

Post-Election Tabulation Audits: A Practical Guide for Beginners and Beyond

October 2024 | Version 1.0



U. S. Election Assistance Commission

633 3rd Street NW, Suite 200 | Washington, DC 20001

www.eac.gov

Table of Contents

Introduction2

What Are Audits?2

Why Is Auditing Election Outcomes Important?.....3

Post-Election Audits vs. Recounts.....5

Checking Election Results Without Recounting All the Ballots6

Organizing and Selecting Ballots for Auditing14

Types of Post-Election Tabulation Audits19

Story of an Audit26

Public Communication About Audits29

Bibliography31



Introduction

Welcome! This is an introductory guide to post-election tabulation audits. It is designed primarily for election officials who want to begin, improve, or better understand post-election audit practices. However, it is not just for election officials. This guide is written in plain language with easy-to-understand examples so that it can be useful to anyone seeking a better understanding of post-election audits. It can help legislators, government officials, candidates, political organizations, journalists, and members of the public learn about current best practices in election audits.

Whether a jurisdiction is looking to implement new audit procedures or is already required to perform certain types of audits, this guide takes a “meet you where you are” approach. It is important to note that work may be constrained by local election laws, available technology, budget, and other factors beyond an election official’s control. Post-election audits can be part of a cycle of continuous improvement, and they do not have to be flawless to be beneficial. All types of audits are valuable, though some types of audits allow stronger claims than others.

There is no “one-size-fits-all” audit. This guide will cover fundamental principles that can be applied to all post-election audits, discuss different types of audits, and describe best practices to strive for over time.

What Are Audits?

In the most general terms, **an audit is an examination of evidence to compare to our expectations.** If the evidence doesn’t match expectations, then the audit may help to identify problems or errors so that they can be corrected.

There are many types of audits, and they are used in a wide variety of industries. A financial audit will examine financial accounts and documents. A security audit will examine security measures and practices. A compliance audit will examine whether required regulations, guidelines, and policies are being followed. Most quality control measures are routine audits to ensure that a product or a service continuously meets expectations.

The field of election administration uses several types of audits. Many election processes can be audited before or after an election, including voter registration lists, procedures, security, and chain of custody. However, this guide is focused on one specific type of audit: **post-election tabulation audits.** These audits provide quality control after all votes are tabulated (counted). The community of election practitioners frequently shortens this name to “post-election audits” or “tabulation audits.” This guide will use these shorter names as well as “audits” to refer to post-election tabulation audits, since they are the sole focus here.



Testing the chemicals in a swimming pool is a type of routine audit.

A post-election tabulation audit reviews votes on cast ballots for one or more election contests and compares them to the expectations set by the tabulated election results.

Post-election audits examine whether election outcomes are supported by the cast ballots. An **election outcome** is the generic term for which candidate(s) or ballot question(s) reportedly won or lost a particular election, regardless of whether it was a landslide or a razor-thin margin. There are a variety of ways to conduct post-election audits. This guide will present the most common choices with some pros and cons to consider.

A typical post-election audit process involves:

1. Determining the audit method to be used
2. Organizing the ballots to be audited to allow for convenient inspection
3. Examining ballots one-by-one or in batches
4. Comparing the results of the audit with the expectations set by the reported results

Some audits have the legal power to escalate procedurally, so that there can be further review and, ultimately, so that any incorrect election outcome can be corrected.

Section Summary

- Audits are routinely used as quality control measures. They examine parts of a process or situation to gather information about the whole process or situation.
- A post-election tabulation audit examines votes on cast ballots and compares them to the expectations set by the tabulated results.
- Audit information that aligns with expectations builds confidence in election outcomes. If it does not align, it can alert auditors to the need for more information or corrective action.

Why Is Auditing Election Outcomes Important?

Post-election audits help ensure election outcomes are correct and build trust in the voting process.



Elections are an essential tool for democracy, but *trust* in elections is what makes democracy work. The legitimacy of an elected official's position in office and the peaceful transfer of power requires public trust that elections are fair and the results are correct.

Trust is strengthened by supporting evidence. If you purchase a house, a car, or anything of much value, you will probably review evidence to confirm you are getting exactly what you intend before you complete the sale. Inspections can reveal problems or issues that need to be addressed, and a clean inspection gives reassurance that major issues weren't overlooked.

Similarly, citizens want *evidence* their votes were counted correctly. Any election can count votes and return decisive outcomes, but **evidence-based elections** can also provide convincing evidence that those outcomes are correct. This affirming evidence helps the public to trust that elections are fair, gives legitimacy to office holders, reassures any candidate who loses an election that it was fairly decided, and helps fight disinformation to the contrary.

Post-election audits also provide quality control and a strong security defense. An audit offers an opportunity to detect and remedy incorrect election outcomes regardless of whether the cause was error, malfunction, or manipulation. If ballots were tabulated by a machine, citizens do not have to implicitly trust the machine or its programming. If ballots were tabulated by people, citizens do not have to implicitly trust those people or the counting process they used. Election administrators work hard to have error-free elections, but errors can still happen from time to time. Audits provide an opportunity to detect errors and, when necessary, remedy them.

Evidence-based elections have three requirements:

1. Collect trustworthy evidence of each voter's intent (ballots)
2. Preserve and protect election evidence (secure chain of custody)
3. Demonstrate that outcomes are consistent with election evidence (post-election audits)

Post-election audits examine ballots that are assumed to reflect voter intent—this means each voter has marked their choices on a durable record, and those records are preserved. When voter intent is captured well and protected, then post-election audits demonstrate that the outcomes are consistent with the will of the people.

Furthermore, helping the public to view and to understand post-election audits increases the transparency of election processes and allows the public to trust the outcome. Evidence-based elections benefit and strengthen a democracy.



Section Summary

- Audits build trust in the election process by demonstrating to the public that the tabulated outcomes are supported by other evidence of voter intent.
- Evidence-based elections collect, protect, and review trustworthy ballots to provide affirmative evidence that elections outcomes are correct.
- Audits rely on trustworthy records and secure chain of custody.

Post-Election Audits vs. Recounts

A post-election audit is not the same as a recount. Most jurisdictions have laws, regulations, and procedures for recounts that are different from those used for post-election audits. Even though they are different, a fundamental understanding of recounting ballots can be helpful for understanding audits.

