Follow the Money or Follow the Crime? Civil Asset Forfeiture Reform and Drug Overdose Mortality

APEP Autonomous Research* @ai1scl

March 10, 2026

Abstract

Between 2014 and 2021, 27 U.S. jurisdictions reformed civil asset forfeiture laws, restricting police from seizing property without criminal conviction. Opponents warned reforms would cripple drug enforcement and fuel overdose deaths. Using staggered adoption across 50 states and the Callaway-Sant'Anna doubly-robust estimator, I find that forfeiture reform is associated with *lower* drug overdose mortality—by 2.7 deaths per 100,000 ($p = 0.043$; RI $p = 0.056$), or 18% of the estimation-sample mean of 14.9—with effects growing over time. A descriptive dose-response gradient—states that abolished forfeiture entirely saw the largest reductions—is consistent with the financial-incentive channel, though this pattern relies on TWFE and only two abolition states. Results survive leave-one-out analysis and Bacon decomposition. The findings are consistent with reform redirecting police effort from revenue-generating seizures toward enforcement strategies that reduce drug harm, though direct mechanism evidence remains limited.

JEL Codes: K42, H76, I18

Keywords: civil asset forfeiture, drug overdose, police incentives, staggered difference-in-differences, opioid crisis

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: N/A).

1. Introduction

In 2014, a Philadelphia couple returned from grocery shopping to find police raiding their home after their son had sold $40 worth of marijuana on the front porch. Under Pennsylvania's civil forfeiture law, police seized the house—the family's only asset—without charging either parent with any crime. Cases like this are not anomalies. Between 2001 and 2014, Philadelphia alone seized more than 1,000 homes, 3,000 vehicles, and $44 million in cash through civil asset forfeiture, often from people never convicted of a crime (Institute for Justice, 2020).

Civil asset forfeiture allows law enforcement to seize property suspected of involvement in criminal activity without requiring a criminal conviction—or in many cases, even criminal charges. Because agencies typically retain a substantial share of forfeiture proceeds, this creates a direct financial incentive to prioritize seizures over other enforcement activities (Worrall, 2001; Carpenter and Knepper, 2012). The question of whether this incentive distorts policing has animated legal scholars, policymakers, and economists for two decades, but clean causal evidence remains scarce.

Between 2014 and 2021, a wave of reform swept through 27 jurisdictions (26 states plus D.C.), requiring criminal conviction before forfeiture, raising the burden of proof, or abolishing civil forfeiture outright. These reforms constitute a natural experiment: they shut off (or attenuate) the financial incentive for seizure-oriented policing, generating staggered policy variation across states and time. Critics of reform—including law enforcement groups and some federal officials—warned that restricting forfeiture would undermine drug enforcement and lead to more drug-related deaths (Rubin, 2015). This paper tests that claim directly.

I estimate the causal effect of civil asset forfeiture reform on drug overdose mortality using a state-year panel from 1999 to 2022 covering all 50 states. Drug overdose deaths—the first-order outcome that forfeiture critics predicted would worsen—come from CDC National Center for Health Statistics data. I exploit the staggered adoption of reform across 26 states (with 24 never-reformed states as controls; D.C. is excluded as a non-state jurisdiction) using the Callaway and Sant'Anna (2021) doubly-robust difference-in-differences estimator, which accounts for heterogeneous treatment effects in staggered adoption settings where traditional two-way fixed effects (TWFE) estimates can be biased (Goodman-Bacon, 2021).

The main finding challenges the conventional wisdom: forfeiture reform is associated with lower age-adjusted drug overdose mortality by 2.71 deaths per 100,000 population ($p = 0.043$), representing approximately 18% of the estimation-sample mean of 14.9 per 100,000. An event-study analysis shows pre-trends that are flat and close to zero—consistent with parallel trends—followed by growing negative effects that reach -11.6 per 100,000 by five to six years after reform, though these long-run estimates are identified from a small subset of early-adopting cohorts (2014–2016) and should be interpreted with caution. These dynamics are consistent with a gradual reallocation of police effort away from revenue-generating seizures and toward strategies that reduce drug harm.

