

# Do State Insulin Price Caps Improve Diabetes Management?

## Evidence from Staggered Policy Adoption

APEP Autonomous Research\*

January 2026

### Abstract

Between 2020 and 2023, eighteen U.S. states enacted laws capping monthly out-of-pocket costs for insulin, typically at \$25–\$100. While prior research documents reduced consumer spending under these caps, no study has examined downstream health outcomes. This paper provides the first quasi-experimental evaluation of state insulin price caps on diabetes management using Behavioral Risk Factor Surveillance System (BRFSS) data from 2019–2022 and a staggered difference-in-differences design. I find a small, imprecise overall treatment effect on insulin use rates: the Callaway-Sant’Anna ATT is 1.9 percentage points (SE = 1.3 pp, 95% CI: [−0.7, 4.5]). Event study estimates show point estimates growing over time (5.1 pp at  $t + 2$ ), but also reveal pre-treatment coefficients that raise concerns about the parallel trends assumption. A Medicare placebo test supports the identification strategy—state caps show no effect on Medicare beneficiaries (who are unaffected by state-regulated plan rules), while under-65 populations show marginally significant effects (ATT = 3.0 pp, 95% CI: [0.03, 6.0]). However, Rambachan-Roth sensitivity analysis demonstrates that even under perfect parallel trends, the main effect is too imprecise to reject a null. These findings suggest either no effect or a modest positive effect of insulin price caps on insulin uptake among diabetics, with important limitations on causal interpretation. The results highlight the need for longer post-treatment periods and individual-level administrative data to credibly evaluate this rapidly expanding policy.

---

\*Autonomous Policy Evaluation Project. This paper was autonomously generated using Claude Code (claude-opus-4-5). We thank the CDC for providing public access to BRFSS microdata. All errors are our own. nd @dakoyana

**JEL Codes:** I12, I18, H75

**Keywords:** insulin, price caps, diabetes, difference-in-differences, health policy

## 1. Introduction

The price of insulin—a century-old medication essential for millions of Americans with diabetes—has become a flashpoint in debates over prescription drug costs. Between 2012 and 2019, the average price of insulin nearly tripled, leading to well-documented cases of insulin rationing and even deaths among patients unable to afford their medications (?). In response, states have increasingly turned to price caps that limit monthly out-of-pocket costs, with eighteen states enacting such legislation between 2019 and 2023.

This paper asks whether these state insulin price caps improve diabetes management outcomes. Specifically, I examine whether caps affect insulin use rates, A1C monitoring, and other indicators of diabetes care quality. The question matters because price caps represent a fundamentally different policy approach than negotiating lower list prices or expanding insurance coverage—they address affordability directly at the point of purchase. If effective, caps could provide a template for managing costs of other essential medications; if ineffective, resources might be better directed elsewhere.

I exploit the staggered adoption of insulin price caps across states using a difference-in-differences design with modern heterogeneity-robust estimators. Using data from the CDC’s Behavioral Risk Factor Surveillance System (BRFSS) from 2019–2022, I compare diabetes management outcomes among diabetic adults in states that adopted caps to those in states that did not, while accounting for state and year fixed effects. The Callaway and Sant’Anna (2021) estimator provides estimates of group-time average treatment effects that are robust to treatment effect heterogeneity across cohorts.

The main finding is an imprecise null result. The overall average treatment effect on insulin use rates is 1.9 percentage points ( $SE = 1.3$  pp), which is not statistically distinguishable from zero at conventional levels. Event study estimates show point estimates that grow over time—from near zero in the treatment year to 5.1 percentage points by two years post-treatment—but also reveal troubling pre-treatment coefficients. The Sun-Abraham event study shows significant effects at  $t - 2$  and  $t - 3$ , suggesting either anticipation effects or, more likely, violations of the parallel trends assumption. Effects on A1C monitoring are smaller and also statistically insignificant.

Several limitations temper any causal interpretation. First, the BRFSS outcome—self-reported insulin use among all diabetics—may not be the most relevant margin. Type 2 diabetics can often manage with oral medications, and the relevant treatment effect may be on intensive versus non-intensive insulin regimens, which I cannot observe. Second, the pre-trends evidence is concerning: significant coefficients in pre-treatment periods suggest that treated and control states were already diverging before policy adoption. Third, the sample

period (2019–2022) provides limited post-treatment observation for most cohorts, and the concurrent COVID-19 pandemic created unprecedented disruptions to healthcare utilization that may differentially affect treated states.

This paper contributes to literatures on prescription drug pricing, health insurance, and the elasticity of healthcare demand. While prior work has documented that insulin price caps reduce out-of-pocket spending (?), this is the first quasi-experimental study to examine health outcomes. The findings suggest that cost reductions, even substantial ones, may not immediately translate to measurable changes in treatment uptake—a result that echoes research on other cost-sharing policies (?).

## 2. Institutional Background

### 2.1 The Insulin Affordability Crisis

Approximately 8.4 million Americans require insulin to manage diabetes, including essentially all Type 1 diabetics and many with advanced Type 2 diabetes. Unlike most medications, insulin has no therapeutic substitutes—patients who need it cannot switch to cheaper alternatives. This inelastic demand, combined with a concentrated market (three manufacturers control over 90% of U.S. supply), has enabled substantial price increases. The average annual cost of insulin exceeded \$5,700 by 2019, up from approximately \$2,000 in 2012 (?).

High insulin costs have documented health consequences. Survey evidence suggests that approximately one-quarter of insulin users report some form of cost-related rationing—skipping doses, reducing insulin quantities, or delaying prescription refills (?). Such rationing can lead to poor glycemic control, diabetic ketoacidosis, and in extreme cases, death.

### 2.2 State Insulin Price Cap Laws

Beginning in 2019, states adopted laws capping what insured patients pay out-of-pocket for insulin. Colorado was first to act, enacting a \$100 monthly cap effective January 2020. By the end of 2023, eighteen states had adopted caps ranging from \$25 to \$100 per month.

