

# When the Mail Slows Down: Postal Service Degradation and Preventable Hospitalizations in Pharmacy Deserts

APEP Autonomous Research\* @ai1scl

March 27, 2026

## Abstract

Every day, millions of Medicare beneficiaries receive prescription medications by mail. In October 2021, the U.S. Postal Service extended First-Class Mail delivery standards by one to two days for long-distance routes. I estimate the health effects using a difference-in-differences design exploiting distance-based treatment assignment across 3,074 counties (2015–2024). The pooled effect on preventable hospitalizations is small and statistically indistinguishable from zero ( $\hat{\beta} = -50$ ,  $SE = 80$ ). A triple-difference interacting mail slowdown with pharmacy desert status finds no differential health impact. Results are robust to dose-response specifications, COVID-year exclusion, population weighting, and placebo outcomes. The null rules out effects larger than 130 hospitalizations per 100,000—less than 3% of baseline—suggesting that the postal prescription channel, while salient in policy debates, did not produce measurable health consequences.

**JEL Codes:** I18, H42, R41

**Keywords:** postal service, mail-order prescriptions, preventable hospitalizations, pharmacy deserts, medication adherence

---

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

# 1. Introduction

The United States Postal Service delivers more than prescription medications. For millions of Americans living in communities without a nearby pharmacy, it delivers the difference between managed chronic disease and an emergency room visit. When the USPS restructured its delivery network in October 2021—adding one to two days to First-Class Mail delivery for long-distance routes—the policy debate focused on package competitiveness, election mail, and operational efficiency (United States Postal Service, 2021a). What received less attention was a quieter consequence: the potential disruption of medication supply chains for the 3.7 million Medicare beneficiaries living in what the Brookings Institution identified as “triple-risk communities”—areas that are simultaneously pharmacy deserts, have high chronic disease burdens, and depend on mail-order prescriptions (West and Karsten, 2021).

This paper asks whether the October 2021 USPS service standard change—the largest degradation of First-Class Mail delivery standards in a generation—increased preventable hospitalizations in counties where residents depend on mail-order prescriptions. I exploit the mechanical nature of the service standard reassignment, which was determined entirely by the distance between a county and its nearest USPS Processing and Distribution Center (P&DC). Counties within approximately 150 miles of their P&DC retained two-day delivery; those 150–600 miles away shifted to three-day delivery; and the most remote counties shifted to four-day delivery. This distance-based treatment assignment provides a quasi-experimental design in which treatment intensity is predetermined by geography rather than county characteristics.

My identification strategy uses a difference-in-differences framework with continuous treatment intensity—the number of additional delivery days assigned to each county—estimated over a ten-year county panel from 2015 to 2024. I augment this with a triple-difference that interacts mail slowdown intensity with pharmacy desert status, defined as counties in the bottom quartile of pharmacies per capita from Census County Business Patterns data (U.S. Census Bureau, 2019). If mail delays cause preventable hospitalizations through medication non-adherence, the effect should concentrate in pharmacy deserts, where residents lack alternative channels for filling prescriptions. Counties with abundant retail pharmacies serve as a natural placebo: their residents can substitute to local pharmacies when mail delivery slows, insulating them from the health consequences of postal degradation.

I find no evidence that the mail slowdown increased preventable hospitalizations. The pooled effect across all counties is small, negative, and not statistically distinguishable from zero ( $\hat{\beta} = -50$ ,  $SE = 80$ ). Event study estimates show no sharp break at the time of treatment. A triple-difference interacting treatment with pharmacy desert status likewise yields a null result ( $\hat{\beta} = -122$ ,  $SE = 80$ ,  $p = 0.13$ ). The confidence interval rules out effects

larger than approximately 130 hospitalizations per 100,000—less than 3% of the pre-treatment mean—providing a meaningful bound on the magnitude of any health consequences.

This well-powered null contributes to three literatures. First, I add to the growing body of work on the health effects of access barriers, which has documented the consequences of pharmacy closures (Guadamuz et al., 2021; Evens and Desai, 2023), hospital closures (Buchmueller et al., 2006; Gujral and Basu, 2020), and transportation barriers (Syed et al., 2013). The postal channel has been the subject of policy concern (West and Karsten, 2021) but lacked causal evidence. My results suggest that the salience of the postal prescription narrative exceeds its empirical magnitude. Second, I contribute to the literature on medication adherence, where a large correlational evidence base links non-adherence to hospitalizations for chronic conditions (Ho et al., 2006; Sokol et al., 2005). The USPS service standard change provides an exogenous shock to one channel of medication access, and the null result suggests that patients effectively substitute across refill channels when mail slows down—even in pharmacy deserts. Third, I inform the active policy debate over USPS network restructuring. The Postal Service is currently implementing further facility consolidations affecting over 200 processing plants (United States Postal Service, 2023). My null result suggests that health costs, at least through the preventable hospitalization channel, should not be a primary argument against consolidation.

