# A Framework for Studying Vote Secrecy in Elections

Carsten Schürmann[1]
IT University of Copenhagen
carsten@itu.dk

## Abstract

In this draft paper, we study the role of vote secrecy in elections. Besides election integrity, vote secrecy is the second defining characteristics of genuine and credible elections. Modern vote casting methods, including internet voting and optical ballot scanners are game changers in terms of understanding and guaranteeing vote secrecy. This is because the concerns of election integrity can no longer be cleanly separated from vote secrecy, as it used to be the case for pen-and-paper election in controlled environments. Many modern election methods resolve this tension by regarding election integrity as more important than vote secrecy. Universal verifiable systems for example, are designed with election integrity in mind; but a universal verifiable election technology may break vote secrecy. To study this phenomenon in a bit more in detail, we present in this draft paper a framework and demonstrate it using three election methods, pen-and-paper voting, internet voting, and optical scan voting.

## Introduction

The two hallmark characteristics of genuine and credible elections are election integrity and vote secrecy. Election integrity refers to the accuracy of the result and that it correctly reflects the voters' intent, whereas vote secrecy refers to ambition to make sure that the voter and only the voter will know what he or she voted for. It is exactly those two properties that render the use of modern election methods (often using Information and Communication Technologies ICT) difficult. Voter registration systems have access to databases identifying eligible voters, and result management systems have access to the results. Voter secrecy entails, that during this process, the identity of a voter must be permanently and irrevocably severed from the vote.

The notion *vote secrecy* can be precisely defined in a technical sense --- it refers to the fact that the relation between voter and vote must not be known to anyone but the voter. It does not mean that the vote per se should be secret -- after all, votes need to be counted. It also does not mean that the the identities of those who voted must be kept secret. Breaking vote secrecy

[1] This paper was made possible by grant NPRP 097-988-1-178, Automated verification of properties of concurrent, distributed and parallel specifications with applications to computer security, from the Qatar National Research Fund (a member of the Qatar Foundation). The statements made herein are solely the responsibility of the author.

is therefore tantamount to revealing the relationship between voter and vote. In the literature there has been much discussion if vote secrecy is a duty or a right, and for good reasons: Even in a paper only election, the use of technology, such as film cameras, or other, enables voters to document the process of voting, and hereby breaking vote secrecy. In the case that it is a duty to preserve vote secrecy, this duty applies to all voters --- and voters caught breaking vote secrecy may be penalized. In contrast, if vote secrecy is a right then this would imply that voters should have the ability to waive this right, and in practice, voters are actually asked to waive this right, if returning ballots by mail, fax, or email. In this draft paper we quantify the effort required by an attacker to break vote secrecy on a global scale vs. directed against individual voters.

In non ICT-based elections (where voters using paper to record their preferences), the fact that paper is not tracked (except perhaps in Great Britain) and that the ballots are shuffled before opening the ballot box protects to a certain degree vote secrecy. In contrast, tracking and shuffling is difficult to imitate in the digital world: Once a computer system “knows” about a voter’s identity and the vote, for example during vote casting using an internet based voting system, it is almost impossible to make the system forget. Forcing the computer to forget, by deleting information is usually not a good strategy, because an insider or a backup service could have copied this information already elsewhere, a software bug prevented the deletion, or a versioning file system may not really delete, but mark this information as deleted, instead.

One way to implement forgetting in internet based voting systems is by using modern cryptography, in particular, mixing networks [1]. Such as network consists of several mixing computers, each of which shuffles the ballot box digitally and rencrypts the vote, which means that it changes the ciphertext (the result of encrypting the votes) without modifying the vote itself. However, this comes also at a price: Although the use of cryptographic protocols can be used as an effective tool to protect vote secrecy, it threatens election integrity: What if a dishonest mixing computer changes the content of the votes? Clearly, this would be very difficult to detect, because from ciphertexts are designed to resemble random noise.

