Information and voting: Evidence from Peru's 2026 presidential election

Marcelo Gallardo[1,2] Nicolas Velarde[1,3] Cristina Gutarra[1]

[1]Pontificia Universidad Católica del Perú (PUCP)
[2]University of California, Berkeley
[3]Centro de Investigación de la Universidad del Pacífico (CIUP)

June 18, 2026

Preliminary draft. This paper is strictly theoretical and empirical, and does not reflect the political views or preferences of the authors.

Outline

Motivation

- In multi-candidate plurality elections with weak party identification, late public signals about candidate **viability** can change voters' minds.
- Peru 2026: extraordinarily fragmented first round — **36 candidates**, weak party brands, a documented logic of rejection (*mal menor*) rather than affiliation.
- On election night the two leading pollsters released *flash electoral* estimates minutes apart:
  - Both agreed: **Fujimori** first.
  - They **disagreed** on second: Datum named **López Aliaga**; Ipsos named **Sánchez**; **Nieto** within the margin of error in both.

Question

What is the causal effect of election-night flash estimates on voting behavior?

Polling dynamics: a fragmented, late-deciding field

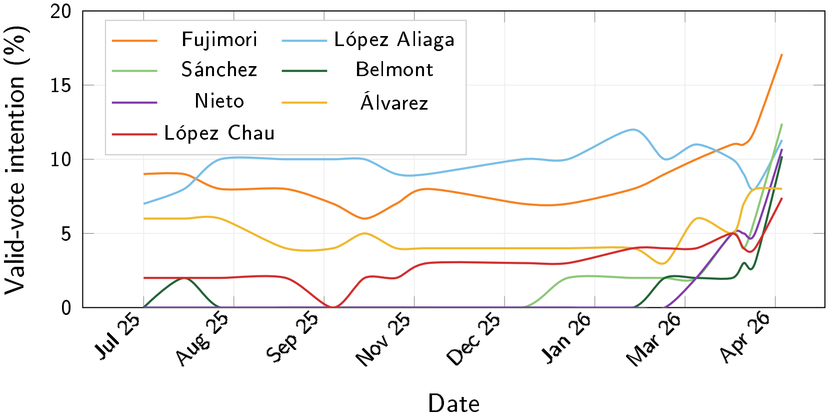

Ipsos Perú regular tracker, Jul 2025–Apr 2026 (valid-vote basis). Only Fujimori is clearly ahead at the close; Sánchez, Nieto, and Belmont surge in the final weeks — the second slot is undecided going into election night. Source: *La Encerrona* poll tracker.

The natural experiment

The logistical failure

- Sunday 12 April 2026: **187 polling tables** across **13 voting centers** (Lima Sur) failed to install — materials were not delivered on time.
- The *Jurado Nacional de Elecciones* (JNE) moved the affected $\approx 55\,000$ electors to **Monday 13 April**.

Treatment

- *Monday* voters cast ballots **after** the Ipsos and Datum flash estimates and the partial count.
- Comparable *Sunday* voters did not.

Identification

Assignment driven by a logistical failure, not by voter characteristics $\Rightarrow$ compare treated vs. matched control *locales de votación*.

Election night: the flash estimates

- No clear front-runner: Fujimori 17.06%, Sánchez 12.04%, López Aliaga 11.90%, Nieto 11.03% — barely a point separating the second runoff slot.
- 18:00, 12 April — Ipsos and Datum release flash estimates:
  - Datum: López Aliaga second (12.8%). Ipsos: Sánchez second (12.1%), Belmont third (11.8%).
  - Nieto within a percentage point of the named candidates in both.
- The estimates **confirmed Fujimori's passage** and rendered **three candidates viable** for the remaining slot, spanning the spectrum:

López Aliaga	Nieto	Sánchez
right	center, off-axis	left

Contributions

1. **Empirical (central).** Exploit the JNE ruling as a natural experiment; identify the causal effect of overnight flash estimates via **cardinality matching** at the polling-place level and a within-group regression estimator.
2. **New dataset.** *Locales de votación* characterized socioeconomically from block-level (*manzana*) INEI Redatam census data merged with ONPE *actas* for 2021 and 2026.
3. **Theory as organizing device.** Bayesian-updating / information-design model (Kamenica–Gentzkow; Blackwell): treatment is a Blackwell improvement of the public experiment, and an identification result equates the matching ATT with its behavioral value.