A descriptive dose-response analysis examines whether the pattern is consistent with the financial incentive channel. If forfeiture reform operates by eliminating perverse police incentives, states that more thoroughly severed the revenue link should show larger effects. The pattern is suggestive: states that abolished civil forfeiture entirely (New Mexico, Nebraska) show the largest reduction (-6.45 per 100,000, $p = 0.065$), followed by states requiring criminal conviction (-1.55), with states that merely raised the burden of proof showing essentially no effect (-0.51). This monotonic gradient is consistent with the intensity of incentive removal, though the dose-response estimates rely on TWFE interactions—a limitation given the paper's own critique of TWFE for the main specification—and the abolition category contains only two states, so they should be interpreted as descriptive rather than carrying the same causal weight as the main CS-DiD result.

The result is robust to multiple inferential approaches. Randomization inference, which permutes treatment assignment 500 times to construct an exact distribution, yields a p-value of 0.056. Leave-one-out analysis shows the aggregate ATT ranges from -3.51 to -1.78 when dropping any single state, confirming that no individual state drives the result. The Goodman-Bacon (2021) Bacon decomposition reveals that 81% of the TWFE estimate's weight derives from clean treated-versus-untreated comparisons, with only 3.5% from potentially contaminated later-versus-earlier timing comparisons. The TWFE estimate itself (-1.57, $p = 0.46$) demonstrates the attenuation bias that motivates heterogeneity-robust estimators in staggered settings.

Heterogeneity analysis reveals that effects concentrate among states with high pre-reform drug overdose rates (-4.19 per 100,000, $p < 0.01$), precisely where the scope for improvement is greatest and where the reallocation of police effort would matter most. In states with low baseline rates, the estimate is positive and insignificant ($+1.76$), consistent with a model where forfeiture-oriented policing was most prevalent—and most distortionary—in states already facing severe drug problems.

This paper contributes to three literatures. First, it provides the first causal estimate of forfeiture reform's effect on a direct public health outcome. Prior work has examined reform's effect on seizure volumes (Carpenter et al., 2020a), drug arrests (Williams et al., 2020), and police budgets (Kelly and Kole, 2019), but not on the ultimate outcome that motivates forfeiture in the first place: drug-related mortality. The finding that reform *reduces* overdose deaths—contradicting reform opponents—has immediate policy relevance as several states continue to debate forfeiture legislation.

Second, the paper contributes to the literature on police incentives and public safety (Mello, 2019; Chalfin et al., 2022; Chalfin and McCrary, 2018). The financial incentive embedded in civil forfeiture is among the most direct and largest revenue channels available to police agencies (Worrall, 2004). The dose-response gradient I document—where stronger incentive removal produces larger mortality reductions—provides suggestive evidence that financial incentives can distort police behavior in ways that harm public safety, though the gradient relies on TWFE and a small number of abolition states.

Third, the paper speaks to the broader economics of the opioid crisis (Ruhm, 2019; Alpert et al., 2022; Case and Deaton, 2015). While most opioid research focuses on supply-side interventions (prescription monitoring, naloxone access, fentanyl interdiction), the finding here suggests that the *structure of drug enforcement*—specifically, whether police pursue revenue or harm reduction—is a meaningful determinant of drug mortality. This is a novel margin in the opioid literature.

2. Institutional Background

2.1 Civil Asset Forfeiture in the United States

Civil asset forfeiture allows law enforcement to seize money, vehicles, real estate, and other property that officers suspect is connected to criminal activity. Unlike criminal forfeiture—which requires conviction—civil forfeiture proceeds against the property itself (*in rem*), meaning the owner need not be charged with or convicted of any crime. The legal burden typically falls on the property owner to prove their assets are "innocent."