In order to alleviate this problem, modern mixing networks emit digital evidence, so called zero-knowledge proofs of knowledge (ZKPKs), that can be independently checked by a third party, for example a scrutineer, a party agent, or even the voter. If all ZKPKs are validated, we can be certain that 1) the original and the shuffled ballot boxes contain the same vote and 2) that the link between voters and votes is permanently severed. But this argument requires at least one mixing servers not to reveal the permutation and it also requires a change in voter behavior, who no longer only vote but are also expected to verify ZKPKs.

For several years now, two other technologies have become popular. Ballot marking devices and optical ballot scanners. Both devices provide mechanisms for voters to check if their intent of whom to vote for is correctly recorded. A ballot marking device is a computer system, usually equipped with a touch screen, where voters record their preferences electronically, and then a paper ballot is printed in clear text and/or a digital barcode. This paper ballot represents the voters intent and (depending on local election laws) may be interpreted as the binding vote.

Optical scan machines, on the other hand work differently. They expect voters to have completed paper ballot forms with little ovals that need to be completely blackened before they represent valid marks. Completed forms are then being fed into an optical scanner that interprets that digital scan of the ballot as a representation of the vote.

The main contribution of this paper is a framework with which allows us, to compare the vote privacy for different election methods, including manual ballot interpretation, paper ballots, internet voting, and optical ballot scanners. The framework takes different factors into account, including (1) the degree to which secrecy of the vote is affected by a particular vote casting method, (2) the legal basis of vote secrecy (3) the assumptions under which vote secrecy can be guaranteed, in which scope and to what extent, (4) the balance between vote secrecy and election integrity for a particular election method and (5) the degree to which the introduction of voter verification mechanisms affects the quality of the overall election.

When studying vote secrecy, one strictly speaking also needs to look at the instruments social media companies such as Facebook, Google, Apple or others have, and how these instruments can be used to predict (statistically) voter behavior and therefore break vote secrecy, without ever setting a foot into a polling place. Although we recognize importance and the impact of such instruments, it is outside the scope of this draft paper, and we will not discuss it further as the methodology to study their effects is orthogonal to the methodology underlying this work.

# A framework for classifying secrecy

## Understanding Vote Secrecy

The notion of vote secrecy needs to examined from two vantage points. First, there is *personal vote secrecy,* which refers to if the vote casting mechanism and the setting in which voting takes place can (if used correctly) protect the relation between the voter and his or her vote. Not all countries provide personal vote secrecy: In Switzerland, for example, voting on a market place by raising one's hand, definitely violates this property. Personal vote secrecy can be guaranteed to a large extent, if voting takes place in a controlled environment that is adequately prepared to provide for vote secrecy. For example, curtains are useful to give voters a private space, which gives adequate protection for pen and paper ballots. If technology is used in polling stations, then vote secrecy depends heavily on which technology is used and it may be that the curtains do not provide enough protection for example to shield electronic magnetic waves from the voting device. If not, an adversary could record those waves and analyze it to determine who voted for what. Attacks of that form become relevant, for example, if electronic ballot markers are used. It is even more difficult to protect personal vote secrecy in uncontrolled environments, such as voting from or over the internet. The device used for casting a vote may be infected will malware or keyloggers, communicating the vote back to the adversary in clear text.

Second, there is a property that we call *group vote secrecy*, which refers to if there are mechanisms in place that protects the vote secrecy of other voters. For example, in smaller voting districts with low voter participation, if there is group of voters that chooses to reveal their intent (for example by filming the voting process and publishing the film), such an action may coincidentally also impede the vote secrecy of others.

In our framework, we consider to which extent a particular voting technology affects personal and group secrecy.

