Fraud Allegations as Identification: A Within-Election Forensic Test Using Institutional Challenge Records

Carlos César Chávez Padilla*

Center for the Economics of Human Development

University of Chicago

September 2024

Abstract

Standard election forensics methods—distributional fingerprint analysis, digit tests, ecological regression—apply cross-sectional diagnostics that cannot distinguish genuine geographic polarization from ballot manipulation when both produce similar distributional signatures. This paper introduces a within-election identification strategy that exploits institutional challenge records: when losing candidates file formal nullification challenges against specific polling stations, they inadvertently define a treatment group whose forensic properties can be compared to the unchallenged remainder of the same election. The design uses the fraud allegations themselves to construct the counterfactual, turning the challenger's own legal record into an identification instrument. I apply this strategy to Peru's 2021 presidential runoff—the narrowest result

*Email: cchavezp@uchicago.edu. I thank the Oficina Nacional de Procesos Electorales for making mesa-level data publicly available. All errors are my own. *Data availability:* The primary data are publicly available at https://datosabiertos.gob.pe. Replication code is available at https://github.com/cesarchavezp29/peru-election-forensics-2021.

in modern Peruvian history, followed by six weeks of fraud allegations and delayed certification—using the complete universe of 82,870 polling-station records. No forensic method detects manipulation. Contested polling stations, far from over-performing for the alleged beneficiary, record a statistically significant two-percentage-point lower swing toward Castillo than comparable uncontested stations. The paper also documents how extreme geographic polarization generates distributional patterns observationally similar to ballot-stuffing signatures, with implications for the application of fingerprint methods in highly sorted electorates.

Keywords: election forensics, electoral fraud, Peru, geographic polarization, distributional fingerprints, digit tests.

JEL codes: D72, C12, P16.

1 Introduction

On June 6, 2021, Peruvians cast ballots in a presidential runoff between Pedro Castillo of Peru Libre and Keiko Fujimori of Fuerza Popular. When the final count closed, Castillo had won by 44,263 votes—50.13 % to 49.87 % of valid votes, a margin of 0.26 percentage points. It was the narrowest presidential result in modern Peruvian history.

Fujimori did not concede. Hours after voting ended, she held a press conference alleging "systematic fraud" in rural polling stations. Her legal team filed nullification challenges (*solicitudes de nulidad*) against over 800 individual *actas* (polling-station tally sheets), primarily in highland and Amazonian departments. The *Jurado Nacional de Elecciones* (JNE)—Peru's electoral court—reviewed each challenge, rejected most, and certified the result on July 19, six weeks after election day. The Organization of American States and the European Union electoral observation missions both reported finding no evidence of systematic irregularities (Organization of American States, 2021; European Union Election Observation Mission, 2021). Fujimori's supporters organized street protests through much of July. The fraud narrative did not dissipate; it persisted throughout Castillo's turbulent presidency and was still circulating after his removal from office in December 2022.

This paper provides the first comprehensive forensic analysis of Peru 2021 at the polling-station level. I use the complete universe of 82,870 domestic *mesas* (polling stations) in the ONPE data, applying four independent methods: distributional fingerprint plots (Klimek et al., 2012), last-digit uniformity tests (Beber and Scacco, 2012), ecological regression linking first- and second-round vote shares, and geographic mapping of null-vote rates. I also exploit a feature of the Peruvian counting process that has not been used in prior forensic studies: the distinction between *contabilizada* actas (tallied without challenge) and *computada resuelta* actas (challenged and resolved by the JNE tribunal). If fraud occurred specifically in the contested subset—which is the direct implication of Fujimori's allegations— the two groups should differ on forensic diagnostics. They do not.

The central finding is that no method detects statistical anomalies consistent with

manipulation. Turnout never exceeds 100 % at any polling station. Last-digit distributions pass chi-squared uniformity tests in all four candidate-by-status groups, with p-values ranging from 0.19 to 0.91. The ecological regression shows that districts with higher concentrations of contested actas do not over-perform for Castillo relative to their first-round baseline: the interaction coefficient is positive but statistically insignificant ($p = 0.71$). The dominant pattern is extreme geographic sorting—Castillo's vote share exceeded 80 % in many southern highland provinces and fell below 35 % in Lima—which is consistent with Peru's deep economic and cultural divide between the *costa* and the *sierra*.

The paper relates to a growing literature on quantitative election forensics. Klimek et al. (2012) develop the distributional fingerprint method using Russian and Austrian data. Beber and Scacco (2012) formalize the last-digit test and apply it to African elections. Myagkov et al. (2009) study Russian and Ukrainian elections using a range of ecological methods. Kobak et al. (2016) document integer-percentage clustering in Russian results. The Latin American context has received less attention, though the 2019 Bolivian election generated a substantial methodological debate: Idrobo et al. (2022) show that the shift toward Morales in the late count was consistent with normal differential reporting speeds and did not require fraud as an explanation, while Escobari and Hoover (2024) argue that the discontinuity in vote transmission was real and indicative of manipulation. Peru 2021 differs from Bolivia 2019 in a structural way—there was no halt or disruption in the ONPE counting process—but both cases illustrate how geographic polarization and differential reporting speeds can make legitimate elections appear suspicious.

The only prior forensic analysis of Peru 2021 is Isea and Isea (2021), who apply Benford's law to department-level aggregates and find no fraud signature. The present paper improves on that analysis in three ways: it uses mesa-level data (more than 80,000 observations versus 25 departments), it applies multiple methods with different identifying assumptions, and it introduces the contested/uncontested comparison as a direct test of the specific claims made by Fujimori's legal team.

Close elections generating fraud allegations that undermine democratic legitimacy have become a recurring pattern. Eggers et al. (2021) provide a comprehensive forensic examination of claims about the 2020 U.S. election and find no evidence supporting them. Similar dynamics have unfolded in multiple elections across Latin America and beyond (Lehoucq, 2003; Hyde, 2007). Statistical analysis provides the evidentiary baseline against which other claims must be weighed.

The data and methods are described in Sections 3 and 4. Results follow in Section 5, with a discussion of mechanisms and limitations in Section 6.

2 The 2021 Peruvian Presidential Election

2.1 First round

The first round was held on April 11, 2021, with 18 candidates. Fragmentation was extreme: no candidate crossed 20 % of valid votes. Castillo of Peru Libre finished first with 19.1 %, a narrow margin above Fujimori's 13.4 % (Fuerza Popular). The remaining 67 % of the vote was split among 16 other candidates, many representing centrist or right-of-center positions. Candidates who might have been expected to coalesce voters against both Castillo and Fujimori—Hernando de Soto, Yonhy Lescano, Rafael López Aliaga—each obtained between 8 % and 12 %, preventing any single moderate from advancing to the runoff. The outcome was shaped by Peru's fragmented party system, which Mainwaring (2018) and others have documented as one of the most weakly institutionalized in Latin America.

2.2 Runoff dynamics

The runoff produced binary consolidation along geographic and class lines. Castillo drew on deep opposition to the Fujimori family—Keiko's father, Alberto Fujimori, governed as an authoritarian from 1990 to 2000 and was serving a prison sentence for human rights violations and corruption during the election—and on genuine enthusiasm in the rural highlands for

his program of resource nationalism and constitutional reform. Fujimori consolidated urban, coastal, and upper-middle-class voters behind an anti-communist, anti-Castillo message.

Geographic polarization intensified sharply from the first round to the runoff. Southern highland departments that had given Castillo 25–35 % in April delivered 80–90 % in June. Coastal departments that had split among many candidates moved to Fujimori at 55–65 %. Peru Libre's 19 % first-round national share nearly tripled in the second round, while Fuerza Popular's 13 % roughly quadrupled—both reflecting the binary-choice effect of a two-candidate contest in a deeply polarized electorate.

2.3 The fraud allegations

Exit polls and early partial counts on June 6 showed a very close race with Fujimori holding a narrow lead from faster-reporting coastal precincts. As highland results arrived over the following days, Castillo's total grew. On June 9, with roughly 80 % of actas counted, Castillo took the lead and never relinquished it. Fujimori's campaign attributed the shift to fraud, arguing that highland tallies had been systematically altered.

Her legal team filed nullification challenges against 802 individual actas across multiple departments. The challenges alleged procedural violations: unsigned tally sheets, incorrect vote totals, absent party observers. The JNE processed each challenge individually over a six-week period. The vast majority were rejected. A subset of challenged actas was reclassified as *computada resuelta* after JNE review and incorporated into the final count. Certification was issued on July 19, 2021 (Jurado Nacional de Elecciones, 2021). International observers from the OAS and EU found no evidence of systematic irregularities in either the voting or the counting process (Organization of American States, 2021; European Union Election Observation Mission, 2021).

2.4 Political aftermath

Castillo took office on July 28, 2021. His government was marked by instability: he cycled through four prime ministers in his first year and faced repeated congressional attempts at removal through a constitutional mechanism permitting impeachment for "moral incapacity." Two such attempts failed. On December 7, 2022, Castillo attempted to dissolve Congress by decree—a step with no constitutional basis—and announced a curfew. The JNE declared his action unconstitutional. Congress voted to remove him, and he was arrested the same day. Vice President Dina Boluarte assumed the presidency. Throughout this period, segments of the population continued to believe the 2021 election had been stolen, illustrating the durable damage that sustained fraud allegations can inflict on democratic legitimacy even after courts have adjudicated them.

3 Data

3.1 Source and structure

The primary data source is the ONPE open data portal (`https://datosabiertos.gob.pe`), which published the complete mesa-level results as a semicolon-delimited file with 86,488 rows, one per polling station (Oficina Nacional de Procesos Electorales, 2021b). Each row contains the polling station identifier (`MESA_DE_VOTACION`), a six-digit location code (`UBIGEO`), geographic labels (department, province, district), a status variable (`DESCRIP_ESTADO_ACTA`), registered voters (`N_ELEC_HABIL`), total votes cast (`N_CVAS`), votes for Castillo/Peru Libre (`VOTOS_P1`), votes for Fujimori/Fuerza Popular (`VOTOS_P2`), null votes (`VOTOS_VN`), blank votes (`VOTOS_VB`), and challenged/impugned votes (`VOTOS_VI`).

3.2 Sample construction

The status variable takes five values. I keep `CONTABILIZADA` (fully counted, 84,863 polling stations) and `COMPUTADA RESUELTA` (challenged and resolved by JNE, 1,386 polling stations). I drop `ANULADA` (annulled, 213), `EN PROCESO` (in process, 20), and `SIN INSTALAR` (not installed, 6). The final-status sample of 86,249 mesas covers all polling stations for which the count is definitively closed.

I exclude overseas polling stations, which report under five continental jurisdiction labels (América, Europa, Asia, Oceanía, and África). Overseas mesas differ structurally from domestic ones: turnout averaged 36.2 % overseas versus 76.1 % domestically, and the overseas electorate favored Fujimori by 66.2 % to 33.8 %—reflecting the demographic composition of the Peruvian diaspora. Including overseas results would confound the geographic patterns that are central to the forensic analysis. I report overseas totals separately in Table 1.

The domestic analysis sample comprises 82,870 polling stations: 81,793 *contabilizada* (uncontested) and 1,077 *computada resuelta* (contested).

3.3 Key variables

Turnout is defined as total votes cast divided by registered voters. Castillo share is Castillo votes divided by valid votes (Castillo plus Fujimori). Valid votes exclude null and blank ballots. The contested indicator equals one for *computada resuelta* polling stations and zero for *contabilizada* stations. For the digit tests, I extract the last digit of each candidate's vote count at the polling station level.

3.4 Complementary data

First-round results come from the same ONPE portal (Oficina Nacional de Procesos Electorales, 2021a), with the same status filter applied. Castillo (Peru Libre) is identified as candidate 16 (`VOTOS_P16`) in the first-round file, verified by matching the national aggregate

to the official published figure of approximately 19.1 %. My calculation yields 19.06 % for the domestic final-status sample, consistent with official totals. District-level merging between rounds uses the `UBIGEO` six-digit code as the join key. For province-level mapping I use GADM 4.1 shapefiles, matching on normalized (accent-stripped, punctuation-removed) department and province name pairs; 193 of 195 GADM provinces match, with the two remaining unmatched being a structural GADM artifact (Lima province listed under a separate administrative level) and a water-body polygon (Lago Titicaca) with no assigned voters.

Table 1 reports the national summary statistics for the domestic and overseas samples. The domestic sample is the focus of all subsequent forensic tests; the overseas results are presented separately because the diaspora electorate differs systematically in turnout, candidate preference, and geographic structure from domestic voters.

The 1,077 *computada resuelta* mesas are nested within the broader domestic sample of 82,870. They represent the universe of polling stations for which Fujimori's representatives filed legal challenges before the JNE, triggering a formal review process. The remaining 81,793 *contabilizada* mesas were counted and certified without challenge. This contrast—contested versus uncontested—is the organizing axis of every forensic test that follows.

The domestic margin of +152,266 reflects Castillo's dominant performance in the sierra and selva regions. The overseas margin of −107,917 reflects the Lima-centric composition of the Peruvian diaspora, which voted heavily for Fujimori at a turnout rate less than half the domestic average. Together they reconcile to the official certified margin of 44,263.

