



www.ElectionDefenseAlliance.org

UBS
Universal Ballot Sampling
To Validate Computerized Vote Counts
In Federal and Statewide Elections

Jonathan D. Simon, JD, and Bruce O'Dell
Election Defense Alliance¹

Introduction

There is a deep and spreading distrust of the computer-based systems of vote tabulation in the United States, and a rapidly growing concern that outcome-determinative electronic vote fraud can easily go unchecked and undetected. While there are proposals being put forth to eliminate or drastically limit the role of proprietary computerized counting systems in elections, these have little prospect of immediate adoption in the vast majority of venues, and it is clear that the upcoming 2006 election (“E2006”) will be highly computerized and that “faith-based voting” on computerized systems will be even more widespread in E2006, and again in E2008, than in any previous election.

The need for non-computerized check mechanisms on the vote counts in E2006 and E2008, in the absence of full hand counting, is common ground for virtually all who are working for honest elections, but the advantages and drawbacks of the various potential check mechanisms remain a source of considerable confusion. Having previously established the fundamental shortcomings of legislation currently pending before Congress,² our purpose in this paper is to introduce and explain the workings of a check mechanism that is both manifestly effective and eminently feasible: **Universal Ballot Sampling (UBS)**.

By drawing and hand-counting a uniform and homogeneous random sample from all of the actual ballots cast in a given federal race, whether for the United States Senate or House of Representatives, we can evaluate the validity of the computerized vote count with great accuracy and confidence. And we can do so with a vastly smaller investment of labor and resources than would be required for the hand counting of every ballot.

¹ See www.ElectionDefenseAlliance.org for further information on Election Defense Alliance and for biographies of the authors.

² See, O'Dell et al, “The HR 550 Election Audit Methodology in U.S. Congressional Elections: Fundamental Shortcomings and Proposed Solutions,” <http://electiondefensealliance.org/files/HR550auditaccuracypaper.pdf> .

Putting The Power Of Statistics To Work

The benefit of statistical sampling lies in the surprisingly strong power of a small part to predict the behavior of a large whole. Although we tend to accept the results of polls and other research based on sampling, it is nevertheless something of a head-scratcher that one can predict with considerable precision the preferences of a nation of 300 million individuals—whether in candidates, policies, or favorite kind of cheese—by questioning a mere 3000, or .001%, of them. This is, however, the case, providing that certain conditions apply and certain procedures are followed.

The laws of statistics tell us that by sampling a small percentage of the ballots cast in a relatively large target venue such as a state or congressional district, we will be able to determine with great confidence whether the otherwise “faith-based” computerized vote counts are accurate or whether significant, potentially outcome-altering “mistabulation” has indeed occurred. As will be detailed below, if the actual ballots are sampled and counted, and if this is done *randomly* and *proportionally throughout the venue*, then the statistical laws governing the accuracy of the sample count will be simple and mathematically straightforward, what we might call “crisp.” If bias or certain bows to convenience enter into the sampling process, our statistical process loses its “crispness,” the cut-and-dried rule of simple equations, and all bets may well be off. Thus, when only certain precincts are selected for auditing, or when exit polls are used as a check mechanism, factors can intrude which complicate the statistical picture, and the laws governing accuracy become “soggy.” Factors such as the “cluster effect,” which comes from sampling in convenient bunches (e.g., selecting only some precincts), or “differential response,” which comes from depending on voluntary voter interviews rather than actual ballots, are difficult if not impossible to quantify with the precision required for a consistently effective and reliable check mechanism. To use the power of statistics most effectively as a check mechanism, it is imperative that we keep the statistics “crisp.”

Assuming access to the actual ballots cast (including absentee and early-voting ballots, as well as those cast at the precinct), there are only two significant factors that will determine how accurate a handcount sample will be relative to the count of the whole: 1) number of ballots sampled, and 2) the randomness of the sample.³ The Universal Ballot Sampling protocol is free of such bugaboos as cluster effect or differential response, and satisfies the “randomness” criteria extremely well. The number of ballots sampled will thus be straightforwardly determined by how accurate we will require our sampling to be. Accuracy, in turn, boils down to two familiar statistical measurements: margin of error and level of confidence.