Model: states and signals

States. Fujimori's first place is uncontested; uncertainty is over the **identity of the runoff entrant**,

$$\Theta = \{\text{RS},\ \text{RLA},\ \text{JN}\}, \quad K \equiv |\Theta| = 3, \quad \mu \in \Delta(\Theta), \quad \mu_0 \text{ the common prior.}$$

Signals. Each pollster's flash is a symmetric η-*noisy* experiment with signal space $S = \Theta$ and error rate η:

$$\pi^{\eta}(s \mid \theta) = \begin{cases} 1-\eta, & s = \theta, \\ \dfrac{\eta}{K-1}, & s \neq \theta, \end{cases} \qquad \eta \in \left[0,\, 1 - \tfrac{1}{K}\right).$$

- η absorbs sampling and non-sampling error and yields a clean Blackwell ordering: $\pi^{\eta_1} \succeq_B \pi^{\eta_2}$ iff $\eta_1 \leq \eta_2$.
- Voters observe **two** public signals (s_D, s_I) (Datum, Ipsos). In 2026 the calls were **discordant**, $s_D^* \neq s_I^*$.

Bayesian updating: one signal

A *signal* is a kernel $\pi : \Theta \to \Delta(S)$; the *posterior* $\mu(\theta \mid s^*)$ is the updated belief that the true state is θ after observing the call s^*. For a realized call $s^* \in \Theta$, Bayes' rule on the η-noisy kernel gives

$$\mu(\theta \mid s^*) = \begin{cases} \dfrac{(1-\eta)\,\mu_0(s^*)}{Z}, & \theta = s^*, \\ \dfrac{\big(\eta/(K-1)\big)\,\mu_0(\theta)}{Z}, & \theta \neq s^*, \end{cases} \qquad Z = (1-\eta)\mu_0(s^*) + \tfrac{\eta}{K-1}\big(1-\mu_0(s^*)\big),$$

where Z is the normalizing constant (the marginal probability of the call s^*). A noisier flash (larger η) pulls the posterior back toward the prior μ_0.

Bayesian updating: two signals and the 2026 posterior

With **two conditionally independent signals** (s_D^*, s_I^*) (Datum, Ipsos) of error rates (η_D, η_I), the joint posterior is

$$\mu(\theta \mid s_D^*, s_I^*) = \frac{\pi^{\eta_D}(s_D^* \mid \theta)\, \pi^{\eta_I}(s_I^* \mid \theta)\, \mu_0(\theta)}{\sum_{\theta' \in \Theta} \pi^{\eta_D}(s_D^* \mid \theta')\, \pi^{\eta_I}(s_I^* \mid \theta')\, \mu_0(\theta')}.$$

2026 discordant calls (illustrative)

$(s_D^*, s_I^*) = (\mathrm{RLA}, \mathrm{RS})$, $(\eta_D, \eta_I) = (0.40, 0.30)$, uniform prior μ_0; likelihoods 0.090, 0.140, 0.030 give

$$\mu(\cdot \mid \mathrm{RLA}, \mathrm{RS}) \approx (0.35_{\,\mathrm{RLA}},\ 0.54_{\,\mathrm{RS}},\ 0.12_{\,\mathrm{JN}}).$$

Discordance **splits** viability mass *within* Θ; candidates outside Θ receive none.