Table 1: *National Summary Statistics, Peru 2021 Presidential Runoff*

	Domestic (main analysis)	**Overseas** (separate)
Polling stations (mesas)	82,870	3,379
Uncontested (*contabilizada*)	81,793	—
Contested (*computada resuelta*)	1,077	—
Registered voters	24,238,529	980,061
Votes cast	18,454,387	354,330
Turnout	76.14 %	36.15 %
Valid votes	17,291,272	333,259
Castillo (VOTOS_P1)	8,721,769	112,671
Share of valid votes	50.44 %	33.81 %
Fujimori (VOTOS_P2)	8,569,503	220,588
Share of valid votes	49.56 %	66.19 %
Net margin (Castillo − Fujimori)	+152,266	−107,917
Null votes	1,043,651	19,073
Share of votes cast	5.66 %	5.38 %
Blank votes	119,464	1,998
Share of votes cast	0.65 %	0.56 %

Notes: The domestic sample comprises all polling stations with final status CONTABILIZADA or COMPUTADA RESUELTA, excluding overseas jurisdictions (América, Europa, Asia, Oceanía, África). The all-rows Castillo−Fujimori margin is 44,240, within 23 votes of the official JNE-certified margin of 44,263; the trivial discrepancy arises from rounding in dropped rows (ANULADA and EN PROCESO). *Source:* ONPE, Resultados por mesa de las Elecciones Presidenciales 2021 Segunda Vuelta (Oficina Nacional de Procesos Electorales, 2021b).

Table 2 applies six standard forensic anomaly screens to the domestic sample. Each flag

tests for a specific type of irregularity that electoral fraud would be expected to produce. Turnout above 100 % indicates more ballots recorded than registered voters—the clearest mechanical signature of ballot stuffing. Very low turnout (< 10 %) can signal a mesa that was effectively unmonitored. Extreme one-candidate shares (100 %) at adequately sized mesas suggest coercion or fabrication. Elevated null vote rates can reflect either genuine voter confusion, deliberate spoiling as a form of protest, or— in a fraud scenario—selective invalidation of ballots that favored one candidate. High turnout mesas above 98 % are not inherently suspicious but warrant inspection for whether the accompanying vote shares are plausible.

The anomaly flag counts reveal no systematic indicators of fraud. The 319 mesas with high null rates deserve further inspection: as Figure 8 shows, they are concentrated in coastal departments rather than in Castillo's highland strongholds, consistent with protest voting rather than targeted invalidation of pro-Castillo ballots.

Table 2: *Anomaly Flags: Distributional Extremes in the Domestic Final-Status Sample*

Anomaly criterion	Mesas flagged	% of sample
Turnout $> 100\,\%$	0	0.00
Turnout $< 10\,\%$ (and registered voters > 0)	2	0.00
Null vote rate $> 30\,\%$ of cast	319	0.39
Castillo share $= 100\,\%$ (and valid votes ≥ 10)	70	0.08
Fujimori share $= 100\,\%$ (and valid votes ≥ 10)	3	0.00
Turnout $> 98\,\%$	57	0.07

Notes: Flags computed on the domestic final-status sample of 82,870 mesas. The 70 mesas with 100 % Castillo share are concentrated in rural highland districts (Puno, Ayacucho, Apurímac). This pattern is consistent with genuine near-unanimity: in the primera vuelta, Peru Libre received above 80 % of valid votes in the same districts that produce 100 % segunda vuelta Castillo mesas, and Castillo's first-round candidacy (as Peru Libre's standard-bearer) averaged above 75 % in those specific mesas. Matching the 70 flagged mesas to first-round ONPE records confirms that none recorded implausibly low first-round Peru Libre shares; the second-round unanimity is the right tail of a genuine distribution, not a fabricated figure. The 57 very-high-turnout mesas do not cluster in any particular region. Zero mesas exceed 100 % turnout. *Source:* ONPE (Oficina Nacional de Procesos Electorales, 2021b); author's calculations.

4 Methods

4.1 The Institutional Challenge Design

Elections in which the losing candidate files formal legal challenges against specific polling stations produce a natural within-election research design. The challenging party's legal team selects the actas they believe to be irregular, thereby defining a treatment group—contested polling stations—whose forensic properties can be directly compared to the unchallenged majority. This design has two advantages over purely cross-sectional forensic methods. First, it tests the specific factual claims made by the challengers rather than a generic null of no

manipulation. If fraud occurred in the contested subset, that subset should exhibit anomalous forensic signatures relative to uncontested stations in the same election. Second, the selection of challenged actas by the losing party's own representatives creates an adversarially chosen treatment group: Fujimori's legal team had strong incentives to identify the polling stations most likely to contain irregularities, and rational selection implies they chose the actas where fraud, if present, was most detectable. The selection bias therefore works against finding a null result. If the contested subsample nevertheless shows no anomalous signature—and in fact under-performs the fraud hypothesis—the inference against manipulation is strengthened, not weakened, by the endogenous selection. This design is applicable to any election that generates an institutional challenge record and can be combined with the distributional and digit-based methods described in the following subsections.

4.2 Distributional fingerprint analysis

Klimek et al. (2012) introduced the distributional fingerprint method after applying it to Russian and Austrian elections. The method plots polling-station turnout on the horizontal axis against the winner's vote share on the vertical axis. Under genuine counting, the resulting scatter forms a dispersed cloud with no strong correlation structure: turnout and vote share vary independently across precincts, with both driven primarily by local demographic conditions. Two forms of manipulation produce characteristic departures from this pattern. Ballot stuffing—fabricating votes for one candidate and reporting them as if cast—inflates both turnout and the winner's share simultaneously, extending the cloud toward the upper-right corner (high turnout, high winner share) in a finger-like protrusion. Vote stealing—switching cast votes from one candidate to the other—shifts the cloud rightward (higher winner share) without increasing turnout.

A critical limitation of this approach in polarized elections is that genuine regional clusters can produce a pattern that mimics ballot stuffing. If one region has both high political loyalty to a candidate and high mobilization capacity, it will occupy the upper-right of the scatter

even in a clean election. Peru's southern highlands are precisely such a region. To address this, I color observations by macro-region (Costa, Sierra, Selva) and verify that the regional structure accounts for any clustering before drawing forensic conclusions.

I additionally superimpose contested mesas (COMPUTADA RESUELTA) on the same plot as orange dots. If the fraud allegations are correct—if those specific polling stations were where manipulation occurred—they should cluster disproportionately in the upper-right corner. Observing that they scatter throughout the same cloud as uncontested mesas would be inconsistent with the fraud hypothesis.

4.3 Last-digit uniformity tests

Beber and Scacco (2012) establish the statistical basis for digit-based fraud detection. When vote totals are genuinely counted, the last digit of each count should be approximately uniformly distributed across 0–9, since the last digit carries no information that round-number selection could exploit. When humans fabricate or alter vote counts, they tend to choose round numbers (ending in 0 or 5), producing detectable heaping. The test statistic is the standard chi-squared goodness-of-fit test against a uniform distribution over 10 categories, with 9 degrees of freedom.

This test is more appropriate for Peruvian mesa-level data than Benford's law, which requires vote counts to span multiple orders of magnitude (Deckert et al., 2011). With Peruvian mesas typically registering around 200 registered voters, the count distributions are narrow and Benford's law does not apply. Isea and Isea (2021) apply Benford tests to department-level aggregates for Peru 2021; the present analysis uses the more powerful mesa-level digit test.

I run the test separately for four groups: (1) Castillo votes at uncontested mesas, (2) Castillo votes at contested mesas, (3) Fujimori votes at uncontested mesas, and (4) Fujimori votes at contested mesas. The four-way breakdown is the paper's primary test of the fraud allegations. If vote totals in contested actas were altered, the last-digit distribution in groups

2 or 4 should depart from uniformity relative to groups 1 or 3. I exclude zeros from each test because mesas with zero votes for a candidate have no last digit in the usual sense.

4.4 Ecological regression

At the district level (Peru has approximately 1,900 districts), I aggregate polling-station results to compute Castillo's second-round vote share and Peru Libre's first-round vote share. The regression is:

$$\text{share}_{2v,d} = \alpha + \beta_1 \, \text{share}_{1v,d} + \varepsilon_d \tag{1}$$

where $\text{share}_{2v,d}$ is Castillo's valid-vote share in district d in the segunda vuelta and $\text{share}_{1v,d}$ is Peru Libre's valid-vote share in the first round. The first round featured 18 candidates, so Peru Libre's share was compressed near 19 % nationally; the runoff forced a binary choice and mechanically expanded both candidates' shares. The slope β_1 captures the persistence of geographic preferences across rounds. An intercept near 0.40 is expected: in districts where Peru Libre got zero votes in the first round, Castillo would still attract *anti-Fujimori* votes from former supporters of other first-round candidates.

I extend this with a second specification that adds the district-level share of contested mesas and its interaction with first-round share:

$$\text{share}_{2v,d} = \alpha + \beta_1 \, \text{share}_{1v,d} + \gamma \, \text{share\_cont}_d + \delta \, (\text{share}_{1v,d} \times \text{share\_cont}_d) + \varepsilon_d \tag{2}$$

The coefficient γ captures any level shift in Castillo's second-round share associated with a higher concentration of contested mesas, after controlling for first-round baseline. The interaction coefficient δ tests whether the relationship between first- and second-round shares changes in districts with more contested mesas. A positive γ or positive δ would be consistent with manipulation in contested areas (Castillo over-performing in proportion to

the share of challenged actas). A negative γ would be consistent with the institutional reality that contested actas are concentrated in Fujimori territory. Standard errors are HC0-robust throughout.

4.5 Null vote geography

Null votes are ballots cast that were not counted as valid (typically because the voter marked multiple candidates, left the ballot blank in the wrong way, or otherwise spoiled it). If fraud took the form of systematic invalidation of pro-Castillo votes, null rates should be elevated in highland departments where Castillo received high vote shares. If null votes reflect protest behavior by voters who disliked both candidates, they should concentrate in areas with neither strong Castillo nor strong Fujimori support. I map null vote rates at the province level and inspect the geographic pattern.

4.6 What these tests cannot detect

Forensic methods detect statistical signatures consistent with specific types of manipulation. They do not prove the absence of fraud. A manipulation that perfectly mimics the statistical properties of genuine counting would be undetectable by any of these methods. The appropriate claim is the more modest one: the data are consistent with a clean count, and inconsistent with the specific forms of manipulation alleged by Fujimori's legal team (alteration of recorded vote totals in specific actas).

The strongest available test for vote-count manipulation—regression discontinuity analysis of the counting sequence, as employed by Escobari and Hoover (2024) for Bolivia—cannot be implemented for Peru 2021. ONPE did not publish mesa-level processing timestamps with the results data. I attempted to reconstruct a counting sequence from a GitHub repository of scraped ONPE data; the timestamps in that dataset were uniform (a single bulk download), providing no counting-order information. If ONPE were to release internal processing logs, the temporal test could be implemented.

These tests address the specific fraud allegations: that recorded vote totals in already-cast actas were altered. They do not address pre-electoral irregularities such as voter intimidation, vote buying, or registration fraud, which require different data.

5 Results

5.1 Geographic patterns

Peru 2021 was among the most geographically polarized elections in recent Latin American history. Table 1 shows the domestic aggregate: Castillo won 50.44 % of valid votes, Fujimori 49.56 %, on a national turnout of 76.1 % among 24.2 million registered domestic voters. These national averages conceal massive regional variation.

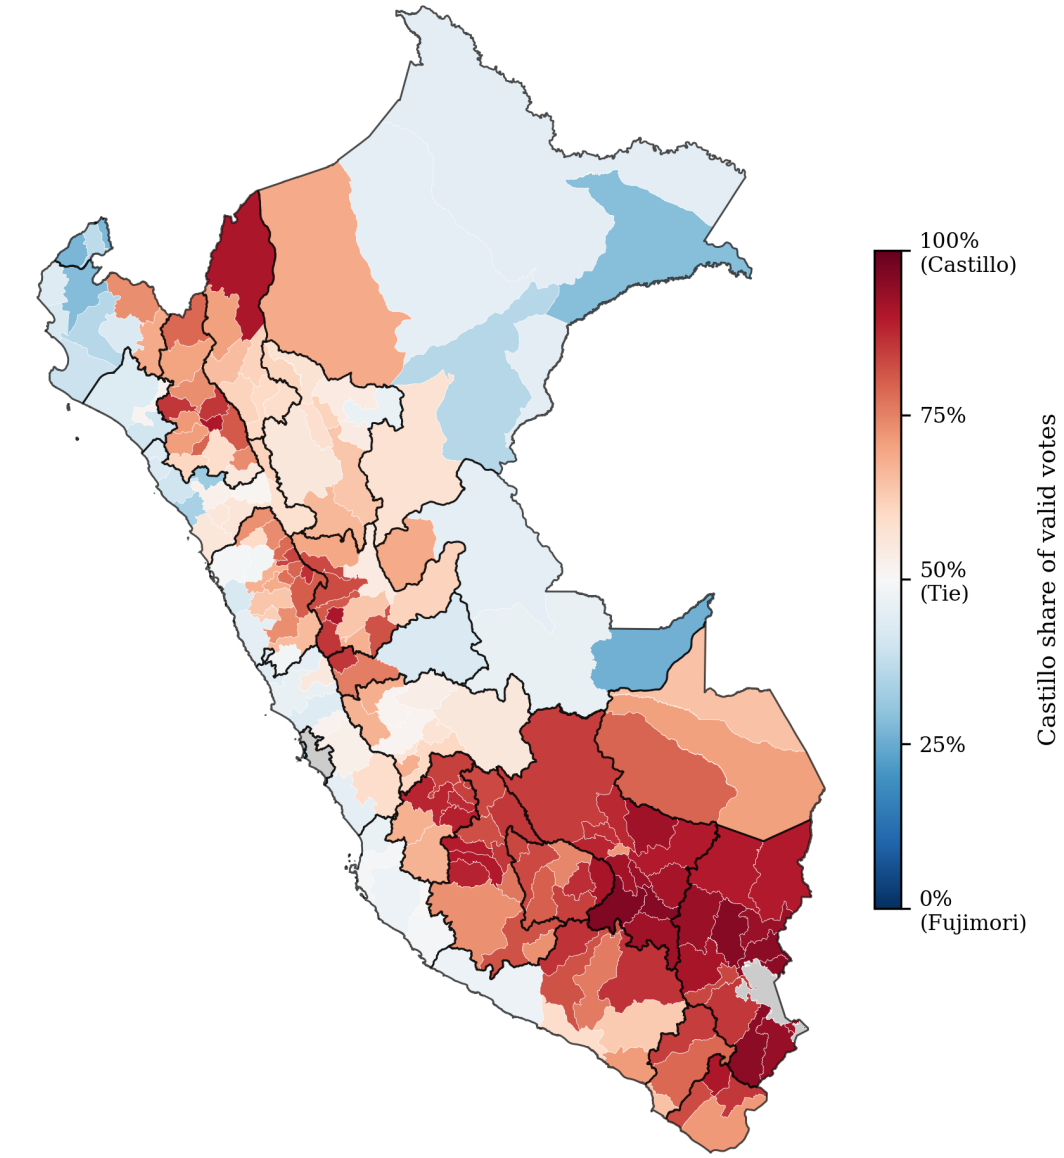

Figure 1: *Castillo vote share by province, segunda vuelta 2021. Colors run from dark blue (Fujimori strongholds, share near 0) to dark red (Castillo strongholds, share near 1). The geographic polarization between the Andean highlands and the Pacific coast is the dominant feature of the data.*