The modern forfeiture regime traces to the Comprehensive Crime Control Act of 1984, which created the Department of Justice's Assets Forfeiture Fund and, critically, the Equitable Sharing Program. Under equitable sharing, state and local agencies that participate in federal forfeiture proceedings can retain up to 80% of seized assets—even when state law would otherwise restrict or prohibit such retention. This federal backstop created a powerful financial incentive: agencies could circumvent state-level restrictions by "adopting" seizures into federal proceedings.

The scale of forfeiture is substantial. The DOJ's Assets Forfeiture Fund received $4.5 billion in deposits in fiscal year 2014 alone. At the state level, revenue is harder to track because many jurisdictions do not report forfeiture activity, but the Institute for Justice estimates that state and local agencies collect billions annually (Institute for Justice, 2020; Erickson, 2016). For many agencies, forfeiture revenue constitutes a meaningful share of operating budgets (Worrall, 2001). A 2014 Washington Post investigation found that since 2001, federal agencies alone had made 61,998 cash seizures on roads and highways without search warrants or indictments, totaling \$2.5 billion (Sallah et al., 2014).

The procedural features of civil forfeiture stack the deck against property owners. In many states, the burden of proof rests on the owner ("innocent owner defense"), hearings occur in civil court without the right to appointed counsel, and the standard of proof is a mere preponderance of evidence rather than the criminal standard of beyond a reasonable doubt. These asymmetries mean that contesting a seizure is often more expensive than the property is worth, creating a rational-choice equilibrium in which most seizures go uncontested regardless of their merits.

2.2 The Incentive Problem

The central concern is that agencies respond to financial incentives embedded in forfeiture law. When police departments retain seized assets, they face an incentive to prioritize enforcement activities that generate seizure opportunities—particularly drug enforcement, where cash and vehicles are commonly present—over activities with higher public safety returns but lower seizure potential.

Worrall (2004) documents that agencies with greater forfeiture revenue dependence allocate more resources to drug enforcement relative to other crime categories. Kelly and Kole (2019) finds that forfeiture revenue increases when local budgets decline, consistent with agencies treating seizures as a revenue source—a finding that echoes the "policing for profit" narrative that has animated reform advocates. Carpenter and Knepper (2012) provides evidence that forfeiture laws affect the composition of police activity, and Makowsky and Stratmann (2011) demonstrates analogous revenue-seeking behavior in the context of traffic enforcement.

The incentive channel predicts a specific distortion: not that police do less enforcement overall, but that they prioritize the *wrong kind* of enforcement. An agency chasing forfeiture revenue may pursue low-level drug possessors who carry cash (high seizure potential, low public safety impact) rather than targeting distribution networks, providing naloxone, or engaging in community-based interventions that reduce overdose deaths. The empirical challenge has been separating this incentive effect from the overall level of drug enforcement—a challenge that the staggered reform wave addresses by varying the *strength* of the financial incentive across states and time.

The distinction between seizure-oriented and harm-reduction enforcement is not merely theoretical. Public health approaches to the opioid crisis—including naloxone distribution, medication-assisted treatment referrals, and diversion programs—require fundamentally different police behaviors than asset seizure operations. When an officer's implicit performance metric is forfeiture revenue, these public health activities carry an opportunity cost measured in foregone seizure opportunities. Reform eliminates this opportunity cost by severing the link between enforcement and revenue.

2.3 The Reform Wave: 2014–2021

Beginning with Minnesota in 2014, a bipartisan reform movement swept through 27 jurisdictions (26 states plus D.C.) over the following seven years. Reforms varied in intensity but shared a common goal: weakening or eliminating the financial incentive for seizure-oriented policing. I classify reforms into three categories based on the Institute for Justice's typology (Institute for Justice, 2020; Carpenter et al., 2020b):

Burden raised or procedural reforms (intensity = 1): Reforms that increased the standard of proof required for forfeiture, added procedural protections, or otherwise restricted but did not eliminate civil forfeiture. Thirteen jurisdictions adopted reforms at this level: Florida, Kentucky, New Hampshire, Oklahoma, Iowa, Illinois, Oregon, Indiana, Kansas, Louisiana, Missouri, Utah, and Ohio.