Model: from utility to vote probabilities

Voter j has type (α_j, β_j, M_j): expressive weight α_j, instrumental weight β_j, and match score $M_j(c) \in [0,1]$, the expressive affinity with candidate c. Utility:

$$U_j(c,\theta) = \underbrace{\alpha_j\, M_j(c)}_{\text{expressive}} + \underbrace{\beta_j\, M_j(c)\, 1\{c \in \mathrm{runoff}(\theta)\}}_{\text{instrumental}} + \varepsilon_{jc}, \qquad \varepsilon_{jc} \sim \text{Gumbel (i.i.d.)}.$$

Perceived viability $q_c(\mu) \in [0,1]$ is the probability the voter assigns to candidate c reaching the runoff: $q_c(\mu) = 1$ for Fujimori (her passage was confirmed), $q_c(\mu) = \mu(c)$ for $c \in \Theta$, and $q_c(\mu) = 0$ otherwise. Writing μ_j^s for voter j's posterior after the signal profile s, the choice probability is logit:

$$\mathbb{P}_j(c \mid s) = \frac{\exp\big(M_j(c)\,[\alpha_j + \beta_j\, q_c(\mu_j^s)]\big)}{\sum_{c'} \exp\big(M_j(c')\,[\alpha_j + \beta_j\, q_{c'}(\mu_j^s)]\big)}.$$

The mechanism: viability–affinity interaction

Viability–affinity interaction

A rise in q_c pulls voter j toward c **in proportion to** $M_j(c)$: a viable but expressively rejected candidate ($M_j(c) \approx 0$) exerts no pull. Formally, $\partial \mathbb{P}_j(c \mid s)/\partial q_c \;\propto\; M_j(c) \;\geq 0$.

The flash is **public**: it raises q_c by the *same* amount for every voter in the location. But the response is not uniform — it is scaled by affinity, so only voters who already find c acceptable actually switch. Aggregating, the gain for c concentrates among voters **already sympathetic to** c, and the location-level effect τ_c is **scaled by the local affinity mass** on c. Same signal, different movement: this is the economic core of the results.

From individual choices to *acta*-level shares

Secret ballot: ONPE records only *acta*-level counts, so the logit is identified from shares. Let N_a be the valid votes in *acta* a and $V_{jc} = 1$ if voter j votes for c (else 0). The **observed vote share** for candidate c in *acta* a is

$$Y_{ac} = \frac{1}{N_a} \sum_{j \in a} V_{jc}, \qquad \mathbb{E}[V_{jc} \mid \mu_j] = \mathbb{P}_j(c \mid \mu_j).$$

Baseline posterior $\mu_j^{(0)} = \mu_0$ (Sunday); post-flash $\mu_j^{(1)} = \mu_j(\cdot \mid s_D^*, s_I^*)$ (Monday). The **individual information effect** is

$$\delta_{jc} := \mathbb{P}_j(c \mid \mu_j^{(1)}) - \mathbb{P}_j(c \mid \mu_j^{(0)}), \qquad \tau_{ac} = \frac{1}{N_a} \sum_{j \in a} \delta_{jc},$$

i.e. within an *acta* the effect is the average of the individual effects.

The estimand and its model counterpart

Averaging τ_{ac} over the treated population gives the **ATT estimand** for candidate c; its model counterpart Δ_c is computable in closed form from the structural choice model:

$$\tau_c := \mathbb{E}\big[\, Y_{ac}(1) - Y_{ac}(0) \mid T_\ell = 1 \,\big] \; = \underbrace{\mathbb{E}\Big[\int \delta_{jc}\, dF_a(j) \;\Big|\; T_\ell = 1\Big]}_{=\Delta_c},$$

where F_a is the within-*acta* distribution of voter types (mass $1/N_a$ on each voter $j \in a$). The unobservable individual effect δ_{jc} reaches the data only through its *acta* average τ_{ac}.

Identification: assumptions

Let H_a be the vector of pre-treatment covariates (*acta* demographics and location structural variables) and $T_\ell \in \{0, 1\}$ the treatment indicator of location ℓ.

- **A1 (conditional independence).** Given the state θ, the Datum and Ipsos flashes factor (used in Prop. 1) — independent samples and weighting; stress-tested next.
- **A2 (selection on observables + overlap).** $Y_{ac}(0) \perp\!\!\!\perp T_\ell \mid H_a$ and $0 < \mathbb{P}(T_\ell{=}1 \mid H_a) < 1$ — justified because displacement followed the contractor's logistical protocol (routes, delivery times), correlated with district socioeconomics but *orthogonal to political preferences*.
- **A3 (SUTVA).** An *acta*'s potential outcomes depend only on its own location's status — no between-location spillovers.