Margin of error (MOE) of a sample refers to the range in which we expect to find the difference between the count of the sample and the count of the whole from which the sample is drawn. And confidence level (CL) tells us how strong our expectation of finding the difference within that range should be. In most research, which by convention uses a CL of 95%, a +/- x%

³ One factor which has negligible impact is the size of the whole. Despite the mind’s natural and intuitive protest that larger wholes require larger samples, once we reach a whole of a certain size (e.g., the size we are dealing with in federal elections), larger elections in fact do not require appreciably larger samples.

MOE means that 95% of the time, or 19 out of 20 times, we expect the count of the sample to fall within x% of the count of the whole. This means that in such a sample, at a CL of 95%, 5% of the time a random sample with an x% MOE will count up *more* than x% off from the whole. That's just the way it is. You don't get 100% certainty. But what *is* certain is that if you ran that sample a million times, the number of times it missed the whole by more than x% would approach 50,000 (5%) very closely. Computer simulations, which we have employed in developing our protocol, are most helpful because they enable one to actually do this and check the results.

The Sampling Protocol

For votecount checking, we propose a MOE of +/- 1% at a CL of 99%. That is we recommend sampling a number of ballots which will yield results within 1% of an *accurately* tabulated official count 99 times out of 100. Thus, if the results of the sample are “off” by more than 1%, we would conclude with 99% certainty that the official count was inaccurate. Or, to look at it the other way, there is only a one in 100 chance that the official count is *accurate* and we're looking at a “false alarm.”

Once we specify the MOE and CL, the statistical law governing random samples tells us with great precision how many ballots we need to sample for a given race.⁴ For a +/- 1% MOE at a CL of 99%, in a Congressional District race with 200,000 votes cast, the magic number is exactly 15,317 ballots.⁵ Because this number varies so little with change in the total number of votes cast, we find that for competitive U.S. House races, where the total number of votes cast varies between about 150,000 and 250,000, we need to sample between 14,936 and 15,556 ballots to achieve a +/- 1% MOE at 99% confidence. Thus it may be seen that if we sample 10% of the total ballots cast in each House race (that is, between about 15,000 and 25,000 ballots), we will achieve 99% (or greater) confidence that our sample results are within 1% of an accurate vote count in all competitive races.

It should be noted that while this analysis of our sample ballots does not appear to tell us *how* inaccurate a “mistabulated” vote count might be, we can easily use our UBS sample results to determine, again with excellent precision, the likelihood that an inaccuracy of any specified size

⁴ There are several good websites that will do the work for you; e.g., www.raosoft.com/samplesize.html. Should you wish to perform your own rough calculations, the simplified but still quite serviceable magic formula boils down to: MOE at 95% CL = 1/square root of the number of ballots sampled (generally referred to as "N"). So, to plug in a few numbers: if you sample 10,000 ballots then $1/\sqrt{N} = 1/100 = 1\%$, and you'd say that your MOE is +/-1%, and you'd expect the sample results to differ from the total tabulated results by less than 1% in 19 out of 20 such elections. If you sampled only 400 ballots, then $1/\sqrt{N} = 1/20 = 5\%$, and you'd say the MOE is +/-5%, and you'd expect the sample results to differ from the total tabulated results by less than 5% in 19 out of 20 such elections. If you sampled 30,000 ballots then $1/\sqrt{N} = 1/173.2 = 0.58\%$, and you'd say your MOE is +/- 0.58%, and you'd expect the sample results to differ from the total tabulated results by less than 0.58% in 19 out of 20 elections.

While each of these examples presumed a CL of 95%, MOE can be easily calculated for *any* CL. To find the MOE at a 99% Confidence Level, for example, just take the MOE numbers above and multiply by 1.29: 10,000 ballots would give you a MOE of +/-1.29%; 400 ballots would give you a MOE of +/-6.45%; 30,000 ballots would give you a MOE of +/-0.75%; all at 99% confidence.

⁵ As pointed out above, this number varies very little with the total number of votes cast. In a Congressional District race with 200,000 votes cast, the number is 15,317; in a large statewide race with 10,000,000 votes cast, the number increases only to 16,560.

has occurred in the vote count. Indeed, once the result of the handcount sample is known, the likelihood of any particular machine count occurring in the absence of error can be calculated with precision. This in turn tells us the likelihood that a given machine count, *or indeed the victory of the purported “winner,”* is legitimate.