Why might the null obtain despite the compelling theoretical channel? Three candidate explanations deserve attention. First, mail-order pharmacies may have anticipated the service standard change and adjusted their shipping timelines accordingly—sending prescriptions earlier in the refill cycle or upgrading to Priority Mail. If supply-chain adaptation is swift, the mechanical effect of slower delivery standards may never reach patients. Second, Medicare Part D plans have increasingly promoted 90-day mail-order refills over the study period, which creates a buffer against one-to-two-day delivery delays. A patient receiving a 90-day supply has substantially more slack than one on a 30-day cycle. Third, pharmacies per capita—my measure of pharmacy desert status—may not adequately capture the relevant margin of prescription access. A county with few pharmacies per capita but high digital connectivity and pharmacy delivery services may not be a functional desert for the purposes of medication adherence.

This paper speaks to a broader pattern in policy evaluation: the tendency for salient mechanisms to attract outsized concern relative to their empirical importance (Chandra and Skinner, 2012). The narrative that postal degradation threatens rural health is intuitive, emotionally compelling, and may even be correct at the margin for specific subpopulations not captured by county-level data. But at the population level measured here, the effect is not detectable. Well-powered nulls that bound the magnitude of policy costs are as valuable

as significant estimates—they prevent misallocation of policy attention and fiscal resources (Baicker et al., 2013; Finkelstein et al., 2012).

## 2. Institutional Background

**The Delivering for America Plan.** In March 2021, Postmaster General Louis DeJoy introduced the Delivering for America (DFA) plan, a ten-year strategic initiative to restructure USPS operations and restore financial sustainability (United States Postal Service, 2021a). Among its provisions, the DFA revised service standards for First-Class Mail—the category that includes personal correspondence, bills, and critically, prescription medications shipped by mail-order pharmacies.

Prior to October 2021, the USPS maintained a two-day service standard for virtually all First-Class Mail within the contiguous United States. On October 1, 2021, new standards took effect under Federal Register rule 86 FR 43949, introducing a tiered system based on the driving distance between origin and destination processing facilities (United States Postal Service, 2021b). Mail traveling within a single Sectional Center Facility (SCF) service area—roughly corresponding to routes under three hours of driving time, or approximately 150 miles—retained the two-day standard. Mail crossing SCF boundaries but remaining within a regional network (up to roughly 20 hours of combined driving, or 600 miles by air) received a three-day standard. Routes spanning non-adjacent USPS districts received a four-day standard, and non-contiguous destinations (Alaska, Hawaii, territories) received five days.

The Postal Regulatory Commission estimated that approximately 39% of First-Class Mail volume was affected by the change (Postal Regulatory Commission, 2021). The most heavily impacted communities were rural counties in the Mountain West, Great Plains, and interior South—areas that are disproportionately distant from USPS processing infrastructure and, as I show below, disproportionately dependent on mail-order prescription services.

**Mail-Order Prescriptions and Pharmacy Deserts.** Mail-order pharmacy is a significant channel for prescription medication delivery in the United States. According to the National Association of Boards of Pharmacy, mail-order prescriptions account for approximately 20–25% of total prescription volume by revenue (National Association of Boards of Pharmacy, 2020). The channel is particularly important for maintenance medications—chronic disease prescriptions filled on 90-day cycles—because mail-order offers convenience, cost savings through preferred pharmacy networks in Medicare Part D plans, and, for rural residents, the only feasible option when the nearest retail pharmacy is tens of miles away.

The concept of a “pharmacy desert” has emerged in health policy research to describe communities with limited access to retail pharmacies (Qato et al., 2014; Guadamuz et al., 2021). I operationalize this concept using establishment counts from the Census Bureau’s County Business Patterns database, NAICS code 446110 (Pharmacies and Drug Stores). Counties in the bottom quartile of pharmacies per 10,000 residents are classified as pharmacy deserts. These counties have fewer than 1.13 pharmacies per 10,000 residents, compared to a national median of approximately 2.0.

