Briefing Paper

Black Swans, Conspiracy Theories, and the Quixotic Search for Fraud: A Look at Hausmann and Rigobón's [1] Analysis of Venezuela's Referendum Vote

By Mark Weisbrot, David Rosnick, and Todd Tucker[2]

September 20, 2004

[1] Hausmann, Ricardo; and Roberto Rigobón. "In Search of the Black Swan: Analysis of the Statistical Evidence of Electoral Fraud in Venezuela". Available at http://ksghome.harvard.edu/~rhausma/new/blackswan03.pdf

Executive Summary

On September 3, economists Ricardo Hausmann of Harvard University's Kennedy School of Government, and Roberto Rigobón of the M.I.T. Sloan School of Management, presented econometric results that the authors maintain are evidence of fraud in Venezuela's August 15 recall referendum. The paper was reported by four major international news outlets and was used to raise doubts about the validity of the referendum among U.S. legislators and policy-makers. It was also used to support claims of fraud by opposition leaders in Venezuela.

In this paper we examine the results presented by Hausmann and Rigobón (available at http://ksghome.harvard.edu/~rhausma/new/blackswan03.pdf) and find that they provide no evidence of fraud. This concurs with the findings of the Carter Center (September 17), showing that the sample selected on August 18 for an audit of the vote that they observed, was indeed a random sample of all voting centers, and that electronic fraud of the type suggested by Hausmann and Rigobón was therefore impossible.

In this referendum voters expressed their preference (YES or NO) with a touch screen voting machine. The machine then printed out a paper ballot with the voter's choice, which voters deposited in a ballot box. The audit of 150 voting centers, observed and certified by the Carter Center and the OAS, found that the paper ballots matched the electronic votes within a 0.1 percent margin.

However, Hausmann and Rigobón put forth a theory of electronic fraud that was consistent with a clean audit. According to their illustrative example, suppose the machines were rigged at 3,000 polling centers, and the remaining 1,580 were randomly selected to be left clean. If the computer program that generated the sample could be fixed to sample only from the clean centers, the electronic votes would match the paper ballots in the audit -- in spite of the fraud.

The authors then present two sets of evidence which they claim indicates that fraud of this type took place.

The main problem with their analysis is that, according to their assumptions, the audited sample of 150 voting sites should reflect the true -- that is, non-fraudulent -- referendum result. Such a large sample provides incontrovertible evidence of the validity of the official results, which were well within the range that would be expected given the results found in the audited sample.

By contrast, the exit poll used by Hausmann and Rigobón, published by the American polling firm Penn, Schoen, Berland & Associates found that 59 percent of voters were in favor of the recall (YES), and 41 percent opposed (NO). This was the opposite of the official results certified by the Carter Center and the Organization of American States, in which voters rejected the recall by a margin of 59 percent (NO) to 41 percent (YES).

But the audited sample had only 41.6 percent YES votes. This paper finds that:

- The chances of getting an audited sample, under Hausmann and Rigobón's assumptions of how it was selected, of 41.6 percent YES, if the true (non-fraudulent) vote were 59 percent YES, are less than one in 28 billion trillion.
- Even if the true vote had the recall barely succeeding with only 50.1 percent YES, the chances of getting an audited sample of 41.6 percent YES are less than one in a million.

This paper also examines the other statistical evidence presented by Hausmann and Rigobón to support their theory of electronic fraud and finds that it is dependent on implausible assumptions. We conclude that the results that they interpret as evidence of fraud most likely stem from a misspecification in their econometric model.

This issue extends beyond Venezuela, where opposition leaders -- including those that control most of the media -- have continued to question the results of the referendum. Most importantly, it has considerable implications for the effectiveness of international monitoring in elections. This was one of the most carefully monitored elections in modern history, with both the Carter Center and the Organization of American States playing a major role.

If this level of monitoring and verification by some of the most experienced election observers in the world, in a case where the election was not even close, cannot produce a credible result, then the whole system of international monitoring would have to be called into question. Fortunately this is not the case, as the statistical evidence of electronic fraud in the referendum has turned out not to be valid.

Introduction

On August 15, 2004 Venezuelan voters went to the polls in record numbers and voted by a margin of 59 to 41 percent to allow President Hugo Chávez Frías to serve out the remaining two and a half years of his term. International observers from the Carter Center and the Organization of American States certified the result and conducted an audit of the vote, in which a sample of the touch screen voting machine results were compared to the paper ballot receipts that voters deposited in ballot boxes[3].

