

Confidence -- What it is and How to achieve it

NIST Symposium on Building Trust and Confidence in Voting Systems

Jim Adler, Founder, VoteHere, Inc.

Maryland, December 10-11 2003

Introduction

The theme of this symposium is Confidence: We all want it – voters, election officials, candidates, and every single person in this room. I am going to focus my discussion today on how we can guarantee confidence in our elections. We don't need to throw out electronic voting machines (DREs) to get this confidence and we don't need to add paper ballots to DREs to get this confidence. There is a third verification option that guarantees this DRE confidence in the face of the most heinous attacks, and it deserves careful consideration.

Last year at around this time, I was speaking at the *National Meeting on Election Reform* in Florida. At that talk, I highlighted the need for verification, both of voter intent and tabulated result. However, what I did not talk about (and what I want to talk about today) is that verification must support a more fundamental, and traditionally ill-defined, goal of *election confidence*. That's the name of this conference and why we're here today.

I will begin with defining and quantifying election confidence. How do traditional machine-counted paper ballots, current DREs, voter-verified paper ballots (VVPB), and receipt verification systems compare in this regard? The current California practice of automatically hand-recounting 1% of a county's precincts bounds the error margin to more than 50% for many races, including federal Congressional races! This means that 50% of the ballots could be cheated and the 1% recount would not detect it with acceptable assurance. We, as an industry and society, can and must do much better than this!

I'll also discuss how any current electronic voting machine can be easily upgraded to support end-to-end verification from voter intent to tabulated result. A low error margin of 0.5% can be achieved even if as few as 1000 voters or election observers participate in the verification process. This

Confidence -- What it is and How to achieve it

NIST Symposium on Building Trust and Confidence in Election Systems, December 10-11 2003

Jim Adler

GUARANTEES, FOR EVERY ELECTION, that ballots were counted-as-intended at a confidence level exceeding 99%.

Figures of Merit

Let's face it; we don't know how to quantify election confidence. Before Election 2000, many believed that elections were perfect (unless, of course, the jurisdiction had a history of electoral fraud). This idyllic belief was shattered in many respects and we, as an industry and society, have struggled with that reality. Without defining and quantifying confidence, we are in an uncomfortable place where we are tempted to manage *perceptions* rather than scientifically provable *realities*.

This is not new. In the mid 17th Century, for example, the Black Plague struck Edinburgh, Scotland and thousands were dying from the disease. The city council was desperate for a solution and was politically pressured to act. So, at one of the town meetings, with no science to support the decision, the council concluded that cats were responsible, and so ordered that every feline in the town be slaughtered. This was a bad move considering that cats made excellent rat catchers, and rats carried the fleas that carried the plague bacteria. As you've already guessed, by killing the cats, the city council caused the rat population to skyrocket along with the plague. The punch line, of course, is that you'd better have a firm grasp on the science that drives an intended outcome.

I don't mean to compare elections to the Black Death, but without applying clear science, we are being tempted into similarly bad policy, and are starting to fall victim to such temptation at both the state and national levels.

In elections, of course, we have the tough task of simultaneously ensuring election integrity and ballot secrecy – the classic audit/secrecy conundrum. However, this difficulty doesn't let us off the hook in delivering elections that can guarantee *real* confidence. Again, I underscore *real, guaranteed confidence* here because you don't need the National Institute of Science and Technology if you're just willing to manage perceptions.

Real confidence is based on well-understood science. We base our medicine on it; we base our polls on it; we base our state regulated lottery and casino games on it. Las Vegas never has a losing year because of it. We set desired confidence levels, acceptable error margins, quantify the odds, and

Confidence -- What it is and How to achieve it

NIST Symposium on Building Trust and Confidence in Election Systems, December 10-11 2003

Jim Adler

make informed, scientific policy decisions. Why don't we do the same for elections? If we did, we could set election policy that guarantees confidence.

Is there a good measure of confidence that we can all agree on? Is there some strawman figure-of-merit? Something like *every election must guarantee an error margin of no more than 0.5% with 99% confidence*. This guarantee can be quantifiably proven for each election in the face of the most heinous threats from insiders and outsiders, hackers and vendors, proven from voter intent to tabulated result.

If we can agree on such a figure-of-merit and guarantee that every election meets or exceeds it, real confidence will rise. Any contested race will have the proof to resolve the protest within an acceptable error margin, in court if necessary. Before we dive into cures, like killing cats, we should define how we measure success. Maybe 0.5% error margin is too high for something as critical as an election. Maybe 99% confidence is too low. That's not for me to say. But, at least, we need a yardstick to measure how good our elections are.

There is some precedent for this. According to California Elections Code, 1% of the precincts is randomly chosen by county elections officials for an automatic "public manual tally" in order to boost confidence that the machine count can be trusted.¹ We'll look at the effectiveness of this in a moment.