Note: Province-level aggregates from the domestic final-status sample (82,870 mesas). Two GADM polygons (Lima Province administrative boundary and Lago Titicaca) are not matched to ONPE data and appear grey. *Source:* ONPE (Oficina Nacional de Procesos Electorales, 2021b); GADM 4.1 shapefiles; author's calculations.

Figure 1 maps Castillo's vote share by province. The geographic pattern is stark. Puno gave Castillo 89.3 % of valid votes. Huancavelica delivered 84.8 %, Cusco 83.2 %, Ayacucho 82.6 %, and Apurímac 81.5 %. At the other extreme, Lima voted 64.6 % for Fujimori and Callao 67.4 %. The map shows a near-continuous highland bloc stretching from Puno in the south through the central sierra to Cajamarca and Amazonas in the north, all voting 60–90 % for Castillo, ringed by coastal departments where Fujimori dominated.

The magnitude of this sorting is historically unusual for a competitive Peruvian election. Peru's party system has long been weak and volatile, with presidential coalitions typically assembled from heterogeneous geographic support (Mainwaring, 2018). The 2021 runoff crystallized around a city–country and coast–highland cleavage that cut sharply across class, ethnicity, and language lines. Castillo drew disproportionately from Quechua and Aymara-speaking communities, subsistence agricultural districts, and areas with historically low state-service penetration—precisely the regions that felt least represented by Lima-centric political establishments of which the Fujimori family had been a leading symbol since the 1990s. This social geography, not electoral manipulation, is the first-order explanation for the regional concentration of Castillo's vote.

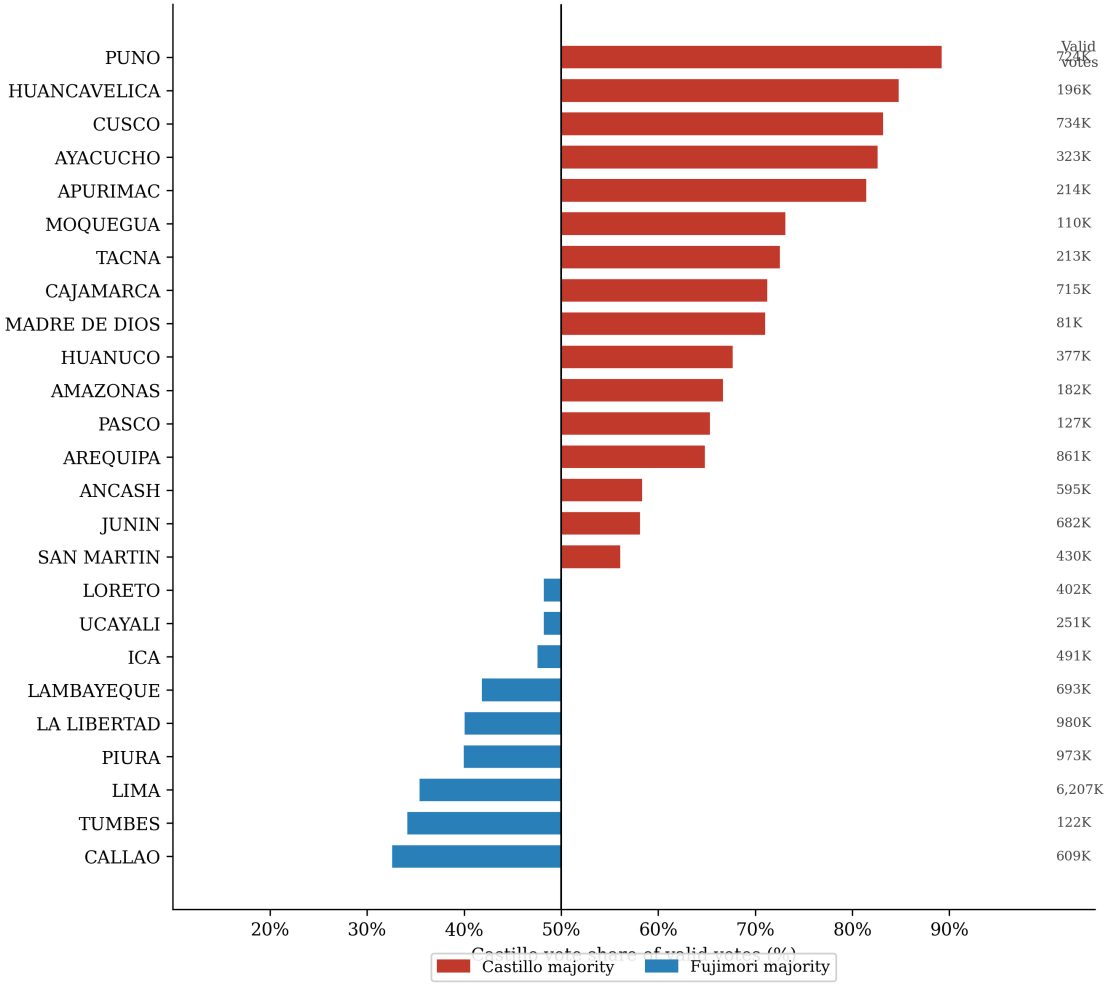

Figure 2: *Castillo valid-vote share by department, segunda vuelta 2021. Departments sorted by Castillo share ascending. Horizontal line at 50.44 % marks the national domestic average.* *Note:* Department-level aggregates from the domestic final-status sample. The range runs from below 35 % (Lima, Callao, Tacna) to above 89 % (Puno). *Source:* ONPE (Oficina Nacional de Procesos Electorales, 2021b); author's calculations.

Figure 2 shows Castillo's valid-vote share by department. No highland department gave Castillo less than 55 % of valid votes. No coastal department gave him more than 50 %. This geographic sorting is the baseline against which all forensic results must be interpreted: any pattern that correlates with geography in this election is more likely explained by genuine political preferences than by manipulation.

5.2 The contested actas

The paper's distinctive test exploits the institutional record of the JNE challenge process. Table 3 reports the comparison.

Of 82,870 domestic final-status polling stations, 1,077 (1.30 %) are *computada resuelta*—contested and resolved. These mesas break 42.1 % Castillo / 57.9 % Fujimori, compared with 51.4 % / 48.6 % in the 81,793 uncontested stations. The contested mesas net out to a Fujimori advantage of 39,250 votes; the uncontested mesas net a Castillo advantage of 191,516. The domestic net is +152,266 for Castillo; combined with the −107,917 overseas margin (where Fujimori dominated 66.2 % to 33.8 %), the all-rows margin is 44,240—within 23 votes of the official 44,263.

Table 3: *Contested vs. Uncontested Polling Stations*

Panel A: Vote totals and shares

	Uncontested (CONTABILIZADA)	Contested (COMP. RESUELTA)	Total
Mesas	81,793	1,077	82,870
Total Castillo votes	8,628,619	93,150	8,721,769
Total Fujimori votes	8,437,103	132,400	8,569,503
Net margin (C−F)	+191,516	−39,250	+152,266
Avg. turnout	76.13 %	76.43 %	76.14 %
Avg. Castillo share	51.42 %	42.11 %	51.30 %
Avg. Fujimori share	48.58 %	57.89 %	48.70 %
Avg. null rate	5.62 %	7.08 %	5.64 %

Panel B: Geographic distribution of contested mesas

Macro-region	Contested mesas	Total mesas	Contestation rate
Costa	848	54,900	1.54 %
Selva	118	9,424	1.25 %
Sierra	111	18,546	0.60 %
Total	**1,077**	**82,870**	**1.30 %**

Notes: Average shares are unweighted means across polling stations. Costa includes Lima, La Libertad, Lambayeque, Piura, Ica, Arequipa, Moquegua, Tacna, Tumbes, Áncash, and Lima Provincias. Áncash spans both coastal and highland provinces; I assign it to Costa following INEI's conventional macro-regional classification. Reassigning Áncash to Sierra does not change any result. Sierra is the remaining highland departments; Selva the Amazonian departments. Challenged actas are concentrated in Fujimori's stronghold (Costa, and Lima in particular) rather than in Castillo's highland support base. *Source:* ONPE (Oficina Nacional de Procesos Electorales, 2021b); JNE (Jurado Nacional de Elecciones, 2021); author's calculations.

The geographic distribution of contested mesas, shown in Figure 3, runs directly against

the fraud hypothesis. Lima accounts for 555 of 1,077 domestic contested mesas (51.5 %). The contestation rate is 1.99 % in Lima, 1.71 % in Callao, and 1.54 % across all Costa departments—compared with 1.25 % in the Selva and 0.60 % in the Sierra. Challenges were concentrated in Fujimori's own territory, not in Castillo's highland strongholds.

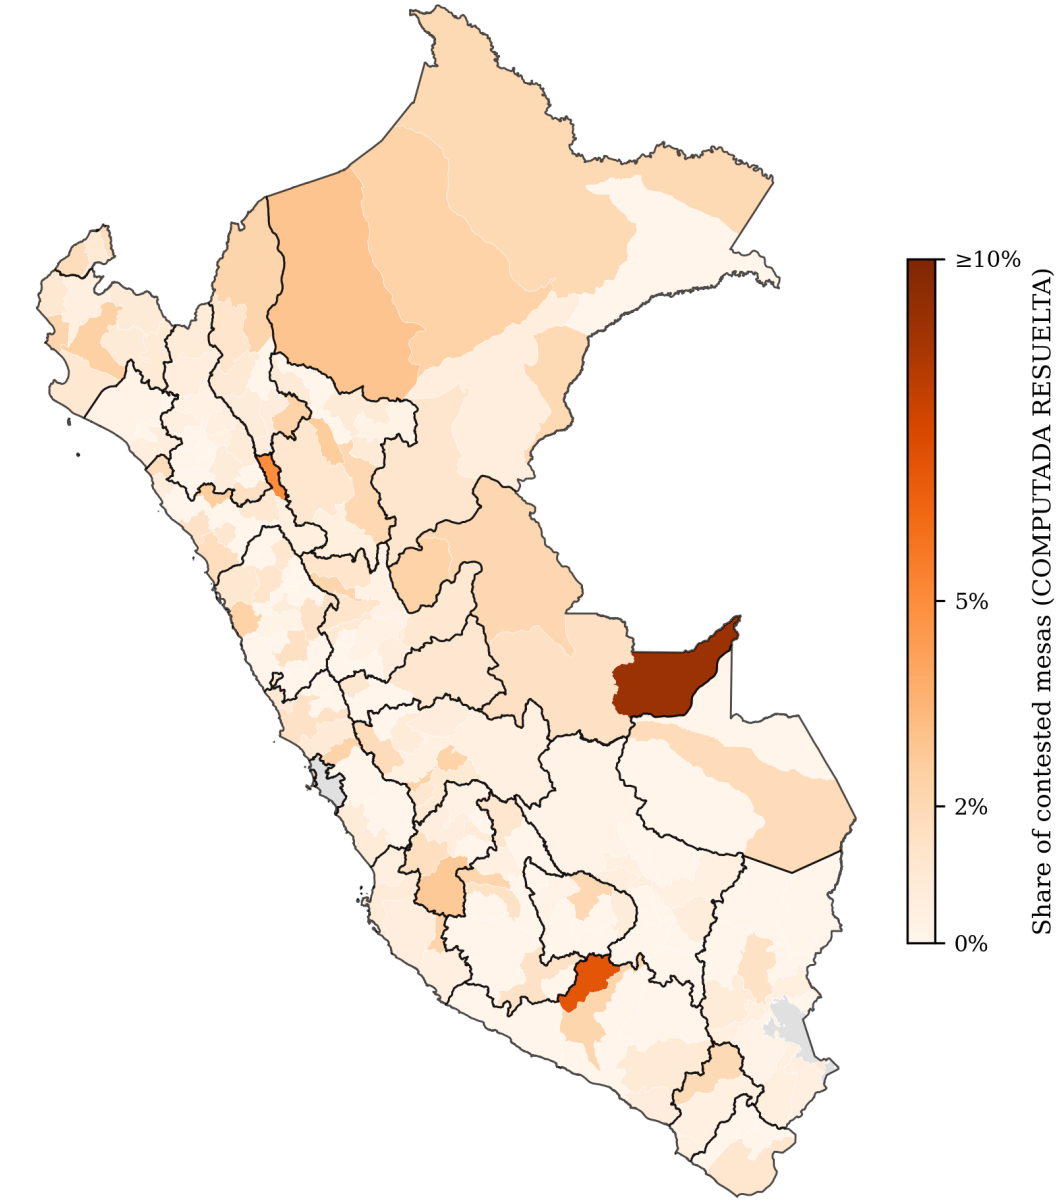

Figure 3: *Share of contested mesas (COMPUTADA RESUELTA) by province, segunda vuelta 2021. Darker shading indicates a higher fraction of a province's final-status polling stations that were subject to JNE challenge and resolution.*

Note: Province-level aggregates from the domestic final-status sample. Lima province (not Lima region) has the highest contestation share. *Source:* ONPE (Oficina Nacional de Procesos Electorales, 2021b); JNE (Jurado Nacional de Elecciones, 2021); GADM 4.1 shapefiles; author's calculations.