Despite the overwhelming margin of the vote, the audit, numerous tests and controls, the certification of the international observers, and the lack of any material evidence of fraud, numerous opposition leaders -- including much of the Venezuelan media -- continue to insist that a massive fraud took place. Opposition exit polls, including one in which voters were interviewed by the opposition group Súmate for the U.S. polling firm Penn, Schoen, Berland & Associates[4] claim to have found the opposite result: the Penn & Schoen exit poll, which reported having canvassed 20,000 people and with a margin of error of less than one percent, alleged that Chávez had been recalled by a margin of 59 to 41 percent. The firm stands by its reported exit poll numbers and insists that the election was stolen[5].

On September 3, economists Ricardo Hausmann of Harvard University's Kennedy School of Government, and Roberto Rigobón of the M.I.T. Sloan School of Management, presented econometric results that the authors maintain are evidence of fraud[6]. The paper was reported by four major international news outlets[7] and used to raise doubts about the validity of the referendum among U.S. legislators and policy-makers. In their conclusion, the authors write, "As [philosopher of science] Karl Popper said when observing 1,000 white swans: this does not prove the accuracy of the thesis that all swans are white. Nevertheless, observing a black swan does allow one to reject it. Paraphrasing Popper, our white swan represents no fraud. The results we obtain make up a black swan."[8].

In this paper we evaluate this evidence and consider whether there is any reason to believe that electronic fraud of the type suggested by Hausmann and Rigobón might have occurred in the recent Venezuelan referendum.

The Alleged Fraud

---

[3] "Audit of Chávez Vote Upholds the Results". Los Angeles Times. Aug. 22, 2004.

[4] The firm states that it was hired by a group of opposition Venezuelan businessmen, but has not revealed the name(s) of its client(s).

[5] "U.S. Poll Firm in Hot Water in Venezuela". Associated Press. Aug. 19, 2004.

[6] Hausmann 2004. P. 36-38. .

[7] Webb-Vidal, Andy. "Chávez opponents face poll losses after recall failure". Financial Times. Sept. 7, 2004. Gunson, Phil. "Still calling vote a fraud, Chávez foes plan challenge". Miami Herald. Sept. 10, 2004. Luhnow, David; and Jose de Cordoba. "Academics' Study Backs Fraud Claim In Chávez Election". Wall Street Journal. Sept. 7, 2004. "Debates and Dilemmas: Venezuela's referendum". The Economist. Sept. 18, 2004.

[8] Hausmann 2004, p. 38.

Before turning to Hausmann and Rigobón's econometric evidence, it is necessary to explain the type of fraud that the authors are talking about, and how it would have to have been carried out. It is not the ordinary type of fraud that occurs in many elections, i.e. ineligible voters with fake i.d.'s, voting more than once, or stuffing of ballot boxes. Rather, the fraud considered in Hausmann and Rigobón's paper is electronic fraud and involves rigging the voting machines that were used in the election, so as to change the votes from Sí (to recall the president) to No.

In this referendum voters expressed their preference (sí or no) with a touch screen voting machine. The machine then printed out a paper ballot with the voter's choice, which voters deposited in a ballot box.

Fraud in this system is thus a difficult and risky enterprise: any rigging of the machines could be caught quite easily by comparison with the paper ballots at any polling center. In an audit after the vote, the National Electoral Council, together with international observers, drew a sample of 150 polling centers and compared the machine results to the paper ballots. The results matched almost perfectly -- with 0.1 percent difference[9] -- a number that is statistically insignificant and easily explainable by the likelihood that some voters might have failed to deposit their paper ballots.

It is of course theoretically possible to stuff the sampled ballot boxes to match rigged machines, as would have been necessary in the previously offered "cap" theory of the fraud. This was the basis of an opposition argument that the voting machines had been programmed to put a limit on the number of "Yes" votes. But this was not plausible, as Jennifer McCoy of the Carter Center explained to those alleging this type of fraud:

> The only way the boxes could have been altered would be for the military—historically the custodians of election material in Venezuela—to have reprogrammed 19,200 voting machines to print out new paper receipts with the proper date, time and serial code and in the proper number of Yes and No votes to match the electronic result, and to have reinserted these into the proper ballot boxes. All of this in garrisons spread across 22 states, between Monday and Wednesday, with nobody revealing the fraud. We considered this to be supremely implausible[10].