## Legislatislative Differences for Vote Secrecy

The respective election laws of most countries live up to the requirements described in the Universal Declaration of Human Rights, paragraph 21 (3). “The will of the people shall be the basis of the authority of government; this will shall be expressed in periodic and genuine elections which shall be by universal and equal suffrage and shall be held by secret vote or by equivalent free voting procedures”. When we look at the details, however, the different laws handle vote secrecy differently. Two main classes exist, some countries treat vote secrecy as a right whereas others treat it as a duty. In the former, a voter has the right to keep the vote secret, meaning that personal vote secrecy is guaranteed, unless the voter actively takes steps to record and make know the way he or she voted. In the later case, the fact of documenting and revealing evidence of one’s vote

In our framework, we consider to which degree, if any, a particular technology can support vote secrecy either as right or as duty.

## Global vs Local Attacks against Vote Secrecy

It is clear that the technology used for vote casting has a direct effect on the attack surface of the electoral process. In paper voting, an adversary trying to learn the relation between voters and votes is limited in staging an attack, because the the adversary’s access to observations in the polling place is limited. In contrast, in internet voting, if the voter register and the vote collection server are under the adversary’s control, the adversary has access to observations on a global scale.

In our framework, we consider the extent to which an adversary can break vote privacy.

## Vote Secrecy vs. Election Integrity

To be sure that an election has integrity it must produce an audit trail. In paper based elections, the actually paper ballots *are* the audit trail. However, an audit trail is only as good as the

mechanisms that check it. Different voting technologies offer different options on how this paper trail can be checked, for example when, i.e. during the election or post election, how, i.e. manually by people or computers, and by whom, i.e. election officials, consultants and experts, scrutineers and agents, or the voters. Such mechanisms are designed primarily to build trust in the election outcome, however, the particular choices also affect vote secrecy. Universally verifiable elections, for example, are designed to involve the voters in checking their part of the audit trail. Voters receive a piece of information pertaining to their vote, and with this information, they can verify that their vote was included in the final count. Some Internet voting systems also produce cryptographic evidence, i.e. zero knowledge proofs of knowledge that prove the correct re-encryption of a ballot box, an operation that changes the appearance of all ciphertexts without changing the votes that are being encrypted, and correct decryption, an operation that decrypts a ballot box into human readable form. Zero knowledge proofs are then shared among all voters via a so called web bulletin boards, so that each voter can check all proofs, basically allowing each and every voter to check the integrity of the result.
Zero-knowledge proofs, technically speaking, do not reveal any knowledge about a particular secret if used correctly. But this does not mean necessarily that vote secrecy is automatically guaranteed. Zero-knowledge proofs support claims of election integrity, but not necessarily of vote secrecy.

In our framework, we consider to which extent the evidence produced by a particular technology supports vote secrecy.

## Verification and Vote Secrecy

Universally verifiable voting protocols rely on voters not only to vote during an election, but also to check the artifacts/evidence that was produced during the election. This appears to be a fundamental break with the traditional voting where voters record their preferences on ballot papers before putting them into the ballot box and then have no other option but to trust that their ballots are included into the count. In universally verifiable systems, voters no longer have to trust, but instead they can verify that the counting process was correct. This apparent improvement, though, comes at a price: the voting and counting process must provide mechanisms to deal with recovering from failure. Examples include what happens if voters claims that they cannot find their votes on the bulletin board, or what happens if the zero-knowledge proof of knowledge checks fail. Universal verifiability does not provide an answer, but instead brings along additional challenges. The mechanisms designed to handle failure of verifications may be exploited by an adversary, for example the losing party, by crying wolf. So this means that the verification mechanisms must be designed in such a way, that it is impossible to claim failure without evidence, and this evidence may, in principle, violate vote secrecy.

In our framework, we consider to what extend a verification mechanisms for various election methods and technologies support vote privacy.

# A Comparative Analysis of Voting Methods

Having the framework in place, we will now consider three different voting methods, and classify them according to the criteria laid in out in the previous section. The three methods are paper voting (as exercised in Denmark, Germany, or the Netherlands), internet voting (for example, as used in Norway 2013), and optical scan machines (as used in the USA). We realize of course, that the different voting methods are very specific to the respective countries that use them, so our findings do not necessarily transferable from one country to another.