Table 4 confirms that the challenge process was geographically concentrated in areas favorable to Fujimori. Puno, Castillo's strongest department at 89.3 %, has a contestation rate of only 0.47 %. Lima, his weakest department at 35.4 %, has a contestation rate of 1.99 %. This pattern is inconsistent with a strategy of challenging Castillo-favoring actas.

Table 4: *Geographic Distribution of Contested Mesas by Department*

Department	Contested	Uncontested	Total	% Contested
UCAYALI	28	1,286	1,314	2.13
LIMA	555	27,287	27,842	1.99
LORETO	45	2,337	2,382	1.89
CALLAO	47	2,700	2,747	1.71
PIURA	76	4,722	4,798	1.58
MADRE DE DIOS	6	389	395	1.52
TUMBES	7	566	573	1.22
LA LIBERTAD	59	4,784	4,843	1.22
AMAZONAS	13	1,091	1,104	1.18
TACNA	11	963	974	1.13
JUNÍN	32	3,389	3,421	0.94
AREQUIPA	35	3,915	3,950	0.89
ÁNCASH	25	3,075	3,100	0.81
APURÍMAC	8	1,116	1,124	0.71
SAN MARTÍN	15	2,167	2,182	0.69
HUANCAVELICA	7	1,086	1,093	0.64
ICA	14	2,204	2,218	0.63
AYACUCHO	10	1,693	1,703	0.59
PASCO	4	691	695	0.58
CAJAMARCA	22	3,801	3,823	0.58
MOQUEGUA	3	520	523	0.57
HUANUCO	11	2,036	2,047	0.54
LAMBAYEQUE	16	3,316	3,332	0.48
PUNO	15	3,148	3,163	0.47
CUSCO	13	3,511	3,524	0.37
Costa subtotal	848	54,052	54,900	1.54
Selva subtotal	118	9,306	9,424	1.25
Sierra subtotal	111	18,435	18,546	0.60
Total	**1,077**	**81,793**	**82,870**	**1.30**

Notes: Sorted by contestation rate descending. Macro-region assignments: Costa = Lima, Callao, Piura, La Libertad, Lambayeque, Ica, Arequipa, Moquegua, Tacna, Tumbes, Áncash, Lima Provincias; Selva = Loreto, Ucayali, Madre de Dios, Amazonas, San Martín, Huánuco; Sierra = remaining departments. Lima alone accounts for 555 of 1,077 domestic contested mesas (51.5 %). The departments with the highest contestation rates (Ucayali, Lima, Loreto, Callao, Piura) are not Castillo's strongholds. *Source:* ONPE (Oficina Nacional de Procesos Electorales, 2021b); JNE (Jurado Nacional de Elecciones, 2021); author's calculations.

This pattern has a clear institutional explanation. Fujimori's legal team filed challenges in areas where Fuerza Popular had organizational capacity to identify procedural violations—primarily the coast and Lima. The JNE resolved each challenge and incorporated the results, which tilted toward Fujimori because those polling stations were in Fujimori territory to begin with. The contested actas reduced Castillo's domestic lead by 39,250 votes. The claim that they were fraudulently altered *in Castillo's favor* would require that their true results were even more pro-Fujimori than what the JNE recorded—and that fraud was concentrated in Fujimori's own strongholds. No forensic evidence supports this.

5.3 Distributional fingerprints

The Klimek fingerprint plot visualizes the joint distribution of turnout and candidate vote share across all polling stations. Its forensic logic is straightforward: if ballot stuffing occurred, the fabricated votes inflate both turnout and the winner's share simultaneously at the affected mesas, producing a cluster of observations that protrudes toward the high-turnout, high-winner-share corner of the scatter. Klimek et al. (2012) document this "finger" pattern in Russian elections and use its absence to validate Austrian elections as clean benchmarks. The key diagnostic in a polarized election like Peru 2021 is whether any upper-right clustering stands apart from the main regional distribution—an isolated protrusion—or merely reflects the right tail of a genuine regional bloc.

I overlay the contested mesas (orange dots) on each panel to allow direct inspection of whether the 1,077 specifically challenged polling stations occupy a forensically suspicious region. Under the fraud hypothesis, those mesas should lie in the high-turnout, high-Castillo-share corner, since the alleged manipulation was intended to shift the aggregate count in Castillo's favor. Their absence from that corner is a direct test of the specific factual claim made by Fujimori's legal team.

Figure 4 shows the Klimek fingerprint plots for both candidates. The scatter is organized by macro-region: Sierra mesas cluster in the upper-right for Castillo (high Castillo share,

moderate-to-high turnout) and in the lower-right for Fujimori; Costa mesas cluster in the lower-right for Castillo and upper-right for Fujimori; Selva mesas are dispersed between the two extremes.

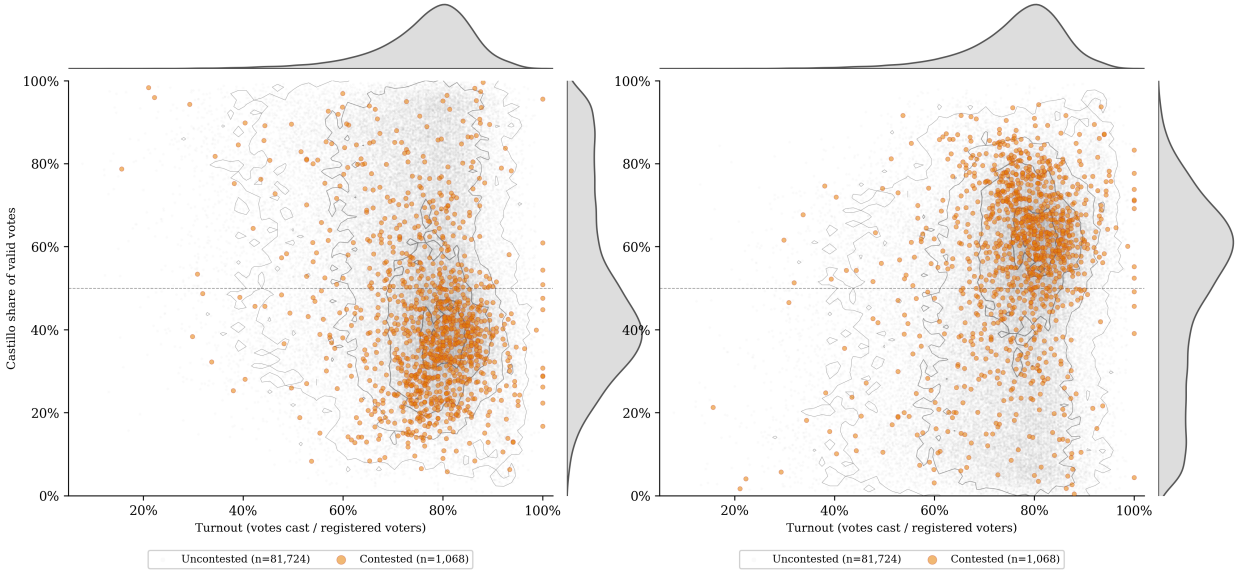

Figure 4: *Distributional fingerprint plots (Klimek et al., 2012). Left panel: Castillo vote share on vertical axis. Right panel: Fujimori vote share on vertical axis. Observations colored by macro-region (Costa = red, Sierra = blue, Selva = green). Orange dots are contested mesas (COMPUTADA RESUELTA). Density contours overlay the full distribution.*

Note: Each dot is one domestic polling station. The Sierra cluster in the upper-right of the Castillo panel reflects genuine geographic polarization, not ballot stuffing; it does not extend as an isolated finger beyond the main distribution. Orange (contested) dots show no concentration in any corner.

Source: ONPE (Oficina Nacional de Procesos Electorales, 2021b); author's calculations.

The key forensic feature is the absence of a "finger"—a protrusion extending toward the $(1.0, 1.0)$ corner that Klimek et al. (2012) identify as the ballot-stuffing signature in their Russian benchmarks. The upper-right region is not empty: there are polling stations with both high turnout (above 90 %) and high Castillo share (above 90 %). But these stations sit *within* the Sierra cluster, which reaches those values throughout. They are not a separate mode or protrusion; they are the right tail of the Sierra distribution. The contested mesas (orange dots) are scattered throughout the cloud with no visible concentration in any corner.

5.4 Last-digit tests

The last-digit test provides the most direct statistical assessment of whether vote totals were fabricated or rounded by human actors. When people invent numbers, they systematically avoid round endings and overuse certain digits, producing a distribution of final digits that departs from uniformity (Beber and Scacco, 2012). The test is particularly well-suited to polling-station data because the counts are small enough (typically 100–250 votes per candidate per mesa) that genuine results have no mechanical reason to cluster on round numbers.

I run four separate chi-squared tests: Castillo and Fujimori votes, each split between contested and uncontested mesas. The four-way breakdown is the key design feature. If the fraud allegations are correct—if vote totals in the challenged actas were altered by party operatives or electoral officials— then the last-digit distribution of the contested subsample should deviate from uniformity relative to the uncontested baseline. Both subsamples should also be compared against the expectation of 10 % per digit under genuine counting. Table 5 reports all four tests; the figure that follows displays the full digit-by-digit distributions.

For Castillo votes at uncontested mesas (n = 81,793), the test statistic is $\chi^2 = 12.44$ ($p = 0.190$). For Fujimori votes at uncontested mesas (n = 81,724, after dropping 69 mesas with zero Fujimori votes), $\chi^2 = 7.84$ ($p = 0.551$). The contested subsamples are smaller ($n \approx$ 1,070 each) and therefore have less statistical power, but both also pass cleanly: $\chi^2 = 4.05$ ($p = 0.908$) for Castillo and $\chi^2 = 4.66$ ($p = 0.863$) for Fujimori. The 0+5 digit share—the primary heaping indicator—ranges from 18.6 % to 20.5 % across the four groups, all within half a percentage point of the expected 20 %.

Table 5: *Last-Digit Uniformity Tests*

Group	n	χ^2 (9 d.f.)	p-value	Share digits 0+5
Castillo, uncontested	81,793	12.44	0.190	20.44 %
Castillo, contested	1,069	4.05	0.908	18.62 %
Fujimori, uncontested	81,724	7.84	0.551	20.11 %
Fujimori, contested	1,071	4.66	0.863	20.54 %
Expected under H_0	—	—	—	20.00 %

Notes: The null hypothesis is that last digits are uniformly distributed over $\{0, 1, \ldots, 9\}$. The test excludes observations with zero votes for the candidate (the last digit is undefined). Zero-exclusions: Castillo uncontested 0, Castillo contested 8, Fujimori uncontested 69, Fujimori contested 6. None of the four tests rejects uniformity at the 10 % level. The contested subsamples ($n \approx 1{,}070$) have approximately 25 % power to detect heaping at the level observed in documented fraud cases (Beber and Scacco, 2012); the uncontested tests have power exceeding 99 % for any practically meaningful departure from uniformity. *Source:* ONPE (Oficina Nacional de Procesos Electorales, 2021b); author's calculations.