This argument was also refuted when computer scientists Aviel Rubin and Adam Stubblefield of Johns Hopkins, and Edward W. Felten of Princeton showed that the

---

[9] "Report on an Analysis of the Representativeness of the Second Audit Sample, and the Correlation between Petition Signers and the Yes Vote in the Aug. 15, 2004 Presidential Recall Referendum in Venezuela". The Carter Center. Sept. 17, 2004. http://www.cartercenter.org/documents/nondatabase/report091604.pdf

[10] McCoy, Jennifer. "What really happened in Venezuela?" The Economist. Sept. 4, 2004.

number of machines that reported the same numbers was within the range that could be expected by random occurrence[11].

Hausmann and Rigobón also reject the theory of the caps. But they put forth another theory of electronic fraud that does not require any ballot box stuffing. To take their hypothetical example: suppose there are 3,000 voting centers where the machines are rigged, and the CNE randomly selects the remaining 1,580 centers to be clean. The electoral authorities then fix the program that randomly selects a sample of 150 centers for auditing, so that the sample is selected from only the 1,580 clean centers. The electronic results from these centers would then match the paper ballots, and the audited sample would show the election to be clean.

The biggest technical problem with this type of fraud -- aside from rigging the machines without anyone finding out[12] -- is to make sure that the sample is chosen from the "clean" polling centers. To do this, the CNE would have to have fooled the Carter Center, OAS, and other international observers into thinking it was randomly selecting a sample to be audited from the total universe of centers, while secretly substituting a program that selected only from the "clean" voting centers[13]. It is worth noting that the sample was chosen in front of a live television audience, as well as the international observers from the Carter Center, the Organization of American States, and another group of European observers[14].

In response to a request by the opposition group Súmate to consider the theory and evidence offered by Hausmann and Rigobón, the Carter Center examined the program that was used to generate the sample. The Center's report (issued Friday, September 17) states:

> The CNE requested a group of university professors to develop a sample generation program for the 2nd audit. The program is written in Pascal for the Delphi environment.
> The program receives a 1 to 8 digit seed. The CNE delivered to the international observers the source code, the executable code, the input file, and the sample. Carter Center experts analyzed the program and concluded:

---

[11] Felten, Edward; and Aviel Rubin and Adam Stubblefield. "Analysis of Voting Data from the Recent Venezuela Referendum". Sept. 1, 2004. Available at http://venezuela-referendum.com/

[12] It is worth noting that two of the five members of the CNE are staunchly pro-opposition. (See Forero, Juan. "Chávez Urges Deference for Electoral Board". The New York Times. Sept. 1, 2003). All of the preparations and execution of the fraud conspiracy would have to have eluded them as well as the international observers.

[13] "Conned in Caracas" was the title of a September 9 piece by the Wall Street Journal Editorial Board, based on the Hausmann and Rigobón paper: "Both the Bush Administration and former President Jimmy Carter were quick to bless the results of last month's Venezuelan recall vote, but it now looks like they were had. A statistical analysis by a pair of economists suggests that the random-sample "audit" results that the Americans trusted weren't random at all."

[14] For a complete report on the auditing of the referendum, see "Audit of the Results of the Presidential Recall Referendum in Venezuela". The Carter Center. Aug. 26, 2004. Available at http://www.cartercenter.org/documents/1820.pdf

1. The program generates exactly the same sample given the same seed.
2. The program generates a different sample given a different seed.
3. The program generates a sample of voting stations (*mesas)* based on the universe of *mesas* that have voting machines.
4. The source code delivered produces the executable file delivered.
5. The input file used to generate the sample is missing only six of 8,147 voting stations (*mesas).* The input file has one missing voting center.
6. The program, when run enough times, includes each *mesa* (voting station) in the sample, and the number of times a given mesa is included in a sample is evenly distributed, indicating the sample generation program is random.

The sample generation program was run 1,020 times. With no exception all of the 8,141 *mesas* appeared at least 14 times in a sample. Not a single mesa was excluded from the sample in the test run[15].

The Carter Center therefore concluded that:

The sample drawing program used Aug. 18 to generate the 2nd audit sample generated a random sample from the universe of all *mesas* (voting stations) with automated voting machines. The sample was not drawn from a group of pre-selected *mesas*[16].