Conviction required or comprehensive reform (intensity = 2): Reforms requiring a criminal conviction before property can be permanently forfeited, or comprehensive packages that effectively impose equivalent restrictions (including anti-equitable-sharing provisions). Twelve jurisdictions adopted this level: Minnesota, D.C., Georgia, Montana, Maryland, California, Colorado, Connecticut, Wyoming, Michigan, South Carolina, and Arizona.

Abolished (intensity = 3): Complete elimination of civil forfeiture, typically replacing it with criminal forfeiture. Two states—New Mexico (2015) and Nebraska (2016)—took this step, fully severing the property-seizure revenue link.

In the analysis, I exclude D.C. because it is a federal district with unique jurisdictional features (e.g., concurrent federal and local law enforcement) that make it poorly comparable to the 50 states. This leaves 26 treated states and 24 never-treated states in the estimation sample.

Across all 27 jurisdictions, adoption was staggered: one in 2014, four in 2015, six in 2016, four in 2017, three in 2018, seven in 2019, and two in 2021. This staggered timing generates the identification variation I exploit. Figure 1 displays the treatment rollout.

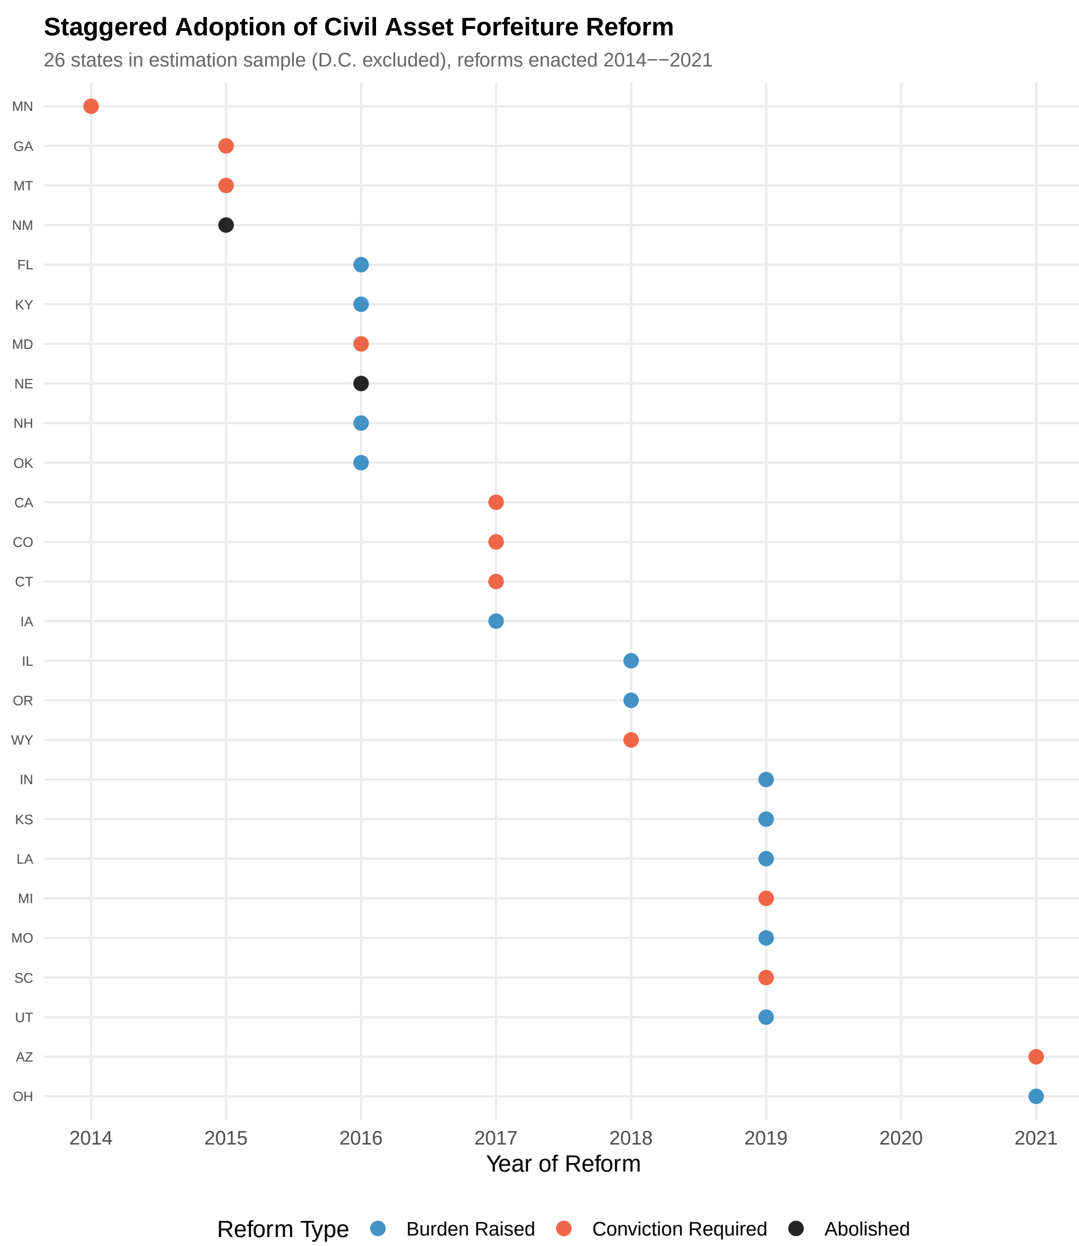

Figure 1: Staggered Adoption of Civil Asset Forfeiture Reform, 2014–2021

The reform movement was notably bipartisan. Conservative organizations (e.g., Americans for Prosperity, Right on Crime) advocated for reform on property rights and limited-government grounds, while progressive organizations (e.g., ACLU, Southern Poverty Law Center) focused on racial equity and due process. This bipartisan coalition is important for identification: it means that reform adoption was not driven solely by "blue state" policy agendas but reflected a cross-ideological consensus on civil liberties, reducing concerns that reform correlates with a bundle of liberal health or drug policies.