For the larger Senate races, the percentage of the total ballots required to be sampled in order to meet the 1% MOE and 99% CL criteria falls off significantly. In a statewide race with 5,000,000 votes cast, for example, the number of sampled ballots needed is 16,533, or only 0.33% of the total. We would need to sample only every 300th ballot in that race, or *three* ballots in a precinct with 900 votes cast, to meet our basic criteria. Needless to say, were we to sample as needed for House races (that is, at a rate of 10%) and simply count votes for both House and Senate (or any other statewide race) on the ballots selected, in such a state with 5,000,000 votes cast our statewide handcount sample would yield a 0.3% MOE at a CL of 99.999%—an extraordinarily precise and powerful verification indeed.

The efficacy of the sampling criteria outlined above has been confirmed by running thousands of simulated elections employing such criteria. Such simulations—made possible, ironically enough, by computer power—are a very powerful tool for examining how proposed protocols will work in the field, because they permit us to replicate both the sampling and vote counting processes under the various conditions that will obtain in the field: variable precinct sizes, variations in partisanship patterns, time-of-day voting preference variations, varying sampling methodologies, etc.

The results of our first run of “elections”⁶ indeed confirm that sampling proportionally at the 10% rate recommended above yields sample counts within 1% of *accurate* vote counts in virtually every congressional district election (better than 99%), however the electoral dynamics are varied, and thus would detect significant *mistabulation* in virtually every case in which it occurred, while validating accurate counts.

The Mechanics Of Drawing The Sample

With this powerful tool at our disposal, we turn our attention to the nuts-and-bolts of the UBS sampling process itself. The mathematical formulas that tell us what results to expect are, as initially pointed out, contingent on the randomness of the sample. If the sample were to depend on the willingness of the responder to participate or were taken in convenient “clusters,” its accuracy would diminish accordingly, often in ways difficult to quantify.⁷ Since we are dealing with actual ballots, differential nonresponse is not a factor, but we need to sample in a way that is truly random and thus ensures that each ballot has an equal chance of being selected, whatever the patterns in which the ballots are cast, and that minimizes opportunity and incentive for any

⁶ See Appendix A for a summary. Additional election simulation results are posted on the Election Defense Alliance website (www.ElectionDefenseAlliance.org) and are also included in O’Dell et al “The HR 550 Election Audit Methodology in U.S. Congressional Elections: Fundamental Shortcomings and Proposed Solutions,” <http://electiondefensealliance.org/files/HR550auditaccuracypaper.pdf> . See Appendix B for the algorithm employed.

⁷ See, for one attempt: Merkle, D. and Edelman, M. “A Review of the 1996 Voter News Service Exit Polls from a Total Survey Error Perspective,” in *Election Polls, the News Media and Democracy*, ed. P.J. Lavrakas, M.W. Traugott, New York: Chatham House, pp. 68 – 72.

intentional bias or other partisan impact in the sampling and counting process. Fortunately these goals are readily achievable.

If all the ballots cast in a given race were gathered in one place and each were assigned a consecutive number from 1 to T, a random sample meeting our criteria could be drawn by allowing a random number generator to select the number of ballots needed (N) by picking N times from the numbers between 1 and T. In reality, however, ballots will be found at the precincts where cast, as well as in additional “piles” in the case of early or absentee voting. While this may initially be perceived as an inconvenience when it comes to drawing a sample, we will see below that it does in fact make possible a level of security and integrity that would not be present were all the ballots aggregated in one place.

In order for ballots to be sampled in the “piles” in which they accrue on Election Day, a protocol must be established that will clearly instruct officials working at the precinct level how to choose the ballots for the sample. We present here a generic protocol for sampling the recommended 10% of the ballots. There would of course be leeway for any jurisdiction adopting our approach to tailor the sampling parameters to its own evaluation of security needs and resources—that is, more or fewer sampled ballots depending on degree of precision desired and investment of labor considered acceptable. We have, however, several sound reasons for proposing, first, that a fixed percentage (rather than a fixed number) of ballots be sampled and, second, that that percentage be 10%.