A full recount re-examines *every ballot* cast in the election. A recount may involve running all ballots through an electronic tabulator again. A recount may involve teams of people who tally the votes by hand. It is common for candidate representatives to play a role in monitoring whether each vote is being attributed fairly and for there to be a resolution process if there is disagreement. A recount produces vote totals which can be compared with the initial vote totals, or which can replace them altogether. Recounts are large projects that often require significant labor, time, and expense.



A team of people recounting votes on ballots by hand.

A post-election audit takes a different approach. It uses *sampling* to examine a portion of the ballots to check if they are consistent with the original vote totals. Different types of audits use different techniques for choosing and examining the sampled ballots. We will discuss sampling and those different techniques in the pages ahead.

A post-election audit is not a partisan process where one or more sides can advocate for interpreting ballots to their preferred candidate's benefit. Auditors are neutral investigators whose objective is to detect possible tabulation errors. Auditors are often election administration staff or volunteers but can include anyone who is authorized and follows set procedures to ensure neutrality and transparency.

Post-election audits often require manual inspection of ballots (the "hand-to-eye" method) but are sometimes performed by electronic tabulators. Manual inspection is done to reduce the chance that an error during initial tabulation will be repeated during the audit. (Even when using different tabulation hardware, the same software and election configuration files are

frequently reused.) Of course, human error during the audit is also possible, and audit procedures are thoughtfully designed to guard against it.

A core benefit of audits is that they can be performed much faster and with less effort than a full recount, because they examine fewer ballots. Most post-election audits in use today will examine 2% of the ballots or fewer, and some types of audits are designed to examine significantly fewer ballots.

This makes post-election audits a powerful tool for routine use in election administration. Recounts can provide exact results with high confidence, but audits can provide evidence that the election outcome is likely correct with significantly less time, labor, and expense. Many contests on a ballot can be audited together with reasonable effort, while a recount of many contests is burdensome and simply may not be possible in the limited time before an election must be certified.

Section Summary

- Recounts and audits are different procedures with differing goals, but both examine evidence from cast ballots to check the outcome.
- Recounts usually review all the ballots in a contest, while audits examine only some of the ballots.
- Recounts validate the exact vote totals, while audits look for problems that are significant enough to have affected the outcome.

Checking Election Results Without Recounting All the Ballots

Post-election audits can give valuable evidence of the trustworthiness of an election when they rely on both good practices for ballot accounting and statistically sound principles for checking the results. The process of checking *only some* of the ballots, while trying to glean information about the outcome represented by *all the ballots*, is known as **sampling**. Audits examine a **sample** of ballots from the full set, rather than counting every ballot.

Sampling is a common quality control process. A swimming pool's chlorine levels are tested with only a small amount of water from the whole pool. Chefs check the taste of their dishes by trying only a small spoonful. Factories can monitor for defects in their products by testing only a few.

A **representative sample** is one whose contents are similar to the overall contents. While we can't guarantee that a sample is representative, we do know some sampling methods are not likely to give us representative samples. For example, if a chef adds salt to one side of a pot and immediately samples from that side, they may incorrectly conclude the dish is too salty because that sample was not representative. This is why chefs stir before tasting! Likewise, if a factory



always tests the first product off the production line, they won't get an accurate picture of what's happening in the whole factory.

Post-election audits can select a *sample* from all the ballots cast in a contest and use it to judge the contents of the whole set. Auditors do not need to examine every ballot. However, auditors should try to select a *representative sample*, as best as they can, to improve their ability to draw conclusions about the whole set of ballots. The good news is that sampling doesn't require checking ballots from every precinct. It only requires that every ballot in the contest has the *opportunity* to be chosen. For example, if auditors choose to sample only ballots from the largest precinct, it doesn't give every ballot in the contest the opportunity to be chosen. This can lead to an inaccurate picture and missed issues just like when a factory only tests products from one part of the production line.



Stirring the pot before sampling a spoonful helps to ensure that a small taste is representative of the whole.

The strength of conclusions drawn from a post-election audit depends on the way samples are taken, so it's worthwhile to explore this concept in more depth.

Sampling Pieces of Candy

Let's suppose that three jars each contain 400 pieces of candy, some purple and some yellow, but we cannot see inside the jars. Our job is to determine which jar contains the most purple candies by taking one small scoop from the top without looking in the jar. We scoop from each jar and get three samples of candy.



Jar A: 10 purple candies

Jar B: 8 purple candies and 2 yellow candies

Jar C: 5 purple candies and 5 yellow candies

We reasonably conclude that Jar A must have the most purple candies. Then the jar contents are revealed.



Surprise! Each jar contains 200 purple and 200 yellow. What happened? Not every jar was mixed well, so a scoop off the top wasn't a representative sample. The yellow candies in the bottom of Jar A had no chance of being scooped up! Likewise, the purple clusters at the top of Jar B made purple more likely to be chosen. Only Jar C, mixed thoroughly, gave a sample whose colors closely represented the entire jar.

It's important to keep in mind that even good sampling techniques can yield a variety of samples. Even from Jar C, choosing only 10 candies could have easily scooped 4 purple and 6 yellow. But it would have been surprising to get all purple in a scoop from Jar C.

Larger samples should yield results closer to the overall contents of the jar. Scooping 100 or 200 candies from Jar C should show almost equal proportions of each. Even in Jar A, which wasn't mixed, larger scoops might have contained at least some yellow candies.

Let's consider another example with two jars of candy. We've been told each jar contains 320 purple candies and 80 yellow candies (80% purple, 20% yellow). The jars are well-mixed, because we know that's important.

We scoop 20 candies from the first jar. We get 2 purple and 18 yellow. Do we believe that the first jar contains 80% purple candy? No. It's not impossible to get this sample, but it seems very strange for a well-mixed jar that's 80% purple.

Next, we scoop 20 candies from the second jar. We get 15 purple and 5 yellow candies. This sample isn't 80% purple either, but it's not so far off that it makes us suspicious. This could be a jar with 80% purple candy – we have no reason to believe it isn't.



Both how well the candies are mixed and the size of the sample are important to good sampling practices. Samples aren't miniature replicas of the entire contents we're examining, but they should let us know when our assumptions are off.

Section Summary

- Sampling some of the contents is useful for drawing conclusions about all the contents.
- A representative sample has contents that are similar to the overall contents.
- The strength of the conclusions drawn from a sample depends on how the sample is taken.
- It is important that all items have an opportunity to be part of the sample.
- Randomly selecting items for a sample gives every item an opportunity to be chosen and allows us to draw reliable conclusions.

