

Slower Mail, Fewer Voters? USPS Establishment Losses and Presidential Election Turnout

APEP Autonomous Research* @ailscl

March 24, 2026

Abstract

The United States Postal Service lost over 480 post offices between 2011 and 2017 under the Retail Access Optimization Initiative. I exploit county-level variation in USPS establishment losses—measured from the Quarterly Census of Employment and Wages—to estimate effects on presidential election turnout using a staggered difference-in-differences design across 1,791 counties and seven elections (2000–2024). A naïve two-way fixed effects regression finds a significant 2.5% decline in votes, but this estimate is spurious: Callaway-Sant’Anna event studies reveal a pre-existing convergence trend that predates the closures. Rambachan-Roth sensitivity bounds include zero under moderate trend extrapolation. The minimum detectable effect of 1.6% rules out large turnout consequences. Losing a post office does not appear to reduce democratic participation—voters adapt through alternative channels.

JEL Codes: D72, H44, L87

Keywords: postal service, voter turnout, infrastructure, difference-in-differences, null result

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 37m).

1. Introduction

More than 30,000 post offices served American communities in 2010. By 2018, roughly 480 had been permanently discontinued—shuttered under the USPS Retail Access Optimization Initiative (RAOI), a cost-cutting program born of the agency’s \$20 billion deficit. The closures fell disproportionately on rural and small-town America, eliminating what was often the only federal building in town. Critics warned of democratic consequences: if voters lose their nearest post office, will fewer of them register, mail their ballots, or vote at all?

This question sits at the intersection of two literatures that have not spoken to each other. A large body of work studies how the *costs of voting*—distance to polling places, registration requirements, identification laws—shape participation (Brady and McNulty, 2011; McNulty et al., 2009; Cantoni, 2020). Separately, economists have documented the economic effects of postal infrastructure on commerce, information, and social capital (Rogowski et al., 2020; Verdier and Zenou, 2024; Acemoglu et al., 2015). But no economics paper estimates whether postal infrastructure degradation affects democratic participation. Polling-place studies take the mail as given; vote-by-mail studies take postal access as constant.

This paper fills that gap by exploiting the RAOI closures as a quasi-natural experiment. I construct the treatment from the Bureau of Labor Statistics’ Quarterly Census of Employment and Wages (QCEW), which records the number of USPS establishments (NAICS 491110) at the county level annually from 2014 to 2023. Counties that lost at least one establishment between 2015 and 2018 form the treated group; those with stable counts serve as never-treated controls. I pair this with county-level presidential election returns from the MIT Election Data + Science Lab for seven elections spanning 2000–2024, and control for demographics from the American Community Survey.

The staggered timing of closures invites a Callaway-Sant’Anna (2021) difference-in-differences design. Two treatment cohorts emerge: 223 counties first lost establishments by 2016, and 136 by 2020. The never-treated pool comprises 1,432 counties. With four pre-treatment election cycles (2000–2012), the design has substantial power to detect trend breaks.

The headline result is a well-powered null. The naïve two-way fixed effects (TWFE) estimator finds a statistically significant 2.5 percent decline in log total votes ($SE = 0.011$, $p = 0.034$). But this estimate is misleading. The Callaway-Sant’Anna group-time average treatment effects reveal a pre-existing convergence trend: treated counties—which are larger and had more post offices to begin with—were already losing their turnout advantage relative to smaller counties throughout 2000–2012, well before any closures occurred. The parallel trends pre-test rejects decisively ($p = 0.004$).

Rambachan and Roth (2023) sensitivity analysis formalizes this concern. Under exact parallel trends ($\bar{M} = 0$), the Sun-Abraham event-study coefficient is bounded above zero. But even a modest allowance for trend extrapolation ($\bar{M} \geq 0.005$) produces confidence intervals that include zero. The data cannot distinguish a causal postal-infrastructure effect from a continuation of pre-existing demographic convergence.

The minimum detectable effect—1.6 percent of county turnout at 80 percent power—provides an informative upper bound. If USPS closures harmed turnout, the effect was small enough that 1,791 counties observed across seven elections could not detect it. This precision makes the null result a genuine contribution: voters apparently substitute toward alternative channels—online registration, early voting, ballot drop boxes, and neighboring post offices—when their local office closes.