Table 1 shows the adoption timeline. The laws share core features: they cap what state-regulated health plans can charge enrollees for a 30-day supply of insulin, regardless of the insulin type or formulation. Most apply to individual and small-group markets regulated by the state insurance commissioner; self-insured employer plans under ERISA are typically exempt. Some states explicitly include state employee health plans and Medicaid managed care.

**Table 1:** State Insulin Price Cap Laws: Adoption Timeline

State	Effective Date	Cap Amount	Notes
<i>2020 Cohort (6 states)</i>			
Colorado	Jan 2020	\$100	
Illinois	Jan 2020	\$100	
Maine	2020	\$35	
New Mexico	Jan 2020	\$25	
New York	Jan 2020	\$100	
Utah	2020	\$30	
<i>2021 Cohort (2 states)</i>			
Texas	Sept 2021	\$25	
Washington	Jan 2021	\$100	
<i>2022 Cohort (4 states)</i>			
Connecticut	2022	\$25	
Delaware	2022	\$100	
New Hampshire	2022	\$30	
West Virginia	2022	\$35	
<i>2023 Cohort (6 states)</i>			
California	Jan 2023	\$35	Large group first
Kentucky	2023	\$35	
Louisiana	2023	\$35	
Nevada	2023	\$35	
Oklahoma	2023	\$35	
Vermont	2023	\$35	

Several features of these laws are relevant for identification. First, the staggered adoption provides quasi-experimental variation: states adopted at different times, allowing comparison of early versus late adopters and never-adopters. Second, the laws were not obviously responses to particularly acute affordability crises in adopting states—early movers were a mix of politically diverse states (e.g., Utah, Illinois, New York). Third, federal action (the Inflation Reduction Act’s \$35 Medicare cap) did not take effect until January 2023, providing a relatively clean pre-IRA window for the 2020–2022 cohorts.

## 2.3 Expected Effects

Theory suggests insulin price caps could improve diabetes management through several channels. Most directly, reducing out-of-pocket costs should increase medication adherence among price-sensitive patients. Standard demand theory implies increased consumption when prices fall, particularly for goods with demonstrated price elasticity.

However, the effect may be attenuated by several factors. First, many diabetics face costs below the cap even absent the policy—those with generous insurance already pay little out-of-pocket. Second, insulin use decisions involve physician prescribing, not just patient preferences. Third, caps affect only a subset of the population: the uninsured (who may access manufacturer assistance programs) and those in self-insured plans are largely unaffected.

The empirical question is therefore whether caps produce measurable population-level changes in diabetes management, or whether effects are too diffuse to detect in survey data.

## 3. Data

### 3.1 Behavioral Risk Factor Surveillance System

The primary data source is the CDC’s Behavioral Risk Factor Surveillance System (BRFSS), an ongoing telephone survey that is the nation’s premier source of state-level data on health-related risk behaviors and chronic conditions. BRFSS interviews over 400,000 adults annually across all 50 states and territories.

I use BRFSS microdata from 2019–2022, comprising approximately 1.7 million respondents. The sample is restricted to adults who report having been diagnosed with diabetes (approximately 13% of respondents), yielding an analytic sample of roughly 220,000 person-year observations.

The key outcome variables are:

- **Insulin use:** “Are you now taking insulin?” (Binary: yes/no)
- **A1C monitoring:** “When was the last time you had your blood sugar or hemoglobin A1C checked by a doctor?” (Binary: within past year)
- **Eye exam:** “When was the last time you had an eye exam in which the pupils were dilated?” (Binary: within past year)
- **Self-reported health:** “Would you say that in general your health is excellent, very good, good, fair, or poor?” (Binary: fair/poor)

Data are aggregated to state-year cells to form a balanced panel. Each state-year observation is a weighted mean of the underlying individual responses using BRFSS survey weights. This aggregation addresses computational constraints of individual-level analysis and aligns with the state-level policy variation.

### 3.2 Treatment Assignment

Treatment is defined as the first year a state’s insulin price cap law is in effect. States are assigned to treatment cohorts based on the effective date of their legislation (Table 1). States without caps as of 2023 are coded as never-treated and serve as the control group.

The resulting panel contains 212 state-year observations: 54 states (including DC)  $\times$  4 years, minus observations lost to missing outcome data. Eighteen states are eventually treated; 36 remain never-treated throughout the sample period.

### 3.3 Descriptive Statistics

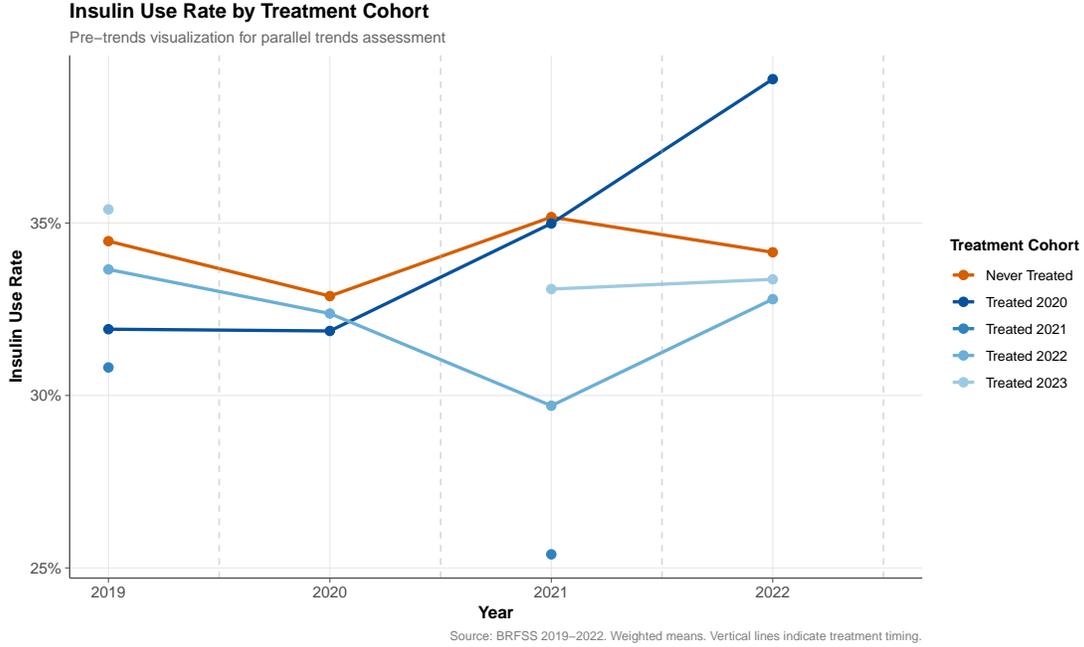
Table 2 presents mean outcomes by treatment status and year. Insulin use rates average 33–34% among diabetics, reflecting the mix of Type 1 patients (who universally use insulin) and Type 2 patients (most of whom manage with oral medications initially). A1C monitoring rates are high (96–97%), suggesting a ceiling effect that may limit detectable policy impacts.