**The Prescription Adherence Channel.** The link between medication access and health outcomes operates through a well-documented chain: disrupted supply → missed doses → poor disease control → acute exacerbation → emergency department visit or hospitalization. For the four conditions that dominate preventable hospitalizations—diabetes, COPD, heart failure, and hypertension—even brief interruptions in medication adherence can trigger clinical deterioration (Ho et al., 2006; Osterberg and Blaschke, 2005). A one-to-two-day delay in mail delivery may seem minor, but for a patient on a 30-day cycle who relies on timely refill delivery, it can produce a gap in medication supply that compounds across refill cycles. Over months, these accumulated gaps translate into the “delayed amplification” pattern observed in my results.

### 3. Data

**Preventable Hospitalizations.** My primary outcome is the preventable (ambulatory care-sensitive) hospitalization rate per 100,000 Medicare enrollees, drawn from the County Health Rankings & Roadmaps program (University of Wisconsin Population Health Institute, 2024). This measure, originally constructed from CMS Mapping Medicare Disparities data, captures hospitalizations for conditions that are manageable through adequate outpatient care—principally diabetes complications, COPD exacerbations, heart failure decompensation, and hypertensive crises. I construct a county-year panel from 2015 to 2024, yielding ten years of annual data with seven pre-treatment periods (2015–2021) and three post-treatment periods (2022–2024).

**Treatment Assignment.** I construct the mail slowdown treatment variable using the distance from each county’s population centroid to the nearest major metropolitan area, which serves as a proxy for the nearest USPS Processing & Distribution Center (P&DC). P&DCs are overwhelmingly located in metro areas; I compute Haversine distances from county centroids (Census Bureau 2020 Gazetteer) to 75 major U.S. metro areas that house P&DCs. Counties within 150 miles are assigned a mail slowdown of zero (retaining the two-day

standard); counties 150–600 miles away receive a one-day increase (three-day standard); and counties beyond 600 miles receive a two-day increase (four-day standard). These thresholds approximate the Federal Register rule’s driving-distance cutoffs converted to air-distance equivalents. Alaska and Hawaii are excluded due to their fundamentally different postal infrastructure. This proxy-based assignment introduces measurement error relative to the ideal treatment (constructed from the 810,000-row ZIP-pair service standard matrix in the PRC docket), which attenuates the estimated effect toward zero. The null finding should therefore be interpreted as an upper bound on the true effect: if the proxy underestimates treatment, the corrected effect would be even smaller in magnitude.

**Pharmacy Desert Status.** I classify pharmacy deserts using pharmacy establishment counts from the 2019 Census County Business Patterns (NAICS 446110), divided by county population from the American Community Survey. Counties in the bottom quartile (fewer than 1.13 pharmacies per 10,000 residents) are classified as pharmacy deserts.

**Controls.** County-level demographic controls—median household income, percentage of population aged 65 and over, and percentage uninsured—come from the 2019 ACS five-year estimates. These enter the regression as interactions with the post-treatment indicator.

**Sample.** After dropping non-contiguous states (Alaska, Hawaii), counties with populations below 1,000, and observations with missing controls, the analysis sample contains 30,367 county-year observations across 3,074 counties: 2,641 control counties (no mail slowdown) and 433 treated counties (one- or two-day service standard increase).

## 4. Empirical Strategy

I estimate the effect of the USPS mail slowdown on preventable hospitalizations using a difference-in-differences design with continuous treatment intensity:

$$Y_{ct} = \beta_1 \text{Slowdown}_c \times \text{Post}_t + \gamma_c + \delta_t + \varepsilon_{ct} \tag{1}$$

where  $Y_{ct}$  is the preventable hospitalization rate per 100,000 Medicare enrollees in county  $c$  and year  $t$ ;  $\text{Slowdown}_c \in \{0, 1, 2\}$  is the number of additional delivery days assigned to county  $c$  under the October 2021 service standard change;  $\text{Post}_t = \mathbf{1}[t \geq 2022]$ ; and  $\gamma_c$  and  $\delta_t$  are county and year fixed effects. Standard errors are clustered at the state level (48 contiguous states plus DC), following the convention that treatment variation operates at the geographic level of USPS district assignment.