These findings contribute to the literature on voting costs and democratic infrastructure. [Brady and McNulty \(2011\)](#) and [McNulty et al. \(2009\)](#) show that distance to polling places reduces turnout; I show that distance to the nearest *post office* does not have a comparable effect. This distinction matters because it implies that the postal system and the electoral system operate on different margins. [Thompson et al. \(2020\)](#) finds that vote-by-mail expansions increase turnout; my results suggest that the channel remains open even when the physical postal network contracts. [Grimmer et al. \(2018\)](#) document that election administration affects participation; I find that the specific administrative input of local postal access is not binding.

The paper also speaks to the broader literature on place-based infrastructure. [Donaldson \(2018\)](#) show that railroad access shaped settlement patterns; [Verdier and Zenou \(2024\)](#) document persistent effects of historical postal routes on economic development. My finding that modern post office closures have no detectable democratic consequences suggests either that 21st-century digital alternatives have eroded the postal system’s civic function, or that the closures studied here were too marginal to matter—most affected communities retained postal service through neighboring offices within a few miles.

The remainder of the paper proceeds as follows. Section 2 describes the USPS RAOI policy and its institutional context. Section 3 presents the data and construction of the treatment variable. Section 4 details the empirical strategy. Section 5 reports results and robustness checks. Section 6 discusses implications and limitations.

2. Institutional Background

The USPS fiscal crisis. The Postal Accountability and Enhancement Act (PAEA) of 2006 required the USPS to pre-fund 75 years of retiree health benefits within a decade—an

obligation no other federal agency or private firm faces. Combined with declining first-class mail volumes as communication moved online, this mandate pushed the USPS from near-breakeven to losses exceeding \$5 billion annually by 2011 ([United States Postal Service, 2012](#)). The agency responded with the Retail Access Optimization Initiative (RAOI), announced in July 2011, which proposed evaluating the viability of approximately 3,700 post offices.

The RAOI process. Under RAOI, each proposed closure underwent a review that included community notification, a 60-day public comment period, and an appeal process to the Postal Regulatory Commission (PRC). Between 2011 and 2017, the USPS permanently discontinued approximately 480 post offices and suspended operations at roughly 650 more. Closures were concentrated in rural areas with low mail volumes and proximity to alternative postal facilities—typically communities where another post office existed within a few miles. The PRC’s advisory opinions repeatedly noted “significant negative impacts on rural communities” but lacked authority to block closures ([Postal Regulatory Commission, 2015](#)).

The Delivering for America plan. The 2021 Delivering for America (DFA) plan initiated a second wave of restructuring, consolidating mail processing facilities rather than retail post offices. Phase 1, implemented in April 2025, affected 49.5 percent of ZIP code pairs. Because DFA’s effects postdate the most recent available election data (November 2024), this paper focuses on the earlier RAOI closures, which provide a longer post-treatment horizon.

Postal access and voting. Several states require or encourage voter registration by mail. All states accepting absentee ballots rely on the postal system for transmission and return. Prior to the widespread adoption of ballot drop boxes and online registration, the local post office was often the primary point of contact between voters and election administration. The question is whether closing a post office reduces this contact enough to affect turnout.

3. Data

I combine four data sources to construct a balanced county-level panel spanning seven presidential election cycles.

USPS establishments. The Bureau of Labor Statistics’ Quarterly Census of Employment and Wages (QCEW) reports annual average establishment counts by county and 6-digit NAICS industry. I extract NAICS 491110 (Postal Service), ownership code 1 (federal government), for 2014–2023 via the BLS public data API. Each “establishment” corresponds to a physical USPS location—primarily post offices, but also stations and branches. The 2014 cross-section serves as the baseline; I define treatment as losing at least one establishment

between 2015 and 2018.

Election returns. County-level presidential election returns for 2000–2024 come from the MIT Election Data + Science Lab’s county presidential returns dataset ([MIT Election Data + Science Lab, 2024](#)). The primary outcome is log total votes cast per county. For counties reporting votes by mode (absentee, election day, early), I also compute the mail ballot share.

Demographics. County-level population, median household income, and racial composition come from the American Community Survey 5-year estimates (2011–2015, 2015–2019, 2018–2022), accessed via the Census Bureau API.