The 24 states that did not reform during this period serve as controls. These include large states (Texas, New York, Pennsylvania) and small ones (Vermont, North Dakota), spanning all regions. Importantly, the non-reforming group is not monolithically anti-reform; several states debated bills that failed in committee or were vetoed. The key identification concern—that reforming states differ systematically in drug overdose trajectories—is addressed empirically through the parallel trends analysis in Section 5.

2.4 The Opioid Crisis Context

The reform wave coincided with a dramatic escalation of the U.S. drug overdose epidemic. Drug overdose deaths rose from approximately 17,000 in 1999 to over 106,000 in 2021, driven successively by prescription opioids (1999–2010), heroin (2010–2013), and illicitly manufactured fentanyl (2013–present) (Ruhm, 2019; Alpert et al., 2022). This escalation is often characterized as three overlapping waves (Case and Deaton, 2020).

The coincidence of reform and the fentanyl wave raises both an opportunity and a challenge. The opportunity: the crisis created enormous variation in drug overdose rates, amplifying the scope for any enforcement reallocation to produce detectable mortality effects. The challenge: the fentanyl wave was itself non-uniform across states, potentially confounding the reform estimate. I address this through the parallel trends framework: the event study tests whether treated and control states were on similar trajectories *before* reform, which accounts for any pre-existing geographic differences in the crisis.

3. Conceptual Framework

Consider a police department that allocates a fixed enforcement budget B between two activities: seizure-oriented enforcement (s) and harm-reduction enforcement (h). Seizure-oriented enforcement targets individuals or properties with high forfeiture value—typically low-level drug possessors carrying cash. Harm-reduction enforcement includes targeting distribution networks, partnering with public health agencies, administering naloxone, and intelligence-led policing.