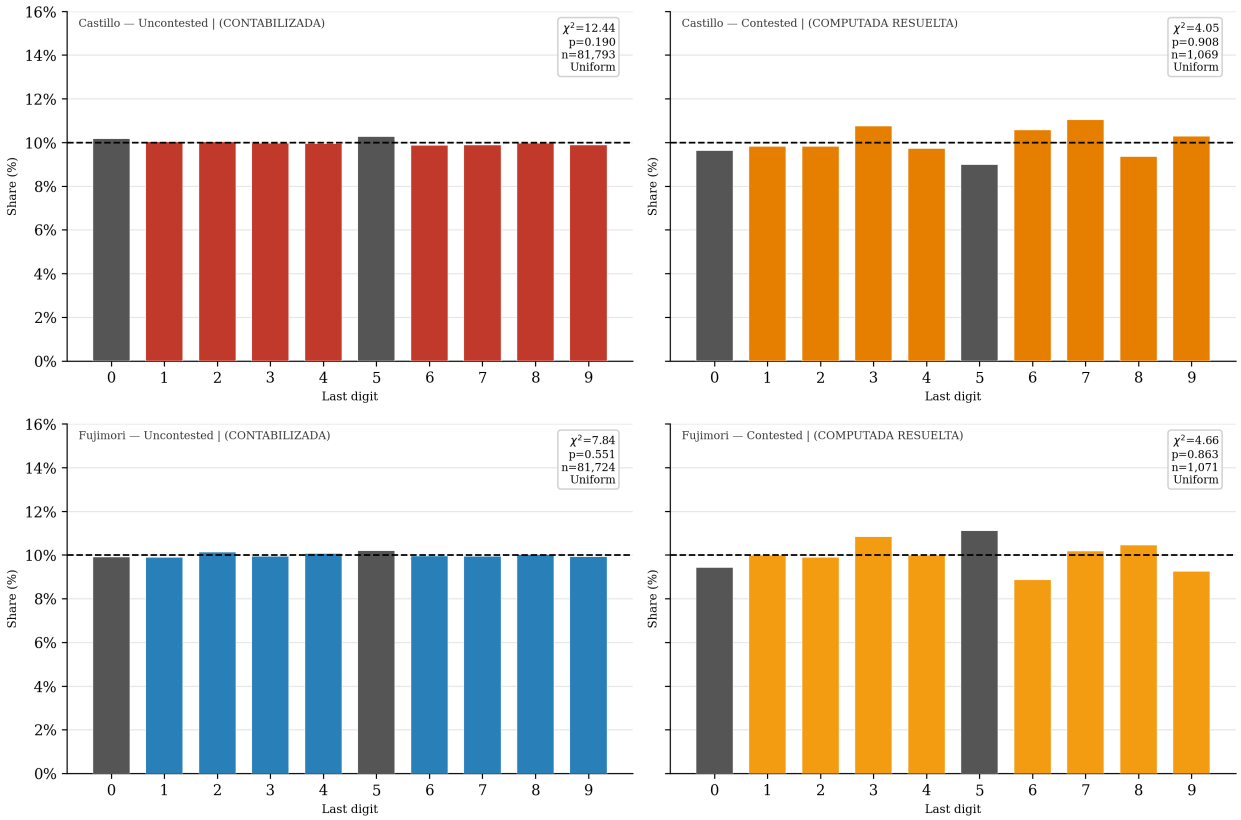

Figure 5: *Last-digit distributions for Castillo (left column) and Fujimori (right column) votes, separated by polling-station status. Top row: uncontested mesas (CONTABILIZADA). Bottom row: contested mesas (COMPUTADA RESUELTA). Dashed horizontal line at 10 % indicates the expected frequency under a uniform distribution.*

Note: Observations with zero votes for the candidate are excluded. Chi-squared p-values: Castillo uncontested 0.190; Castillo contested 0.908; Fujimori uncontested 0.551; Fujimori contested 0.863.

Source: ONPE (Oficina Nacional de Procesos Electorales, 2021b); author's calculations.

Figure 5 plots the full digit distributions. Each of the four panels shows a near-uniform bar chart with no discernible heaping at 0 or 5. The contested and uncontested distributions are visually indistinguishable. If vote totals in contested actas had been fabricated or altered by human actors, last-digit heaping would be the expected signature. No such pattern appears.

5.5 Ecological regression

The ecological regression exploits the two-round structure of the election to test whether contested mesas distort the district-level relationship between first- and second-round vote

shares. Under a genuine election, the fraction of mesas that were challenged and reviewed by the JNE should have no predictive power for Castillo's second-round performance once geographic preferences—captured by first-round results—are held constant. A positive coefficient on the share of contested mesas would suggest that those polling stations systematically boosted Castillo's total beyond what his local support base would predict.

Specification 1 establishes the baseline relationship. Peru Libre received a national first-round share of roughly 19 %, since that vote was divided among 18 candidates; the runoff forced consolidation around Castillo or Fujimori. The regression estimates how strongly a district's first-round preference for Peru Libre translated into second-round support for Castillo. Specification 2 then adds the share of a district's final-status mesas that are *computada resuelta* and its interaction with first-round share. Table 6 reports both specifications.

Specification 1 regresses Castillo's second-round district share on Peru Libre's first-round district share across $N = 1{,}873$ districts with matched data. The estimated relationship is:

Table 6: *Ecological Regression: First-Round and Second-Round Vote Shares at the District Level*

	Specification 1	Specification 2
	Eq. (1)	Eq. (2)
Constant	0.392***	0.393***
	(0.006)	(0.006)
Primera share (share_{1v})	0.817***	0.815***
	(0.012)	(0.013)
Share contested ($\text{share}_{\text{cont}}$)		−0.113
		(0.253)
Primera × contested		0.231
		(0.614)
N	1,873	1,873
R^2	0.708	0.708

Notes: Dependent variable: Castillo valid-vote share in the segunda vuelta, aggregated to the district level. Primera share is Peru Libre's valid-vote share in the primera vuelta (candidate 16, VOTOS_P16). Share contested is the fraction of a district's final-status segunda vuelta mesas that are *computada resuelta*. HC0-robust standard errors in parentheses. *** $p < 0.001$. Sample restricted to districts with at least 10 valid votes in the segunda vuelta and positive Peru Libre votes in the primera vuelta. Districts with at least one contested mesa: 314 (16.8 % of sample). *Source:* ONPE (Oficina Nacional de Procesos Electorales, 2021b,a); author's calculations.

$$\widehat{\text{share}}_{2v} = 0.392 + 0.817 \times \text{share}_{1v}, \quad R^2 = 0.708.$$

The slope of 0.817 is less than 1: in districts where Peru Libre was strong in the first round, Castillo gained proportionally but not one-for-one, reflecting the fact that many additional votes came from former supporters of other first-round candidates rather than

simple mobilization of the existing Peru Libre base. The intercept of 0.39 confirms that Castillo drew substantial anti-Fujimori votes even in districts where Peru Libre had minimal first-round presence. The R^2 of 0.71 indicates that first-round geography explains most of the variation in second-round outcomes.

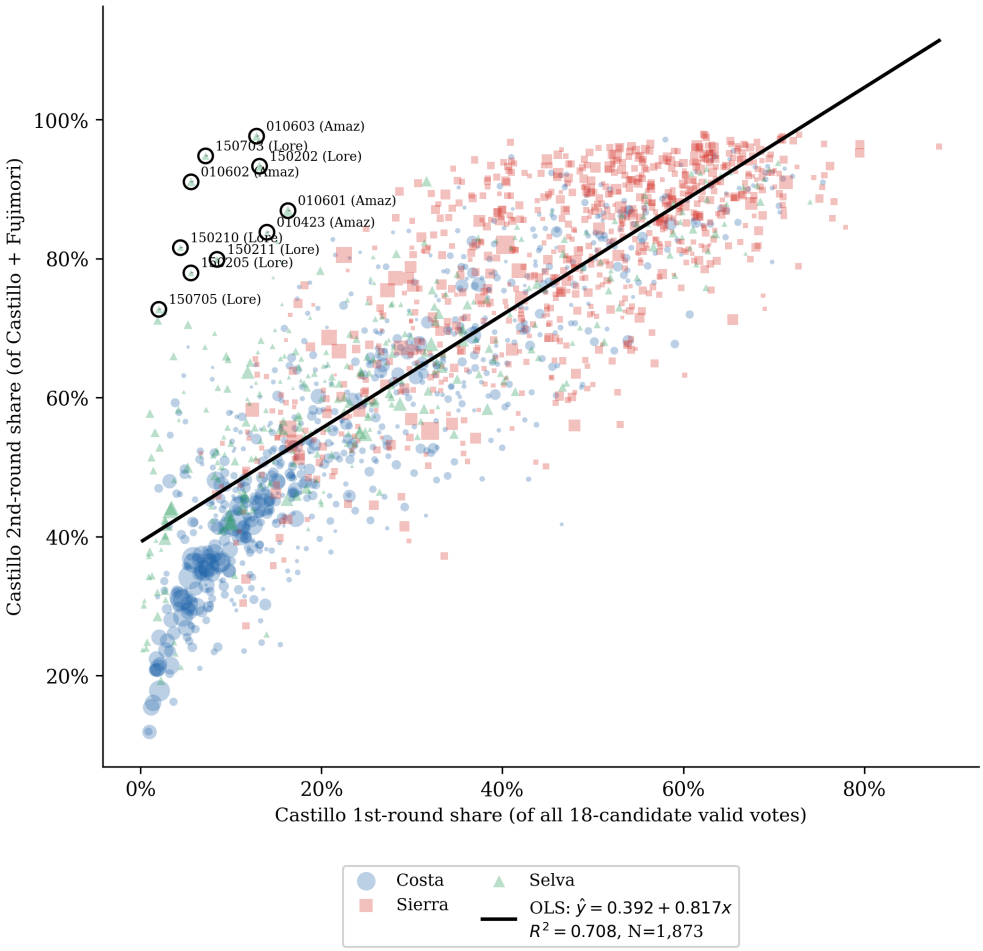

Figure 6: *First-round (primera vuelta) vs. second-round (segunda vuelta) Castillo/Peru Libre vote share at the district level. Each point is one of 1,873 districts. The fitted line is from Specification 1 of Table 6. Orange points are districts with any contested (COMPUTADA RESUELTA) mesas.*

Note: Districts with fewer than 10 valid segunda vuelta votes or zero primera vuelta Peru Libre votes are excluded. The top-10 positive-residual districts are in Loreto and Amazonas, which have high geographic isolation; none are in the departments identified in Fujimori's fraud allegations.

Source: ONPE (Oficina Nacional de Procesos Electorales, 2021b,a); author's calculations.

Figure 6 displays the district-level scatter with the fitted regression line. The relationship

is tight and linear, with no obvious outliers or breaks. The districts with the largest positive residuals (where Castillo most over-performed relative to first-round prediction) are concentrated in Loreto and Amazonas departments—remote Amazonian areas with historically variable turnout and high geographic isolation, not departments identified in the fraud allegations.

Specification 2 adds the district share of contested mesas (share_cont_d) and its interaction with first-round share. Neither coefficient is statistically significant: $\hat{\gamma} = -0.113$ (SE $= 0.253$, $p = 0.656$) and $\hat{\delta} = 0.231$ (SE $= 0.614$, $p = 0.707$). The R^2 is unchanged at 0.708. Districts with a higher concentration of contested mesas do not over-perform for Castillo after conditioning on first-round baseline.

Two features of Specification 2 deserve emphasis. First, the point estimate $\hat{\gamma} = -0.113$ is negative: if anything, a higher share of contested mesas is associated with a *lower* Castillo second-round share, conditional on first-round support. This is the opposite direction from what ballot-stuffing fraud would predict. The estimate is far from conventional significance thresholds, so no causal interpretation is warranted, but the sign alone rules out a large, positive fraud effect of the kind alleged.

Second, the variable share_cont_d is not exogenous: Fujimori's legal team selected which actas to challenge, and that selection was presumably driven by their own assessment of which mesas contained exploitable procedural errors. This endogeneity does not undermine the forensic conclusion—it strengthens it. Fujimori's representatives chose the actas they believed were most irregular; even in those districts, there is no detectable Castillo over-performance. The selection bias works in favor of finding a positive coefficient, yet the estimate is negative. An instrumental-variables correction would, if anything, move $\hat{\gamma}$ further from the fraud hypothesis.

Appendix A extends this analysis to the mesa level. Matching each polling station across both rounds yields a panel of 79,672 mesas, of which 1,004 are contested. The mean swing from Peru Libre's first-round share to Castillo's second-round share is +31.3 pp—exactly

what binary consolidation of anti-Fujimori votes predicts mechanically. Contested mesas exhibit a mean swing of +29.3 pp, *lower* than the +31.3 pp uncontested average (Welch $t = -6.0$, $p = 2.6 \times 10^{-9}$). The mesa-level OLS residual is -0.020 for contested mesas, confirming that the fraud hypothesis—which requires a positive residual in precisely those mesas—is directly rejected.

5.6 Turnout patterns

Figure 7 shows the turnout distributions for both rounds. Average domestic turnout in the segunda vuelta was 76.1 %, down from the primera vuelta. This directional finding runs against what ballot stuffing would produce: adding fabricated votes to tallies would inflate recorded turnout, generating a distribution shifted toward the high end. The segunda vuelta distribution is unimodal and roughly symmetric, with no secondary peak near 100 % that would suggest mass turnout manipulation.

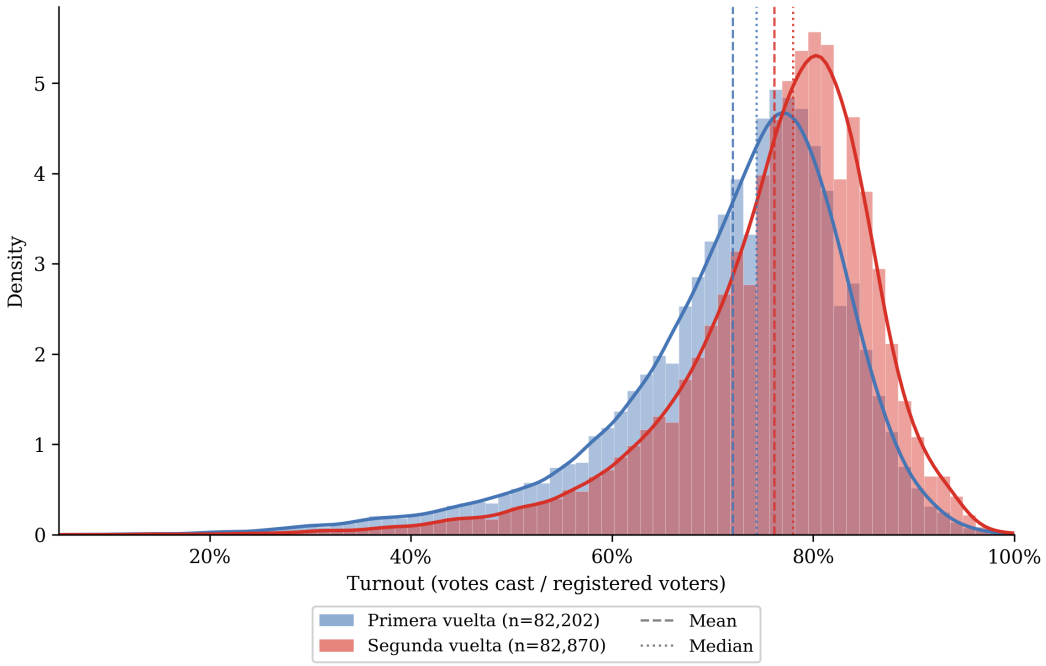

Figure 7: *Turnout distribution by polling station, primera vuelta (April 2021) and segunda vuelta (June 2021). Vertical lines mark the respective medians.*

Note: Domestic final-status samples for each round, with turnout clipped at 1.05 for display purposes.