Sampling Ballots

Let's look at how this concept of sampling can apply to election audits with some simple examples using made-up election data.

In the table below, suppose each row represents one precinct. Each precinct has exactly 10 ballots, so there are 100 ballots total. We'll have a Purple Party and a Yellow Party, and ballots cast for each party candidate are represented by P and Y.

Example 1											Purple	Yellow
Precinct A	P	P	P	P	P	P	P	P	P	P	10	0
Precinct B	Y	P	P	Y	Y	P	Y	P	Y	Y	4	6
Precinct C	Y	P	P	Y	Y	P	P	Y	Y	Y	4	6
Precinct D	P	P	P	P	P	P	P	P	P	P	10	0
Precinct E	P	P	P	P	Y	P	P	P	P	P	9	1
Precinct F	Y	P	P	Y	Y	Y	P	Y	P	Y	4	6
Precinct G	Y	Y	Y	Y	Y	P	P	P	Y	P	4	6
Precinct H	Y	P	P	P	P	P	P	P	P	P	9	1
Precinct I	Y	P	Y	Y	P	Y	P	Y	Y	Y	3	7
Precinct J	P	P	P	Y	P	P	P	P	P	P	9	1
											66	34

In this election, Purple is the clear winner: 66 votes to 34 votes. Suppose we want to check 10 ballots to help verify that Purple won this election. If we take the first ballot from each precinct, shown in the first column, we'll see 4 Purple ballots and 6 Yellow ballots. This doesn't accurately represent the overall election results or even show the correct winner!

What if instead we check one precinct (row) at random? Notice that the totals columns show that Yellow won 5 out of 10 precincts, so choosing a precinct at random will show Yellow winning 50% of the time, even though Purple won the overall election.

Why don't the samples we selected look like the overall results? Because we did not choose *representative samples*. We saw this occur in the first group of candy examples when we didn't give every piece of candy an opportunity to be chosen by mixing up the candies first. While we can't physically mix up the ballots, we can make sure we are sampling from all the ballots cast in the contest and that we aren't just choosing ballots from the top or from only one precinct when the precincts are so different.

To sample effectively in this example, we need to *randomly* select 10 ballots from the entire pool of 100 ballots cast in this election. When we do, we'll be more likely to see sample results closer to the actual election results. It is of course still possible to draw 10 ballots that don't show Purple as the winner. In fact, it's possible to draw 10 that are all Yellow (there are 34 of them after all), but that's unlikely if we're sampling randomly. Most samples of 10 ballots chosen randomly should contain a majority of Purple ballots, even if the percentage of Purple ballots is not exactly 66%.

Let's try another example where the election results are very close.



A birthday cake may be 10% frosting and 90% cake, but if you've ever watched the birthday child select a piece, they have a knack for cutting a corner piece that is 50% frosting. This doesn't mean the whole cake is half frosting, only that they haven't chosen a representative piece.

Example 2											Purple	Yellow
Precinct A	P	Y	Y	Y	P	Y	P	P	P	P	6	4
Precinct B	P	P	P	Y	Y	Y	P	P	P	Y	6	4
Precinct C	P	Y	Y	Y	Y	Y	P	Y	P	Y	3	7
Precinct D	Y	Y	Y	Y	Y	Y	Y	Y	Y	P	1	9
Precinct E	P	Y	P	Y	P	Y	P	Y	P	Y	5	5
Precinct F	P	P	P	P	P	P	P	Y	Y	Y	7	3
Precinct G	P	P	Y	Y	Y	P	Y	Y	P	P	5	5
Precinct H	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	0	10
Precinct I	P	P	P	P	P	P	P	P	P	P	10	0
Precinct J	Y	Y	P	P	P	Y	Y	P	P	P	6	4
											49	51

In this election, Yellow barely beat Purple: 51 votes to 49 votes. Suppose we sample 10 ballots for our audit.

We know now we need to draw 10 ballots at random. Because this election is so close, even when our sample mirrors the election results, it can't show 51% Yellow and 49% Purple with only 10 ballots. Instead, our samples will be close to 50/50. We should still see at least 5 Yellow ballots in most random samples, but the chance of getting a sample that shows Purple winning is higher, as is the chance of getting a sample that shows a tie. That's not because we're sampling badly. It's because the election is so close.

In elections this close, looking at results from small random samples is not as helpful for getting a clear picture of the overall results. We need to look at more than 10 ballots in this situation. Some audits are designed to look at ballots in rounds to add more ballots to the sample when needed to draw reliable conclusions.

What does it mean when our sample disagrees with the reported outcome?

We explored ways that inadequate sampling can lead to misleading results. Is it possible for a correctly drawn sample to show a different result than the reported outcomes even when those outcomes are correct? Absolutely—we saw in the candy example that a correctly drawn sample can still differ from the overall contents. However, a sample that is very different should raise our suspicions.

When choosing ballots to audit randomly, it's again *technically* possible that our sample could include any of the ballots, but samples that show a different winner should be very, very unlikely, especially with large samples. So unlikely, in fact, that we might doubt the reported winner is correct, just like we doubted the first candy jar in the second example had 80% purple candy.

In an the election example where the Purple Party won 70/30, a sample of 10 ballots with only Yellow Party ballots would be very strange. It would even be strange for a sample of 10 ballots to have more than 5 Yellow Party ballots in such a wide-margin election.

The goal of sampling ballots isn't to use the sample to tell us what the actual margin is. Samples help us detect issues that could have affected the accuracy of the originally reported winner. Let's return to the example where the Purple Party won with 66 out of 100 ballots. Audits are designed to check the report that Purple won by asking, "Is it possible that Yellow actually won or tied this election?" A random sample with 6 Purple ballots and 4 Yellow ballots means the answer is "probably not" because it would be strange to see that sample if Yellow had won. Thus, we agree to keep the reported results that Purple won because the sample did not show evidence to the contrary.

The audits described in this section are only one way to use the audited ballots to check for potential issues. Later, we'll examine other audit methods. Regardless of the audit method, random selection helps give confidence that the audit didn't miss any issues that could have affected the reported outcome.



Product reviews are a sample of customer experiences, but not necessarily a representative sample. When a product gets positive or negative reviews, it matters whether those reviews come from 5 people or 5,000 people.