As we have demonstrated above, a 10% sample ensures +/- 1% MOE at 99% confidence (or better) in any competitive House or Senate race. The most obvious advantage of preordaining a 10% sample is that the task at the precinct level becomes the simple one of randomly choosing one out of every ten ballots for hand counting. Were a jurisdiction to mandate instead that a fixed *number* (e.g., 15,000) of ballots be sampled, it would remain first to be determined how many total ballots were cast in the jurisdiction, then what percentage of that total 15,000 would be, and finally it would remain for the precinct level workers to apply that percentage to their ballot pile. This approach is not only more subject to errors in communication and execution, but more time-consuming as well. With 10% sampling, the ballots are ready for sampling as soon as the polls close and the ballots are gathered from their secure storage compartments, thereby greatly minimizing chain-of-custody concerns and speeding the certification process.

Although it is not essential that all ballots be of the same format, it is imperative that every ballot in a precinct have an equal chance of being selected. Hence we wish, wherever possible, to avoid small “sub-piles,” and recommend that all ballots at a given precinct be aggregated before sampling. Optical scan ballots—and absentee, disabled-voter, and early-voting ballots which are produced in the op-scan format—would greatly expedite and simplify the sampling process. It is possible that DRE (touchscreen) paper trails could be rendered suitable for aggregation and sampling, but if the vote records were limited to a thermal roll, this would be quite problematic.⁸

⁸ One major problem with this so-called DRE paper trail is that, where the voter does not mark his or her own ballot, the machine’s paper record may not reflect the voter’s intent, and it is left up to individual voter verification efforts to detect this, hence introducing the potentially very distorting factor of voter verification rate, as well as dramatically slowing the voting process. In noting the superiority of op-scans over DREs when it comes to sampling, we may further point out an obvious cost and/or usage advantage: a single op-scan device may service an entire precinct

Obviously, where DREs without *any* paper record are deployed, no effective venue-wide handcount sample can be carried out. In such cases, individual machine function tests would constitute the only safeguard. But spot-testing of machines is neither rigorous nor vigorous in practice. Even when functional tests are performed there are inherent technical limitations which make any point-in-time testing scheme powerless to safeguard electronic voting transactions⁹. If DRE equipment is to be used at all it must minimally produce a voter-verified paper audit trail; but we strongly recommend *replacement* of DRE vote-tallying systems by voter-marked paper ballots.

Once all the ballots associated with a given precinct are aggregated, it remains only to assign each one a number (this can be done with an inexpensive consecutive stamper or by hand) and draw the sample. To do this, a simple random number generator can be programmed with the total number of ballots in the pile and the portion to be chosen (10%).¹⁰ It will then output a random sample of ballot numbers to be chosen, amounting to 10% of the total.¹¹ The indicated ballots are then selected, the federal races hand counted and tallied. The tallies are publicly posted and submitted to the central counting location for the county or jurisdiction, where they are again publicly posted and combined into a jurisdiction-wide tally.

This UBS tally, which represents a random sampling of the jurisdiction's ballots, will be the reference and check for the official vote count. If the total vote percentage for the winner of a race differs from the handcount sample percentage for that candidate by greater than 1%, significant mistabulation is presumed because the chances are at least 99% that it has occurred. Further scrutiny—either in the form of a second sampling or, if so mandated in the jurisdiction, by a full hand recount—is triggered.

without causing any back-up at the polls; the same can not be said for a DRE. With the op-scan system, voters read and mark their ballots (the time-consuming part) away from the machine in a booth or cubby, and many may do so simultaneously, taking as much time as they please without holding up the line; they then bring their ballots over to the op-scan device, where a ballot is fed in and read in a matter of a few seconds. With the DRE, the machine is occupied while the voter reads, considers, chooses, and verifies. Slow voters at the DREs cause major bottlenecks and the long lines that embarrassed many DRE jurisdictions in 2004.

⁹ See <http://electiondefensealliance.org/auditability> on inherent limitations of point-in-time testing of voting equipment

¹⁰ Alternately, an interval method could be used, in which the percentage of total ballots to be sampled determines the interval and only one random number is chosen. Thus, if we were sampling 10% of the ballots, we would in effect be choosing every 10th ballot. The interval (k) would be 10 and the random number chosen between 1 and 10. If it were, for example, four (4), this protocol would select ballots 4, 14, 34, 54, 74, 94 . . . until the ballots were exhausted. This selection process could be done at the precinct or the central level. Although this procedure is used in some polling applications, it may introduce by its very regularity a slight distortion of randomness, given that the vote patterns themselves have some tendency to cluster when numbered consecutively as cast. The chief advantage is that no random number generator is necessary—ten slips of paper in a hat will do fine.