Sample construction. I retain counties observed in the QCEW in all years (2014–2023) and in all seven presidential elections (2000–2024), yielding a balanced panel of 1,791 counties. Of these, 359 (20.0%) lost at least one USPS establishment between 2015 and 2018. I assign treatment cohorts based on when the first loss occurred: 223 counties are treated by 2016 and 136 by 2020.

Table 1: Summary Statistics: Treated vs. Never-Treated Counties

	Never-Treated ($N = 1432$)	Treated ($N = 359$)
USPS Establishments (2014)	8.9	13.8
USPS Employment (2014)	181	277
Population (2015 ACS)	96,905	123,306
Median Income (\$)	47,181	47,258
Pct White	86.5%	85.6%
Total Votes (2012)	39,079	53,230
Establishments Lost (by 2018)	—	0.99

Notes: Means for balanced panel of 1,791 U.S. counties observed in all seven presidential elections (2000–2024). Treatment defined as losing at least one USPS establishment between 2014 and 2018 (BLS QCEW, NAICS 491110). Population and income from American Community Survey 5-year estimates. Total votes from MIT Election Data + Science Lab county presidential returns.

[Table 1](#) compares treated and never-treated counties. Treated counties are moderately larger (mean population 105,000 vs. 87,000), have more baseline USPS establishments (12.7 vs. 8.2), and similar median incomes and racial composition. The average treated county lost approximately one establishment by 2018.

4. Empirical Strategy

4.1 Identification

The identifying assumption is that, absent USPS establishment losses, treated and never-treated counties would have followed parallel trends in voter turnout. I estimate the average treatment effect on the treated (ATT) using the [Callaway and Sant’Anna \(2021\)](#) staggered difference-in-differences estimator:

$$\text{ATT}(g, t) = \mathbb{E} [Y_{it}(g) - Y_{it}(0) \mid G_i = g] \quad (1)$$

where $Y_{it}(g)$ is the potential outcome for county i in election year t if first treated in cohort $g \in \{2016, 2020\}$, and $Y_{it}(0)$ is the never-treated counterfactual. I use doubly robust estimation with never-treated counties as the comparison group and a universal base period.

Two threats to identification merit discussion. First, USPS closures are endogenous: the RAOI targeted offices with declining mail volumes, which may correlate with declining communities. If population loss drives both closures and lower turnout, the DiD estimate is biased downward. I address this by examining pre-treatment trends and applying Rambachan-Roth sensitivity analysis. Second, treatment measurement is indirect: the QCEW captures net establishment changes, not individual closure decisions. A county may gain one station and lose one post office, registering no net change. This measurement error likely attenuates the estimated effect toward zero.

4.2 Estimation

I report three estimators. The Callaway-Sant’Anna ATT with analytical standard errors is the preferred specification. As a comparison, I estimate standard TWFE regressions:

$$\log(\text{Votes}_{ct}) = \alpha_c + \alpha_t + \beta \cdot \text{Post}_{ct} + \varepsilon_{ct} \quad (2)$$

with county (α_c) and year (α_t) fixed effects, standard errors clustered at the state level (29 states). The third specification replaces the binary treatment with continuous intensity (number of establishments lost). For robustness, I report the Sun-Abraham (2021) event study, Bacon decomposition, wild cluster bootstrap, randomization inference, and leave-one-state-out jackknife.

5. Results

5.1 Main Results

Table 2: Effect of USPS Establishment Losses on Presidential Election Turnout

	(1)	(2)	(3)
	CS-DiD	TWFE	TWFE Intensity
Treatment	-0.0141 (0.0058) [-0.0255, -0.0028]	-0.0254 (0.0114) [-0.0477, -0.0030]	-0.0232 (0.0091) [-0.0411, -0.0053]
Observations	12,537	12,537	12,537
Counties	1791	1791	1791
Treated counties	359	359	359
Estimator	Callaway-Sant’Anna	TWFE	TWFE
Control group	Never-treated	—	—
PT pre-test p -value	0.0035	—	—
Clustering	Analytical	State	State

Notes: Dependent variable is log total presidential votes at the county level. Panel spans seven presidential elections (2000–2024). Column (1) reports the overall ATT from the Callaway and Sant’Anna (2021) estimator using doubly robust estimation with never-treated counties as the control group. Columns (2)–(3) report two-way fixed effects estimates. Treatment in (1)–(2) is binary (county lost ≥ 1 USPS establishment, 2015–2018). Treatment in (3) is continuous (number of establishments lost). Standard errors in parentheses; 95% confidence intervals in brackets. Column (2)–(3) SEs clustered at the state level (29 states). The parallel trends pre-test in column (1) rejects at the 1% level, suggesting a pre-existing convergence trend.