Given this evidence, the Carter Center's conclusion appears to be the only logical conclusion. It follows that the theory of electronic fraud put forth by Hausmann and Rigobón, and supported by others in Venezuela, appears to be logistically impossible.

The Econometric Evidence

Ignoring the technical difficulties in conceiving of how such a fraud might have taken place, let us turn to the econometric evidence presented by Hausmann and Rigobón. They provide two pieces of evidence; we will look at the second one first. In this part of their paper,[17] the authors use a regression analysis to test whether the sample drawn for the audit is truly a random sample. Without going into the mathematics of the model, the authors use a regression to test whether the "Sí" votes in the August 15 signatures gathered for the recall (in November-December 2003) are related differently to the audited sample as compared to how they are related to the overall universe of polling centers[18]. The theory is that if the audited sample is truly a random sample of the polling

---

[15] The Carter Center. Sept. 17, 2004. P. 6.
[16] The Carter Center. Sept. 17, 2004. P. 7.
[17] Pages 28-35.
[18] The regression is:
Log SI = Constant + Log FIRMA + D * FIRMA + Log Electores Reafirmazo + D * Log Electores Reafirmazo + Log Electores Nuevos + D * Log Electores Nuevos + Log Electores No Votantes + D * Log Electores no Votantes + D
Where SI = the number of YES votes in the referendum;
FIRMA = the number of signatures (in November-December 2003, on the recall petition);

centers, then the relationship between the signatures and the "Sí" vote count in the audited sample should not be significantly different from that in the rest of the 4580 voting centers.

The authors find that this relationship is significantly different, and interpret this as evidence of fraud.

But there is a very serious problem with this analysis. Let us return to the example offered by Hausmann and Rigobón in their paper, which describes the fraud that their regression model is here attempting to detect. Say there are 3,000 voting centers that are rigged, and 1,580 left "clean" for the CNE and observers to draw their sample from. How is this division to be made, and the sample drawn? Hausmann and Rigobón assume that the both of these selections are drawn randomly[19]. If that is the case, we would expect that the sample would show a very different proportion of YES votes than the total count. In other words, if the real vote was, as Penn, Schoen, Berland & Associates allege, -- 59-41 YES, instead of the opposite (59-41) NO, the official count -- then the sample should reflect that.

We can do a statistical test to see how likely is it that the audited sample, which Hausmann and Rigobón allege was randomly drawn from "clean" voting machines, came from a universe of voters that actually voted for the recall.

Table 1: Referendum Results*

	YES	NO	Percent YES
National Total	3,584,835	4,917,279	42.16
Audited Sample	145,785	204,640	41.60

Source: Carter Center (September 17, 2004)
* for centers with electronic voting

Electores Reafirmazo = the number of registered voters at the time of the signature gathering;
Electores Nuevos = the number of new voters (registered after the signature drive);
Electores no votantes = number of registered voters that did not vote;
and D is a dummy variable that takes on the value of 1 for voting centers that are in the audited sample and zero for non-audited centers.
The authors report a coefficient of .105 (significant at the .01 level) for the variable D * FIRMA, indicating that the elasticity of YES votes with respect to signatures is 10.5 percent higher for audited than for non-audited centers. The authors interpret this as evidence that the sample was not was not a random sample of the universe of voting centers. (Hausmann and Rigobón, p. 34)

[19] "To give an example, suppose that out of the 4,580 automated precincts used in the election, 3,000 precincts were altered but the rest were not. *Let us further suppose that the unaltered 1,580 precincts were picked at random.* This implies that they would represent a balanced sample of the country from a regional and social point of view." Hausmann 2004. P. 28.

Table 1 shows the percentage of YES and NO votes in the audited sample and from the total universe of machines[20].

As can be seen from the table, the number of YES votes for the sample (41.6 percent) is very close to the number for the overall universe (42.2 percent).

Table 2: Probability of Occurrence of the Actual Audited Sample Under Various Assumptions Regarding the True (Non-Fraudulent) Percentage "YES" Votes in the National Referendum

National vote percentage YES	Audit vote percentage YES	Approximate chance of lower audited percentage
50.10	41.60	1 in 1,400,000
55.00	41.60	1 in 73,000,000,000,000
59.00	41.60	1 in 28,000,000,000,000,000,000,000

Source: CNE, Carter Center, and authors' calculations (see Appendix)

Table 2 shows the probability of finding a sample with 41.6 percent YES votes under various assumptions for the true mean of the universe from which it came. For example, if the Penn & Schoen exit poll reflected the correct (non-fraudulent) total (59 percent YES, 41 percent NO), then the probability of the audited sample showing only 41.6 percent YES would be less than one in 28 billion trillion.