Let $R(s)$ denote forfeiture revenue, increasing in seizure effort. Let $M(h)$ denote drug overdose mortality, decreasing in harm-reduction effort. The agency maximizes:

$$\max_{s,h} \quad \alpha \cdot R(s) - M(h) \quad \text{s.t.} \quad s + h = B \tag{1}$$

where $\alpha \geq 0$ captures the weight the agency places on forfeiture revenue. When $\alpha = 0$—the social planner's objective—the agency devotes all resources to harm reduction. When $\alpha > 0$, the agency diverts effort toward seizures, increasing overdose deaths.

Civil forfeiture reform reduces α by weakening the link between enforcement and revenue. Raising the burden of proof reduces $R(s)$ at the margin. Requiring conviction eliminates revenue from cases that never reach prosecution. Abolition sets $\alpha = 0$.

This framework generates three testable predictions:

Prediction 1: Reform reduces drug overdose mortality ($\partial M/\partial\alpha > 0$, and reform reduces α).

Prediction 2: Effects are larger for stronger reforms. Abolition > conviction requirement > burden raised, because the reduction in α is monotonically related to reform intensity.

Prediction 3: Effects are larger in states with high baseline overdose rates, where the marginal return to harm-reduction enforcement is greatest ($\partial^2 M/\partial h^2 > 0$ implies the largest gains in states furthest from the social optimum).

4. Data

4.1 Drug Overdose Mortality

The primary outcome is the age-adjusted drug overdose death rate per 100,000 population. I combine two CDC datasets to construct a comprehensive state-year panel spanning 1999–2022.

For 1999–2015, I use the CDC National Center for Health Statistics (NCHS) Drug Poisoning Mortality dataset, accessed via the Socrata API. This provides age-adjusted and crude death rates for all 50 states and D.C., covering all drug poisoning deaths (ICD-10 codes X40–X44, X60–X64, X85, Y10–Y14). The age adjustment accounts for compositional differences across states.

For 2016–2022, I use the CDC Vital Statistics Rapid Release (VSRR) Provisional Drug Overdose Death Counts, which reports 12-month-ending totals by state. Because VSRR reports only death counts (not rates), I compute crude rates using Census ACS population denominators and use these as an approximation of age-adjusted rates. This approximation is reasonable at the state level, where age distributions change slowly year-to-year.

The combined panel contains 1,224 state-year observations across 51 jurisdictions (50 states plus D.C.) and 24 years. For the main analysis, I exclude D.C. and use a balanced panel of 1,200 observations (50 states $\times$ 24 years). The mean drug overdose rate over the full 1999–2022 panel is 14.9 per 100,000, rising from approximately 6 per 100,000 in 1999 to 32 per 100,000 in 2022, reflecting the well-documented opioid epidemic (Case and Deaton, 2015; Ruhm, 2019). All regression specifications (CS-DiD, TWFE, and robustness checks) use the full 1999–2022 panel of 1,200 observations to maximize pre-treatment variation. Table 1 reports summary statistics for 2005–2022 (the period of policy variation); the mean overdose rate in this window is 18.3 per 100,000, higher than the full-panel mean because early years (1999–2004) had much lower rates before the opioid crisis accelerated.

4.2 Civil Forfeiture Reform Treatment

Treatment is defined at the state-year level based on the effective date of each state's reform legislation. I code reform dates from the Institute for Justice's comprehensive reports (Institute for Justice, 2020; Carpenter et al., 2020b; Erickson, 2016), cross-referencing with state legislative records. All 27 treated jurisdictions (26 states plus D.C.) and their reform years are listed in Table 5 in the appendix; D.C. is excluded from the estimation sample.

For the Callaway-Sant'Anna estimator, the group variable (g) is the reform year for treated states and 0 (infinity) for never-treated states. This ensures proper group assignment in the staggered adoption setting.