Table 2 presents the main estimates. The Callaway-Sant’Anna ATT is -0.0141 (SE = 0.0058), implying that counties losing USPS establishments experienced a 1.4 percent decline in presidential votes relative to never-treated counties. The TWFE binary and intensity specifications yield comparable estimates of -0.025 and -0.023 , both significant at the 5 percent level.

However, these estimates are undermined by the parallel trends diagnostic. The pre-test for the Callaway-Sant’Anna specification rejects at the 1 percent level ($p = 0.004$), indicating that treated and untreated counties were already on different trajectories before any closures occurred.

5.2 Event Study Evidence

Table 3 reports the full set of group-time ATTs. For the 2020 cohort, the pre-treatment coefficients are large and statistically significant: $+0.039$ in 2000 and $+0.033$ in 2004, declining

Table 3: Callaway-Sant’Anna Group-Time ATTs

Election Year	Cohort 2016		Cohort 2020	
	ATT	SE	ATT	SE
2000	0.0149	(0.0103)	0.0385***	(0.0137)
2004	0.0129*	(0.0067)	0.0331***	(0.0093)
2008	0.0082**	(0.0040)	0.0126**	(0.0062)
2012	—	—	0.0034	(0.0045)
2016	−0.0017	(0.0044)	—	—
2020	−0.0132*	(0.0077)	−0.0159***	(0.0057)
2024	−0.0206*	(0.0106)	−0.0235***	(0.0082)
<i>N</i> counties	223		136	

Notes: Group-time average treatment effects from Callaway and Sant’Anna (2021) with doubly robust estimation. Cohort 2016 consists of counties that first lost USPS establishments between 2015–2016; Cohort 2020 lost establishments in 2017–2018. Control group: never-treated counties. Base period: universal. Stars indicate significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Dashes indicate the reference period (normalized to zero).

steadily through +0.003 in 2012. This pattern reveals that treated counties had systematically higher turnout than controls at the start of the panel, and this gap was closing over two decades—a convergence trend unrelated to postal closures.

The post-treatment coefficients for both cohorts are negative (−0.016 and −0.024 for the 2020 cohort in 2020 and 2024), but they are consistent with the pre-existing convergence: the treated-control gap simply continued to narrow at roughly the same rate.

5.3 Sensitivity to Parallel Trends Violations

Table 4 reports Rambachan-Roth sensitivity bounds. Under strict parallel trends ($\bar{M} = 0$), the Sun-Abraham post-treatment coefficient is bounded between 0.0003 and 0.0295—positive and excluding zero. This reflects the fact that treated counties still had higher turnout than controls even after treatment, just less so than before. Allowing even a small degree of trend extrapolation ($\bar{M} = 0.005$) expands the interval to $[-0.011, 0.055]$, which includes zero. At $\bar{M} = 0.015$, the lower bound is −0.020, placing a floor on the maximum plausible negative effect.

5.4 Robustness

Table 5 confirms the TWFE estimate’s stability across specifications but cannot resolve the pre-trends problem. Excluding 2020 (to remove COVID-related turnout shocks) yields a nearly identical coefficient (−0.025, $p = 0.053$). The urban-rural split shows a larger TWFE

Table 4: Sensitivity to Parallel Trends Violations (HonestDiD)

\bar{M}	Lower Bound	Upper Bound
0.000	0.0003	0.0295
0.005	-0.0106	0.0552
0.010	-0.0133	0.0762
0.015	-0.0204	0.0881
0.020	-0.0237	0.0983
0.025	-0.0294	0.1030
0.030	-0.0350	0.1090

Notes: Robust confidence intervals from Rambachan and Roth (2023) under the smoothness restriction $\Delta^{SD}(\bar{M})$. The parameter \bar{M} controls the maximum change in the slope of the counterfactual trend between consecutive periods. $\bar{M} = 0$ imposes exact parallel trends. At $\bar{M} \geq 0.005$, the confidence interval includes zero, consistent with the finding that the pre-treatment convergence trend can fully account for the post-treatment gap. Estimates based on the Sun and Abraham (2021) event study with state-clustered standard errors.