Source: ONPE (Oficina Nacional de Procesos Electorales, 2021b,a); author's calculations.

The anomaly flag counts in Table 2 are consistent with this picture. Zero mesas exceed 100 % turnout. Only 2 mesas fall below 10 % turnout (both are small rural mesas with a single registered-voter discrepancy). There are 57 mesas with turnout above 98 %—none show the accompanying perfect or near-perfect vote shares that would mark ballot stuffing.

The cross-round comparison in Figure 7 is also informative about the plausibility of large-scale fraud. The segunda vuelta attracted roughly the same absolute number of voters as the primera vuelta, but the registered electorate had grown slightly, producing a marginally lower official turnout rate. Had a substantial number of fraudulent votes been added to the segunda vuelta count, the turnout distribution would shift rightward relative to the primera vuelta. The distributions are nearly superimposed, with the segunda vuelta median slightly lower. This directional evidence is inconsistent with net ballot stuffing at scale.

5.7 Null vote geography

Null votes—ballots cast but invalidated at the polling station during counting—offer a complementary window into potential manipulation. If Castillo's opponents or complicit electoral officials systematically invalidated valid pro-Castillo ballots by marking them as null, the null rate would be elevated in Castillo's strongholds: the southern highland departments of Puno, Ayacucho, Cusco, Apurímac, and Huancavelica. This form of fraud is distinct from ballot stuffing or vote-count alteration; it requires collusion at the mesa level and would leave a geographic fingerprint precisely where Castillo received his largest majorities.

Nationally, 1,043,651 null votes were cast, representing 5.66 % of total domestic votes (Table 1). The provincial distribution of this rate is mapped in Figure 8. Three features are forensically relevant. First, the highest null rates are concentrated in Lima's metropolitan provinces, the northern coastal provinces of Piura and La Libertad, and the Lambayeque region—all Fujimori strongholds, not Castillo's. Second, Puno, the department with the highest Castillo share (89.3 %), has a null rate near or below the national average. Third, the 319 mesas flagged for null rates above 30 % in Table 2 are geographically clustered on the coast, where protest voting against both candidates was highest among urban, educated voters. None of these patterns is consistent with targeted invalidation of pro-Castillo votes.

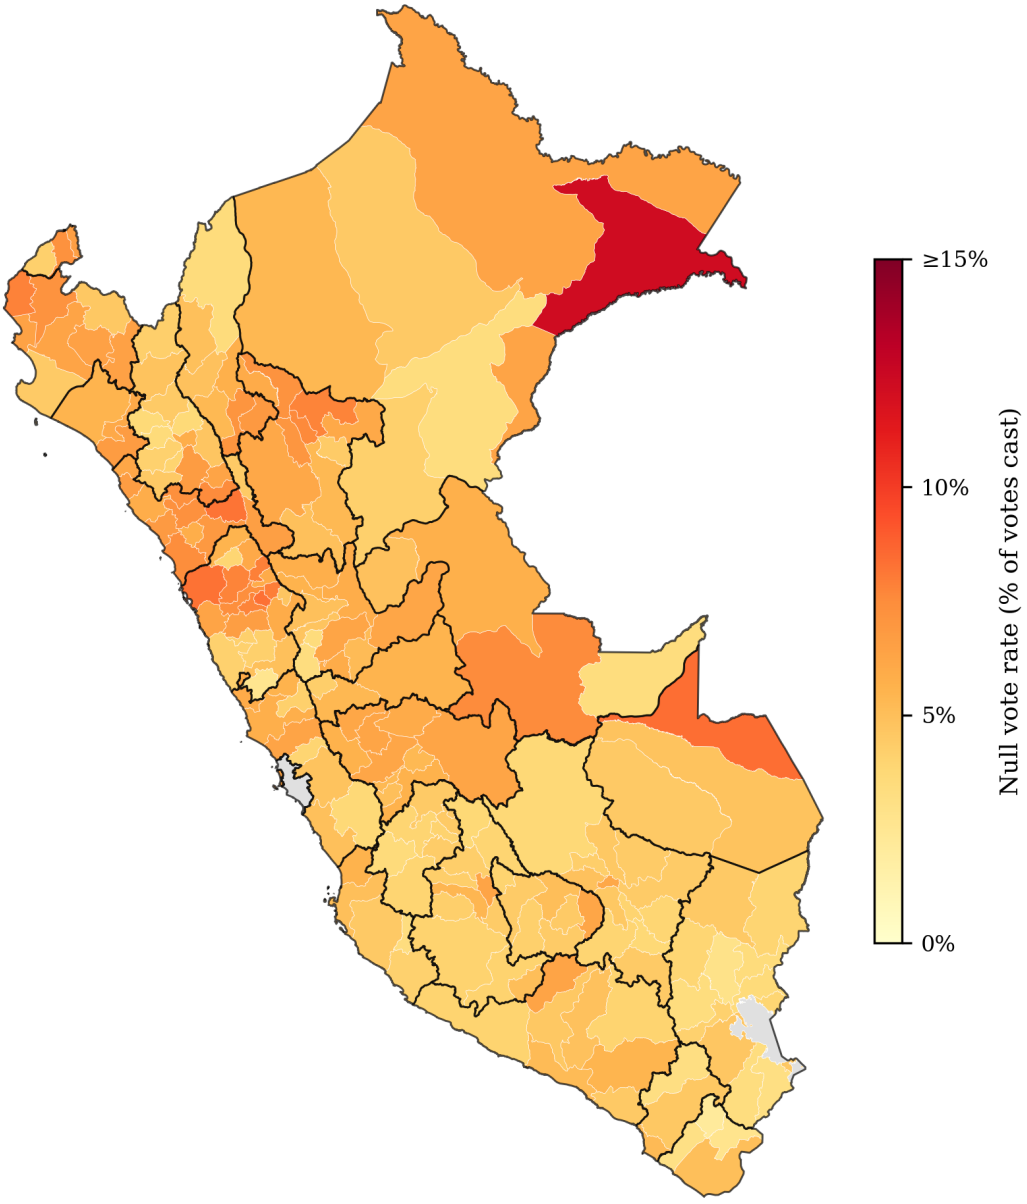

Figure 8: *Null vote rate by province (null votes as a share of total votes cast), segunda vuelta 2021. Darker shading indicates higher null rates.*

Note: Province-level aggregates from the domestic final-status sample. Null votes (VOTOS_VN) include all ballots invalidated at the polling station during counting. *Source:* ONPE (Oficina Nacional de Procesos Electorales, 2021b); GADM 4.1 shapefiles; author's calculations.

High null rates concentrate in coastal and metropolitan areas: Lima (the capital), Piura,

La Libertad, and Lambayeque all show elevated null vote rates. These are Fujimori's strongest departments. If fraud had taken the form of systematic invalidation of pro-Castillo votes, null rates would be elevated in Castillo's highland strongholds. The observed pattern points instead to protest voting by urban Limeños who disliked both candidates—a plausible behavioral response given the polarizing choice on offer.

6 Discussion

6.1 Why fraud allegations gained traction

Three factors explain why Fujimori's fraud narrative resonated despite the absence of supporting evidence. First, the margin was historically narrow. At 44,263 votes out of 17 million cast, any irregularity in a few hundred polling stations could, in principle, have been decisive. Close elections generate fraud allegations in many countries (Lehoucq, 2003), and Peru was no exception.

Second, the geographic pattern of results appeared anomalous to urban observers unfamiliar with the political geography of the Peruvian sierra. Vote shares of 80–90 % for a single candidate in polling stations with near-universal turnout look like fraud to someone who has never visited these regions and does not know their voting history. In Andean communities with strong collective identities and long histories of political mobilization, near-unanimous results are not unusual and do not require manipulation.

The mesa-level swing analysis in Appendix A bears directly on this perception. Of the 116 mesas recording a swing of ≥ 80 percentage points between rounds, only 2 are contested. Lima and Amazonas account for 88 % of extreme-swing mesas—not the southern highland departments of Puno, Ayacucho, or Cusco that anchored the fraud allegations. These extreme swings reflect high anti-Fujimori consolidation from former Sagasti and Toledo voters who had near-zero Peru Libre support in the first round, not manipulation: the mesas with the most striking two-round swings were almost entirely unchallenged by Fujimori's own legal

team.

Third, and most consequential for the fraud narrative, the running count shifted toward Castillo over several days as highland actas arrived later than urban ones. This is the "early-count mirage" documented by Idrobo et al. (2022) for Bolivia 2019: urban precincts count faster, early returns favor the urban candidate (Fujimori), and as rural results arrive the lead shifts. The shift is a mechanical consequence of differential reporting speeds in a geographically polarized electorate. It is not suspicious, but it is visible, and it was interpreted by Fujimori's supporters as evidence that votes were being added to Castillo's total after the initial count.

A fourth factor was Fujimori's personal legal exposure. She faced ongoing criminal prosecution for corruption and illegal campaign financing at the time of the election. A Castillo presidency increased the probability that she would be convicted and imprisoned. This gave her a rational personal incentive to contest the result beyond what the evidence warranted, independent of any genuine belief that fraud had occurred.

6.2 Comparison with Bolivia 2019

The Bolivia 2019 election offers the closest methodological parallel. In Bolivia, the OAS alleged fraud based on a discontinuity in the TREP (preliminary vote-count) transmission: the real-time count was suspended for roughly 24 hours and resumed with a composition that favored Evo Morales. Escobari and Hoover (2024) argue that this discontinuity was real and consistent with manipulation. Idrobo et al. (2022) counter that the late-arriving vote share in the resumed count was nearly identical to the composition of the geographic areas still unreported at suspension time, and that no manipulation is needed to explain the shift.

Peru 2021 differs from Bolivia 2019 in one structural respect: there was no counting halt. The ONPE system reported results continuously and in real time. The shift toward Castillo over several days was gradual and uninterrupted, not a jump around an unexplained stoppage. The regression-discontinuity strategy that anchors the Bolivia debate has no application in

Peru.

The temporal test that would provide the strongest identification—regressing vote shares on processing order within the count, as in Escobari and Hoover (2024)— cannot be implemented for Peru 2021 because ONPE did not publish processing timestamps. This is the primary data limitation of the present analysis. If ONPE were to release the internal log recording when each acta was processed, the temporal analysis could be carried out and would substantially strengthen the conclusions. As matters stand, the available evidence from cross-sectional forensic methods uniformly fails to detect manipulation.

6.3 Limitations

Three limitations apply. First, the absence of processing timestamps precludes temporal analysis. This is the strongest available identification strategy in the forensics literature (Escobari and Hoover, 2024) and its unavailability is a genuine gap.

Second, forensic tests are designed to detect specific statistical signatures. Absence of anomalies does not prove absence of fraud. A sophisticated and well-coordinated manipulation that altered vote totals in statistically consistent ways—maintaining uniform digit distributions, not exceeding 100 % turnout, preserving first-to-second-round relationships—would not be detected by any of the tests in this paper. The appropriate inference is that the data are inconsistent with the specific, unsophisticated forms of manipulation that were alleged.

Third, the analysis covers recorded vote totals only. It cannot speak to pre-electoral irregularities—voter intimidation, vote buying, biased access to media—which require different data and methods (Callen and Long, 2015).

6.4 Implications

For Peru's electoral institutions: ONPE should publish processing timestamps with future election data. This is a low-cost improvement that would enable temporal forensic analysis and allow electoral authorities to preempt fraud narratives in real time. The timestamps

presumably exist in ONPE's internal systems; publishing them alongside the vote counts would transform the evidentiary basis available to independent researchers.

For the election forensics literature: geographic polarization is a first-order empirical challenge for distributional methods. Klimek et al. (2012) developed their fingerprint method using Russia and Austria, neither of which exhibits the degree of regional sorting present in Peru. When a country has provinces voting 89 % for one candidate and 67 % for the other, the fingerprint scatter will show regional clusters in the corners that a naive application of the method might flag as fraud. Conditioning on region is essential, and even then, the method distinguishes polarization from ballot stuffing only imperfectly. Combining multiple methods—as this paper does— provides a more credible basis for inference.

For democratic governance: the gap between forensic evidence and public perception is a problem that statistics cannot solve on its own. Multiple independent tests find no manipulation in Peru 2021; the OAS and EU missions found no irregularities; the JNE reviewed each challenged acta individually and certified the result. Yet a substantial fraction of the electorate believed the election was stolen. This pattern—unfounded fraud allegations undermining democratic legitimacy in close, polarized elections—has recurred across different institutional contexts: the United States in 2020 (Eggers et al., 2021), Bolivia in 2019 (Idrobo et al., 2022; Escobari and Hoover, 2024), and Kenya in 2017. Rapid communication of forensic evidence by credible institutions, produced quickly enough to compete with the fraud narrative as it forms, is an institutional response worth developing.