Identification: are the two polls independent?

A1 requires independence *conditional on* the true entrant θ. Datum and Ipsos draw separate samples with different field teams and weighting, so their sampling noise is independent; what they share is the *electorate* — which is the signal θ, exactly what A1 conditions out.

- The 2026 calls were **discordant** ($s_D^* \neq s_I^*$): identical methods would have produced the same call. Discordance is direct evidence the errors did not co-move.
- Their methods are similar but not identical, so some residual error correlation is plausible. That would make the two flashes carry *less* independent information, splitting viability mass *more* diffusely — which reinforces, not reverses, the sign predictions. A robustness check allows a correlation parameter ρ in the joint kernel.

Identification: what we cannot measure

Treatment is a bundle

The flash does not arrive alone: Monday voters also saw live media coverage, news commentary, and the ONPE partial-count platform. We identify the effect of the Monday *information environment* as a whole, not the flash in isolation.

- These channels are neither separately observed nor separable, so the estimand is the joint effect of the post-flash environment — the honest object the design delivers.
- This does not threaten identification (A2 still holds: displacement was logistical, not political), but it bounds interpretation: "flash estimates" is shorthand for the overnight information shock they anchored.

Identification: the central result

Proposition 2 (Identification of τ_c)

Under A2–A3, with control regression $m_0(h) := \mathbb{E}[Y_{ac} \mid H_a = h,\ T_\ell = 0]$ (the average control share at covariate profile h),

$$\tau_c \;=\; \underbrace{\mathbb{E}[Y_{ac} \mid T_\ell = 1] - \mathbb{E}[m_0(H_a) \mid T_\ell = 1]}_{\text{observable: the matching estimand}} \;=\; \underbrace{\Delta_c}_{\text{structural}}\;.$$

Two edges close the triangle:

- **(I) structural identity** $\tau_c = \Delta_c$ — definitional, from how the model maps beliefs to shares; it lets the measured ATT be read as the behavioral value of the added information.
- **(II) population identification** — selection-on-observables identifies the treated counterfactual $\mathbb{E}[Y_{ac}(0) \mid T_\ell = 1]$ from matched controls.

Treatment as a change of experiment: theory $\leftrightarrow$ data

A *Sunday* voter decides under the **null experiment** $\pi^{\varnothing}$ (no flash, posterior $= \mu_0$); a *Monday* voter under the **flash experiment** $\pi^F = \pi^{\eta_D} \otimes \pi^{\eta_I}$ (posterior $\mu^{(1)}$). The null garbles every experiment, so $\pi^F \succeq_B \pi^{\varnothing}$ — strict once $\eta_D, \eta_I < 1 - 1/K$. Treatment is a **Blackwell improvement**, and τ_c is its behavioral value:

$$\underbrace{\pi^F \succeq_B \pi^{\varnothing}}_{\text{information design}} \quad \Rightarrow \quad \underbrace{\delta_{jc} = \mathbb{P}_j(c \mid \mu_j^{(1)}) - \mathbb{P}_j(c \mid \mu_j^{(0)})}_{\text{updating / choice}}$$

$$\Rightarrow \underbrace{\tau_c = \mathbb{E}\left[\int \delta_{jc}\, dF_a \mid T_\ell = 1\right] = \Delta_c}_{\text{structural estimand}} = \underbrace{\mathbb{E}[Y_{ac} \mid T_\ell = 1] - \mathbb{E}[m_0(H_a) \mid T_\ell = 1]}_{\text{matching estimator}}.$$

What the theory fixes, what the data delivers

Blackwell fixes the *value* of the extra information, not its *direction*: which c gains comes from the updating arrow, scaled by $M_j(c)$. National viability (experiment level) can diverge from local gain (updating level) — the Sánchez case.

Hypotheses (for the treated Lima Sur sample)

Within the right-leaning Lima sample, viable gains rank by **match**, not viability alone: Nieto (off-axis, "clean") > López Aliaga (right-aligned, tied to the congressional pact) > Sánchez (left-anchored, provincial base).