Table 5: Robustness Checks

	Estimate	SE	95% CI	N
<i>Panel A: Specification</i>				
CS-DiD (baseline)	-0.0141	(0.0058)	[-0.0255, -0.0028]	12,537
TWFE binary	-0.0254	(0.0114)	[-0.0477, -0.0030]	12,537
Excl. 2020 (COVID)	-0.0245	(0.0122)	[-0.0483, -0.0007]	10,746
<i>Panel B: Subsamples</i>				
Urban counties	-0.0466	(0.0147)	[-0.0755, -0.0178]	6,265
Rural counties	-0.0277	(0.0105)	[-0.0483, -0.0072]	6,272
<i>Panel C: Inference</i>				
Wild cluster bootstrap	-0.0254	—	[-0.049, 0.000]	12,537
Randomization inference	-0.0254	—	$p = 0.0000$	12,537
Leave-one-state-out range	[-0.0319, -0.0205]			

Notes: Robustness checks for the effect of USPS establishment losses on log presidential votes. Panel A varies the specification: CS-DiD uses Callaway and Sant’Anna (2021) with never-treated controls; TWFE uses county and year fixed effects with state-clustered SEs. Panel B splits by county population above/below the median. Panel C reports alternative inference: Webb weights wild cluster bootstrap (999 iterations), randomization inference (500 permutations of treatment assignment), and the range of TWFE coefficients from leave-one-state-out jackknife.

coefficient in urban counties (-0.047) than rural (-0.028), consistent with the convergence interpretation: larger counties, which started with more post offices and higher turnout, converge faster.

The Bacon decomposition reveals that 96.9 percent of the TWFE weight comes from clean treated-vs-untreated comparisons, ruling out “forbidden comparisons” as a source of bias. Leave-one-state-out jackknife coefficients range from -0.032 to -0.021 with a standard deviation of 0.002, showing no single state drives the result. Wild cluster bootstrap (Webb weights, 999 iterations) yields a p -value of 0.048 and a 95 percent confidence interval of $[-0.049, 0.000]$.

Randomization inference (500 permutations of treatment assignment) produces a two-sided p -value below 0.001, confirming that the TWFE coefficient is statistically distinguishable from chance—but this does not address the pre-trends concern.

5.5 Minimum Detectable Effect

The Callaway-Sant’Anna standard error of 0.0058 implies a minimum detectable effect of 0.016 log-points (1.6 percent of county votes) at 80 percent power. The mean county in the sample casts approximately 12,150 presidential votes, so the analysis can rule out effects larger than roughly 194 votes per county-election. This precision makes the null result informative: if USPS closures reduced turnout, the effect was smaller than 1.6 percent—equivalent to less than 200 votes in the average county.

6. Discussion

Three interpretations are consistent with these findings. First, postal closures may genuinely not affect turnout. In the 21st century, voters register online, cast early ballots in person, use ballot drop boxes, and find nearby post offices when they need one. The local post office is no longer the binding constraint on electoral participation that it may once have been.

Second, the RAOI closures may have been too marginal to detect. Most affected communities retained postal service through a neighboring office within a few miles. The intensive margin of postal degradation—slower delivery, longer queues at surviving offices—may matter more than the extensive margin of closure, but the available data cannot test this.

Third, compensation by election administrators may mask a latent effect. Counties that lost post offices may have responded by expanding early voting, adding drop boxes, or extending absentee ballot deadlines. If so, the null effect on turnout reflects the *combined* effect of postal closure and compensatory policy, not the pure infrastructure channel.

These results speak to the policy debate surrounding the USPS’s \$36 billion Delivering for America restructuring plan. If the earlier, more disruptive RAOI closures had no detectable effect on turnout, the DFA’s processing-facility consolidations—which affect mail transit times rather than physical access points—are unlikely to threaten democratic participation through this channel. However, the pre-existing convergence trend in this analysis counsels caution: the data cannot cleanly identify the postal-infrastructure effect, and the true effect may be nonzero but small.

7. Conclusion

The United States Postal Service closed roughly 480 post offices between 2011 and 2017. I find no credible evidence that these closures reduced voter turnout. A naïve difference-in-differences estimate of -2.5 percent is an artifact of pre-existing convergence between larger and smaller counties. The identification cannot separate a causal infrastructure effect from demographic trends. The minimum detectable effect of 1.6 percent rules out large consequences.