7 Conclusion

I apply four independent forensic methods to the complete universe of polling-station records from Peru's 2021 presidential runoff and find no evidence of vote manipulation. Turnout is below 100 % everywhere. Last-digit distributions pass chi-squared uniformity tests for all four candidate-by-status groups, with p-values above 0.19. The ecological regression finds

that districts with higher concentrations of contested actas do not over-perform for Castillo after controlling for first-round baseline ($p = 0.71$). The contested actas—the specific subset targeted by Fujimori's legal team— are concentrated in Fujimori's coastal strongholds, favor Fujimori by a wide margin, and show forensic signatures identical to uncontested actas.

The paper makes three contributions. It provides the first comprehensive mesa-level forensic analysis of Peru 2021, improving on the state-level Benford tests of Isea and Isea (2021) in both statistical power and methodological range. It introduces the contested/uncontested institutional distinction as a within-election forensic diagnostic—a strategy that uses the fraud allegations themselves to define the treatment group. And it documents how extreme geographic polarization can generate distributional patterns that appear anomalous to non-specialists while being entirely consistent with genuine political preferences.

One data improvement would substantially strengthen future analyses of Peruvian elections: processing timestamps for each acta, recording when each polling-station tally was uploaded to the ONPE system. These data likely already exist within ONPE's infrastructure. Publishing them as a standard part of the post-election data release would enable the temporal identification strategy that provides the sharpest available test and would give electoral institutions a powerful real-time tool for responding to fraud allegations before they take hold.

References

Beber, B. and Scacco, A. (2012). What the numbers say: A digit-based test for election fraud. *Political Analysis*, 20(2):211–234.

Callen, M. and Long, J. D. (2015). Institutional corruption and election fraud: Evidence from a field experiment in Afghanistan. *American Political Science Review*, 109(3):468–484.

Deckert, J., Myagkov, M., and Ordeshook, P. C. (2011). Benford's Law and the Detection of Election Fraud. *Political Analysis*, 19(3):245–268.

Eggers, A. C., Garro, H., and Grimmer, J. (2021). No evidence for systematic voter fraud: A guide to statistical claims about the 2020 election. *Proceedings of the National Academy of Sciences*, 118(45):e2103619118.

Escobari, D. and Hoover, G. A. (2024). Late-arriving votes and electoral fraud: A natural experiment and regression discontinuity evidence from Bolivia. *World Development*, 173:106418.

European Union Election Observation Mission (2021). Final report: Peru presidential and congressional elections 2021. Technical report, European External Action Service, Brussels.

Hyde, S. D. (2007). The observer effect in international politics: Evidence from a natural experiment. *World Politics*, 60(1):37–63.

Idrobo, N., Kronick, D., and Rodríguez, F. (2022). Do shifts in late-counted votes signal fraud? evidence from bolivia. *Journal of Politics*, 84(2):1128–1132.

Isea, J. and Isea, R. (2021). Was there voter fraud in the 2021 Peru presidential elections? *Journal of Model Based Research*, 1(4):1–5.

Jurado Nacional de Elecciones (2021). Resoluciones sobre actas observadas e impugnadas, elecciones generales 2021 segunda vuelta presidencial. Technical report, Jurado Nacional de Elecciones del Perú, Lima, Perú.

Klimek, P., Yegorov, Y., Hanel, R., and Thurner, S. (2012). Statistical detection of systematic election irregularities. *Proceedings of the National Academy of Sciences*, 109(41):16469–16473.

Kobak, D., Shpilkin, S., and Pshenichnikov, M. S. (2016). Integer percentages as electoral falsification fingerprints. *Annals of Applied Statistics*, 10(1):54–73.

Lehoucq, F. (2003). Electoral fraud: Causes, types, and consequences. *Annual Review of Political Science*, 6:233–256.

Mainwaring, S., editor (2018). *Party Systems in Latin America: Institutionalization, Decay, and Collapse*. Cambridge University Press, Cambridge.

Myagkov, M., Ordeshook, P. C., and Shakin, D. (2009). *The Forensics of Election Fraud: Russia and Ukraine*. Cambridge University Press, Cambridge.

Oficina Nacional de Procesos Electorales (2021a). Resultados por mesa de las Elecciones Presidenciales 2021 Primera Vuelta. Datos abiertos, https://datosabiertos.gob.pe. File: Resultados_1ra_vuelta_Version_PCM.csv.

Oficina Nacional de Procesos Electorales (2021b). Resultados por mesa de las Elecciones Presidenciales 2021 Segunda Vuelta. Datos abiertos, https://datosabiertos.gob.pe. File: Resultados_2da_vuelta_Version_PCM.csv.

Organization of American States (2021). Preliminary report of the OAS electoral observation mission to the general elections of Peru, second round. Technical report, OAS Electoral Observation Mission, Washington, D.C.

A Mesa-Level Swing Analysis

This appendix extends the district-level ecological regression in Section 5 to the polling-station (mesa) level. The swing analysis is not the primary forensic test in this paper— that role belongs to the contested/uncontested comparison in Table 3 and Table 6—but it provides a complementary and more granular examination of whether contested mesas deviate from expected vote patterns.

A.1 Data Construction

I match the primera vuelta (April 2021) and segunda vuelta (June 2021) polling-station files released by ONPE, joining on `MESA_DE_VOTACION`. Both rounds are filtered to final-status mesas (*contabilizada* and *computada resuelta*) in domestic jurisdictions, following the same sample restrictions applied in the main analysis.

For each mesa i, I define:

$$s_i^{1v} = \frac{\texttt{VOTOS\_P16}_i}{\texttt{N\_CVAS}_i - \texttt{VOTOS\_VN}_i - \texttt{VOTOS\_VB}_i} \tag{3}$$

$$s_i^{2v} = \frac{\texttt{VOTOS\_P1}_i}{\texttt{VOTOS\_P1}_i + \texttt{VOTOS\_P2}_i} \tag{4}$$

$$\text{swing}_i = s_i^{2v} - s_i^{1v} \tag{5}$$

where s_i^{1v} is Peru Libre's share of valid first-round votes and s_i^{2v} is Castillo's share of valid second-round votes. The denominator in equation (4) uses only Castillo and Fujimori votes because the second round was a strictly binary contest with no other valid candidates.

The matched panel contains $N = 79{,}672$ polling stations, of which 1,004 (1.26 %) are *computada resuelta* (contested). Mesas with missing vote counts or shares outside $[0, 1]$ in either round are excluded (less than 0.5 % of matched observations).

A.2 Swing Distribution

The expected swing under pure binary consolidation is approximately +31 pp: Peru Libre's national first-round share was 19.1 %, and Castillo won 50.4 % of second-round valid votes, implying a mechanical consolidation of roughly 31 percentage points from former supporters of other first-round candidates.

Figure A1 shows the full distribution of mesa-level swing. The distribution is unimodal, approximately normal, and centered very close to the mechanically predicted value. Negative swings are possible at individual mesas where, for example, first-round Peru Libre support was concentrated but Castillo won fewer second-round votes than Peru Libre received in the first round—a pattern consistent with some Peru Libre voters defecting to Fujimori in the runoff.

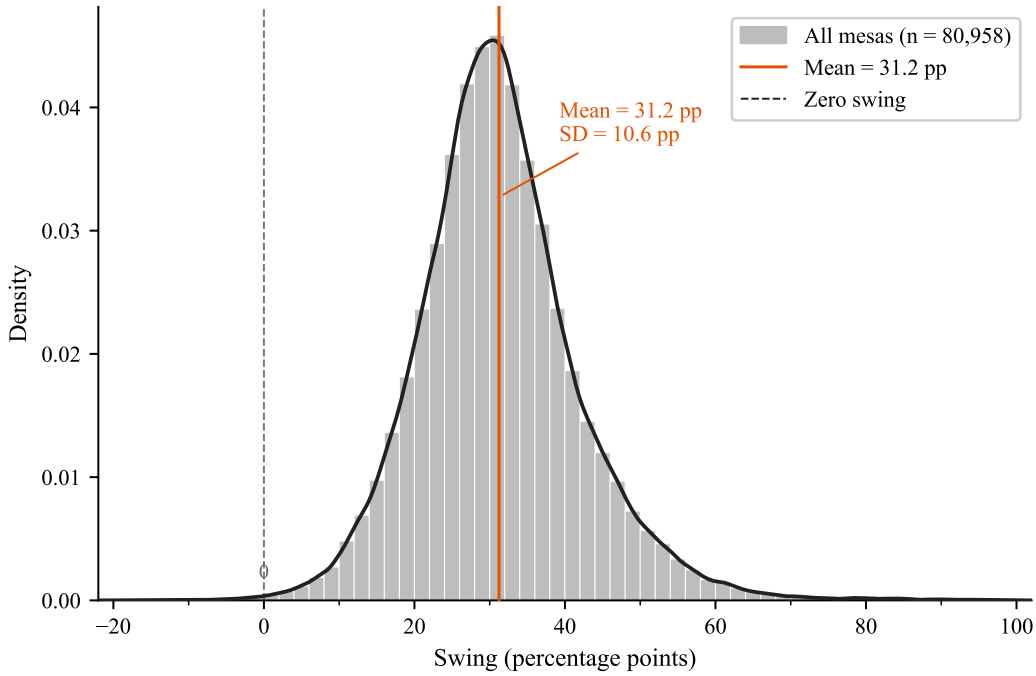

Figure A1: *Distribution of mesa-level swing between rounds. Swing is defined as Castillo's valid-vote share in the second round minus Peru Libre's valid-vote share in the first round, expressed in percentage points. The histogram uses 2 pp bins. The solid orange vertical line marks the sample mean (+31.3 pp); the dashed black line marks zero swing. The curve is a Gaussian kernel density estimate. Source: ONPE (Oficina Nacional de Procesos Electorales, 2021b,a); author's calculations.*

Table A1 reports the summary statistics. The mean (+31.3 pp) and median (+30.6 pp) are nearly identical, confirming approximate symmetry. The 10th percentile is +19.0 pp and the 90th is +44.4 pp, indicating that the overwhelming majority of polling stations—even those with below-median swings—show a positive consolidation shift.

Table A1: *Summary Statistics: Mesa-Level Swing*

Statistic	Value
Mean	+31.3 pp
Median	+30.6 pp
Std. deviation	10.5 pp
10th percentile	+19.0 pp
90th percentile	+44.4 pp
Minimum	−18.6 pp
Maximum	+98.2 pp
N	79,672

Notes: Unit of observation: domestic final-status polling station matched across both rounds. Swing = Castillo second-round valid-vote share minus Peru Libre first-round valid-vote share, in percentage points. Negative values indicate polling stations where Peru Libre's first-round share exceeded Castillo's second-round share. *Source:* ONPE (Oficina Nacional de Procesos Electorales, 2021b,a); author's calculations.

A.3 Contested versus Uncontested Mesas

The fraud hypothesis has a specific implication for the swing distribution: contested mesas—those specifically selected by Fujimori's representatives as allegedly irregular—should have *higher* swings toward Castillo than comparable uncontested mesas if their totals were manipulated upward. Figure A2 and Figure A3 test this directly.

Figure A2 plots each mesa's first-round Peru Libre share against its second-round Castillo share. Contested mesas (orange) are overlaid on the cloud of uncontested mesas (grey).

Visually, contested mesas are scattered throughout the distribution with no concentration in the upper-right corner (high first-round and high second-round share) that the fraud hypothesis would predict.

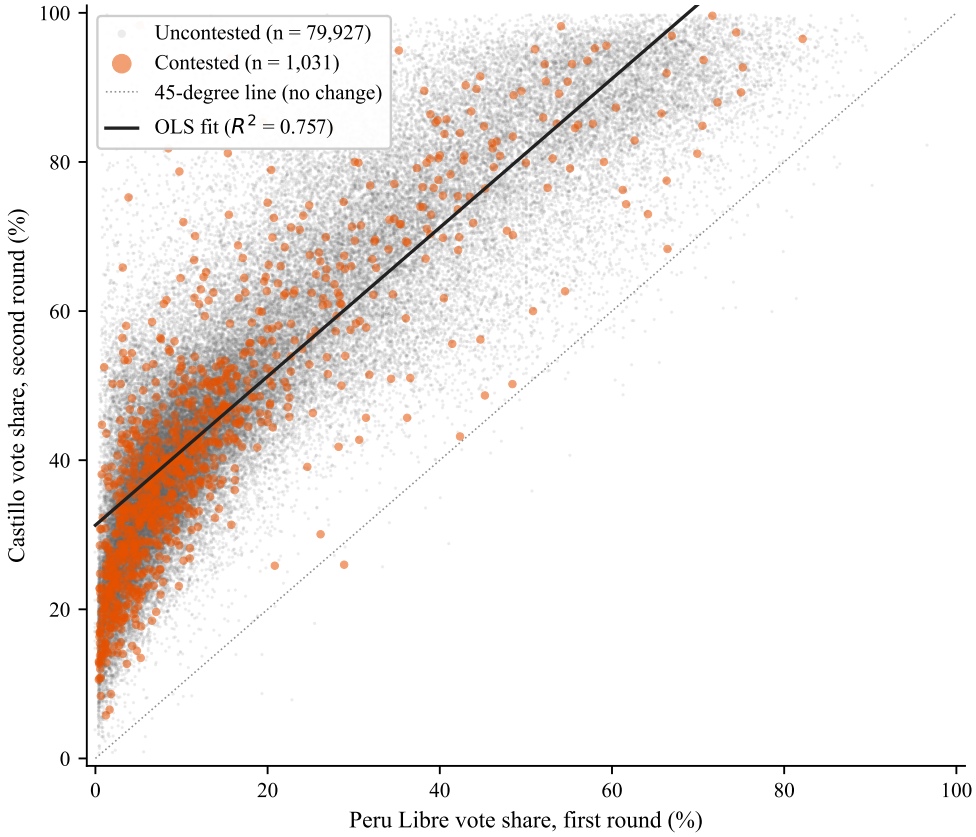

Figure A2: *Mesa-level vote shares: first round vs. second round. Each point is one of 79,672 matched polling stations. Grey points are uncontested (*contabilizada*, $n = 78{,}668$); orange points are contested (*computada resuelta*, $n = 1{,}004$). The dotted diagonal is the 45-degree line (no change between rounds). The solid line is the OLS fit from a regression of second-round Castillo share on first-round Peru Libre share ($R^2 = 0.759$). Source: ONPE (Oficina Nacional de Procesos Electorales, 2021b,a); author's calculations.*

Figure A3 shows the full swing distributions separately for uncontested and contested mesas. The contested distribution is shifted *leftward*: contested mesas have a lower mean swing (29.3 pp) than uncontested mesas (31.3 pp), a difference of -2.0 pp that is highly statistically significant (Welch $t = -6.0$, $p = 2.6 \times 10^{-9}$).