If the opposition had won by a smaller margin, 55 to 45 percent, the probability of the audited sample having only 41.6 percent YES is less than one in 72 trillion.

Finally, if Chávez had barely been recalled, with 50.1 YES to 49.9 NO, the chances of drawing a sample of voting centers with the audited sample's percentage of 41.6 percent YES is still miniscule, at less than one in a million.

The theory that there was electronic fraud that could have even come close to affecting the result of the election is therefore not plausible, under Hausmann and Rigobón's assumptions about how it could have taken place[21].

There is another possibility that would allow for the observed vote count in the audited sample. The "clean" centers and audited sample could have been selected, not randomly as Hausmann and Rigobón suggest, but from very pro-government voting centers. In this way the audited sample could show a yes vote similar to a fraudulent total vote, even if the real (non-altered) vote count had a much higher proportion of YES votes. But in this case the audited sample would not appear to be representative of the

---

[20] The Carter Center. Sept. 17, 2004. P. 5. The percentages for the totals (58.4 NO to 41.6 YES) differ slightly from the totals for the whole country (59.2 NO to 40.7 YES), because these include only the centers that used machines. About 1.4 million voters (out of 9.8 million total) used only paper ballots.

[21] Of course it is possible that just a handful of machines were rigged, increasing the government's margin of victory only slightly. But this would be a rare crime indeed: lacking not only opportunity and evidence, but also motive.

electorate; it would have to look, by other measures, very pro-Chávez -- especially if there were significant fraud in the rigged voting centers.

One obvious measure is the percentage of votes for Hugo Chávez[22] in the 2000 election, in the centers that were selected to be audited in the 2004 referendum[23]. By this measure, the audited centers do not appear significantly different from the rest of the electorate. The audited centers show a total of 61.8 percent for Chávez, as opposed to 61.4 percent for the country. As shown in Appendix below, this is well within the margin of sampling error.

In light of these results, it is difficult to interpret Hausmann and Rigobón's one regression as significant evidence of fraud. It is much more likely to be the result of a spurious correlation or a misspecification in the model that they used[24].

Results Using Exit Poll Data

Hausmann and Rigobón also claim to find evidence of electronic fraud by means of another statistical test. As in the analysis of the sample, they use a statistical model in which the intention of the voters -- an unobservable variable -- is measured, with some error, by different events. These include exit polls conducted on August 15, and the petition signatures. Both of these data are assumed to be imperfect measures of the voters' intentions, and even possibly biased.

But the authors assume that the errors in the measures of voter's intentions are not correlated with each other[25]. In other words, although we would expect (and the authors

---

[22] The numbers show the percent of votes received by Chávez as between him and the candidate who finished second Francisco Arias Cardenas (with most of the remaining votes). To see the national level results, see "Elecciones 30 de julio de 2000: Presidente de la República – Total Votos a nivel Nacional y por Entidad Federal". Consejo Nacional Electoral (Venezuela).
http://www.cne.gov.ve/estadisticas/e015.pdf
[23] The actual audit was conducted on 150 of these 196 centers.

[24] It is worth noting that this regression (see footnote 18) is only one of many regressions that could have been run to test whether the sample was significantly different from the rest of the voting centers in its relationship to the signatures. This is true not only because of the choice of control variables but also because there were 4,580 voting centers in the referendum, but only 2,700 centers for gathering signatures. Since the latter did not map directly to the voting centers, this would allow for many possible regressions of the referendum vote on the signatures, with many different data sets and/or control variables. Hausmann and Rigobón also take 1,000 random samples from the total number of polling centers and run the same regression, finding a significant result (at the 1 percent level, the same as in the regression in footnote 18) in less than one percent of the regressions. But this does nothing to validate their regression results in footnote 18; if the significant coefficient stemmed from a spurious correlation, then we would expect exactly what they found for the 1,000 regressions run on different samples with the same variables.