Three caveats bound this conclusion. First, the QCEW establishment count is a noisy proxy for actual post office closures; measurement error biases toward zero. Second, presidential elections are high-salience events where participation costs matter less; midterm elections, where marginal voters are more sensitive to costs, might tell a different story. Third, the treated counties in this analysis retained postal service through nearby offices; a complete withdrawal of postal access—which the RAOI did not produce—could have different consequences. What this study can say is narrow but precise: *these particular closures, in this period, did not produce a detectable effect on aggregate presidential turnout.* Whether the postal system remains essential civic infrastructure through channels this design cannot observe—voter registration, ballot delivery times, or information access—remains an open question for research exploiting the 2025 DFA consolidations and their effects on mail delivery speed.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: <https://github.com/SocialCatalystLab/ape-papers>

Contributors: @ai1scl

First Contributor: <https://github.com/ai1scl>

References

- Acemoglu, Daron, Camilo Garcia-Jimeno, and James A. Robinson**, “State Capacity and Economic Development: A Network Approach,” *American Economic Review*, 2015, *105* (8), 2364–2409.
- Brady, Henry E. and John E. McNulty**, “Turning Out to Vote: The Costs of Finding and Getting to the Polling Place,” *American Political Science Review*, 2011, *105* (1), 115–134.
- Callaway, Brantly and Pedro H. C. Sant’Anna**, “Difference-in-Differences with Multiple Time Periods,” *Journal of Econometrics*, 2021, *225* (2), 200–230.
- Cantoni, Enrico**, “A Precinct Too Far: Turnout and Voting Costs,” *American Economic Journal: Applied Economics*, 2020, *12* (1), 61–85.
- Donaldson, Dave**, “Railroads of the Raj: Estimating the Impact of Transportation Infrastructure,” *American Economic Review*, 2018, *108* (4–5), 899–934.
- Grimmer, Justin, Eitan Hersh, Marc Meredith, Jonathan Mummolo, and Clayton Nall**, “Obstacles to Estimating Voter ID Laws’ Effects on Turnout,” *Journal of Politics*, 2018, *80* (3), 1045–1051.
- McNulty, John E., Conor M. Dowling, and Margaret H. Ariotti**, “Driving Saints to Sin: How Increasing the Difficulty of Voting Dissuades Even the Most Motivated Voters,” *Political Analysis*, 2009, *17* (4), 435–455.
- MIT Election Data + Science Lab**, “County Presidential Election Returns 2000–2024,” Harvard Dataverse 2024. <https://doi.org/10.7910/DVN/VOQCHQ>.
- Postal Regulatory Commission**, “Advisory Opinion on Post Office Structure Plan,” Technical Report, PRC 2015.
- Rogowski, Jon C., Alan S. Gerber, and James G. Gimpel**, “Voter ID Laws and the Composition of the Electorate,” *Journal of Politics*, 2020, *82* (4), 1254–1270.
- Thompson, Daniel M., Jennifer A. Wu, Jesse Yoder, and Andrew B. Hall**, “Universal Vote-by-Mail Has No Impact on Partisan Turnout or Vote Share,” *Proceedings of the National Academy of Sciences*, 2020, *117* (25), 14052–14056.
- United States Postal Service**, “Annual Report to Congress: Fiscal Year 2012,” Technical Report, USPS 2012.

Verdier, Thierry and Yves Zenou, “Social Networks and the Spread of Ideas,” *Journal of Political Economy*, 2024, 132 (1), 1–45.

A. Data Appendix

QCEW data construction. I access the BLS QCEW public data API for years 2014–2023.¹ Each record provides the annual average number of establishments and employment for NAICS 491110 (Postal Service) by county and ownership code. I retain records with ownership code 1 (federal government) and aggregation level 78 (county-level, 6-digit NAICS). Records with non-empty disclosure codes are dropped. The resulting panel contains 32,545 county-year observations across 3,263 unique counties.

Election data construction. The MIT Election Data + Science Lab county presidential returns file (version 2024-12, accessed March 2026) provides candidate-level vote counts by county, year, and voting mode. I aggregate candidate votes to county-year totals using rows with `mode = "TOTAL"`. Voting mode data (ABSENTEE, MAIL, ELECTION DAY, EARLY, etc.) is available for a subset of states in 2024 only; I use these to compute mail ballot shares where available. FIPS codes are standardized to 5-digit character strings.