Now let's see how three potential solutions (current DREs, VVPB, and ballot receipt verification) measure up.

How do DRE's Compare?

First, how do DREs compare to our 0.5% error margin. Well, the bad news is that there is no current way to really tell what the error margin is for a DRE since there is no end-to-end Election Day test from voter intent to tabulated result. DREs can fall victim to undetected changes in voter intent,

¹ California Elections Code, Section 15360. During the official canvass of every election in which a voting system is used, the official conducting the election shall conduct a public manual tally of the ballots tabulated by those devices cast in 1 percent of the precincts chosen at random by the elections official. If 1 percent of the precincts should be less than one whole precinct, the tally shall be conducted in one precinct chosen at random by the elections official. In addition to the 1 percent count, the elections official shall, for each race not included in the initial group of precincts, count one additional precinct. The manual tally shall apply only to the race not previously counted. Additional precincts for the manual tally may be selected at the discretion of the elections official.

Confidence -- What it is and How to achieve it

NIST Symposium on Building Trust and Confidence in Election Systems, December 10-11 2003

Jim Adler

whether accidental or malicious. Of course, much is done procedurally to ensure confidence such as pre-election L&A (logic and accuracy) tests; post-election L&A tests; physical security; hardening to external and internal attacks; and security reviews. But, as many of my fellow panelists have pointed out, there are holes in all this apparatus that could corrupt voter intent or election results.

The conclusion, and what I believe fuels the DRE security debate, is that current DRE error margin is *unknown* to a guaranteed level of confidence. Confidence is merely asserted, based on assumptions of “good” system design and election procedures. Since much of today’s election apparatus is currently under scrutiny and doubt, in the absence of real guarantees, confidence undoubtedly suffers.

How do Paper Ballots Compare?

Next, how do machine-counted paper ballots compare to our 0.5% error margin? For this, it’s instructive to look more closely at how effectiveness of California’s 1% automatic precinct hand-recount.

Just to be clear, a precinct recount compares a hand-count of the ballots to the machine count. As part of California Election Code 15360, at least 1% of the precincts are randomly chosen by the local election official and hand-recounted. More precincts are included for down ticket races and at the discretion of the local election official.

For sake of discussion, let’s assume the ballot contains a statewide race, like Governor or US Senator, and a more localized countywide race, like member of US Congress. Assuming a county with 500 precincts (the average California county has about 450 precincts), 5 precincts would be chosen for hand-recount of the Congressional race. For the Senate race, 260 precincts would be chosen across the roughly 26,000 California precincts.

Well, it turns out that the 1% precinct hand-recount narrows the error margin to $\pm 2.3\%$ for our Senate race, which is not bad but clearly no better than the confidence you’d have in a typical poll. To put this in context, 2.3% error margin equates to more than 200,000 votes in a California statewide race where roughly 9 million votes are cast. The jaw dropper, however, is the Congressional race where the error margin is a whopping 60%, or 105,000 votes in a Congressional district where

Confidence -- What it is and How to achieve it

NIST Symposium on Building Trust and Confidence in Election Systems, December 10-11 2003

Jim Adler

175,000 votes are cast. Tripling the number of hand-counted precincts would only lower the error margin to about 30%. And to reach our 0.5% error margin target in a Congressional race, roughly 1000 precincts would require hand-counting in order to achieve 0.5% error margin.

These conclusions are based on a white paper by VoteHere's Chief Scientist, Andy Neff, which details the mathematics behind this analysis.²

How does VVPB Compare?

It has been proposed that adding a voter-verified paper ballot (VVPB) to a DRE will improve confidence. One of the more compelling arguments for the VVPB is that it provides a way to conduct a hand-recount like that used in California. Given the analysis I just went through, the error margin *does* drop from an unknown level with current DREs. But it doesn't improve much at all, dropping to merely 60%. Given that the entire purpose of the VVPB is to ensure election confidence, this analysis shows that the VVPB is really just a fig leaf on election confidence. It helps the perception with false confidence but really just papers over the real problem.

So, if you hand-recount the paper ballots, you get a small increase in confidence. However, we've now created two unintended consequences that may dwarf the original problem. First, we've introduced two, potentially contradictory ballot boxes – one paper, the other electronic. Second, there is a question of when to count the paper ballots. Let's deal with those in turn.

Proponents say that the paper ballot is the ballot of record and would legally supercede any electronic count. This is debatable given the dozens of election fraud cases perpetrated with paper ballots throughout our history (I have about a dozen right here over the last 50 years). But for sake of argument, let's concede this – the paper ballot is the ballot of record. What political fall-out would descend when one candidate wins the electronic count and another candidate wins the paper count? What happens when it is highly suspected, but maybe not proven, that the paper ballots were mistakenly lost or maliciously changed? Do the electronic ballots count then? It's these ambiguous cases that will fuel future election controversies resulting in divisive politics and further erosion of confidence.