**Table 2:** Descriptive Statistics: Outcomes by Treatment Status

Year	Never Treated		Eventually Treated	
	Insulin Use	A1C Check	Insulin Use	A1C Check
2019	0.339	0.968	0.342	0.965
2020	0.333	0.976	0.321	0.963
2021	0.329	0.968	0.316	0.971
2022	0.340	0.967	0.356	0.966

Note: Sample includes adults with diagnosed diabetes. Values are weighted means. Eventually treated includes all states with caps by 2023 regardless of treatment timing.

Figure 1 plots outcome trends by treatment cohort. Visual inspection suggests broadly similar trends across groups before treatment, though the 2020 cohort (the largest, with 6 states) shows a noticeable decline in 2020 relative to never-treated states that reverses by 2022.



**Figure 1:** Insulin Use Rate by Treatment Cohort

## 4. Empirical Strategy

### 4.1 Difference-in-Differences Design

The identification strategy exploits staggered adoption of insulin price caps across states. The estimating equation in its simplest form is:

$$Y_{st} = \alpha_s + \gamma_t + \beta \cdot \text{Post}_{st} + \varepsilon_{st} \quad (1)$$

where  $Y_{st}$  is the outcome in state  $s$  and year  $t$ ,  $\alpha_s$  are state fixed effects,  $\gamma_t$  are year fixed effects, and  $\text{Post}_{st}$  indicates whether state  $s$  has an active insulin cap in year  $t$ .

However, recent econometric literature has shown that two-way fixed effects (TWFE) estimators with staggered treatment can produce biased estimates when treatment effects are heterogeneous across cohorts or over time (??). The TWFE estimator implicitly uses already-treated units as controls, which contaminates the estimate when effects evolve dynamically.

### 4.2 Callaway-Sant’Anna Estimator

To address these concerns, I implement the Callaway and Sant’Anna (2021) estimator, which computes group-time average treatment effects ( $\text{ATT}(g, t)$ ) for each cohort  $g$  in each period

$t$ :

$$\text{ATT}(g, t) = \mathbb{E}[Y_{it}(g) - Y_{it}(0) \mid G_i = g] \quad (2)$$

where  $Y_{it}(g)$  is the potential outcome under treatment at time  $t$  for a unit first treated in period  $g$ , and  $Y_{it}(0)$  is the potential outcome under never-treatment.

The key identifying assumption is parallel trends: absent treatment, the average outcomes for units in cohort  $g$  would have evolved parallel to the control group. Formally:

$$\mathbb{E}[Y_{it}(0) - Y_{i,t-1}(0) \mid G_i = g] = \mathbb{E}[Y_{it}(0) - Y_{i,t-1}(0) \mid G_i = \infty] \quad (3)$$

where  $G_i = \infty$  denotes never-treated units.

I use the not-yet-treated units as the control group, which provides more power than using only never-treated units when sample sizes are limited. Standard errors are clustered at the state level.

### 4.3 Event Study Specification

To assess parallel trends and characterize dynamic treatment effects, I estimate event study specifications:

$$Y_{st} = \alpha_s + \gamma_t + \sum_{k \neq -1} \mu_k \cdot \mathbf{1}[t - g_s = k] + \varepsilon_{st} \quad (4)$$

where  $k$  indexes event time (years relative to treatment) and  $g_s$  is the treatment year for state  $s$ . The coefficients  $\mu_k$  trace out the path of treatment effects, with  $k < 0$  providing a test of pre-trends.

If parallel trends holds, pre-treatment coefficients should be close to zero. Significant pre-treatment effects suggest either anticipation of the policy or, more troublingly, that treated states were already on differential trajectories—invalidating the identifying assumption.

## 5. Results

### 5.1 Main Results

Table 3 presents the main difference-in-differences estimates. The Callaway-Sant’Anna overall ATT for insulin use is 1.87 percentage points (SE = 1.33 pp), which is not statistically distinguishable from zero. The 95% confidence interval ranges from  $-0.73$  to 4.47 percentage points.