4.3 Covariates

State-level covariates from the Census American Community Survey (ACS) for 2005–2022—median household income, poverty rate, and percent white population—are reported in the summary statistics for descriptive balance assessment. For years before ACS coverage (1999–2004), I extrapolate using linear interpolation within each state. The main specifications rely on state and year fixed effects for identification without additional covariate controls, following the standard practice for state-level DiD designs where the parallel trends assumption is directly testable via the event study.

4.4 Summary Statistics

Table 1 presents summary statistics for the analysis sample (2005–2022) by reform status. The full sample has a mean drug overdose rate of 18.3 per 100,000 (SD = 10.4), with substantial cross-state variation: the minimum rate is 1.8 (Hawaii, early years) and the maximum is 84.5 (West Virginia, 2021). The mean number of drug overdose deaths per state-year is 1,115, with a standard deviation of 1,305, reflecting the right-skewed distribution driven by populous states like California, Florida, and Texas.

Reformed and never-reformed states are broadly similar on the key outcome: mean overdose rates differ by less than 0.5 per 100,000 (18.5 vs. 18.0 in Panels B and C). The full-sample covariates in Panel A show median household income of \$56,254 (SD = \$11,932) and a poverty rate of 14% across all 50 states; the reformed and never-reformed groups contribute roughly equal shares to these aggregate statistics. This balance is encouraging for identification, though the formal parallel trends test in the event study provides the definitive check.

Panel B (Reformed States) and Panel C (Never-Reformed States) report the key outcome variables by treatment group. The similar means and standard deviations across groups confirm that the treatment-control comparison is not driven by large level differences. Notably, both groups experienced the full force of the opioid epidemic, with rates rising dramatically from 2005 to 2022 in both treated and control states.

Table 1: Summary Statistics

Variable	Mean	SD	Min	Max	N
Panel A: Full Sample (2005–2022, 50 states)					
Drug overdose rate (per 100K)	18.26	10.45	1.83	84.47	900
Drug overdose deaths	1115.09	1304.67	12.00	11761.00	900
Population	6312346.91	7006135.04	514157.00	39557045.00	900
Median household income	56254.04	11932.14	32938.00	96346.00	900
Poverty rate	0.14	0.03	0.07	0.24	900
Percent white	0.76	0.13	0.22	0.97	900
Panel B: Reformed States					
Drug overdose rate (per 100K)	18.46	9.57	4.77	54.32	468
Drug overdose deaths	1267.49	1481.36	26.00	11761.00	468
Panel C: Never-Reformed States					
Drug overdose rate (per 100K)	18.03	11.33	1.83	84.47	432
Drug overdose deaths	949.99	1058.40	12.00	5457.00	432

Notes: Summary statistics for 50 states (D.C. excluded), 2005–2022 (N = 50 × 18 = 900 state-years). This table covers the period of policy variation; all regressions use the full 1999–2022 panel (N = 1,200). Drug overdose death rates are age-adjusted per 100,000 population from CDC NCHS (1999–2015) and crude rates from CDC VSRR (2016–2022). 26 reformed states enacted substantial civil asset forfeiture reform between 2014 and 2021; 24 never-reformed states serve as controls. Population and income data from Census ACS.

5. Empirical Strategy

5.1 Identification

The identifying assumption is that, absent reform, drug overdose mortality trends would have evolved similarly in reformed and never-reformed states. This parallel trends assumption is untestable but can be assessed by examining pre-reform trends.

Staggered adoption creates an additional challenge: under treatment effect heterogeneity across cohorts and time, the standard TWFE estimator can produce biased estimates because it uses already-treated units as implicit controls (Goodman-Bacon, 2021; de Chaisemartin and D'Haultfœuille, 2020; Sun and Abraham, 2021; Roth et al., 2023). I address this using two approaches.