² C. Andrew Neff, *Election Confidence: A Comparison of Methodologies and Their Relative Effectiveness at Achieving It*, VoteHere, November 2003.

Confidence -- What it is and How to achieve it

NIST Symposium on Building Trust and Confidence in Election Systems, December 10-11 2003

Jim Adler

The other issue concerns when to count the voter-verified ballots? If the ballots are only counted on a close election, you create a procedural way to cheat by a large amount. It comes under the time-honored axiom of dictators: “if you’re going to tell a lie, make it a big one.” On the other hand, if you only count the ballots on landslides, you create an onerous administrative burden by hand-counting many precincts. Actually, both of these strategies are poor – with our 0.5% error margin figure-of-merit, we now know have to hand-count 1000 precincts anyway.

The VVPB add-on is really just a band-aid for the ills of unknown DRE confidence. This cure adds little real confidence and threatens to institutionalize ambiguity in our election systems.

The solution: Ballot Receipt Verification

The good news is that the same math that mandates 1000 hand-counted precincts to achieve a 0.5% error margin is the same math that allows 1000 verified ballots to achieve the same error margin. By verifying individual ballots, instead of precincts, this high level of confidence can be achieved with relatively few voters participating – like 1000 out of 500,000. This is the punch line so let me say it again: verifying 1000 *ballots* instead of 1000 *precincts* can achieve our error margin of 0.5% with 99% confidence. In a typical county of, say, 500 precincts, that’s two verifications per Election Day; and since any voter can choose to verify, the more verifications you do, the higher the more confidence you get.

So how can we accomplish end-to-end verification without violating ballot secrecy? To explain how this could work, let’s try this “thought” experiment. For the moment, ignore the secret ballot requirement and allow the voter to leave the polling place with a copy of their ballot. Stay with me here. After the election, all the ballots used for counting would be published. The voter would compare their ballot copy with the published ballots and make sure that what he/she intended is what got counted. Any cheating would be detected. Any voter can do it, and if more than 1000 voters did it, the error margin would be below our target of 0.5%.

Now to iron out the wrinkle of the secret ballot; voters, of course, can’t be allowed to leave the polling place with proof of how they voted. This is the classic audit/secrecy conundrum that has confounded elections since the Australian secret ballot was introduced in 1856. During a recent election, I noticed that a DRE county had the following answer in their Election Day Q&A:

Confidence -- What it is and How to achieve it

NIST Symposium on Building Trust and Confidence in Election Systems, December 10-11 2003

Jim Adler

“No system can guarantee to a voter through some verification process that their vote was counted as it was intended. To provide such a system would eliminate the secrecy of the ballot.”

This isn't true anymore. This conundrum has been solved.

To accommodate the secret ballot, the verifying voter is given a receipt with codes that represent their chosen candidates. The voter is shown which code corresponds to which candidate in the privacy of the polling place. The voter takes the coded receipt with them. The codes are unique for each ballot and candidate and so a voter cannot *prove* to anyone how they voted because only the voter knows which code matches which candidate. After tabulation, the codes are generated from the ballots in the ballot box and published so the voter can verify what he/she intended actually got counted.

Granted that this is a bit more complicated for those voters that choose to verify; but only 1000 voters *need* to go through this process in a county of 500,000 voters to achieve our 0.5% error margin. These voters safeguard the election for the rest of us that don't bother to verify. Voters that don't choose to verify are not encumbered at all.

The receipt verification approach provides three real benefits:

- 1) It is essential for smaller races, like Congressional districts, where recounting precincts with VVPB does not provide an acceptable confidence level.
- 2) It is much simpler and cheaper for the same level of confidence than recounting precincts.
- 3) Only a small number of ballots need be verified to guarantee a large confidence in the election as a whole.
- 4) The most heinous, malicious attempts to cheat anywhere in the system, from the voter's intent to tabulated result, would be detected with *guaranteed* confidence.

Conclusion

I think we can all agree that confidence in our elections is critical. That confidence comes from verification. Using the common science I've discussed today, election verification can provide real, guaranteed confidence. Assertions and rhetoric on all sides of this debate only provide false confidence. We've provided one way to objectively measure and prove election confidence.

Confidence -- What it is and How to achieve it

NIST Symposium on Building Trust and Confidence in Election Systems, December 10-11 2003

Jim Adler

The knee jerk to paper ballots is not based on any scientific measure of confidence. Actually, the science tells us that the error margin for paper ballots is many times larger than that of a common poll and can undermine the very confidence we're trying to improve.

The solution exists. The ballot receipt verification methods I've discussed preserve the secret ballot and provide a *guarantee* that elections were not cheated.

Just like the best science, I don't ask you to believe me. The proof is in the math; the proof is in the verified ballots. The proof is in the data.

Thank you for your attention.