PHYS.ORG

Using randomness to protect election integrity

April 10 2017, by Eugene Vorobeychik

Credit: AI-generated image (disclaimer)

Democratic societies depend on trust in elections and their results. Throughout the 2016 presidential election, and since President Trump's inauguration, allegations of Russian involvement in the U.S. presidential campaign have raised concerns about how vulnerable American elections are to hacking or other types of interference.

Various investigations – involving congressional committees, the FBI and the intelligence community – are underway, seeking to understand what happened and how. There are many potential problems with elections: Voters can be individually coerced or bribed into changing their votes; the public can be misled about important facts, causing them to draw inaccurate conclusions that affect their votes; and the physical – and electronic – process of voting can itself be hacked.

Without conducting a full, vote-by-vote manual recount, which is impossible because many voting machines leave no paper trail, how can we be sure an election was conducted fairly and not interfered with?

My research, as a scholar of game theory applied to computer security, has highlighted how combining two approaches can help solve this vital problem. First, my collaborators and I use game theory to think like an attacker – imagining that we want to influence the outcome of an election and determining the best way to do so. Then, we use our expertise in computer security – including an understanding of the value of randomness – to inform our design of an audit process that maximizes our chances of catching someone conducting that kind of attack.

Sampling ballots to recount

An important way to ensure public confidence in electoral results is to audit the machines' vote counts. This is best done by checking the numbers each machine reports at the end of an election against paper records made in real time as voters cast their ballots throughout the day. But even if every machine did keep a paper record – and many don't – doing a simultaneous manual count could cost tens of millions, or even billions, of dollars.

It's much more efficient – and just as mathematically accurate – to conduct a selective audit, examining a small sample of the voting results to identify evidence of tampering. But that leaves open the question of which districts to audit.

Thinking like an attacker

Just as the best place to look for evidence of a crime is at the place the incident happened, the best election districts to audit are also the places that might be the most attractive for an attacker to target.

But how do we identify which ones these are? Could hacking one large district, such as the state of California, have the same overall effect as hacking three or four smaller ones, such as Delaware, Vermont, Wyoming and Idaho? What if it's more difficult to hack a single big target, and easier to hack the smaller ones?

Game theory can help us with this problem. In 1992, the first rigorous, though highly theoretical, study was published in the field that would come to be known as "election control." In essence, the paper's authors investigated how difficult it would be for a malicious party to change an election outcome.

The specific type of difficulty they looked at was not how much money such an effort might cost, nor how many people or how much time would be required. Rather, they looked at the computational burden involved, attempting to calculate which votes would need to be manipulated to change an election's outcome. They identified several factors that might affect how hard it would be to influence an election, such as the nature of the influence (for example, adding candidates or voters to the election) and different types of voting systems, like those used in different countries.

Since that initial work, many researchers have investigated variations on the general theme, such as targeting specific people's votes or influencing groups of voters rather than individuals.

We use that type of approach to think like an attacker, seeking to identify the districts that are most vulnerable to influences in ways that would deliver an attacker's desired outcome.

Considering the likelihood of detection

But it's not enough just to identify those districts that are the best targets for attack. A hacker's goal is to make a difference in the election outcome without being detected. In the U.S. at the moment, election audits don't happen very often, if ever. So the attacker could pick any of the best targets to determine the outcome.

However, our work assumed that election audits happened regularly, and that their existence was public knowledge. So an attacker would have to pick districts that were both vulnerable to attack and where auditors would be unlikely to look.

If both attackers and election officials have equal expertise at evaluating the hackability of election districts and choosing which to audit, their analyses will be the same. They'll identify the most vulnerable districts and suggest the hackers hack there – and tell the auditors to audit those same districts. But that creates a real quandary: a smart attacker will decide to attack somewhere other than where the auditors are looking. And the auditors will realize the attackers will shift targets, and search elsewhere for evidence of outside influence.

Again, game theory can help. It's an ideal method for analyzing situations where every decision on one side influences the other's moves, in an apparently endless loop. Game theoretic analysis shows that while the loop may be repetitive, it is not endless. There is a set of districts that are vulnerable enough to attack to be worth considering for an audit, and a group of districts that are not sufficiently vulnerable to attack to be worth either attacking or auditing. This helps narrow the field for both auditors and attackers.

Nevertheless, it may still not be enough to make an audit plan. In elections involving large numbers of districts, like 50 states, or 435 Congressional districts, attackers may be able to efficiently influence the outcome in dozens of possible ways. And auditors may not have the resources to check them all.

Enter randomness

But if we randomly pick which districts to audit, the added unpredictability makes the attacker's choice considerably more difficult, and less certain of success.

What we end up with is an audit plan that is admittedly challenging to compute, but also imposes high calculation burdens on an attacker seeking to evade detection, and reduces or eliminates the impact of nearly all attacks. That not only makes an attacker less interested in trying to change an election result, but provides high public confidence that anyone who did would be found out.

We evaluated this method by looking at the results of the 2002 French presidential election and the 2016 U.S. presidential election results of the 10 largest districts in Michigan. We found that our method identified the districts in those elections that could have the greatest effects on the overall election outcome.

And we found that our use of randomness improved our ability to select which among those districts we should audit to maximize use of limited auditing resources. By using this combined approach, we were able to design an audit plan that could significantly reduce the likelihood of a successful – and undetected – attack on those elections.

To improve public trust in election results, we need to be sure we're looking in all the right places for evidence of tampering. Randomness can make the auditors' jobs easier – and an attacker's task much harder.

This article was originally published on The Conversation. Read the original article.

Provided by The Conversation

Citation: Using randomness to protect election integrity (2017, April 10) retrieved 26 April 2024 from https://phys.org/news/2017-04-randomness-election.html