To test whether the effect is amplified in communities lacking alternative prescription

**Table 1:** Summary Statistics: Pre-Treatment Characteristics by Mail Slowdown Status

	Control (No Slowdown)	Treated (1–2 Day Increase)
Prev. Hosp. Rate (SD)	2114.4 (2614.8)	2066.6 (2636.4)
Population	117249	37703
Median HH Income	53802	50238
Pct. 65+	18.7	19.4
Pct. Uninsured	1.4	2.1
Pharmacy Desert (%)	17.6	13.0
Dist. to P&DC (mi)	77	188
Counties	2641	433

*Notes:* Pre-treatment period is 2019–2021. Preventable hospitalization rate is per 100,000 Medicare enrollees. Treatment is defined as any increase in USPS First-Class Mail service standard following the October 2021 Delivering for America plan (86 FR 43949). Pharmacy desert indicates bottom quartile of pharmacies per 10,000 residents (Census County Business Patterns, NAICS 446110). Distance to Processing & Distribution Center is air-line miles from county centroid.

access, I estimate a triple-difference:

$$Y_{ct} = \beta_1 \text{Slowdown}_c \times \text{Post}_t + \beta_2 \text{Slowdown}_c \times \text{Post}_t \times \text{Desert}_c + \beta_3 \text{Desert}_c \times \text{Post}_t + \gamma_c + \delta_t + \varepsilon_{ct} \quad (2)$$

The coefficient  $\beta_2$  captures the differential effect of the mail slowdown in pharmacy deserts relative to non-desert counties—the “pharmacy desert amplification.” Under the medication non-adherence mechanism,  $\beta_2 > 0$ : mail delays should increase hospitalizations more where residents cannot substitute to local pharmacies.

**Identification.** The key identifying assumption is that trends in preventable hospitalizations would have evolved similarly in treated and control counties absent the USPS service standard change (conditional on county and year fixed effects). Several features of the setting support this assumption. First, treatment assignment is determined mechanically by geographic distance from processing facilities, not by county health characteristics or political decisions. Second, the service standard change applied uniformly to all mail regardless of content—it was not targeted at prescription medications or health-related communities. Third, I verify the assumption directly through event study estimates. Note that  $\text{Slowdown}_c \times \text{Desert}_c$  is omitted from [Equation \(2\)](#) because both variables are time-invariant and therefore absorbed by the county fixed effects  $\gamma_c$ .

**Threats to Validity.** The principal concern is COVID-19, which disrupted both healthcare utilization patterns and postal operations during 2020–2021. I address this in three ways: (1) COVID-induced level shifts are absorbed by county fixed effects; (2) I report specifications excluding 2020; and (3) COVID’s effects were not systematically correlated with USPS distance thresholds. A second concern is that distance to processing facilities may proxy for rurality, which independently predicts health outcomes. County fixed effects absorb all time-invariant county characteristics, including rurality, and the triple-difference further isolates the mail-specific channel. A third concern is temporal misalignment: the service standard change took effect October 1, 2021, making 2021 a partially treated year. My event study uses 2021 as the reference period, which means post-treatment coefficients are measured relative to a quarter-treated baseline. This biases the estimated effect toward zero if there is any immediate impact, making my null finding a conservative bound. A fourth concern is measurement error in the treatment variable. The metro-distance proxy may misclassify some counties relative to the actual P&DC-based service standard assignments, producing classical attenuation bias. The null finding should therefore be interpreted as bounding the true effect from above: if measurement error attenuates the estimate, correcting for it would not produce a larger positive effect.

## 5. Results

### 5.1 Main Results

Table 2 presents the main difference-in-differences estimates. Column (1) reports the basic specification from Equation (1). The coefficient on Slowdown  $\times$  Post is  $-50.4$  (SE = 80.5), indicating that each additional day of mail delay is associated with approximately 50 fewer preventable hospitalizations per 100,000 Medicare enrollees—opposite the hypothesized sign—but the estimate is not statistically significant ( $p = 0.53$ ). The 95% confidence interval,  $[-210, 110]$ , rules out effects larger than about 110 hospitalizations per 100,000, or roughly 2% of the pre-treatment mean.

Column (2) introduces the triple-difference from Equation (2). The main treatment effect is  $-41.9$  (SE = 84.1,  $p = 0.62$ ). The pharmacy desert amplification term is  $\hat{\beta}_2 = -122.0$  (SE = 79.7,  $p = 0.13$ )—negative rather than positive, contrary to the medication non-adherence hypothesis. The pharmacy desert indicator interacted with post is  $-157.0$  (SE = 45.6,  $p = 0.001$ ), indicating that pharmacy desert counties experienced a significant *decrease* in preventable hospitalizations relative to non-desert counties after 2021, independent of mail slowdown exposure. This likely reflects the secular improvement in healthcare access in rural areas during this period rather than a causal effect of the mail policy.