¹¹ It will of course be noted that 10% of any total number of ballots *not* divisible by ten will leave a remainder—that is, to choose an integral number of ballots will require either slight *over-* or *undersampling*. As long as either “over” or “under” is chosen and followed consistently at all precincts, no significant distortion will be introduced. We recommend the slight oversample, since it ensures that at least 10% of ballots are selected and the threshold MOE and CL criteria exceeded. So, for example, if a precinct aggregated 536 ballots, 54 would be sampled; if 73 ballots, eight would be sampled. Alternately, traditional rounding can be employed and 10% of the rounded number of ballots can be sampled. In the above examples 536 would now be rounded up to 540 and 54 ballots would be sampled; but 73 would be rounded down to 70 and seven ballots sampled. Either method works without aggregate distortion, as long as it is applied consistently.

Universal Ballot Sampling is Effectively “Ungameable”

It would certainly be reasonable to inquire what protections ensure the fairness and honesty of the UBS handcount sample itself. Can it, as well as the computerized tally, be gamed by a determined rigger? The answer is that the handcount sample is intrinsically extremely difficult, if not impossible, to game *in a way that would falsely validate a rigged election rather than triggering even more penetrating levels of scrutiny.*

This is the case because, while the requirement for rigging an election is simply to add *enough* votes to your candidate to produce a victory, the requirement for gaming the UBS handcount sample is to add rather precisely the *correct* number of sampled votes to match the rigged overall vote tally. That is, to rig an election one has a great deal of numerical leeway in adding, subtracting, and switching votes—as long as it adds up to a win—and it’s all happening out of public view. To rig the UBS handcount, on the other hand, one has very little leeway—having to “hit” the rigged “official” vote count within 1% or else trigger additional and deeper scrutiny—and it’s happening very much in public view. Decentralized, precinct-level, sampling and tallying make the gamer’s task all the more difficult.

Let’s consider an example. The backers of congressional Candidate “A” have gleaned from pre-election polling that “A” is trailing by single digits. They have obtained access to the computerized central tabulators through one of many security vulnerabilities that have been identified. They decide to give “A” a net gain of 20,000 votes, calculating that that will be sufficient to reverse the legitimate result and hand “A” the victory. What do they know at this point, going into the UBS handcount sample? They know that if the sample is done honestly it will show a massive discrepancy and trigger a full handcount or some other form of deeper scrutiny. So they must find a way to shift ballots in the sample. But how? And how many?

If they know, or can approximate, the total turnout, they should have an idea of the percentage by which they have altered the actual vote, and therefore of the percentage by which they would need to alter the UBS handcount to match. But—even assuming the unlikely case that their “idea” was good enough to “hit the number” and match the rigged vote total within 1%—how would they go about it? They would either have to convey instructions to “their” people working at every precinct to try to slip a small number of extra “A” ballots into the sample, or they would have to try to add a large number of “A” ballots to the sample at a few locations or perhaps centrally. If the first approach, that’s a whole battalion of people required to be in on the game (as opposed to a single programmer or hacker), and potential exposure at every precinct in the venue. If the second approach, even a modest observation effort—certainly in “stronghold” precincts and at the central counting location where observers will be camped—will detect the scam.

Successful rigging of the UBS handcount sample would take both the kind of labor-intensive and easily detected “broad-based conspiracy” that is decidedly *not* necessary for rigging computer-tallied elections—and a fair amount of luck, to boot. Attempts to rig the UBS handcount sample will have an extremely high likelihood of backfiring—instead triggering a full handcount or other anti-fraud provisions adopted by the jurisdiction. The strong incentive for all election officials of all parties will be to perform a truly random and accurate handcount sample which, better than 99

times out of 100, will match an honestly and accurately tabulated vote within 1% and send everyone home early.

Conclusion

The UBS protocol is designed to restore voter confidence in the electoral system. It achieves this purpose by marshalling the power of very basic statistics and ensuring that humans participate only in ways in which their partisanship can have little or no impact on the verification decision. Thus, whatever specific parameters are ultimately chosen for the sampling protocol, the verification decision itself is dictated by a fixed, pre-selected standard rather than the discretionary *opinion* of election officials or political appointees, as provided for by such legislative proposals as HR 811, currently pending.¹²

The UBS protocol is also designed to be within the capabilities and resources of election administrators and their citizen helpers in every demographic setting in the United States. It is simple by design, easily mastered, and unvarying from one election to the next. Execution at the precinct level is straightforward, involves no math beyond basic counting, and can be completed on election night within an hour or two of the closing of the polls.