Section Summary

- An audit can meaningfully check election results without examining all the ballots in the election.
- Choosing ballots to audit based on convenience or with procedures that are not random can lead to an audit that does not give adequate information about the reported outcome.
- Choosing a random sample from all ballots cast in a contest ensures the information is useful in either supporting the reported outcome or identifying cause for concern.
- Every ballot in that contest should be available for selection.
- If a sample does not support the reported outcome, it could be because of incorrect sampling methods or choosing too small of a sample. If a large, randomly selected sample does not support the outcome, it indicates there may be a problem and further investigation is needed.

Organizing and Selecting Ballots for Auditing

Ballot Organization

The first step in any audit is to organize the ballots to be audited. An audit needs to account for all the ballots in the election prior to sampling. Every ballot must be available to be chosen as part of the random sample, whether auditing precincts, batches, or individual ballots.

Election administrators typically organize the cast ballots by their source and want to keep them that way. A box or bundle may be kept for each precinct, each day of early voting, or for ballots received by mail, drop box, or other means. Audits do not need to change the organization of ballots, though ballot storage processes can be designed to facilitate auditing. This might include organizing and labeling the ballot containers with the precinct, tabulator, and number of ballots.

It is easy to see how precincts or batches can be selected at random when ballots are well organized in this fashion. A good ballot accounting system can facilitate randomly selecting individual ballots as well. Auditors just need a list of all ballots in the election regardless of whether they are stored in boxes, envelopes, or other containers. A **ballot manifest** is a list that documents how many ballots a voting system scanned and where those ballots are located.

As an example, suppose we have a contest that is in five precincts, and every ballot cast in those precincts contains that contest. Let's suppose our five precincts have 268 ballots, organized into several ballot containers as shown in the table below.



	Ballot Container	Ballots in Container
Precinct 1	1	40
Precinct 1	2	53
Precinct 1	3	11
Precinct 2	1	10
Precinct 3	1	45
Precinct 3	2	33
Precinct 4	1	34
Precinct 5	1	42
		268

Using this ballot accounting information, we can assign a number to each ballot, in order, and add a column to the table showing which numbered ballots are in which container. Now we have a ballot manifest. The ballots can all stay in their containers, and don't even need to be in one location.

	Ballot Container	Ballots in Container	Ballot Numbers
Precinct 1	1	40	1-40
Precinct 1	2	53	41-93
Precinct 1	3	11	94-104
Precinct 2	1	10	105-114
Precinct 3	1	45	115-159
Precinct 3	2	33	160-192
Precinct 4	1	34	193-226
Precinct 5	1	42	227-268
		268	1-268

Ballot Selection

We have seen that random selection is important to effective auditing. Previous examples showed the issues with choosing ballots from just one precinct and with choosing the first few ballots from each precinct. Random selection of precincts, batches, or ballots gives a better picture of the overall election. If precincts and batches are well organized and numbered, then auditors can draw random numbers to choose the associated precincts or batches to audit.

When considering auditing individual ballots at random, we might imagine all the ballots being put in a large jar, shaken, and then drawn out like pieces of candy. Of course, we can't do that, but a ballot manifest and well-organized ballot storage makes it possible to select even individual ballots at random. The key is that the ballot manifest allows us to select *numbers* randomly (and then find the corresponding ballots) instead of having to shake up an imaginary giant jar of ballots.

Generating truly random numbers can be done with a physical process that gives every number an equal opportunity to be chosen. In this example, we could put the numbers 1-268 on identical slips of paper inside a container, shake them up, and draw numbers out individually. Computer software can generate random numbers, but the generated list of numbers is only truly random if the software is given a random starting point, called a **seed**. In many audits, an election official or other participants will roll a 10-sided die several times to get a series of digits that can be entered into software as the seed. This ensures that the list of numbers created is not predictable by knowing how the computer’s hardware and software works since no one can know what the dice will roll ahead of time.



Let’s suppose the software gives us the following list of random numbers from 1-268. What does that mean for our audit?

Random #
57
234
66
260
135
22
53
255
132
153

These are the numbers that correspond to the ballots to be audited. But how do we find ballot 57? We can find it by referencing the ballot manifest. The ballot manifest showed that ballot 57 is in Precinct 1, Container 2. This container starts with ballot 41, so we count on from 41 to get to ballot 57, selecting the 17th ballot in that container. In practice, software can do these calculations and give the precise ballot locations as shown in the table below.

Random #	Corresponding Ballot
57	Precinct 1, Container 2, Ballot #17
234	Precinct 5, Container 1, Ballot #8
66	Precinct 1, Container 2, Ballot #26
260	Precinct 5, Container 1, Ballot #34
135	Precinct 3, Container 2, Ballot #21
22	Precinct 1, Container 1, Ballot #22
53	Precinct 1, Container 2, Ballot #13
255	Precinct 5, Container 1, Ballot #29
132	Precinct 3, Container 1, Ballot #18
153	Precinct 3, Container 1, Ballot # 39

Notice that Precinct 1 has more ballots selected than other precincts and precincts 2 and 4 had no ballots selected at all. It’s important to remember that Precinct 1 had more ballots than other precincts overall, so in a random sample it’s not unusual for more to come from that precinct. It’s also not required that a ballot come from every precinct, only that every ballot is eligible to be selected. In a small sample, some precincts may not be included. If we intervene in the sampling process and insist that one ballot from every precinct be chosen, slightly more sophisticated statistical principles are needed. One way to adjust the sampling method without using more sophisticated methods is to increase the size of the sample, which increases the likelihood of sampling ballots from more precincts.

Notice that the ballots can remain in their containers, and even in their precincts, while the ballot manifest is being generated. The ballots can even be audited in different locations provided that proper procedures for observation and recording are followed. In practice, for efficiency and ballot control, election officials may choose to bring all ballots from an audited contest to one location before choosing the sample of ballots to audit. For a statewide contest, however, election officials may audit the indicated ballots in their own jurisdictions and send results to a central location.

Section Summary

- A ballot manifest lists all ballots in the contest being audited and records their physical location. This ensures every ballot in the contest has an opportunity to be part of the audit.
- All ballots can remain in their containers, and even in several locations, during the preparations for an audit. The ballot containers only need to be opened to examine the ballots selected for the audit.
- Software can generate a random list of numbers that correspond to ballots to audit, but a physical random process (like a dice roll) is needed to give the software a seed. This way no one can predict the audited ballots based on the software's programming.