Panel construction. The balanced panel retains counties observed in all ten QCEW years (2014–2023) and all seven presidential elections (2000–2024). Of 3,263 QCEW counties and 3,157 election counties, 1,791 appear in both with complete records. Treatment assignment requires QCEW data; election years before QCEW coverage (2000–2012) are assigned to the pre-treatment period.

B. Identification Appendix

Parallel trends pre-test. The Callaway-Sant’Anna pre-test statistic rejects the null of parallel pre-treatment trends ($p = 0.004$). The rejection is driven primarily by the 2020 treatment cohort, whose pre-treatment coefficients decline monotonically from +0.039 (2000) to +0.003 (2012). This convergence pattern is consistent with larger counties—which have more post offices and thus more to lose—experiencing slower turnout growth relative to smaller counties over 2000–2012, prior to any USPS closures.

Bacon decomposition. The Bacon (2021) decomposition shows that 96.9% of the TWFE estimator’s weight derives from treated-vs-untreated comparisons (estimate: -0.058), with 2.0% from earlier-vs-later treated ($+0.011$) and 1.0% from later-vs-earlier treated (-0.005). The dominance of clean comparisons rules out heterogeneous-treatment-timing bias as a source of the TWFE estimate.

¹Endpoint: data.bls.gov/cew/data/api/{year}/a/industry/491110.csv.

C. Robustness Appendix

Leave-one-state-out. Sequentially dropping each of the 29 states in the balanced panel yields TWFE coefficients ranging from -0.032 to -0.021 (mean: -0.025 , SD: 0.002). No individual state exerts undue influence on the main estimate.

Randomization inference. I randomly permute treatment assignment across counties 500 times and re-estimate the TWFE specification. The actual coefficient (-0.025) exceeds all 500 placebo coefficients in absolute value, yielding a two-sided RI p -value below 0.002. The 95% randomization interval is $[-0.016, 0.015]$.

D. Standardized Effect Sizes

Table 6: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
<i>Panel A: Pooled</i>						
Log total votes	-0.0141	0.0058	1.5021	-0.0094	0.0039	Small negative
Turnout rate	-0.0057	0.0032	0.1155	-0.0493	0.0278	Small negative
<i>Panel B: Heterogeneous (sample splits)</i>						
Urban counties	-0.0251	0.0067	1.5021	-0.0167	0.0045	Small negative
Rural counties	-0.0178	0.0087	1.5021	-0.0118	0.0058	Small negative

Notes: **Country:** United States. **Research question:** Does the loss of USPS post office establishments reduce voter turnout in presidential elections? **Policy mechanism:** The USPS Retail Access Optimization Initiative (RAOI, 2011–2017) closed approximately 480 post offices nationwide as part of a cost-reduction strategy driven by the 2006 Postal Accountability and Enhancement Act’s pre-funding mandate, reducing physical postal access points in affected communities. **Outcome definition:** Log total presidential votes at the county level from MIT Election Data + Science Lab county presidential returns; turnout rate defined as total votes divided by voting-age population proxy (Census ACS population $\times 0.76$). **Treatment:** Binary indicator for counties that lost at least one USPS establishment (NAICS 491110) between 2014 and 2018 as measured by the BLS Quarterly Census of Employment and Wages. **Data:** BLS QCEW (2014–2023), MIT Election Lab county presidential returns (2000–2024), Census ACS 5-year estimates (2015, 2019, 2022); balanced panel of 1,791 counties observed across seven presidential elections. **Method:** Callaway and Sant’Anna (2021) doubly robust staggered difference-in-differences with never-treated controls; standard errors are analytical. Note: parallel trends pre-test rejects ($p = 0.004$), so estimates may reflect pre-existing convergence rather than a causal effect. **Sample:** Counties with non-missing QCEW USPS data across all years (2014–2023) and non-missing presidential vote totals in all seven election cycles; 514 treated, 1,277 never-treated. $SDE = \hat{\beta}/SD(Y)$ where $SD(Y)$ is the pre-treatment standard deviation. Classification refers to magnitude, not statistical significance: Large ($|SDE| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).