For this very modest investment of citizen effort, the UBS delivers a validation protocol that, in Congressional District elections, will detect outcome-determinative fraud in virtually every case without sounding a “false alarm.” In larger statewide elections—and when electing a President of the United States—its performance is “off the charts:” a truly airtight validation protocol.

American democracy deserves such a guarantee of the integrity of its elections, and American citizens are prepared to support it, demand it, and make it work.

¹²See, O’Dell et al, “The HR 550 Election Audit Methodology in U.S. Congressional Elections: Fundamental Shortcomings and Proposed Solutions,” <http://electiondefensealliance.org/files/HR550auditaccuracypaper.pdf> .

APPENDIX A: Empirical measurements of the accuracy of the UBS audit protocol

Category of UBS audit	UBS Audit %	Average UBS Sample Size	Number of simulated elections	Smallest election in sample	Mean election size	Largest election in sample	Observed MOE < 1%	Observed Confidence Level
Random election size	2.0	16,239	10,000	64	811,950	1,977,835	9,913	99.1300%
US Congressional District	10.0	17,765	10,000	85,783	177,651	296,039	9,973	99.7300%
Large state (about 10 million votes)	0.2	20,003	2,267	7,703,634	10,001,658	12,225,039	2,263	99.8236%
Small state (about 1 million votes)	2.0	20,019	9,999	741,863	1,000,953	1,252,239	9,981	99.8200%
US Congressional District	10.0	22,503	2,997	200,024	225,034	250,000	2,993	99.8665%
Large state (about 10 million votes)	0.5	50,055	2,179	7,680,772	10,011,177	12,174,809	2,179	100.0000%

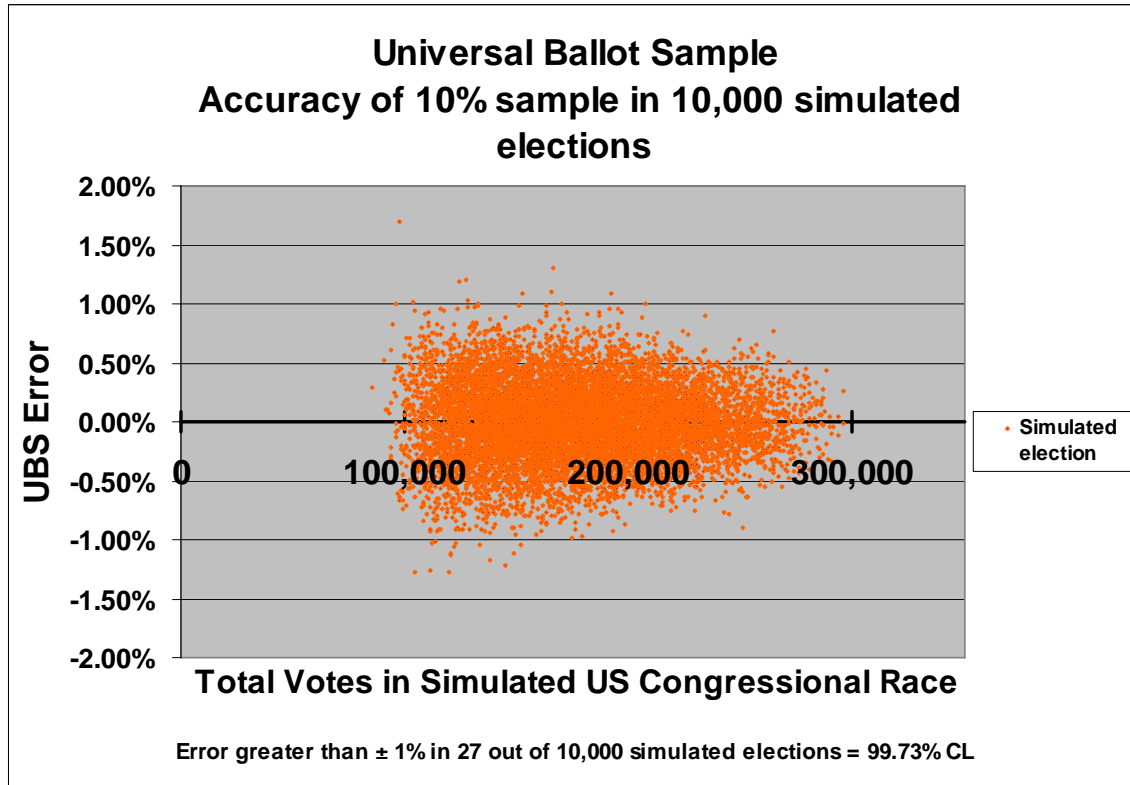
This table summarizes the results of a computer program which measured the accuracy of the UBS protocol over a total of 37,442 simulated elections. The precinct size, number of precincts and precinct partisanship varied between simulated elections. Ballots were sampled in each precinct at random.