H1. Nieto gains the most: $\tau_{\mathrm{JN}} > 0$, largest among viable.

H2. López Aliaga intermediate: $0 < \tau_{\mathrm{RLA}} < \tau_{\mathrm{JN}}$, $\tau_{\mathrm{RLA}} \geq \tau_{\mathrm{RS}}$.

H3. Sánchez gains least among viable, despite national viability.

H4. Fujimori unmoved ($\tau_{\mathrm{KF}} \approx 0$, at most weakly positive): her entry is confirmed in every state.

H5. Non-viable candidates lose: $\tau_c \leq 0$ (clearest: Álvarez, López Chau; Belmont also declines).

Adding-up: $\sum_c Y_{ac} = 1 \Rightarrow \sum_c \tau_c = 0$.

Data

Sample: Lima Sur — 354 voting locations, 4930 *actas*; 13 treated, 341 control pool.

INEI (Redatam census)

- Georeferenced at the block (*manzana*) level.
- Population, migration, language, education, health insurance, services, housing.

ONPE

- *Acta*-level results, 2021 and 2026 general elections.
- 2026 electoral roll: demographic composition of registered voters.

Linkage. For each location ℓ, a 1 km buffer; census blocks $B(\ell)$ within it aggregated by population weights (pop_b, v_b the value at block b):

$$V_\ell = \frac{\sum_{b \in B(\ell)} \mathrm{pop}_b \, v_b}{\sum_{b \in B(\ell)} \mathrm{pop}_b}.$$

Consistent with *Elige tu local*: voters select polling places near their residence.

Cardinality matching: the balance program

With only **13 treated clusters**, propensity-score models are unstable. Cardinality matching instead fixes the balance target *ex ante* and maximizes the retained control sample subject to it.

Let $\mathcal{C}_0$ be the control pool, $q_\ell \in \{0,1\}$ the selection indicator, n_ℓ the *actas* at location ℓ, and $\mathcal{C}^\star = \{\ell : q_\ell = 1\}$:

$$\max_{\{q_\ell\}} \sum_{\ell \in \mathcal{C}_0} q_\ell \, n_\ell \qquad \text{s.t.} \qquad \left|\mathrm{SMD}_j\right| \le \delta_j, \quad j = 1, \dots, J, \qquad \mathrm{SMD}_j = \frac{\bar{H}_{j,T} - \bar{H}_{j,\mathcal{C}^\star}}{s_j},$$

where $\bar{H}_{j,T}$, $\bar{H}_{j,\mathcal{C}^\star}$ are treated and matched-control means of covariate j and s_j its pooled SD.

- Strict per-covariate rule: $\delta_j = 0.10$ for all $J = 9$ covariates, no exceptions; fine balance on the district marginal. Solved as an integer program in Python.

Cardinality matching: the estimation pipeline

Stage 2 (grouping). For each treated location we rank controls by **Mahalanobis distance** $d(\ell, \ell') = \sqrt{(x_\ell - x_{\ell'})^\top \Sigma^{-1} (x_\ell - x_{\ell'})}$ — covariate distance rescaled by Σ so correlated covariates are not double-counted — keep the nearest, and drop any with $d > 3$ (the caliper). Group sizes range 3–10.

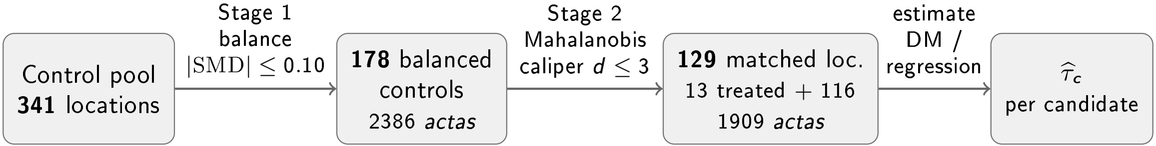

Matching is at the *location* level; outcomes are estimated at the *acta* level.