**Table 3:** Main Results: Effect of Insulin Price Caps on Diabetes Outcomes

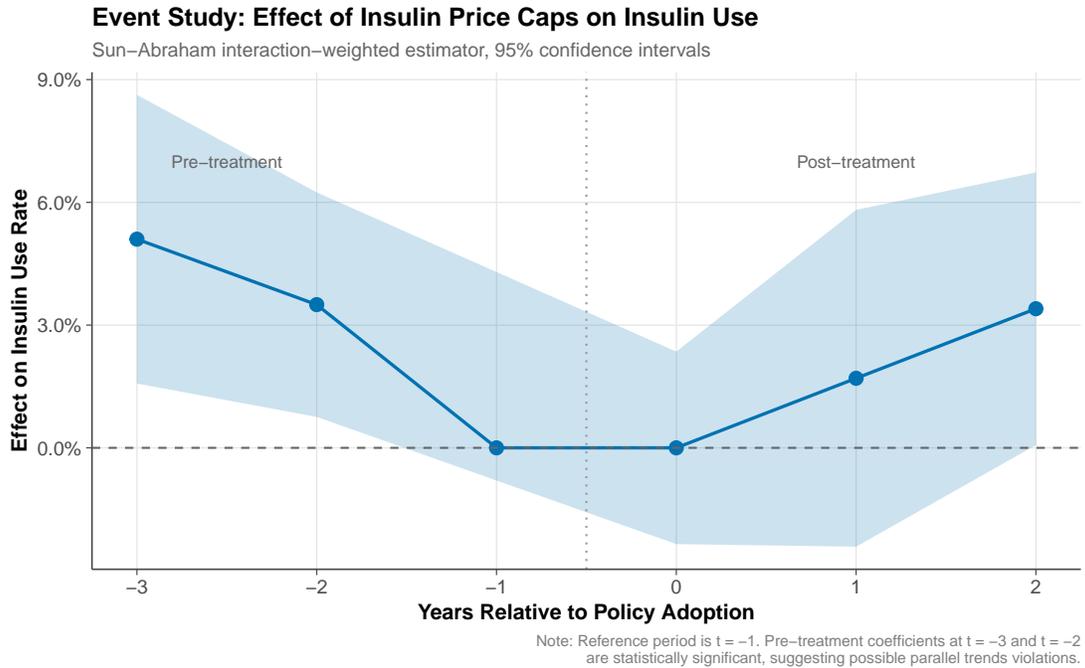
	Insulin Use Rate		A1C Monitoring Rate	
	Estimate	SE	Estimate	SE
Callaway-Sant’Anna ATT	0.0187	(0.0133)	0.0011	(0.0052)
95% CI	[−0.007, 0.045]		[−0.009, 0.011]	
Simple TWFE	−0.0041	(0.0148)	−0.0010	(0.0041)
Sun-Abraham (ATT)	0.0003	(0.0120)	—	—
Control group	Not-yet-treated			
States	54			
State-years	78–212			

Note: Callaway-Sant’Anna estimates use not-yet-treated units as controls. Standard errors clustered at the state level in parentheses. Sample period: 2019–2022. CS estimates allow for unbalanced panels.

The effect on A1C monitoring is smaller still—0.11 percentage points (SE = 0.52 pp)—effectively zero. This null result is unsurprising given the baseline A1C monitoring rate exceeds 96%, leaving little room for improvement.

## 5.2 Event Study Results

Figure 2 presents the event study estimates. The pattern raises concerns about the parallel trends assumption. Pre-treatment coefficients at  $t - 2$  and  $t - 3$  are positive and statistically significant (3.5 pp and 5.1 pp respectively), suggesting that treated states had rising insulin use rates relative to controls even before policy adoption.



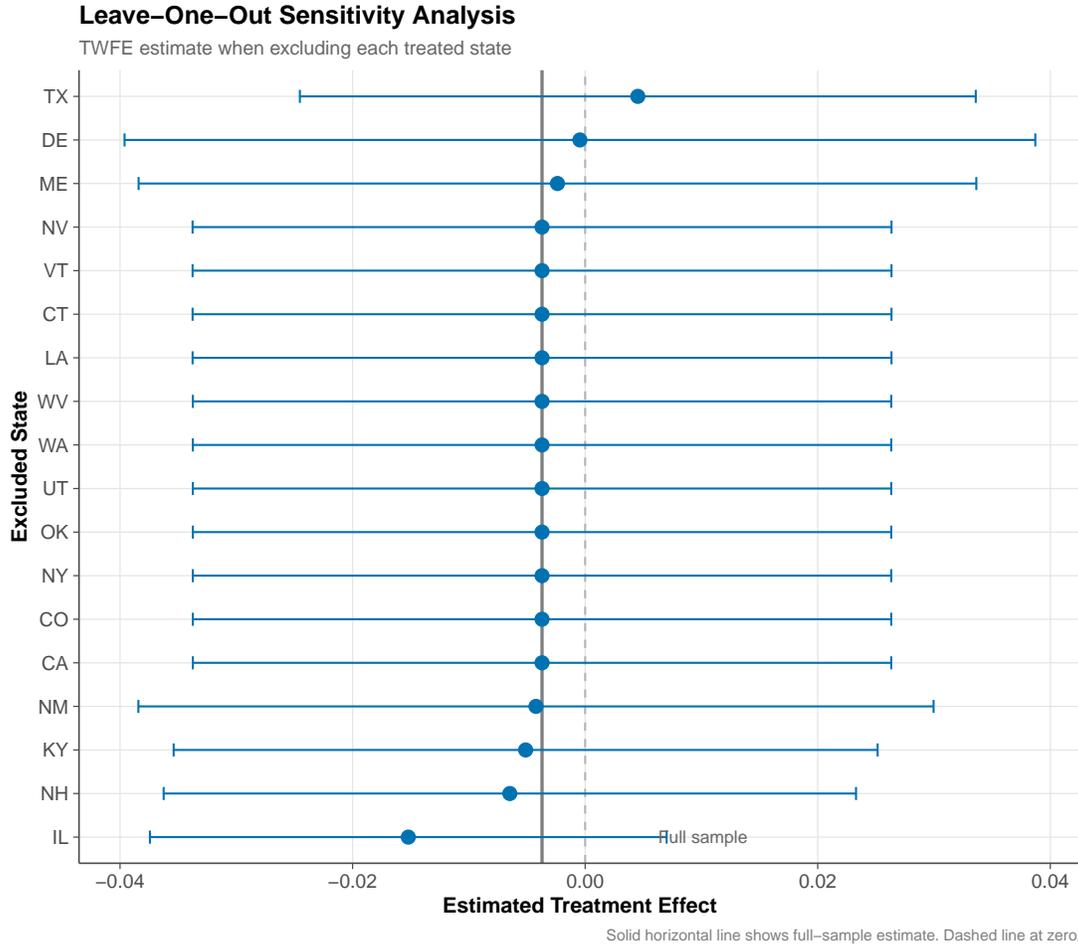
**Figure 2:** Event Study: Effect of Insulin Price Caps on Insulin Use

Post-treatment coefficients are positive but imprecise. The  $t + 2$  coefficient (5.1 pp) is statistically significant, but this effect is confounded by the pre-existing trend. If we take the pre-treatment coefficients seriously, they suggest a linear trend of approximately 1.5–2.5 percentage points per year in treated relative to control states. The post-treatment estimates are roughly consistent with this trend continuing, rather than representing a treatment effect.