Although the sizes of the simulated elections varied widely, the key number is the “Average UBS Sample Size”. All of the average sample sizes were greater than 15,317 (see discussion in the section “Sampling Protocol” above), and the resulting accuracy – in a simulation approximating field conditions – clearly matches theoretical expectations.

The results for the “Large State” category show how even a small UBS sampling rate (0.5%) produces extraordinarily accurate predictions of the simulated statewide vote. Of the 2,179 simulated elections with an average of about 10,000,000 votes cast, *none* were outside the $\pm 1\%$ margin of error.

The accuracy of the UBS protocol in estimating the outcome of US Congressional races is illustrated on the next page.

For the 10,000 simulated “US Congressional District” elections, we see that the highest UBS sampling rate (10%) produces samples easily large enough to ensure 99% confidence of $\pm 1\%$ accuracy. This is illustrated in the following chart:



Appendix B: Description of the Universal Ballot Sampling algorithm

The alternative UBS election verification protocol is simple. Given a paper optical scan ballot or voter-verified paper trail,¹³ select and handcount a random sample of paper vote records from *every precinct*, and then compare those aggregated totals to the official electronic tally.

Basic statistics tells us that if the random sample sizes are adjusted properly, it will be possible to detect even a small amount of deliberate fraud or inadvertent error that causes the official electronic tally to deviate from the paper audit trail, with a very high degree of confidence.

One goal was to measure the size of the handcount sample that would allow us to be 99% certain that the result of handcounting a random sample of the paper trail was within 1% of the actual results. Proving we can achieve this level of accuracy would help ensure that we would not undertake the effort of a full manual recount of all ballots unless there was what amounts to 100-to-1 odds that there was a serious problem with the official electronic count.

The algorithm implementing a random selection of a percentage of the votes in every precinct is straightforward.

Set up a new handcount sampling:

Initialize the working memory that will contain the handcount sampling totals

Conduct the handcount sampling:

For each precinct in the simulated election

Calculate the number of ballots to be sampled in the precinct
(equal to the sampling percentage times the number of cast ballots)
For each ballot to be sampled

Calculate x = a unique random number between 1 and the number of cast ballots in the precinct

Examine the x^{th} ballot in the precinct

If the x^{th} ballot is marked for Candidate A

Increment Candidate A's audit total

If the x^{th} ballot is marked for Candidate B

Increment Candidate B's audit total

Report the handcount sampling results:

When all precincts have been sampled, print out the handcount sampling totals for Candidate A and B

Compare the handcount sampling totals for Candidate A and B with the actual totals for Candidate A and B to assess the accuracy of the handcount sample

If the difference between the handcount sample and the actual, unaltered tally is within 1% of the vote, consider the handcount test a success

If the difference between the handcount sample and the actual, unaltered tally is greater than 1% of the vote, consider the handcount test a failure

¹³ The optical scan ballot is highly preferable to a mere "paper trail" for a number of reasons: it is marked by the voter and does not depend on a separate step of voter verification; it is far easier to number and count than most paper trail formats such as thermal rolls; it is not dependent on printer functions highly subject to breakdown, etc.

Technical notes

The results reported here were generated by an election simulation program written in Microsoft Visual Basic.NET; that language was chosen simply for ease of rapid programming. The algorithms are described in sufficient detail to enable creation of a functionally-equivalent simulation in any alternative language of choice.

Many aspects of the simulation are governed by pseudorandom number generators. No programming language has a perfect random number generator, but the results of this simulation show very small deviations from perfect randomness.