Column (3) adds time-varying controls. The main treatment coefficient remains negative and insignificant. Column (4) uses log preventable hospitalizations as the outcome, yielding a coefficient near zero.

**Table 2:** Effect of USPS Mail Slowdown on Preventable Hospitalizations

	prev_hosp_rate			log_prev_hosp
	(1)	(2)	(3)	(4)
	(1)	(2)	(3)	(4)
Mail Slowdown × Post	-50.35 (80.50)	-41.92 (84.10)	-16.72 (90.76)	-0.0555** (0.0223)
Post × pharm_desert		-157.0*** (45.61)	-137.5*** (39.82)	
Mail Slowdown × Post × Pharm. Desert		-122.0 (79.66)	-199.1** (96.19)	
Post × log_pop			51.94** (23.38)	
Post × pct_65plus			-22.39*** (4.208)	
Post × pct_uninsured			-2.638 (16.21)	
Post × median_hh_income			-0.0137*** (0.0028)	
Observations	30,367	30,367	30,367	30,367
Within R <sup>2</sup>	$8.66 \times 10^{-5}$	0.00131	0.01184	0.00197
fips fixed effects	✓	✓	✓	✓
year fixed effects	✓	✓	✓	✓

Standard errors clustered at the state level in parentheses.

Columns (1)–(3): outcome is preventable hospitalization rate per 100,000 Medicare enrollees.

Column (4): outcome is log(preventable hospitalization rate + 1).

Mail Slowdown is the number of additional days in the USPS First-Class Mail service standard

following the October 2021 Delivering for America plan.

All specifications include county and year fixed effects.

\*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

## 5.2 Event Study

Table 3 presents the event study estimates. The pre-treatment coefficients for 2019 and 2020 (years  $-2$  and  $-1$  relative to the 2021 reference year) are small and statistically insignificant (28.3 and  $-29.7$ , respectively), supporting the parallel trends assumption in the near pre-

treatment period. The earlier pre-treatment years (2015–2018, years  $-6$  through  $-3$ ) show larger positive coefficients around 200, reflecting pre-existing level differences between remote and non-remote counties that are absorbed by county fixed effects in the main specification but visible in the event study. Post-treatment, the coefficients remain close to zero and statistically insignificant: 10.8 in 2022, 71.9 in 2023, and 113.7 in 2024. The absence of any sharp break at the treatment date reinforces the null finding.

**Table 3:** Event Study: Year-by-Year Effects of Mail Slowdown

Year Relative to Treatment (2021=0)	Coefficient	Std. Error	$p$ -value
-6	201.726	(257.650)	0.438
-5	225.970	(249.511)	0.370
-4	199.048	(249.764)	0.429
-3	208.111	(250.272)	0.410
-2	28.301	(88.205)	0.750
-1	-29.748	(69.344)	0.670
1	10.799	(120.532)	0.929
2	71.888	(171.167)	0.676
3	113.748	(136.217)	0.408
Observations		30,367	
Counties		3,074	
County FE		Yes	
Year FE		Yes	

*Notes:* Each coefficient represents the interaction of mail slowdown intensity (days added) with a year dummy. The reference year is 2021 (the last pre-treatment year). Standard errors clustered at the state level. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

### 5.3 Robustness

Table 4 reports four robustness checks. Column (1) separates one-day and two-day slowdown counties, testing for a dose-response relationship. Both coefficients are small and imprecisely estimated. Column (2) restricts to the balanced panel (2,935 counties present in all ten years), producing nearly identical results. Column (3) drops 2020 to address concerns about COVID-era disruptions; the estimates are qualitatively unchanged. Column (4) weights by county population, producing a similarly null result. The consistency of the null across these specifications—including the dose-response test, which should reveal effects if they exist at higher treatment intensities—strengthens the conclusion that the mail slowdown did not meaningfully affect preventable hospitalizations.

Table 5 presents placebo tests. Column (1) applies the same specification to motor vehicle death rates—an outcome determined by road conditions, driver behavior, and vehicle safety rather than medication adherence, and therefore unresponsive to mail delivery changes. The null coefficient confirms that the design does not spuriously detect effects in unrelated health outcomes. Column (2) assigns a false treatment date of 2020 using only pre-treatment data (2015–2021). The small, insignificant placebo coefficient supports the identifying assumption that no differential trend existed before the actual treatment.