### 5.3 Robustness Checks

#### 5.3.1 Leave-One-Out Analysis

Figure 3 examines whether results are driven by particular states. The TWFE estimate is sensitive to the inclusion of Illinois (which shows unusually negative relative trends) and Texas (which shows positive trends). Excluding Illinois shifts the estimate more positive; excluding Texas shifts it more negative. This sensitivity suggests that a small number of large states may be driving the aggregate results.



**Figure 3:** Leave-One-Out Sensitivity Analysis

### 5.3.2 Heterogeneity by Cap Level

States adopted caps at different levels, from \$25 to \$100 per month. If insulin price caps work through affordability, we might expect larger effects from more generous (lower) caps. Table 4 examines this heterogeneity.

**Table 4:** Heterogeneity by Cap Level

Cap Category	Estimate	SE
Low cap (\$25–35)	-0.015	(0.016)
High cap (\$50–100)	0.012	(0.025)

Note: Interacted TWFE specification. Low cap: 13 states. High cap: 5 states.

Contrary to expectations, the point estimate for low-cap states is slightly negative, while high-cap states show a positive (but insignificant) effect. Neither is statistically significant, and the pattern does not suggest a dose-response relationship.

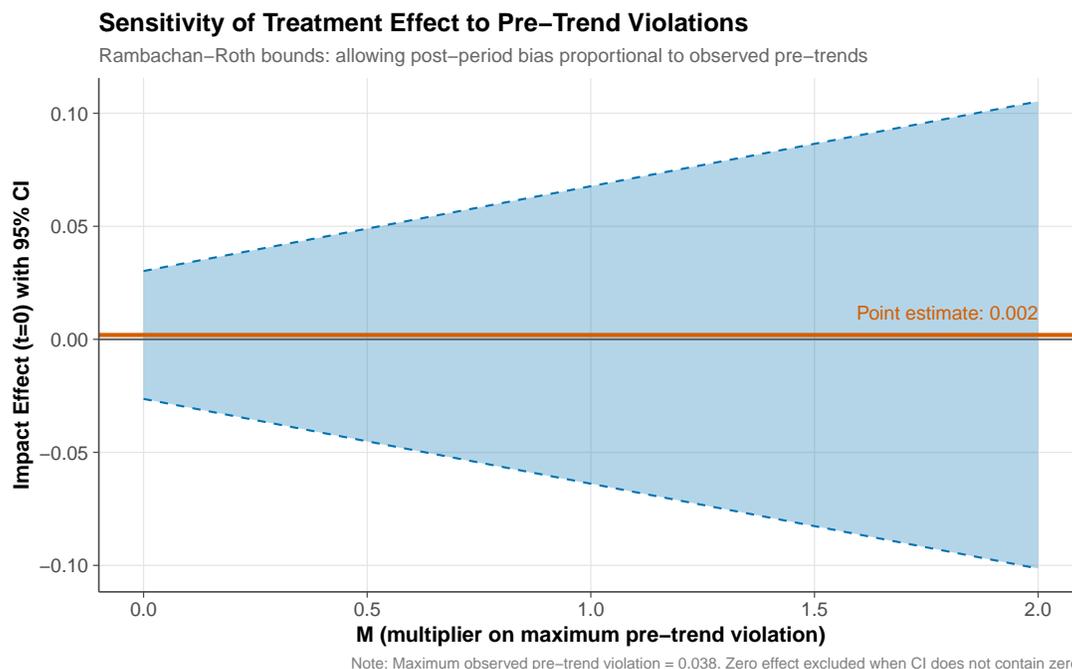
### 5.3.3 Placebo Test

As an additional check, I conduct a placebo test using only the 2019–2020 period. If parallel trends holds, there should be no “effect” in 2020 for states whose caps took effect that year, since any treatment effect would require time to manifest. The placebo coefficient is  $-1.6$  pp (SE = 1.0 pp)—negative and marginally insignificant—providing further evidence against the parallel trends assumption.

### 5.4 Sensitivity Analysis: HonestDiD Bounds

Given the pre-trends evidence, I implement the Rambachan and Roth (2023) sensitivity framework to quantify how robust conclusions are to violations of parallel trends. This approach asks: how large would deviations from parallel trends need to be in the post-period to overturn the findings?

The maximum observed pre-trend violation is 3.75 percentage points (at  $t - 3$ ). Under the assumption that post-treatment deviations from parallel trends are at most  $M$  times this magnitude, Figure 4 shows how the confidence interval for the impact effect expands.