The following table summarizes our findings.

| | Paper Voting (Denmark) | Internet Voting (Norway) | Optical Scan (USA) |
|---|---|---|---|
| Personal VS | Yes. Controlled environment | No | Partially. Controlled environment |
| Group VS | In large polling stations currently yes, may change if there are other techniques to broadcast. In small polling stations no. | No. The return code SMS on a mobile phone together with the voter card can be used as evidence how someone voted. | |
| VS as a right | Yes. | No. (1) The Norwegian does not guarantee everlasting privacy, (2) Once the election key is assembled VS may be compromised. | Yes. |
| VS as a duty | Difficult to enforce. | No. | Difficult to enforce. |
| Local attacks against VS | No. | Yes. The election private key can be used to decrypt all votes in the database, and then to learn who voted for what. | Yes. Election officials can see completed ballot papers rejected by the optical scan machine. |
| Global attacks against | No. | Yes. The election private | No. |

| | | | |
|---|---|---|---|
| VS | | key can be used to decrypt all votes in the database, and then to learn who voted for what. | |
| Election Integrity compared to VS | Two stages: (1) vote secrecy is established the moment ballot boxes are opened. (2) the process is organized to guarantee vote integrity. | Vote secrecy and election integrity concerns are intertwined. The Norwegian election system prioritizes election integrity over VS. | Optical scan machines interpret paper ballots, but they keep the paper ballots as a paper trail that may be audited. The technical reality pose a real challenge for vote secrecy. Vote secrecy and election integrity concerns are intertwined. |
| Verification mechanisms | N/A | Unclear | Yes. The it is the verification mechanism that breaks vote secrecy. |

## Case Study: Paper Voting in Denmark

Paper voting systems are in generally considered privacy preserving, if administered correctly. Voting takes place in controlled environments, curtains give voters privacy to record their votes in secret. In Denmark, after election law, [3] paragraph 45(2), has the voter the right to cast a vote in secret, and the election officials must prepare the space in such a way that this is guaranteed. There are, of course, a few exceptions. In the UK, for example, voters also vote in secret (check), but the ballot forms are numbered that in case the court orders a vote to be revealed, election officials can identify retroactively the vote of particular individuals. We should mention that vote secrecy in Denmark and other countries using paper ballots relies on the fact that empty ballot forms cannot be distinguished. This assumption, however, is no longer true [2]. Each ballot paper, when magnified, exposes a different surface structure. If the election commission were to record this information, it is actually possible to break vote secrecy when only theoretical, because it could be correlated to the registration logs modern voter registration systems create when recording a timestamp with the arrival of each voter. Further advances of information technology may invalidate other common assumptions as well.

As voting is a right in Denmark, it would be possible, although most likely not practical that all voters take pictures/movies of the completing and casting their ballot papers and post them on social media. As it is known to the authorities who is registered and who did vote in a particular polling station, the publication of pictures/movies begins to infringe on the group vote privacy of

other voters. It is not known, if in Denmark or elsewhere, digital recording of ballots is a systemic problem, probably not, but in the future, it might.

In every polling station in Denmark, the ballot boxes used are made of cardboard and resemble moving boxes. They are not transparent. But they are being assembled in the presence of a group of witnesses (an event that is open to the public) when the election opens. This way, the emptiness of the ballot box is confirmed, while guaranteeing that it is extremely difficult, if not impossible to link vote and voters. Therefore the risk of a local attack against vote secrecy is extremely low, and the risk of a global attack even lower.

The Danish election law (as many other election laws in Europe and other countries using pen and paper election) disentangles the concern of vote secrecy from election integrity, Vote secrecy is basically guaranteed when the ballot boxes are opened and the votes are dumped into a big pile onto a counting table.

With vote secrecy guaranteed, the rest of the process is organized to ensure election integrity. This includes that the ballots are counted more than once and by two different groups of people. Indeed, there are many checks that are being conducted, and such checks can fail. However, all of these integrity checks can no longer affect vote secrecy.

