz, M. Korthals, M. Schermer, and T. Swierstra. Ethics in a technological culture: A proposal for a pragmatist approach. In J. Keulartz, M. Korthals, M. Schermer, and T. Swierstra, editors, *Pragmatist ethics for a technological culture*, chapter 1, pages 3–21. Kluwer Academic Publishers, 2002.
15. T. Kohno, A. Stubblefield, A. D. Rubin, and D. S. Wallach. Analysis of an electronic voting system. In *Proceedings of the 2004 IEEE Symposium on Security and Privacy*, 2004.
16. G. Lowe. Breaking and fixing the Needham-Schroeder public key protocol using FDR. In *Tools and algorithms for the contruction and analysis of systems*, volume 1055 of *Lecture notes in computer science*, pages 147–166. Springer, 1996.
17. N. Luhmann. *Trust and power: two works by Niklas Luhmann.* Wiley, Chichester, 1979.
18. N. Luhmann. Familiarity, confidence, trust: problems and alternatives. In D. Gambetta, editor, *Trust: Making and breaking of cooperative relations.* Basil Blackwell, Oxford, 1988.
19. D.P. Moynihan. Building secure elections: E-voting, security and systems theory. *Public administration review*, 64(5), 2004.
20. 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.
21. 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.
22. A.M. Oostveen and P. Van den Besselaar. Security as belief: user's perceptions on the security of electronic voting systems. In A. Prosser and R. Krimmer, editors, *Electronic Voting in Europe: Technology, Law, Politics and Society*, volume P-47 of *Lecture Notes in Informatics*, pages 73–82. Gesellschaft für Informatik, Bonn, 2004.
23. 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.
24. B. Randell and P.Y.A. Ryan. Voting technologies and trust. Technical Report CS-TR-911, School of Computing Science, University of Newcastle upon Tyne, 2005.
25. R. Riedl. Rethinking trust and confidence in european e-government: Linking the public sector with post-modern society. In *Proceedings of I3E 2004*, 2004.
26. A. Riera and P. Brown. Bringing confidence to electronic voting. *Electronic Journal of e-Government*, 1(1):43–50, 2003.
27. A.D. Rubin. Security considerations for remote electronic voting. *Communications of the ACM*, 45(12):39–44, 2002.
28. B. Shneiderman. Designing trust into online experiences. *Communications of the ACM*, 43(12):57–59, 2000.
29. P.P.C.C. Verbeek. *What things do: Philosophical Reflections on Technology, Agency, and Design.* Pennsylvania State University Press, 2005.
30. A. Xenakis and A. Macintosh. Procedural security and social acceptance in e-voting. In *Proceedings of the 38th Hawaii International Conference on System Sciences (HICSS'05)*, 2005.