Types of Post-Election Tabulation Audits

The last two sections described how a random sample of ballots from an election contest can give information about the results of the overall contest using basic statistical principles. Researchers in statistics, political science, and other fields have developed methods for conducting post-election audits that use this and similar sampling methods. Their work continues to evolve with research. (See Bibliography for a selection of early and recent academic papers on the subject.) Today, most post-election audits share the common goal of verifying election outcomes by looking at a portion of the cast ballots, but the procedures and implementations vary.¹

Post-election audits fall into two broad categories: traditional and risk-limiting. Traditional audits usually attempt to answer the question, “Did the vote tabulators function correctly?” Risk-limiting audits attempt to answer the question, “Is the reported outcome correct?” These questions are certainly related since we can only trust the reported outcome if there were not tabulation errors significant enough to affect the outcome. However, there are differences in how each type of audit is conducted and in the information the audit provides.

This section will examine several types of post-election tabulation audits and the benefits and drawbacks of each. Election officials may be constrained by time, resources, laws, and regulations that affect how they conduct post-election audits. Understanding the different types of post-election tabulation audits can help election officials make an informed choice based on their situation and constraints.

¹ U.S. Election Assistance Commission, “Election Audits across the United States,” October 6, 2021, 5-8

https://www.eac.gov/sites/default/files/bestpractices/Election_Audits_Across_the_United_States.pdf.



Traditional Audits

Traditional audits are typically designed to answer the question, “Did the vote tabulators function correctly?” They examine “a pre-determined number of ballots, voting precincts, or devices and compare reported results from voting systems to the paper ballot records”.² The audited ballots may be chosen randomly or according to a set procedure.

Common types of traditional audits:

- Fixed number or percentage of precincts
- Fixed number or percentage of vote tabulators
- Fixed number or percentage of ballots

Audit regulations vary by state but most traditional audits sample between 1–5% of precincts. Sample sizes may be determined by state law or at the discretion of an election official, and frequently there are nuances to the size or selection in a particular state. A traditional audit may have human auditors or machines count votes from any number of contests on the ballots. The goal is to determine whether there are discrepancies between the original count and the audit count and to decide what next steps to take, if any.

Benefits

Traditional audits have a predictable workload. This allows election officials to more reliably plan and allocate resources to the audit, even when auditing several contests on the ballot. Comparing audit results with the original tabulation of the same set of ballots can detect tabulation issues and could detect widespread issues if batches or precincts are chosen randomly. When a traditional audit chooses ballots randomly from all ballots in the contest and the contest margin is not too close, it could also give useful information about the trustworthiness of the reported outcome.

Drawbacks

Traditional audits lack the flexibility to adapt for wide- or close-margin contests, because the number of audited ballots is set ahead of time. Some discrepancies may not affect the outcome in a wide-margin contest, but in a close-margin contest even a few discrepancies could matter. The predetermined number of audited ballots could also be larger than needed for confidence in the outcome causing election officials to use more resources than necessary. Traditional audits that don’t select ballots randomly or examine enough ballots may not give an accurate overall picture of the contest and could miss widespread issues.

² *Election Audits Across the United States*, page 10.



Things to Keep in Mind

- Traditional audits can give useful information about machines sampled but may not give reliable information about the overall contest outcome.
- No matter how many ballots or machines are sampled, using random selection is important to have reliable information about machine performance or outcomes.
- Traditional audits of close-margin contests often need more ballots than the requirements specify to check the outcome. Wide-margin contests could often be checked with fewer.
- Results of individual precincts often cannot be generalized to overall contest results even though they may provide valuable information.

Risk-Limiting Audits

Risk-limiting audits, or RLAs, build on traditional audit methods by focusing on the question, “Is the reported outcome correct?” rather than on the exact count of any set of ballots. RLAs check for tabulation errors that are significant enough to have affected the outcome. Thus, the number of ballots audited depends on the margin of the contest. Risk-limiting audits require that all ballots from the audited contest be available for auditing, be selected randomly, and be examined by human auditors.

Risk-limiting audits proceed in rounds. They start by checking a few ballots, and based on those results, determine whether more ballots need to be audited. We saw in the examples with jars of candy that small samples can sometimes show strange results just by virtue of being small. Large random samples should be more representative of the overall situation. Thus, RLAs start small and grow as needed, until the audit has checked enough ballots that it is unlikely it missed any outcome-affecting issues.

An RLA will continue to audit ballots, round after round, until there is enough evidence to indicate that counting all remaining ballots would very likely show the same election outcome. It is not unlike an election administrator offering to keep recounting ballots until all the candidates are convinced that further counting is not worthwhile.

Of course, anytime we stop short of counting *all* the ballots, there is a risk that we stopped too early, and a full recount would show a different outcome. We want to keep that chance small. This is the basis for the **risk limit** of an RLA. A 5% risk limit means that when the audit stops, there is *at most* a 5% chance that a full recount would have arrived at a different outcome.

In practice, an RLA doesn’t usually *feel like* a full recount that ends early because it may end with only a small fraction of the ballots examined. This is because the statistical principles at work are designed for both efficiency and accuracy. Unlike traditional audits, RLAs take into



account the margin of the reported results and require more evidence for close-margin contests than for wide-margin contests. If audit results do not provide adequate evidence to support the outcomes, the audit will continue to a full hand recount of the ballots to determine the correct outcome. A sample alone cannot determine the outcome; it only indicates when further investigation is needed.

We will look at three types of RLAs: batch-level comparison, ballot polling, and ballot-level comparison.

Batch-Level Comparison Risk-Limiting Audits

Batch-level comparison RLAs randomly select *batches* of ballots and compare the audited results to the original tabulated results for the same batches. They are like traditional audits that look for discrepancies between the audited results and the tabulation results. A batch-level comparison RLA is trying to determine if discrepancies are significant enough that the reported outcome could be incorrect. Some discrepancies, such as additional votes for the winner or offsetting differences between several batches, might not change the reported outcome, even though they could indicate issues with the original tabulators. Batch-level RLAs additionally require that batches of ballots be randomly selected from all batches of ballots with the audited contest. Because batches will likely vary in size, RLAs can also take into account the size of each batch when determining how many batches to audit.

Benefits

Batch-level comparison RLAs share many of the same practices of traditional tabulation audits, so jurisdictions using traditional audits may find them to be an easy path to RLAs. It is usually faster for a jurisdiction to retrieve whole batches of ballots rather than a series of individual ballots. When each audited batch is associated with a particular machine, election officials can check the performance of those machines while also collecting evidence to validate the overall outcome. Discrepancies in the audit and machine totals can be explained in an audit report.