[25] In terms of their model,

Ei = a*Xi +epsi
Si = b*Xi +etai

10

assume) a correlation between the signatures and the exit polls, they assume that in the absence of fraud there is no such relation between the way the signatures and the exit polls, as measures of the (unobservable) intention of the voters, actually differ from this intention, and therefore the vote itself.

The full model is explained in detail and is derived mathematically in the appendices of the paper, and will not be reproduced here. The basic implication of their model, given the assumption, is that if both the exit polls and the signatures differ from the referendum vote in a similar manner then this is the result of fraud[26]. The authors find such a correlation, and interpret this as evidence of fraud.

But their result in this section depends on a crucial assumption that is difficult to justify in this situation, given the sources of the exit poll data: that the errors in the measurement from the exit poll and those from the signatures are uncorrelated. While the signature gathering was subject to controls and international monitoring, and thus can be taken as official data, the same cannot be said about the exit poll data. The exit poll data provided by Súmate was reported by Penn, Schoen, Berland & Associates, and showed the opposition to have won the referendum by 59 percent (Yes) to 41 percent (No). The other exit poll data used in this analysis was provided by the opposition group Primero Justicia, and showed the opposition winning with 62 percent of the vote. These data are not just measured with error but highly implausible. We have no idea how they were gathered or if there was fraud involved in their collection. As such, it is entirely possible that the error term for the exit polls would be correlated with the error term for the signatures. If this is the case, the empirical estimates of the Hausmann and Rigobón model in this section would say nothing about fraud in the election; rather they would simply be a product of how the exit poll data was collected.

---

where Ei is the exit poll data, Si is the signature data gathered in Nov/Dec 2003, Xi is the voter's intention and epsi and etai are error terms. It is these two error terms, not Ei and Si, which are assumed to be uncorrelated.

[26] In Hausmann and Rigobón's model they derive the equation

cov (psi1, psi2)_IV = var(Fi) + f*(1/a-c1iv)cov(Ei,Si) + f^2var(Si),

where psi1 and psi2 are the residuals of regressions of the voting results (Vi) on Ei and Si plus other control variables. As noted by the IV attached to the left hand term, the authors use instrumental variable estimation, for reasons explained in their appendix 1. The term c1iv is the instrumental variable estimator of the coefficient on Ei of the first IV regression. Fi and f are fraud variables associated with the machine rigging; if both are zero the covariance of these residuals would be zero. The authors find a positive covariance when they test this equation; they conclude that this positive covariance is evidence of fraud.

This is one of several steps in which the assumption that the error terms (espi and etai, see footnote beta) are necessary for the authors' result; on this assumption the authors are able to use Ei as an instrument for Sí, and vice versa. But even ignoring the IV estimation, if these error terms are positively correlated, , the right hand side of the equation for this covariance would have additional positive terms. This would give a positive correlation between the residuals -- cov(psi1, psi2) > 0 even if f=o (i.e. the absence of fraud). Therefore, Hausmann and Rigobón's conclusion that there is fraud, in this part of the paper, depends on the questionable assumption that the error terms for the signatures and the exit polls are uncorrelated.

Indeed it is highly unusual to be asked to question the results of an election in which so many controls and monitoring procedures were in place, on the basis of implausible exit poll data that was provided (and gathered) by political activists with no verifiable controls or monitoring. Although Hausmann and Rigobón's analysis does not require this data to be accurate, it does require that its errors be uncorrelated with those of the signatures, something that cannot be assumed without any verifiable knowledge or observation of where the data came from. It is also unusual that the authors used only this opposition data, and ignored other exit poll data that more closely predicted the official results of the election. For example, exit polling by the American polling firm Evans/ McDonough Company, Inc. polled 53,045 voters and found a result of 55% NO to 45% YES[27].

There are many ways in which the model in this section could have been mis-specified even if the political activists who gathered the exit poll data had done their best to deliver an honest and reliable result[28].

But it must be emphasized that there is no need to determine exactly how the model in Hausmann and Rigobón's analysis may have been mis-specified. The fact remains that their theory of how the fraud could have taken place is untenable, as the Carter Center has demonstrated. And the audited sample, which has been shown clearly to be a random sample of the entire universe of voting machines, matched the electronic results almost exactly. On this basis we can safely reject Hausmann and Rigobón's econometric evidence.

Conclusion

Conspiracy theories abound, and of course it is impossible to disprove them. There are tens of millions of people throughout the world who remain convinced that the massacres of September 11 were orchestrated by the Bush Administration. A best-selling book in France maintains that the no airplane ever hit the Pentagon[29]. There are detailed web sites where the evidence is marshaled, and arguments spun.