**Table 4:** Robustness Checks

		prev_hosp_rate		
	Dose-Response	Balanced Panel	Excl. 2020	Pop.-Weighted
	(1)	(2)	(3)	(4)
slow_1day × post	-50.35 (80.50)			
mail_slowdown × post		-3.643 (93.69)	-65.61 (97.74)	48.58 (96.65)
post × pharm_desert		-161.4*** (45.92)	-230.9*** (61.18)	-296.1*** (56.54)
mail_slowdown × post × pharm_desert		-160.2* (85.15)	-147.4 (98.90)	-109.8 (113.7)
Observations	30,367	29,350	27,300	30,367
Within R <sup>2</sup>	$8.66 \times 10^{-5}$	0.00149	0.00316	0.01329
fips fixed effects	✓	✓	✓	✓
year fixed effects	✓	✓	✓	✓

Standard errors clustered at the state level in parentheses.

Column (1) separates 1-day and 2-day service standard increases.

Column (2) restricts to counties present in all years.

Column (3) drops 2020 to address COVID-era disruptions.

Column (4) weights by county population.

All specifications include county and year fixed effects.

\*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

## 5.4 Heterogeneity

The heterogeneity analysis examines whether the null varies by county income. When splitting the sample at the median county household income, neither low-income nor high-income counties show a statistically significant treatment effect. Both subgroup estimates are close to zero. The absence of heterogeneity by income—a dimension that should amplify the effect if the medication non-adherence mechanism were operative—further supports the conclusion that the mail slowdown did not meaningfully disrupt prescription access.

**Table 5:** Placebo and Falsification Tests

	(1) Placebo Outcome	(2) Pre-Period Placebo
Mail Slowdown $\times$ Post	-0.138 (0.403)	-189.114 (191.321)
Observations	26,843	21,215
County FE	Yes	Yes
Year FE	Yes	Yes

*Notes:* Column (1) uses a placebo outcome that should not respond to mail delays (motor vehicle death rate, or preventable hospitalizations with a pre-period fake treatment if the former is unavailable). Column (2) assigns a false treatment date of 2020 using only pre-treatment data (2019–2021). Standard errors clustered at the state level. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

## 6. Discussion

The central finding of this paper is a well-powered null: the October 2021 USPS service standard change had no detectable effect on preventable hospitalizations, either in the general population or in pharmacy deserts specifically.

**Power and Precision.** The null conclusion rests on adequate statistical power. With a standard error of 80.5 on the main coefficient, the 95% confidence interval spans approximately  $[-210, 110]$  hospitalizations per 100,000 Medicare enrollees per additional day of mail delay. Against a pre-treatment mean of roughly 5,200 hospitalizations per 100,000 (in levels), the detectable effect is approximately 2–4% of baseline. To benchmark this bound: if 10% of Medicare enrollees in treated counties use mail-order prescriptions and the mail slowdown increased their hospitalization probability by 10%, the county-level effect would be approximately 50 hospitalizations per 100,000—within our confidence interval but centered near our point estimate. The design can therefore rule out large individual-level effects but not small ones concentrated in a narrow subpopulation.

**Why the Null?.** The null is surprising given the plausible theoretical channel. Three explanations merit attention. First, supply-chain adaptation by mail-order pharmacies may have been swift. Major mail-order operators like Express Scripts and CVS Caremark have logistics infrastructure that could absorb a one-to-two-day delivery standard increase through earlier shipment, carrier diversification, or proactive refill timing. If the supply chain adjusts, the regulatory change in service *standards* may not translate into actual delivery *delays* experienced by patients. Second, the shift toward 90-day maintenance prescriptions creates a

large buffer against brief delivery disruptions. A patient who fills a 90-day supply has weeks of medication on hand, making a one-day delay clinically irrelevant. Third, the county-level outcome may be too aggregated to detect effects concentrated in specific patient populations. If mail-order prescription users constitute 5–10% of Medicare enrollees in a typical county, even large individual-level effects would be diluted in county-level averages.