Drawbacks

Batch-level comparison RLAs require jurisdictions to maintain strict ballot accounting procedures, so that each machine's tabulation totals correspond exactly to a physical batch of ballots. The workload cannot be predicted before the election, but auditors reviewing batches can expect to review more total ballots than with RLAs that select individual ballots. In practice, some batches may be very large, particularly when a scanner processes several days' worth of absentee or early-voting ballots and tabulates them as one batch. These very large batches can also be more likely to be selected for audit, increasing the likely workload.



Things to Keep in Mind

- Traditional audits and batch-level comparison RLAs both review batches, but they ask different questions. RLAs ask, “Is the reported outcome correct?” and not, “Did the vote tabulators function correctly?”
- Comparing batches requires auditors to have tabulated result totals that correspond exactly to each batch of ballots.
- Large batches of ballots, like those from absentee or early voting, can be more likely to be selected and create a large workload.
- When batch sizes are manageable, comparison audit workloads can be smaller than polling risk-limiting audits.

Ballot Polling Risk-Limiting Audits

A ballot polling RLA is an audit that counts votes from a random sample of *individual* ballots to determine if those vote totals are reasonable given the reported outcome. Like all RLAs, it examines ballots until there is enough evidence to conclude a full recount would show the same outcome.

Ballot polling RLAs use sampling like we saw in the “How to Select Ballots to Audit” section with the Purple vs. Yellow contests. For example, if Purple really won the overall election in Example 1, then we would expect a reasonably-sized random sample to contain a majority of Purple ballots, even if the sample margin (the ratio of purple to yellow ballots) is not exactly the same as the reported margin. On the other hand, we would not expect a reasonably-sized random sample to contain a majority of Yellow ballots, though small samples could. If larger samples continued to have a majority of Yellow ballots, it could indicate that the reported outcome was incorrect.

Benefits

Ballot polling RLAs can give confidence in election outcomes without a full recount, and in many cases with a relatively small number of ballots. The EAC’s *Election Audits Across the United States* explains that “[T]hese audits require minimal set-up costs, can be conducted independent of voting system data, and offer an efficient way to audit contests with 10% or greater margins.”³ The workload for the audit is dependent on the margin of the election, so election officials can focus resources on close contests and use fewer resources for wider margin contests. Ballot polling RLAs may be easier to explain to voters than other risk-limiting audit methods.

³ *Election Audits Across the United States*, page 15.



Drawbacks

Ballot polling RLAs are impractical in very close margin elections because even reasonably sized samples can show varying results. When many individual ballots need to be audited, the workload required to select and examine them is often more burdensome than a hand recount of all ballots. The total number of audited ballots is determined in part by what the previously audited ballots reveal, therefore election officials cannot completely predict the resources needed for an audit ahead of time.

Things to Keep in Mind

- All ballots for the audited contest must be available to be selected in the audit.
- Ballots must be selected randomly for the risk calculations to be accurate.
- The audited ballots do not need to show the same margin as the overall election, but they should show the same outcome.
- This audit method can be impractical for very close margins, and the workload may be larger than desirable for other small margins.

Ballot-Level Comparison Risk-Limiting Audits

Ballot-level comparison RLAs randomly select *individual ballots* and compare each of those ballots to the original tabulation records. Ballots are selected exactly as they are in a ballot-polling RLA, but auditors make a direct comparison to the recorded interpretation of each ballot, also known as a **cast vote record** (see box at right). Jurisdictions can only implement ballot-level comparison RLAs if they have the ability to match an individual ballot with the data about its interpretation during the original tabulation. This is typically enabled by printing identification numbers on each scanned ballot that correspond to the cast-vote records generated by a machine tabulator, but it is also possible to keep the ballots in the exact order they were scanned without printed identifiers. Like batch-level comparison RLAs, a ballot-level comparison RLA is trying to determine if discrepancies are widespread enough that the reported outcome could be incorrect.

A cast vote record is an “[a]rchival, tabulatable record of a set of contest selections produced by a single voter as interpreted by the voting system.”⁴

Benefits

Ballot-level comparison RLAs typically examine significantly fewer ballots than polling audits or batch-comparison audits to achieve the same risk limit, which reduces the workload for retrieving and reviewing ballots. Like batch-comparison RLAs, ballot-level comparison RLAs incorporate the additional information gained by directly comparing ballots to the original interpretation. They also give more details about tabulation issues because discrepancies are tied to exact ballots and tabulators. Any discrepancies between the audit and machine



interpretations of a ballot can be explained in an audit report. Ballot-level comparison RLAs are well-suited to jurisdictions that centrally tabulate ballots and have the ability to add identifiers to ballots.

Drawbacks

Ballot-level comparison RLAs require the machines tabulating results to be able (and allowed) to produce cast-vote records for individual ballots. There can be privacy concerns if the identifiers on the ballots could reveal the identity of any voter. While the total number of ballots is usually expected to be small, workload can't be completely predicted ahead of the audit.

Things to Keep in Mind

- Ballot-level comparison RLAs can only be conducted when the voting system can produce cast vote records for individual ballots.
- Individual ballot comparisons must consider privacy concerns and regulations that protect voter privacy.
- Ballot-level comparison audits can be a good choice for jurisdictions with central tabulation.
- Ballot-level comparison RLA workloads are often significantly smaller than polling RLAs but can't be completely predicted ahead of time.

Summary of Types of Post-Election Audits

- Traditional and risk-limiting audits can both give useful information and help verify election outcomes.
- Traditional audits usually attempt to answer the question, "Did the vote tabulators function correctly?" Risk-limiting audits attempt to answer the question, "Is the reported outcome correct?"
- Traditional audits predetermine the quantity of ballots to audit. Risk-limiting audits are responsive to the margin of a contest and the information given by the ballots audited.
- Traditional audits may not examine enough ballots to draw a useful conclusion about the election outcomes, or they may examine more ballots than are needed to reach a strong conclusion. Risk-limiting audits allow audits to escalate through additional rounds when needed but have a less predictable workload as a result.
- For any type of audit, selecting ballots at random from all the ballots in the contest greatly improves the ability to draw reliable conclusions from the audit.