---

[27] Although this firm was hired by CITGO -- owned by the state oil company, PDVSA -- it is a reputable firm and it made its methodology transparent and available. It is difficult to see why this data would be ignored while only opposition-supplied data was used in the Hausmann/ Rigobón analysis.

[28] For example, we know that there was fraud in the signature gathering process, as 375,000 signatures were disqualified (in addition to the more than 800,000 sent to be "repaired.") The disqualified signatures included dead people, children, foreigners, etc. At the same time, the people who conducted the exit poll could have biased results, depending on who would be willing to answer their questions. It is likely that the pro-Chávez areas would have the highest incidence of refusals to answer; and it is entirely possible that these areas would have the largest errors in the signature gathering process, since these areas would present greater opportunities for fraud in the signature gathering (because of the relative scarcity of signers). This is just one example of a number of possibilities that could lead to a correlation between the error terms for the signatures and the exit polls.

[29] Riding, Alan. "Sept. 11 as Right-Wing U.S. Plot: Conspiracy Theory Sells in France". New York Times. June 21, 2002. This article discusses the book "L'Effroyable Imposture" or "The Horrifying Fraud" by French author Thierry Meyssan.

The theory that the Venezuelan referendum was stolen is different from other implausible conspiracy theories in that so long as most of the Venezuelan media is controlled by the conspiracy theorists[30], it will maintain a significant base of support.

But Hausmann and Rigobón's paper provides no credible evidence of fraud in the Venezuelan elections. There is no "black swan" here; maybe a white duck that got stuck in an oil slick.

It would be best if Venezuela could get beyond this referendum, and of course the more democratically oriented members of the opposition would like to accept the results and move on. The issue has implications beyond Venezuela as well -- most importantly for the effectiveness of international monitoring in elections. This was one of the most carefully monitored elections in modern history, with both the Carter Center and the Organization of American States playing a major role. If this level of monitoring and verification by some of the most experienced election observers in the world, in a case where the election was not even close, cannot produce a credible result, then the whole system of international monitoring would have to be called into question.

Fortunately this is not the case. The audit went smoothly; and if the vote had been close, or there had been any evidence of electronic fraud, the Venezuelan electoral authorities or the international observers could have asked for a complete count of the paper ballots. They did not do so, because there was no evidence that the result was in any way wrong. That remains the case today; and by now it is unlikely that even a full audit would convince many of the doubters, since they would maintain that sufficient time had elapsed for all the ballot boxes to be stuffed so as to match the machines.

As American pollsters hired by opposing sides in this election pointed out, every reliable pre-election poll predicted that the recall effort would fail. The most recent and comprehensive polls before the vote predicted the actual margin of victory very closely[31]. With a huge turnout of poor people and first-time voters, which favored the government, the results were even more predictable on Election Day. All the available evidence except for highly questionable exit poll data supplied by opposition activists points in the same direction. There is still no reason to question the results.

---

[30] In Venezuela most of the broadcast and print media is controlled and used as a political tool by the opposition to the government. This is important with respect to this discussion, because conspiracy theories such as those required for the electoral fraud discussed in this paper, that would not gain a large following in most other democracies, sometimes do so in Venezuela.

[31] "Venezuela Recall: Analysis of Pre-Election Polling". Evans / McDonough Company, Inc. Sept. 2004.

Appendix

Suppose, as in the hypothetical example provided by Hausmann and Rigobón, there was widespread voting machine fraud in the election, but with clean centers placed randomly throughout the country, so that the demographics at the clean centers represent the demographics of the country as a whole. Suppose also that the audit was performed on a set of centers chosen randomly from among only the clean centers. Then the results of the vote in the audited centers would, within some margin of error, reflect the election results across the nation had there been no fraud. This appendix serves to estimate the likely margin of error in the results in the audited centers with respect to the entire population.

We can then find the probability of getting the result that was found in the audited sample (41.6 percent "sí"), under various assumptions for the true vote tally for all votes. That is, under the assumption that the true total vote tally was different from the reported vote tally due to fraud. If the sample were simply a random sample of voters taken from the clean centers, then this would be a simple calculation, based on a binomial distribution. It would be analogous to calculating the probability of getting, e.g., 72 "heads" from 100 tosses of a fair coin (where the population mean is assumed to be 50).