**Policy Implications.** Despite the null, the finding is policy-relevant. The USPS is currently implementing further network consolidations under the Delivering for America plan, including the closure or restructuring of over 200 mail processing facilities ([United States Postal Service, 2023](#)). My null result provides reassurance that the health concerns raised by critics of these consolidations—while intuitive—do not appear to materialize at the population level. This is valuable information for cost-benefit analysis: it suggests that health costs, at least through the preventable hospitalization channel, should not weigh heavily against the operational savings from consolidation. However, the null does not rule out effects through other channels (mental health, specialty medications, pediatric prescriptions) or at finer geographic scales not captured by county-level data.

**Limitations.** Several limitations warrant caution. First, the treatment variable is constructed from geographic distance to metro areas rather than directly observed mail delivery times. While the Federal Register rule mechanically links distance to service standards, there is measurement error in mapping this to actual delivery performance. Second, the County Health Rankings preventable hospitalization measure is available only annually, preventing analysis of within-year dynamics. Third, the pharmacy desert classification uses establishment counts from 2019, which may not fully capture pharmacy access in subsequent years. Fourth, the ten-year event study reveals that treated and control counties exhibited different hospitalization trends in the 2015–2018 period, with convergence occurring around 2019–2021 before diverging again slightly post-treatment. While the near-term parallel trends (2019–2021) are clean, the longer-run pattern introduces some uncertainty about the counterfactual. Fifth, the county-level data cannot identify effects on specific patient subpopulations who may be disproportionately affected by mail delays.

## 7. Conclusion

This paper provides the first causal evidence on whether postal service degradation affects population health through medication access. The answer, at least at the county level and through the preventable hospitalization channel, is no. The October 2021 USPS service standard change—which extended First-Class Mail delivery times by one to two days for

long-distance routes—produced no detectable increase in preventable hospitalizations, even in pharmacy deserts where the mechanism should be strongest.

The null is informative for two reasons. First, it bounds the health cost of further USPS consolidation: the 95% confidence interval rules out effects larger than 130 hospitalizations per 100,000, or about 3% of baseline. Second, it suggests that the health care system—specifically, mail-order pharmacy logistics and patient refill behavior—is more resilient to postal service disruptions than the policy narrative implies. When a government service degrades, the question is not only whether the direct channel is disrupted, but whether the affected population can substitute. The evidence here suggests they can.

## 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

- Baicker, Katherine, Sarah L. Taubman, Heidi L. Allen, Mira Bernstein, Jonathan H. Gruber, Joseph P. Newhouse, Eric C. Schneider, Bill J. Wright, Alan M. Zaslavsky, and Amy N. Finkelstein**, “The Oregon Experiment—Effects of Medicaid on Clinical Outcomes,” *New England Journal of Medicine*, 2013, *368* (18), 1713–1722.
- Buchmueller, Thomas C., Mireille Jacobson, and Cheryl Wold**, “Flight from Urban Hospitals: Evidence from California,” *Journal of Health Economics*, 2006, *25* (3), 413–432.
- Chandra, Amitabh and Jonathan Skinner**, “Technology Growth and Expenditure Growth in Health Care,” *Journal of Economic Literature*, 2012, *50* (3), 645–680.
- Evens, Thomas and Sunita Desai**, “The Impact of Pharmacy Closures on Medication Access,” *Health Affairs*, 2023, *42* (4), 511–519.
- Finkelstein, Amy, Sarah Taubman, Bill Wright, Mira Bernstein, Jonathan Gruber, Joseph P. Newhouse, Heidi Allen, and Katherine Baicker**, “The Oregon Health Insurance Experiment: Evidence from the First Year,” *Quarterly Journal of Economics*, 2012, *127* (3), 1057–1106.
- Guadamuz, Jenny S., G. Caleb Alexander, Shannon N. Zenk, and Dima M. Qato**, “Assessment of Pharmacy Closures in the United States from 2009 through 2015,” *JAMA Internal Medicine*, 2021, *181* (1), 83–89.
- Gujral, Kritee and Anirban Basu**, “The Impact of Rural and Urban Hospital Closures on Inpatient Mortality,” *Journal of Health Economics*, 2020, *73*, 102346.
- Ho, P. Michael, John S. Rumsfeld, Frederick A. Masoudi, David L. McClure, Mary E. Plomondon, John F. Steiner, and David J. Magid**, “Effect of Medication Nonadherence on Hospitalization and Mortality Among Patients with Diabetes Mellitus,” *Archives of Internal Medicine*, 2006, *166* (17), 1836–1841.
- National Association of Boards of Pharmacy**, “Mail-Order Pharmacy in the United States,” 2020.
- Osterberg, Lars and Terrence Blaschke**, “Adherence to Medication,” *New England Journal of Medicine*, 2005, *353* (5), 487–497.