Story of an Audit

This guide has covered the essential mechanics of post-election audits in the previous sections. In this section, we'll see how an audit unfolds, step by step, in the fictional jurisdiction of **Electionville**, where an election for **mayor** with two candidates has just concluded. This example is designed to illustrate overall goals and concepts; certain practices and processes may not apply in all situations.

Traditional Audits

This is not the first post-election audit in Electionville. In fact, Electionville has been performing traditional audits for several decades. Each of Electionville's 50 polling places has an electronic tabulator to tally ballots on Election Day. The main election office also has one electronic tabulator that is used to tally ballots returned by mail.

Electionville recognizes the importance of having chain-of-custody procedures to protect the paper ballots that will be relied on as trustworthy evidence in any post-election audit. Poll workers move ballots from a tabulator's ballot box into one or more canvas bags, secure the bags with numbered, plastic, tamper-evident seals, and log the numbers for each seal used. A set of designated poll workers deliver the canvas bags of ballots to the main election office, along with a removable memory card from each tabulator and other election paperwork. The election staff document the receipt of each canvas bag and place them into sealed containers in a secure storage area.

To maintain the chain-of-custody, the election staff organizes the ballots for each polling place. Even in a high-turnout election, no polling place has more ballots than will fit in a single box. They label each box with the polling place information and the total number of ballots inside. They also keep a list of the boxes and the number of ballots inside each one.

After each election, the election staff randomly select batches of ballots to audit by drawing numbers representing each of the 50 polling places and the batch of mail-in ballots out of a coffee can.

The ballots selected are retrieved from storage and kept in batches as they are examined by teams of workers. Each team tallies the votes on a batch of ballots by hand-counting the ballots according to their hand-count procedures. The audited batch totals are compared to the original batch totals. Occasionally, they have found small discrepancies in the totals, but none have been significant enough to cause concern about the outcome. After each successful audit, Electionville reports to the public that a routine audit of the ballots cast showed that the electronic tabulators functioned as expected and no major issues were found.

Risk-Limiting Audits

This year, Electionville is going to pilot a risk-limiting audit in place of the traditional audit they know well. The pilot will benefit from many of these existing audit procedures.



Before the Risk-Limiting Audit

The story starts when the polls close on Election Day. All eligible ballots have been cast. The election staff is hard at work. When the polls close, the ballots in each polling place have already been counted and the totals compiled. The staff's first priority is to receive the paper ballots and to import results information being returned from those polling places into the election management system, just like they have always done.

Electionville also has ballots that were returned by mail and have been processed and scanned on a central count scanner over the past several days. There are also a small number of provisional ballots to review, and a few overseas ballots allowed by the Uniformed and Overseas Citizens Absentee Voting Act (UOCAVA) may still arrive. All these ballots are kept secure and organized as they are processed.

Once every eligible ballot has been counted and all subtotals have been aggregated together, Electionville announces preliminary results in the contest for mayor. These results are obviously important to the candidates, but they are important to our audit story because these *election outcomes* set the expectations for the post-election audit. There is now a reported winner, and the risk-limiting audit will check whether the evidence supports that conclusion.

Electionville has decided that this audit will be a **ballot polling risk-limiting audit**. Electionville's precinct tabulators are not able to support ballot-comparison RLAs, because they don't print identifiers on ballots or keep the ballots stored in the exact order they were tabulated. They decided against a batch-comparison RLA, because the mail-in ballots are in one extra-large batch, which would likely lead to a difficult workload. The audit will examine the **contest for mayor** with a **risk-limit of 10%**. That means that, when the audit concludes, the chance that it overlooked an outcome-affecting issue will be less than 10%.

The audit will take place in a spacious warehouse next to the election office and the public and media have been invited to attend. The ballots will be organized in boxes by precinct just like in past years and will be kept secure until they are brought to the warehouse the day of the audit. This year the election staff needs to create a ballot manifest detailing the inventory of ballots before the audit. It is helpful that the boxes were well-labeled when they were filled, so the staff does not need to count the number of ballots inside each one. They take extra care to be sure that *all* eligible ballots—ballots from polling places, mail-in ballots, UOCAVA, and provisional ballots—are included in the list so that the audit will sample from the full set of ballots in the Mayoral contest. With a complete ballot manifest, they are ready for the audit to begin.

The Day of the Risk-Limiting Audit

The Director of Elections begins the audit with an overview of the audit process. They explain that just like in years past, human auditors will examine evidence by interpreting the cast ballots according to their hand-count rules. However, this year there will be two key differences. They



are not committing to audit a fixed 2% of the ballots. Instead, they'll look at ballots until enough evidence has been gathered to feel confident that no more ballots need to be examined to accept the reported outcome. They'll also select individual ballots from all cast ballots instead of randomly selecting precincts to recount. A projector and screen have been set up to display audit activity so that everyone can follow the progress of the audit.

The Director explains that Electionville will use audit software to manage the automated and computational parts of the audit. The election results, the ballot manifest, and the risk limit have been entered into the software. The Director explains that the audit depends on choosing a random sample of ballots for review and reminds everyone that these will be drawn individually from across all ballots cast in the election. Ten members of the audience take turns rolling a 10-sided die to generate a 10-digit random number (or seed) for the software to use. The news media takes video and photos as each roll is announced, recorded, and projected on the screen. The numbers rolled are: 3, 1, 1, 2, 8, 3, 4, 7, 0, 5.

The software now has everything it needs to produce a list of the randomly selected ballots for auditing. The workers have been organized into seven teams, so the list of ballots is broken up into seven pages. Each team is given their first-round assignment for ballot selections. The teams take their lists over to their ballot review tables, where the teams will work together to retrieve, examine, record, and return ballots that correspond to the numbers in their list.

At table 4, the first ballot to be retrieved is the 30th ballot from Polling Place Q. A set of workers goes into the ballot storage area, locates the box for Polling Place Q, opens it, and finds the 30th ballot. The workers insert a piece of paper as a placeholder at that position while the ballot is out for auditing. The workers close the box and take the ballot back to the review table. The vote for Mayor on the ballot is examined and recorded on paper by the audit team. Then, the workers return the ballot to the original box at its original position. The voters' choices from the audited ballots are also entered into the audit software as the audit progresses.

This same process is repeated by the workers at all seven tables until the entire first round of ballots has been examined, recorded, and entered into the audit software. The software runs calculations based on the number of votes for each candidate found in the audited ballots and the reported margin. The software reports that the risk limit has not been met yet. This does not indicate that the outcome is in doubt, just that the audit needs to make the sample larger to draw useful conclusions.