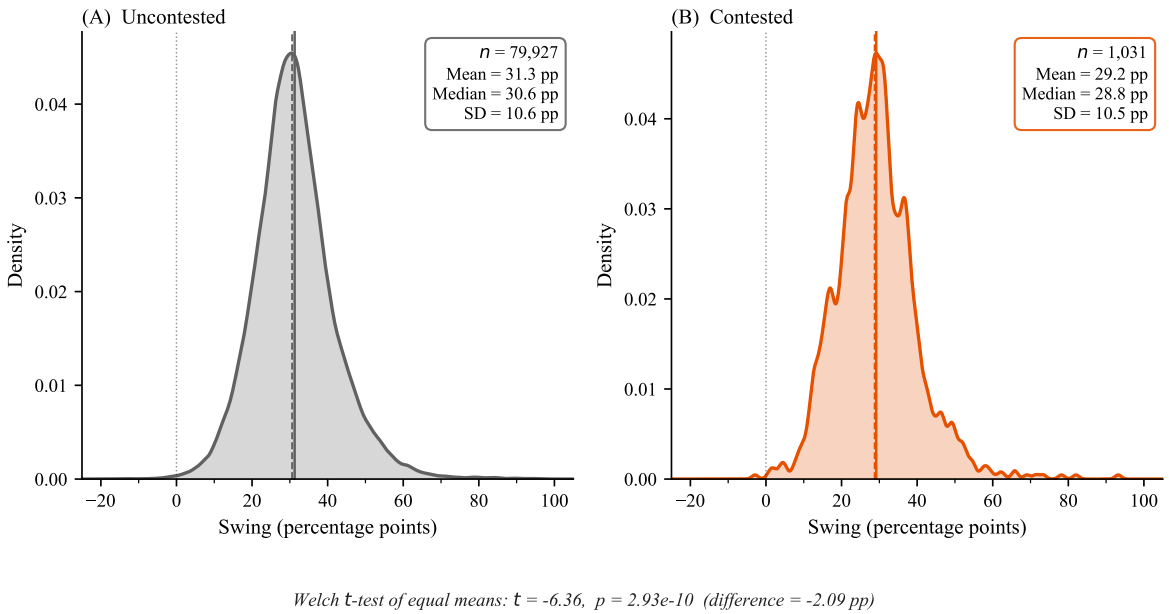

Figure A3: *Distribution of mesa-level swing by contestation status. Panel (A) shows the 78,668 uncontested mesas; Panel (B) shows the 1,004 contested mesas. Solid vertical lines mark the group mean; dashed lines mark the median. Contested mesas average +29.3 pp, compared with +31.3 pp for uncontested mesas. Source: ONPE (Oficina Nacional de Procesos Electorales, 2021b,a); author's calculations.*

Table A2 reports the comparison numerically, including mean residuals from the mesa-level OLS regression $s_i^{2v} = \alpha + \beta\, s_i^{1v} + \varepsilon_i$.

Table A2: *Swing and OLS Residuals by Contestation Status*

Group	N	Mean swing	Median swing	Mean residual
Uncontested	78,668	+31.3 pp	+30.6 pp	+0.000
Contested	1,004	+29.3 pp	+28.8 pp	−0.020
Difference	—	−2.0 pp	−1.8 pp	−0.020

Notes: Mean residual from mesa-level OLS regression of Castillo second-round share on Peru Libre first-round share ($R^2 = 0.759$; $N = 79{,}672$). A negative residual indicates that the mesa's Castillo second-round share is *below* what its first-round geography predicts. Welch two-sample t-test of equal means: $t = -6.0$, $p = 2.6 \times 10^{-9}$. KS test: $D = 0.095$, $p = 3.4 \times 10^{-8}$. *Source:* ONPE (Oficina Nacional de Procesos Electorales, 2021b,a); author's calculations.

The direction and magnitude of the contested–uncontested gap is important to interpret correctly. Fujimori's legal team *selected* the contested mesas—they chose actas they believed contained procedural errors that favored Castillo. Under the fraud hypothesis, these mesas should have *inflated* Castillo totals; the residual should be positive. Instead, contested mesas have a mean residual of -0.020: Castillo under-performs his geographic baseline by 2 percentage points in precisely the mesas that Fujimori's team identified as suspicious. This pattern is inconsistent with the fraud allegations and consistent with random selection noise in the challenge process.

A.4 Extreme Swing Analysis

A natural question is whether the upper tail of the swing distribution— mesas with very large positive swings—contains contested polling stations. If fraud concentrated in a subset of extreme-swing mesas, those mesas should appear disproportionately in the challenged sample.

Figure A4 shows the top-10 departments by count of mesas with swing ≥ 80 pp. Table A3 summarizes two anomaly criteria.

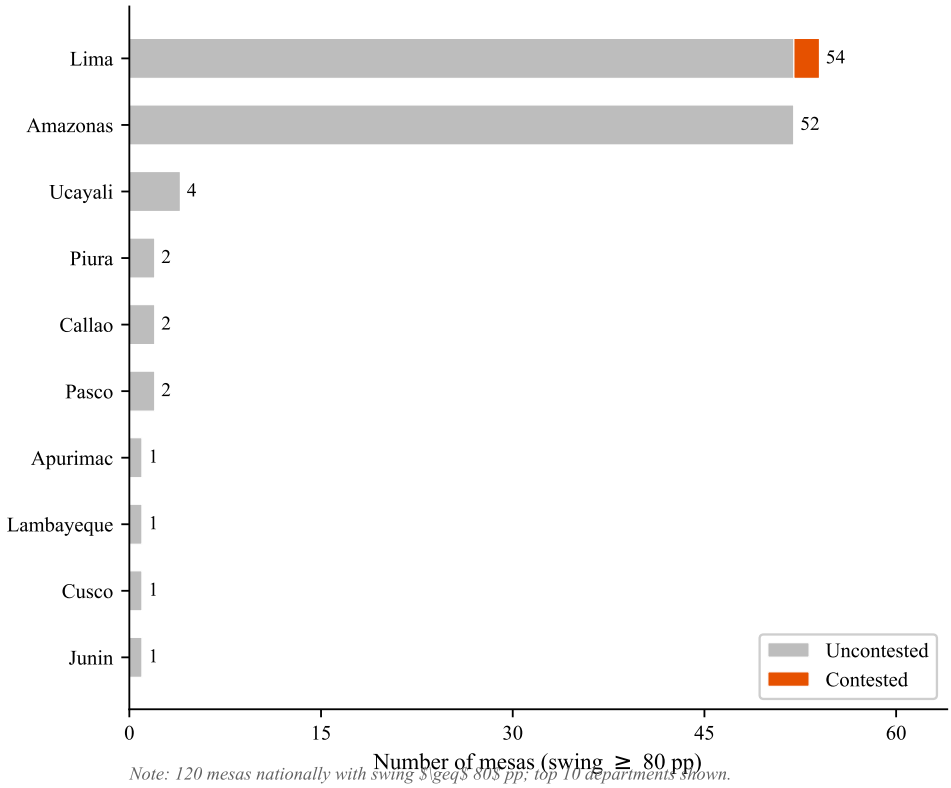

Note: 120 mesas nationally with swing $\geq$ 80 pp; top 10 departments shown.

Figure A4: *Departments with the most mesas recording a swing of* ≥ 80 *percentage points (top 10 of 25 departments). Grey bars are uncontested; orange bars are contested. Of 116 mesas nationwide with swing* ≥ 80 *pp, only 2 are contested. Source: ONPE (Oficina Nacional de Procesos Electorales, 2021b,a); author's calculations.*

Table A3: *Extreme-Swing Mesas: Two Anomaly Criteria*

Criterion	Total	Contested	% Contested	Top departments
PL $< 10\,\%$ in R1 & Castillo $> 90\,\%$ in R2	55	1	1.8 %	Lima (28), Amazonas (26)
Swing ≥ 80 pp	116	2	1.7 %	Lima (52), Amazonas (50)

Notes: "PL $< 10\,\%$ in R1" denotes mesas where Peru Libre's first-round valid-vote share was below 10 %; "Castillo $> 90\,\%$ in R2" denotes mesas where Castillo's second-round valid-vote share exceeded 90 %. These criteria are designed to identify the mesas most consistent with the specific fraud allegation: that previously non-Castillo mesas were altered to show high Castillo support. The share of contested mesas in both extreme categories (1.7–1.8 %) is below their overall share of the panel (1.26 %, i.e., 1,004 of 79,672), inconsistent with selective targeting of extreme-swing mesas for manipulation. Lima and Amazonas are not Castillo highland strongholds; their extreme swings reflect high anti-Fujimori consolidation from former Sagasti and Toledo voters with near-zero first-round Peru Libre support. *Source:* ONPE (Oficina Nacional de Procesos Electorales, 2021b,a); author's calculations.

The contested share among extreme-swing mesas (1.7–1.8 %) is *lower* than the contested share in the full panel (1.3 %). Extreme swings are concentrated in Lima and Amazonas—neither of which is a Castillo highland stronghold—and arise from high anti-Fujimori vote consolidation in urban areas where Peru Libre had minimal first-round presence. There is no geographic or distributional fingerprint of selective manipulation in these mesas.

A.5 Regional Heterogeneity

Table A4 breaks the contested–uncontested comparison down by macro-region (Costa, Sierra, Selva).

Table A4: *Mean Swing and Contested–Uncontested Difference by Macro-Region*

Region	Contested	Uncontested	Difference	t-statistic	p-value	N
Costa	+28.6 pp	+31.4 pp	−2.8 pp	−7.4	< 0.001	49,308
Sierra	+29.1 pp	+30.1 pp	−1.0 pp	−1.3	0.187	26,161
Selva	+29.8 pp	+30.4 pp	−0.6 pp	−0.6	0.550	4,203
All	+29.3 pp	+31.3 pp	−2.0 pp	−6.0	< 0.001	79,672

Notes: Welch two-sample t-tests within each macro-region. The Costa difference is the largest in absolute terms and most precisely estimated, reflecting both the greater number of contested mesas on the coast (848 of 1,004, or 84.5 %) and their higher geographic concentration. The Sierra and Selva differences are also negative but do not reach conventional significance levels, consistent with the smaller contested subsamples in those regions (111 and 45 mesas, respectively). Macro-region assignments follow INEI's conventional classification. *Source:* ONPE (Oficina Nacional de Procesos Electorales, 2021b,a); author's calculations.

The negative contested–uncontested gap holds in all three macro-regions. It is largest and most precisely estimated in the Costa, where challenged actas were most concentrated (848 of 1,004). In the Sierra—the region where Castillo dominated most heavily and where fraud allegations were implicitly directed—the contested–uncontested difference is −1.0 pp and statistically indistinguishable from zero. If manipulation had targeted highland actas to inflate Castillo's totals, the Sierra difference should be large and positive. It is not.

A.6 Interpretation

The mesa-level swing analysis adds four pieces of evidence to the main-paper forensic tests.

First, the mean swing (+31.3 pp) is almost exactly what binary consolidation predicts mechanically from the two rounds' national vote shares. There is no excess swing requiring a manipulation explanation.

Second, contested mesas have a *lower* swing than uncontested mesas by approximately 2 pp ($t = -6.0$, $p < 10^{-9}$). The residual from a mesa-level OLS is negative (−0.020) for

contested mesas: they under-perform their geographic baseline. This is the opposite direction from what inflation of Castillo's totals in contested actas would produce.

Third, extreme-swing mesas are concentrated in Lima and Amazonas—not in the Sierra departments where Castillo's majorities were largest. These mesas are almost entirely uncontested. Their large swings reflect genuine consolidation of anti-Fujimori votes from non-Peru Libre first-round candidates, not manipulation.

Fourth, the regional breakdown shows that the contested–uncontested gap is negative in every region, including the Sierra, where the fraud allegations were implicitly focused. The finding is not an artifact of regional composition.

Taken together, these results reinforce the conclusion of the main analysis: the contested actas do not bear the statistical fingerprints of manipulation.