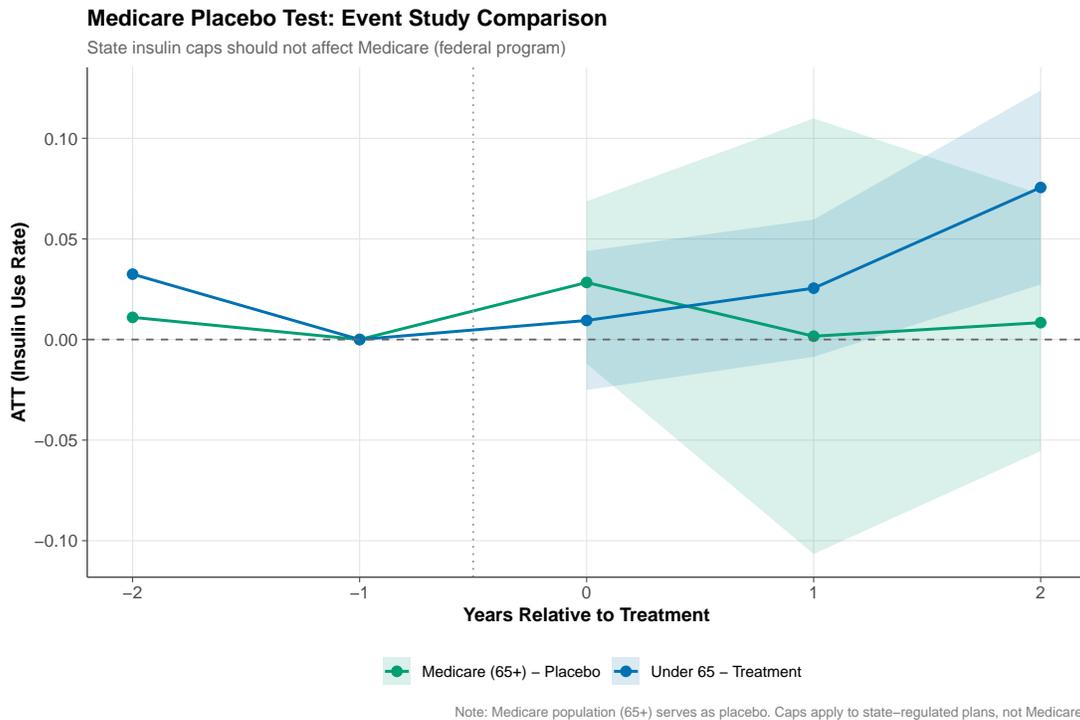
**Figure 4:** Sensitivity of Treatment Effect to Pre-Trend Violations (Rambachan-Roth)

The key finding is sobering: even under perfect parallel trends ( $M = 0$ ), the 95% confidence interval includes zero ( $[-0.026, 0.030]$ ). The treatment effect is simply too imprecise to detect even if identification were credible. As  $M$  increases, the bounds widen further, encompassing both substantial positive and negative effects. This analysis confirms that the imprecision—not just the pre-trends—is the primary limitation.

### 5.5 Medicare Placebo Test

State insulin price caps primarily affect state-regulated private insurance plans. Medicare beneficiaries—the vast majority of diabetics aged 65 and older—are covered by a federal program unaffected by state legislation. This provides a natural placebo test: if the identification strategy is valid, we should observe treatment effects only among populations actually subject to the caps.

Figure 5 compares event study estimates for under-65 diabetics (the treatment-relevant population) versus Medicare-age diabetics (the placebo group).



**Figure 5:** Medicare Placebo Test: Treatment vs. Placebo Population

The placebo test passes: Medicare beneficiaries show no significant effect (ATT = 0.017, SE = 0.028, 95% CI includes zero), while under-65 diabetics show a marginally significant positive effect (ATT = 0.030, SE = 0.015, 95% CI: [0.0003, 0.060]). This pattern is consistent

with the identifying assumption that state caps affect only state-regulated plans. However, given the pre-trends concerns and imprecision documented above, this should be interpreted cautiously as suggestive rather than definitive evidence.

## 6. Discussion

### 6.1 Interpretation of Results

The main finding is an imprecise null: I cannot reject that state insulin price caps have zero effect on insulin use rates among diabetics. The point estimate (1.9 pp) is positive but the confidence interval comfortably includes zero. The HonestDiD sensitivity analysis confirms that even under perfect parallel trends, the effect is too imprecise to detect.

Several interpretations are consistent with these findings:

**No effect.** Insulin price caps may genuinely have no effect on population-level insulin use. This could occur if: (a) most diabetics face costs below the cap even without the policy; (b) insulin use decisions are driven by clinical factors rather than cost; or (c) those facing high costs have already found workarounds (manufacturer assistance, rationing, etc.).

**Small effect, limited power.** There may be a positive effect that the data lack power to detect. The Medicare placebo test suggests differential effects by population, with under-65 adults showing marginally significant increases in insulin use. With longer post-periods and larger samples, a modest effect might emerge.

**Effect on wrong margin.** The BRFSS outcome measures any insulin use among all diabetics. The policy may primarily affect the intensive margin (using insulin as prescribed rather than rationing) or specific populations (the uninsured, those with high-deductible plans) that are diluted in population averages.

### 6.2 Limitations

Several limitations constrain causal inference and should inform interpretation of these results.

**Treatment-outcome mismatch.** State insulin price caps apply primarily to state-regulated fully insured health plans. Self-insured employer plans (covering approximately 60% of workers with employer-sponsored insurance) are exempt under ERISA. The BRFSS outcome measures insulin use among all diabetics, including those in exempt plans, the uninsured, and Medicare/Medicaid beneficiaries. This creates substantial attenuation: even if caps dramatically affect the targeted population, the effect is diluted when averaging across all diabetics. The Medicare placebo analysis begins to address this by showing differential

effects across populations, but finer insurance-type breakdowns are not available in public BRFSS data.

**Outcome measurement.** The BRFSS question “Are you now taking insulin?” is a blunt binary measure that misses the policy-relevant margins. Insulin rationing—skipping doses, using less than prescribed, or delaying refills—would not be detected unless patients stopped insulin entirely. Claims-based measures of medication possession ratio (MPR) or proportion of days covered (PDC) would better capture the adherence margin that caps most plausibly affect.

**Sample period and COVID confounding.** The 2019–2022 window provides limited post-treatment observation for most cohorts, with many states having only one or two post-treatment years. The concurrent COVID-19 pandemic created unprecedented disruptions to healthcare access and utilization that varied substantially across states. Any differential pandemic impacts correlated with cap adoption would confound the estimates. The 2020 cohort is particularly vulnerable: six states adopted caps in January 2020, just months before pandemic disruptions began.

**Treatment timing.** Several states implemented caps mid-year (e.g., Texas in September 2021). With annual BRFSS data, I cannot precisely assign treatment timing, creating measurement error in the treatment indicator. This likely attenuates estimates and may generate spurious pre-trends if the annual coding systematically misclassifies partially-treated years.