Round two begins. Seven new lists of ballots are generated, and the teams repeat the process. The workers are now familiar with the audit procedures and the second round is noticeably faster than the first.

At the end of round two, the software runs its calculations. It finds that the ballots sampled during the audit have satisfied the 10% risk limit. The audit stops because there is less than a 10% chance it missed an outcome-affecting issue. The Director of Elections explains that even if



the audit continued examining ballots, it's unlikely that it would find a different outcome than the reported outcome and so the public can feel confident stopping the audit here. The Director then announces that the audit has concluded, and the reported results can move on to the certification process. Everyone leaves the audit except for the election staff, who secure the ballot boxes and move them into long-term storage.

After the Risk-Limiting Audit

In the coming days, the election result will be certified, and the election officials will publish an audit report. An audit report will describe what happened during the audit, explain why election officials are confident that the outcome is correct, and explain any issues that were encountered. The audit report will also explain that the audit was intended to verify the outcome, so it is normal and expected for the margin of victory of the audited ballots to be different than the overall reported margin. In a few months, the new Mayor of Electionville will take office. No matter who they voted for, the citizens of Electionville will have reason to trust that the election outcome was decided correctly, based on the evidence from the audit.

Public Communication About Audits

Officials need to communicate audit processes and results to the media, policy makers, and voters as part of building public trust in elections. The public is looking for evidence that the election outcomes are trustworthy. Audits provide that evidence. Talking about post-election audits clearly and accurately is an important step in building public trust.

It can be tricky to clearly communicate what the audit results mean. Election officials cannot expect the public to be familiar with the mechanics of audits. However, many parts of a post-election audit can be performed in publicly observable ways, helping to build understanding. In general, the more election officials talk about post-election audits as a routine part of the election process, the more comfortable the public can become with how elections are verified. In addition to inviting the public to observe an audit, officials can publish press releases, information about audit processes and timelines, and a report explaining the audit results.

Below are some suggestions for effectively communicating audit processes and results to the public.

- State the type of audit and explain the process, including how ballots are chosen and safeguarded throughout the audit.
- Explain ahead of time that some discrepancies are normal and don't necessarily cause concern. Auditors may interpret a ballot differently than a machine tabulator, for example, if marks are outside the lines. Randomly selected ballots should confirm that the reported winner is correct but are not expected to match the exact margin.



- Describe each step and what the election auditors are recording for each audited batch or ballot. If discrepancies are discovered, explain how they will be handled, including processes and policies for further investigation.
- Take care not to overstate what the audit confirms. An audit cannot *guarantee* that an election result is correct, but an audit that does not raise concerns gives trustworthy evidence that the outcome is correct.
- Share the final results of the audit. Include the number of ballots audited and the results from those audited ballots.
- Written audit reports can include mathematical calculations like a risk-limit, measured-risk, or p-value which can be helpful to some experts. Be careful about making more generalized statements about what these numbers indicate.



Bibliography

Fuller, Benjamin, Abigail Harrison, and Alexander Russell. “Adaptive Risk-Limiting Comparison Audits.” In *2023 IEEE Symposium on Security and Privacy (SP)*, 3314–31. San Francisco, CA, 2023. <https://ieeexplore.ieee.org/document/10179424>.

Johnson, Kenneth C. “Election Certification by Statistical Audit of Voter-Verified Paper Ballots.” *ssrn.com*, October 31, 2004. <https://ssrn.com/abstract=640943>.

Lindeman, Mark, and Philip B. Stark. “A Gentle Introduction to Risk-Limiting Audits.” *IEEE Security & Privacy* Volume 10, Issue 5 (2012): 42–49. <https://doi.org/10.1109/MSP.2012.56>.

Lindeman, Mark, Philip B. Stark, and Vincent S. Yates. “BRAVO: Ballot-Polling Risk-Limiting Audits to Verify Outcomes.” In *2012 Electronic Voting Technology Workshop/Workshop on Trustworthy Elections (EVT/WOTE '12)*, 2012. reprint: <https://www.usenix.org/system/files/conference/ewtwote12/ewtwote12-final27.pdf>.

National Institute of Standards and Technology. “NIST Election Glossary Draft Version 1.0.” *pages.nist.gov*. Accessed May 19, 2024. <https://pages.nist.gov/ElectionGlossary/#cast-vote-record>.

Ottoboni, Kellie, Philip B. Stark, Mark A. Lindeman, and Neal McBurnett. “Risk-Limiting Audits by Stratified Union-Intersection Tests of Elections (SUITE).” In *Krimmer, R., et al. Electronic Voting. E-Vote-ID 2018. Lecture Notes in Computer Science*, 11143:174–88. Springer, Cham, 2018. https://doi.org/10.1007/978-3-030-00419-4_12.

Stark, Philip B. “ALPHA: Audit That Learns from Previously Hand-Audited Ballots.” *The Annals of Applied Statistics* 17, No. 1 (March 2023). <https://doi.org/10.1214/22-aos1646>.

———. “Conservative Statistical Post-Election Audits.” *The Annals of Applied Statistics* 2, No. 2 (June 2008): 550–81. <http://www.jstor.org/stable/30244217>.

———. “Overstatement-Net-Equivalent Risk-Limiting Audit: ONEAudit.” In *Essex, A., et al. Financial Cryptography and Data Security. FC 2023 International Workshops. Lecture*



Notes in Computer Science, 13953:63–78. Springer, Cham, 2023.

https://doi.org/10.1007/978-3-031-48806-1_5.

———. “Sets of Half-Average Nulls Generate Risk-Limiting Audits: SHANGRLA.” In *Bernhard, M., et al. Financial Cryptography and Data Security. FC 2020. Lecture Notes in Computer Science*, 12063:319–36. Springer, Cham, 2020. https://doi.org/10.1007/978-3-030-54455-3_23.

U.S. Election Assistance Commission. “Election Audits Across the United States,” October 6, 2021.

https://www.eac.gov/sites/default/files/bestpractices/Election_Audits_Across_the_United_States.pdf.

Zagórski, Filip, Grant McClearn, Sarah Morin, Neal McBurnett, and Poorvi L. Vora. “Minerva- an Efficient Risk-Limiting Ballot Polling Audit.” In *USENIX Security Symposium*, 2021.

<https://www.usenix.org/conference/usenixsecurity21/presentation/zagorski>.