## Case Study: Internet Voting in Norway

Next, we consider the internet voting system that was used in Norway in 2011 and 2013 in selected municipalities. (See for example [4]) The system was developed by the voting technology vendor Scytl, and all documentation and source codes are publically available for inspection. The voting process proceeds as follows. Each voter receives a voter registration card in due time before the internet election commences. Besides personal information, the voter card also contains a list of parties and numbers, called *return codes* that voters will be able to use that their vote was correctly recorded as intended. The return codes are unique for each voter.

The voting activity takes place in an uncontrolled environment: voters vote from the comfort of their homes on devices of their choice. For the 2011 election, internet voters had to install a program, the *vote client* on their computers, which forced internet voters to use a Mac or PC. In contrast, in 2013 this software was replaced by a Javascript web client that could be executed in any web browser on virtually any device. Regarding vote secrecy, the main challenge is to find a technical solution to the fact that the device that the voter uses for voting will know both, information about the voters' identity and the vote. Simple keylogger programs and viruses running in the background could steal this information,record it, and share it. Also a programming bug in the vote client software or by a cyberattack against the voting technology vendor or third party websites used in the voting solution [5], might break vote secrecy. It is difficult to detect that such an attack has taken place, after all, it is only data that was stolen.

Modern viruses can be programmed in such a way that they delete themselves after such a privacy attack was successful, meaning that it can be very difficult to find any traces of malfeasance.

Once a vote has been cast, an encrypted version of the vote together with the voter id will be stored in a database maintained by the election commission. This is necessary, as Norwegian voters can re-vote at a polling station with pen and paper, overriding all previous votes cast through the Internet. During election night, after polling stations have closed, but before the votes are counted, the vote database will be cleansed (invalidated votes will be removed). The only defense to protect the vote is the cryptography used to encrypt the vote. Cryptography works analogous to a safe deposit box: Whoever has the (cryptographic) key, learns the vote and can therefore break vote secrecy. This is inherent problem with many of the currently available internet voting system: Vote secrecy is reduced to key management.

In order to mitigate the risk of one person having the election key, the Norwegian system (as well as others) use a technique called threshold cryptography, where the secret election key is either constructed from several key shares (during the decryption ceremony of internet votes) or a generated key is being split into several shares and distributed to the different stakeholders before the election. No matter which way this cookie crumbles, the existence of the secret election key, necessary to decrypt the election result, threatens vote secrecy.

During the Norwegian decryption ceremony in 2013, the secret election key was reconstructed from shares that were previously given to the members of the Internet Voting committee, and recorded on a USB stick, which was than protected by physical means. Although it was unclear what happened to this stick after the observers left, this key could be used to decrypt votes stored in the uncleansed database that contains also voter id information.

The other threat against vote secrecy is the weakness of cryptography itself. The Norwegian system does not guarantee everlasting privacy. No cryptography scheme is 100% secure, no matter the strength of the underlying crypto system. With better and faster computer hardware, smarter hardware, and possibly even quantum computers, we must prepare for the fact that all crypto will eventually be broken. The content of the digital ballot box (if a copy was kept after the election) is bound to be insecure, eventually.

One may be concerned that the open source code policy required by the Norwegian election commission may also pose a threat to vote secrecy. This is not the case. The security of the Norwegian Internet system is derived from the hardness of a mathematical problem --- the discrete logarithm problem --- and not on the source code itself. It is true, however, that programming mistakes can threaten security, and this actually was the case during the 2013 election. Then, the system uses randomized encryption, which means that the ciphertexts (encrypted votes) will always look different, even if two voters voted for the same candidate. To achieve this, the voting client needs to have access to a good source of randomness. Unfortunately, a programming mistake in the voting client prevented exactly this, with the effect

that a substantial number of votes for the same party looked the same in the database. This allowed anyone with access to the database to deduce precisely who voted for whom without decrypting the votes first. Although this was a programming mistake, the fact that an attacker could in principle inject any kind of code into the voting client (using other vulnerabilities of routers and servers), illustrates how bridle vote secrecy is in internet voting systems.