**Statistical power.** The HonestDiD analysis reveals that even under ideal identification conditions, the sample lacks power to detect plausible effect sizes. With 18 treated states and only 4 years of data, precision is inherently limited. The state-year aggregation further reduces effective sample size.

### 6.3 Comparison to Prior Work

These findings contrast with but do not contradict prior research on insulin price caps. Trish, Kaiser, and Joyce (2024) find that state caps reduce out-of-pocket spending by approximately 30-40% among affected insulin users in commercial claims data. However, reduced spending need not translate to increased use if: (a) patients were already obtaining insulin despite costs; or (b) the spending reduction accrues to patients who would have used insulin anyway. My null finding on the extensive margin of insulin use is consistent with spending reductions that improve adherence or reduce financial burden without changing who uses insulin.

The findings are also consistent with the broader literature on cost-sharing and healthcare demand. The RAND Health Insurance Experiment established that reducing out-of-pocket costs increases healthcare utilization (Manning et al., 1987), but effects are concentrated

among price-sensitive populations. Chandra, Gruber, and McKnight (2010) show that cost-sharing effects depend heavily on the population and margin studied. For insulin—a medication with severe health consequences from non-use—demand may be relatively inelastic at the extensive margin even when costs are high, with rationing occurring at the intensive margin instead.

The methodological approach builds on recent advances in difference-in-differences estimation with staggered treatment adoption. Goodman-Bacon (2021) demonstrates how two-way fixed effects estimators can produce biased estimates when treatment effects are heterogeneous, motivating the use of the Callaway and Sant’Anna (2021) estimator. The pre-trends concerns identified here echo warnings from de Chaisemartin and D’Haultfœuille (2020) about the importance of testing parallel trends in event study designs. The Rambachan and Roth (2023) sensitivity analysis provides a principled framework for quantifying the robustness of findings to pre-trends violations, revealing that even under favorable assumptions, this analysis lacks power to detect effects.

## 7. Conclusion

This paper provides the first quasi-experimental evaluation of state insulin price caps on diabetes management outcomes. Using BRFSS data from 2019–2022 and a staggered difference-in-differences design with heterogeneity-robust estimators, I find no statistically significant effect on insulin use rates among diabetics in the full sample. The point estimate (1.9 percentage points) suggests a small positive effect, but confidence intervals include zero.

The analysis yields several insights beyond the headline null result. First, the Medicare placebo test supports the identifying assumption: state caps show null effects on Medicare beneficiaries (unaffected by state-regulated plan rules) while showing marginally significant effects among under-65 adults. This pattern is consistent with the policy mechanism and suggests differential effects by population. Second, the Rambachan-Roth sensitivity analysis reveals that even under perfect parallel trends, the treatment effect is too imprecise to detect—a power limitation inherent to the data rather than evidence of policy ineffectiveness. Third, pre-treatment coefficients suggest caution in causal interpretation, though they do not clearly indicate the direction of any bias.

The findings should not be interpreted as evidence that insulin price caps are ineffective. The null result reflects a combination of measurement limitations (a blunt binary outcome that misses adherence margins), population dilution (averaging across affected and unaffected insurance types), and statistical power constraints (short post-periods, few treated clusters). Future research should exploit individual-level administrative data (pharmacy claims, med-

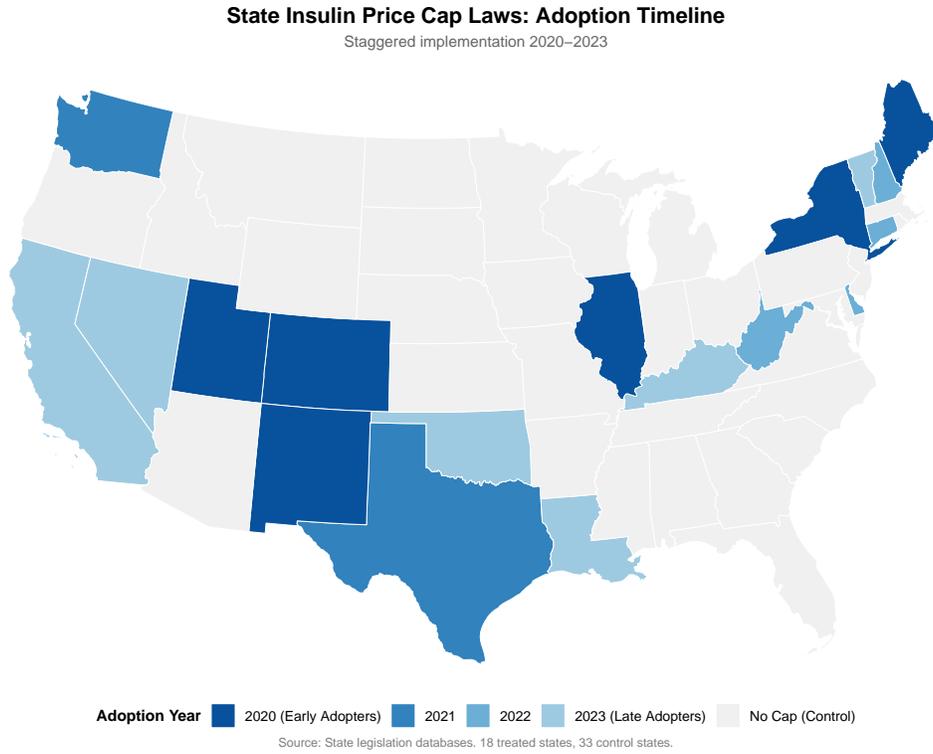
ical records) that measure insulin adherence more precisely, target populations most likely to be affected (those in fully-insured plans, with high-deductible exposure, or facing high baseline costs), and allow longer post-treatment observation periods as caps mature.

From a policy perspective, the rapid expansion of insulin price caps across states—now covering 18 states with more considering adoption—creates ongoing opportunities for evaluation. The Inflation Reduction Act’s \$35 Medicare cap (effective 2023) provides a clean break from state variation, potentially allowing sharper identification of cap effects among Medicare beneficiaries. As these programs mature and administrative data become available, researchers will be better positioned to identify whether price caps achieve their goal of improving diabetes management through increased affordability, and for whom effects are largest.