The procedure, step by step

From raw data to $\widehat{\tau}_c$

1. **Covariates.** For each location ℓ build H_ℓ: census variables (1 km *manzana* buffer) plus 2021 vote shares — nine covariates in all.
2. **Stage 1 (balance).** Solve the integer program: pick the controls that *maximize retained actas* subject to $|\mathrm{SMD}_j| \leq 0.10$ for every covariate. This fixes balance *before* looking at outcomes. $\Rightarrow$ 178 controls.
3. **Stage 2 (grouping).** For each treated location, rank the selected controls by Mahalanobis distance, keep the nearest, and drop any with $d > 3$. $\Rightarrow$ 13 groups, 116 controls (129 locations).
4. **Estimate.** Within each group take the treated−control *acta*-mean gap; average across groups with *acta* weights $\Rightarrow \widehat{\tau}_c^{\mathrm{DM}}$. A group-fixed-effects regression gives $\widehat{\tau}_c^{\mathrm{R}}$ as a check.
5. **Inference.** Leave-one-treated-location-out and within-group permutation tests, given only 13 treated clusters.

Two estimators on the matched sample

1. Acta-weighted difference in means (DM). Within group i, contrast treated and control *acta* means $\hat{\Delta}_{i,c} = \bar{Y}_{i,c}^{T} - \bar{Y}_{i,c}^{C}$; weight by the treated *acta* share $w_i = n_i / N_T$ (n_i treated *actas* in group i, $N_T = \sum_i n_i$):

$$\widehat{\tau}_c^{\mathrm{DM}} = \sum_{i=1}^{13} w_i \, \hat{\Delta}_{i,c}, \qquad \sum_i w_i = 1.$$

2. Regression-adjusted ATT. Pooled OLS over *actas* a, with treatment indicator $T_a \equiv T_{\ell(a)}$, group fixed effects γ_i, and *acta*-level demographic controls X_a:

$$Y_{a,c} = \tau_c \, T_a + \gamma_{i(a)} + X_a^{\top} \lambda_c + u_{a,c}, \qquad \text{cluster-robust SE at the location level.}$$

Inference caveat: 13 treated clusters is a known limitation; within-group permutation tests under the sharp null are the preferred inferential base.

Covariate balance

Covariate	Treated mean	Control (all)	Control (matched)	SMD before	SMD after
Socioeconomic stratum	2.207	2.599	2.179	−0.629	0.045
Indigenous-language share	0.110	0.091	0.112	0.473	−0.028
Higher-education share	0.249	0.312	0.254	−0.733	−0.064
Internet-access share	0.361	0.453	0.372	−0.786	−0.100
2021 left vote share	0.254	0.247	0.257	0.201	−0.095
2021 right vote share	0.279	0.301	0.279	−0.347	0.013
2021 Fujimori vote share	0.169	0.163	0.173	0.167	−0.099
Female-voter share	0.492	0.499	0.495	−0.205	−0.098
Voter age (group index)	2.755	2.855	2.772	−0.193	−0.032

Largest $|\mathrm{SMD}|$ falls from 0.79 to 0.100; every covariate clears the $\delta_j = 0.10$ tolerance.

Main results: treatment effects on vote shares

Candidate	$\widehat{\tau}_c^{\mathrm{DM}}$	$\widehat{\tau}_c^{\mathrm{R}}$
Keiko Fujimori (KF)	0.026*** (0.007)	0.026*** (0.004)
Roberto Sánchez (RS)	0.027*** (0.006)	0.027*** (0.003)
Rafael López Aliaga (RLA)	0.051*** (0.007)	0.051*** (0.004)
Jorge Nieto (JN)	0.173*** (0.008)	0.173*** (0.004)
Ricardo Belmont (RB)	−0.031*** (0.004)	−0.031*** (0.003)
Carlos Álvarez (CA)	−0.068*** (0.003)	−0.068*** (0.002)
Alfonso López Chau (LCH)	−0.044*** (0.001)	−0.044*** (0.001)

Share units; cluster-robust SEs at the location level. $N = 1909$ *actas*; 129 matched locations. $^{***}\ p < 0.01$.

DM and regression-adjusted estimators agree in sign and magnitude for every candidate: the global balance carries the identification.