The way how the return code generator and the zero knowledge proof of knowledge checking is an integral part of the internet voting process, it also shows clearly, that the Norwegian have prioritized election integrity over vote secrecy. This is not uncommon observation, which also applies to the Estonian Internet voting system and also others.

Lastly, we turn our attention to the return code mechanism. An adversary may consider crying wolf, and claim that their return code doesn't match. It is unclear, how the Norwegian election commision would react, and what kind of information would have to enter the court proceedings.

## Case Study: Voting on Optical Scan Machines in USA

The United States uses a plethora of voting equipment in their elections. As of lately, most of the old paperless DRE machines have been or will shortly be replaced by optical scan machines. During an election observation visit to the US during the 2016 presidential election, we made several observations that we describe here. In the US, running elections is a responsibility of each county. Some of the US States see voting as a right, others see it as a duty.

An optical scan machine is a machine that scan paper ballots and interprets them digitally. The process works as follows. A voter completes a ballot form by filling in ovals with a pen. In this, it is not much different from a typical pen and paper election, except that the rules of how a ballot is expected to be completed, are much more stringent than in standard pen-and paper elections (due to limits in ballot interpretation). After completing the ballot form, the voter will slip the ballot into a privacy sleeve, which he or she will then carry (unfolded) to a optical scan ballot box. The ballot will then be inserted into a scanning device, ideally directly out of the privacy sleeve. The system will scan the ballot and run a recognition software on it to determine the digital representation of the ballot, i.e. the voter preferences. If the ballot is successfully scanned it drops into a container below the scanner (out of reach of the voters), the ballot box. But not all ballots scans succeed: If an oval is not filled out completely, an oval is crossed out, or even more than the correct number of ovals are filled, the optical scan machine will reject the ballot and spit it out like a parking fee machine, that does not recognize a $20 bill as valid: usually, inserting the same ballot into the same optical scan machine over and over does not change the result.

We can consider an optical scan machine as a verification mechanism that validates, if the ballot inserted was completed in a computer readable way. The privacy sleeve protects the information on a ballot form from praying eyes. However, what happens, if the ballot was not

scanned? In this situation, the voter most likely turns to an election official and explains him or her the situation, using the completed ballot as evidence. Almost certainly, the election official will learn how the voter voted. Therefore, any optical scan technology comes with a built-in vote secrecy attack. If an election official wants to know who a voter voted, mark the ballot form in such a way that it will be rejected by the optical scan machine, before handing it to the voter. The voter will complete the ballot, and try to scan it. Since it will be rejected, the voter will reveal his vote preferences to the closest election official. As this attack only works locally in one polling station, we consider this a local attack against vote secrecy. It won't scale easily to other/all polling stations.

## Conclusion

In this paper, we present a framework for analyzing vote secrecy for various voting methods. We then apply the framework in the context of three elections that we have observed, including paper voting in Denmark, Internet voting in Denmark, and the use of optical scan machines in the US. Our findings are that commonly the tension between election integrity and vote secrecy is often resolved by strengthening election integrity at the expense of vote secrecy.

## Bibliography

[1] S. Bayer, J Groth. Efficient Zero-Knowledge Argument for Correctness of a Shue, pp. 263–280. Springer Berlin Heidelberg, Berlin, Heidelberg (2012).

[2] Joseph A. Calandrino, William Clarkson, and Edward W. Felten, Some Consequences of Paper Fingerprinting for Elections, USENIX (2009).

[3] Danish Election Law. The Parliamentary Electoral System in Denmark. Translation published by Folketinget, (2011).

[4] Expert Study Mission Report. The Carter Center. Internet Voting Pilot: Norway's 2013 Parliamentary Elections, (2013).

[5] J. Alex Halderman, Vanessa Teague. The New South Wales iVote System: Security Failures and Verification Flaws in a Live Online Election, Proceedings of E-Voting and Identity, Springer Verlag, (2015).