But in this case is it slightly more complicated, because the audited sample is not a sample of voters, but a sample of voting centers. Because each voting center has a different proportion of "sí" and "no" votes, the margin of error for a sample of voting centers will be larger than the margin of error for a sample of voters. Also, we will have to construct a distribution of centers for the universe of centers, in order to estimate this margin of error. We can estimate the margin of error based on a stochastic model of the vote. In constructing this model, we will have to assume certain distributional properties of the universe of voting centers.

Statistics of the audited centers

The data listed 200 centers on the audit list. Four of these centers had no recall vote listed, leaving 196 available for analysis. These centers varied in the total number of valid votes cast as well as the percentage of votes for and against recall.

Let the number of 'sí' votes at the ith center be denoted by S_i and the number of 'no' votes be denoted by N_i. Then the log of the turnout $\ln T_i \equiv \ln\left(S_i + N_i\right)$ turns out to be roughly distributed as normal with mean 7.6 and standard deviation 0.7. The log of the vote ratio $\ln R_i = \ln\left(S_i/N_i\right)$ is also normally distributed with mean –0.4 and standard deviation 0.8. With a Pearson's r of only 0.009, the correlation between $\ln T$ and $\ln R$ is insignificant. This is consistent with independence of the two variables.

Stochastic model of the audited centers

Suppose the voting centers across the entire country have clean turnouts and vote ratios independent and lognormally distributed with means and standard deviations as in the audited sample. Then we can randomly generate 150 sample centers and estimate the audited vote. Repeating the process 1,000 times and computing the standard deviation of the proportion of 'sí' votes produces an estimate of the margin of error. As a result of this process, we estimate the audited sample to have a standard deviation in the proportion of 'sí' votes of only 1.8%. By implication, there is a 95% chance that there an audit of 150 voting centers would result in a vote proportion of 41%, plus or minus 3.5%.

Implications for the recall referendum

Let us assume a margin of error due to center sampling based on the above stochastic model. Let us also assume that 41.0% of the national vote was in favor of recall. Then there is a 37% chance that an audited sample would consist of more than 41.6% of votes in favor. That is, the audit results are consistent with the reported national results.

Let us now assume that the national vote was actually 59% in favor of recall -- as in the Penn, Schoen, Berland & Associates / Súmate exit poll used by Hausmann and Rigobón -- but this was not the reported result due to fraud. Despite any such fraud in the centers where machines were rigged, the probability of getting the observed total in the audited sample of no more than 41.6% in favor is approximately 1 in 28 thousand billion billion.

Rather than 59%, let us assume that the national vote in favor of recall was only 50.1%. That is, assume that the opposition barely won recall vote, but most of the machines were fixed to show Chávez winning by 57.8-42.2 percent. Then the chances of getting the observed result in the clean, audited sample of only 41.6% in favor would be less than 1 in a million.

Table A: Election Odds

National vote percentage	Audit vote percentage	Difference in Standard Deviations	Approximate chance of lower audited percentage
41.02	41.60	0.33	2 in 3
50.10	41.60	-4.83	1 in 1,400,000
55.00	41.60	-7.61	1 in 73,000,000,000,000
59.00	41.60	-9.88	1 in 28,000,000,000,000,000,000,000

Source: CNE, Carter Center, and authors' calculations

We can therefore conclude that the probability of drawing the observed audited sample, under the circumstances described in Hausmann and Rigobón's hypothetical electronic fraud scenario, is infinitesimally small, if there were enough fraud to affect the outcome of the election.

There is one other possibility of electronic fraud, in a slightly different scenario than that described by Hausmann and Rigobón. As noted in the text, another scenario could be the case in which the audited sample, and/or the clean centers are not chosen randomly from the universe of centers. In this case a sample that is very disproportionately pro-government could be chosen, so that the vote tally in the audited sample matches the recorded (fraudulent) vote total for the universe of centers, even though the true mean for the universe is much lower than reported due to fraud. But in this case the audited sample would appear pro-government by other measures. And it does not: to see this, we can look at how the voting centers subject to the audit in 2004 voted in the 2000 election. As between Chávez and Arias (the second place finisher), Chávez received 61.8 of the vote in these audited centers; in the country overall he received 61.4 percent. This difference is far smaller than the margin of error of 3.5 percent for the sample of audited centers.