The effect at a glance

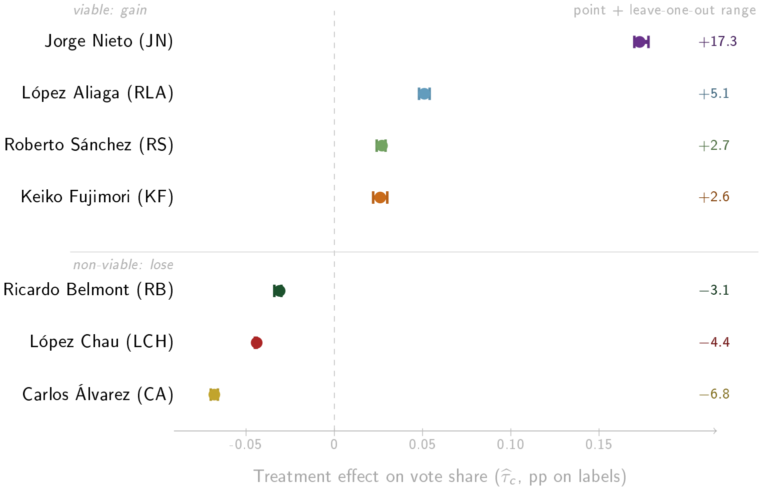

Acta-weighted ATT; caps span the thirteen leave-one-treated-location-out reruns. $N = 1909$ *actas*, 129 matched locations.

Reading the magnitudes

1. **Viable gain, non-viable lose** — exactly the sign pattern of H1–H3, H5.
2. **Fujimori**: small but significant +2.6 pp — weak reinforcement (H4).
3. **Reallocation concentrated on Nieto**: +17.3 pp, three to seven times any other gain.

Within-viable ordering: Nieto > López Aliaga > Sánchez

Tracks **affinity, not national viability**. Nieto (off-axis "clean" option, highest match) takes the bulk; López Aliaga's match is cut by his tie to the congressional pact (+5.1); Sánchez — the eventual runoff entrant — is expressively most distant in this right-leaning sample and gains least (+2.7).

The seven leading effects sum to +0.134: the displaced share is drawn from the long tail of minor candidates (adding-up, H5).

Robustness

- **Location-weighted estimand.** Equal weight per treated location: signs and significance unchanged (Nieto $+0.176$, López Aliaga $+0.054$, Sánchez $+0.025$).
- **Leave-one-treated-location-out.** Thirteen reruns: no sign reversal, no loss of significance at 5%; Nieto stays within $[0.170, 0.178]$; each viable candidate's contrast positive in ≥ 11 of 13 groups.
- **Balance tolerance.** Re-solving at $\delta \in \{0.05, 0.10, 0.15\}$ moves estimates negligibly; all remain significant at 1%.

Conclusion

- A late **viability signal** causally reallocates votes — every effect significant at 1% and robust to leave-one-out, weighting, and tolerance checks.
- The reallocation runs through the **local distribution of expressive affinities**, not the signal alone: q_c moves uniformly, but the vote it moves is scaled by $M_j(c)$.
- A candidate can advance nationally on a signal yet move few local votes: **viability travels with the signal; its reallocation is governed by local affinity** (the Sánchez case).
- First-order relevance for Peru: institutional instability, fragmentation, and political volatility make voters unusually responsive to viability signals.

Caveats and work in progress

This is a rushed first pass

The election was **barely two months ago** (April 2026); the georeferenced census and *acta*-level files arrived late, leaving little time before this draft. The findings are stable but the inference layer is still being finalized.

- **Rosenbaum sensitivity (1:k).** The current sign-flip enumeration is valid only for 1:1 pairs. The correct version uses per-location scores, a worst-case within-set tilt under Γ, and exact combination across the 13 independent groups (k varies 3–10).
- **Heterogeneity.** Effects by a Right–Left index (interaction terms) — not yet in the results files.
- **Few-cluster inference.** Wild cluster bootstrap / randomization inference to complement cluster-robust SEs with only 13 treated clusters.

Thank you

marcelogallardob21@berkeley.edu
na.velardef@up.edu.pe
cristina.gutarra@pucp.edu.pe