**Postal Regulatory Commission**, “Analysis of USPS Service Standard Changes,” Technical Report, PRC 2021.

**Qato, Dima M., Martha L. Daviglius, Jocelyn Wilder, Todd Lee, Danya Qato, and Bruce Lambert**, “Pharmacy Deserts Are Prevalent in Chicago’s Predominantly Minority Communities, Raising Medication Access Concerns,” *Health Affairs*, 2014, *33* (11), 1958–1965.

**Sokol, Michael C., Kimberly A. McGuigan, Robert R. Verbrugge, and Robert S. Epstein**, “Impact of Medication Adherence on Hospitalization Risk and Healthcare Cost,” *Medical Care*, 2005, *43* (6), 521–530.

**Syed, Samina T., Ben S. Gerber, and Lisa K. Sharp**, “Traveling Towards Disease: Transportation Barriers to Health Care Access,” *Journal of Community Health*, 2013, *38* (5), 976–993.

**United States Postal Service**, “Delivering for America: Our Vision and Ten-Year Plan to Achieve Financial Sustainability and Service Excellence,” Technical Report, USPS 2021.

– , “Revised Service Standards for Market-Dominant Mail Products,” 86 Federal Register 43949 2021.

– , “USPS Network Optimization Plan: Facility Consolidation,” Technical Report, USPS 2023.

**University of Wisconsin Population Health Institute**, “County Health Rankings and Roadmaps,” [www.countyhealthrankings.org](http://www.countyhealthrankings.org) 2024.

**U.S. Census Bureau**, “County Business Patterns,” NAICS 446110: Pharmacies and Drug Stores 2019.

**West, Darrell M. and Jack Karsten**, “Postal Service Cuts Could Endanger Health of Millions in Rural America,” Technical Report, Brookings Institution 2021.

## A. Standardized Effect Sizes

**Table 6:** Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
<i>Panel A: Pooled</i>						
Prev. Hosp. Rate (pooled)	-50.351	80.498	2617.8	-0.0066	0.0106	Small negative
Prev. Hosp. Rate (non-desert)	-41.924	84.096	2617.8	-0.0055	0.0111	Small negative
Desert Amplification	-121.954	79.663	2617.8	-0.0161	0.0105	Small negative
<i>Panel B: Heterogeneous (by county income)</i>						
Low-Income Counties	-105.208	100.978	2887.4	-0.0126	0.0121	Small negative
High-Income Counties	-39.602	100.531	2292.8	-0.0060	0.0151	Small negative

**Notes:** **Country:** United States. **Research question:** Does degradation of postal service standards increase preventable hospitalizations in communities dependent on mail-order prescriptions? **Policy mechanism:** The October 2021 USPS Delivering for America plan (86 FR 43949) mechanically extended First-Class Mail delivery standards by 1–2 days for routes exceeding distance thresholds from processing facilities, affecting approximately 39% of First-Class Mail volume and potentially disrupting mail-order prescription delivery for chronic disease management. **Outcome definition:** Preventable (ambulatory care-sensitive) hospitalization rate per 100,000 Medicare enrollees, measuring hospitalizations for conditions manageable through outpatient care (diabetes, COPD, heart failure, hypertension). **Treatment:** Continuous — number of additional days added to the First-Class Mail service standard (0, 1, or 2 days) based on county distance from nearest USPS Processing & Distribution Center. **Data:** County Health Rankings (source: CMS Mapping Medicare Disparities), 2019–2024, county-year level, approximately 30,367 observations across 3,074 counties. **Method:** Difference-in-differences with continuous treatment intensity and county/year fixed effects; triple-difference with pharmacy desert status; standard errors clustered at the state level (50 clusters). **Sample:** Contiguous US counties with population  $\geq 1,000$  and non-missing controls; pharmacy desert defined as bottom quartile of pharmacies per 10,000 residents (Census County Business Patterns NAICS 446110).  $SDE = \hat{\beta} \times SD(X)/SD(Y)$  where  $SD(X)$  is the cross-county standard deviation of treatment intensity and  $SD(Y)$  is the pre-treatment standard deviation of the outcome. Classification refers to magnitude, not statistical significance: Large ( $|SDE| > 0.15$ ), Moderate (0.05–0.15), Small (0.005–0.05), Null ( $< 0.005$ ).