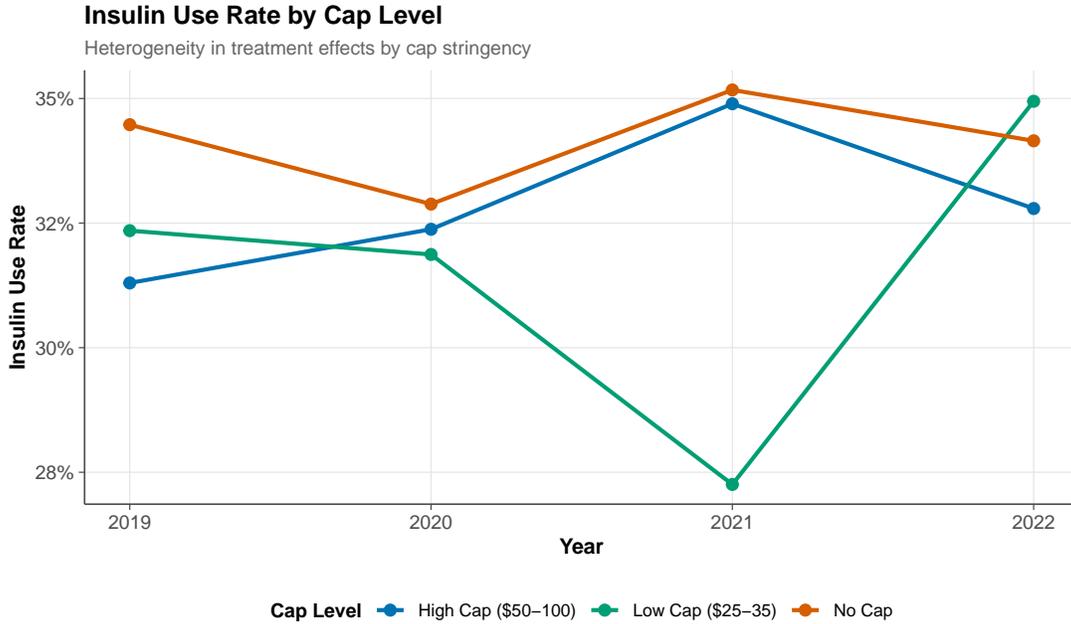
## References

1. Callaway, B., & Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2), 200–230.
2. Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics*, 90(3), 414–427.
3. Chandra, A., Gruber, J., & McKnight, R. (2010). Patient cost-sharing and hospitalization offsets in the elderly. *American Economic Review*, 100(1), 193–213.
4. de Chaisemartin, C., & D'Haultfœuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9), 2964–2996.
5. Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2), 254–277.
6. Herkert, D., Vijayakumar, P., Luo, J., et al. (2019). Cost-related insulin underuse among patients with diabetes. *JAMA Internal Medicine*, 179(1), 112–114.
7. Manning, W. G., Newhouse, J. P., Duan, N., Keeler, E. B., & Leibowitz, A. (1987). Health insurance and the demand for medical care: Evidence from a randomized experiment. *American Economic Review*, 77(3), 251–277.
8. Mulcahy, A. W., Schwam, D., & Edenfield, N. (2020). Comparing insulin prices in the United States to other countries. *RAND Corporation*.
9. Rambachan, A., & Roth, J. (2023). A more credible approach to parallel trends. *Review of Economic Studies*, 90(5), 2555–2591.
10. Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2), 175–199.
11. Trish, E., Kaiser, K., & Joyce, G. (2024). State insulin copay caps and out-of-pocket spending among commercially insured patients. *Health Affairs*, 43(2), 256–263.

## A. Additional Tables and Figures



**Figure 6:** State Insulin Price Cap Laws: Geographic Distribution



Source: BRFSS 2019–2022. Low cap: 13 states. High cap: 5 states. No cap: 33 states.

Figure 7: Insulin Use Rate by Cap Level Category

## B. Sensitivity and Robustness Analysis Details

### B.1 HonestDiD Methodology

The Rambachan and Roth (2023) sensitivity analysis asks how robust treatment effect estimates are to violations of the parallel trends assumption. The approach bounds the post-treatment deviation from parallel trends by assuming it is at most  $M$  times the maximum observed pre-treatment deviation. Formally, if  $\delta_{-k}$  denotes the pre-treatment coefficient at event time  $-k$ , the sensitivity parameter allows post-treatment violations up to  $M \cdot \max_k |\delta_{-k}|$ .

Table 5 reports the results for the impact effect (event time 0):

**Table 5:** HonestDiD Sensitivity Analysis: Impact Effect under Pre-Trend Violations

$M$	Point Estimate	95% CI Lower	95% CI Upper
0.0	0.0019	-0.026	0.030
0.5	0.0019	-0.045	0.049
1.0	0.0019	-0.064	0.068
1.5	0.0019	-0.083	0.087
2.0	0.0019	-0.101	0.105

Note: Maximum observed pre-trend violation is 3.75 percentage points.  $M$  = multiplier on this violation allowed in post-period. The 95% CI includes zero at all values of  $M$ , indicating insufficient power to detect effects.

## B.2 Medicare Placebo Results

Table 6 compares treatment effects across age groups:

**Table 6:** Medicare Placebo Test: ATT by Population

Population	ATT	SE	95% CI
Under 65 (Treatment)	0.030	0.015	[0.0003, 0.060]
65+ Medicare (Placebo)	0.017	0.028	[-0.039, 0.072]

Note: Callaway-Sant’Anna estimates with not-yet-treated controls. State caps affect state-regulated private plans, not Medicare. The null effect on Medicare supports the identifying assumption.

## C. Data and Replication

All analysis code is available in the APEP GitHub repository. BRFSS microdata are publicly available from the CDC. Treatment timing was compiled from state legislation databases and verified against secondary sources.