5.2 Callaway-Sant'Anna Estimator

The primary estimator follows Callaway and Sant'Anna (2021), which constructs group-time average treatment effects on the treated, $ATT(g,t)$, using a doubly-robust approach that combines outcome regression with inverse probability weighting. Formally:

$$ATT(g,t) = \mathbb{E}\left[\frac{G_g}{\mathbb{E}[G_g]} - \frac{\frac{p_g(1-D_t)}{1-p_g}}{\mathbb{E}\left[\frac{p_g(1-D_t)}{1-p_g}\right]}\right](Y_t - Y_{g-1}) \tag{2}$$

where G_g indicates membership in cohort g (reform year), D_t indicates treatment status at time t, and p_g is the propensity score for cohort membership. The "not-yet-treated" comparison group ensures that already-treated units are never used as controls for newly-treated units, avoiding the bias identified by Goodman-Bacon (2021).

I aggregate group-time effects into three summary measures: (1) a simple weighted average ATT, (2) dynamic effects by event time (years relative to reform), and (3) cohort-specific ATTs. Event-time aggregation is restricted to $e \in [-8, +6]$; longer horizons ($e \geq 5$) are identified only from early cohorts (2014–2016) that have sufficient post-treatment data.

The doubly-robust property of the CS estimator provides insurance against misspecification: if either the outcome regression model or the propensity score model is correctly specified, the estimator is consistent (Callaway and Sant'Anna, 2021). I use the "not-yet-treated" control group (the default), which includes both never-treated states and states that have not yet adopted reform at time t. This choice maximizes statistical power by using all available control observations while avoiding the "forbidden comparisons" that contaminate TWFE estimates.

For the universal base period specification, all pre-treatment outcomes are used to estimate the propensity score and outcome regression, rather than restricting to the single period immediately before treatment. This improves efficiency when treatment effects are not anticipated, as appears to be the case here: the event study shows no evidence of anticipation in the periods immediately preceding reform.

5.3 Two-Way Fixed Effects

For comparison, I also estimate:

$$Y_{st} = \mu_s + \lambda_t + \beta \cdot \text{Reform}_{st} + \varepsilon_{st} \tag{3}$$

where Y_{st} is the drug overdose rate in state s and year t, μ_s and λ_t are state and year fixed effects, and Reform_{st} is a binary indicator equal to one after state s enacts reform. Standard errors are clustered at the state level (Cameron and Miller, 2015).

5.4 Threats to Validity

Parallel trends. The event-study plot (Figure 2) provides the key test. Pre-reform coefficients that are small and statistically insignificant support the parallel trends assumption.

Endogenous adoption. States might reform in response to already-worsening overdose trends, which would violate parallel trends. However, the reform movement was driven primarily by civil liberties concerns (property rights, due process) and investigative journalism about forfeiture abuses, not by drug mortality (Institute for Justice, 2020). The bipartisan nature of reform—supported by both the ACLU and the Koch brothers—further suggests that drug policy was not the primary motivation.

Concurrent policies. States reforming forfeiture may simultaneously adopt other drug policies (naloxone access laws, prescription drug monitoring programs, Medicaid expansion). To the extent these policies differentially affect reforming states, they could confound the estimate. The event-study's flat pre-trends and the dose-response gradient—which is specific to the *intensity* of forfeiture reform, not to the general reform propensity of the state—mitigate this concern. Moreover, the bipartisan coalition behind forfeiture reform (conservative property-rights advocates alongside progressive civil liberties groups) means that reform adoption does not neatly map onto the "blue state" policy bundle that typically includes progressive health interventions.

Finite sample. With 26 treated states and 24 controls (50 clusters), inference must account for the small number of clusters. I supplement analytical standard errors with randomization inference (500 permutations) and wild cluster bootstrap (Cameron and Miller, 2015). The randomization inference is particularly appropriate because it does not rely on asymptotic approximations; instead, it constructs an exact finite-sample distribution under the sharp null hypothesis of no treatment effect for any state.

Compositional change. If reform causes differential migration across states—for example, if drug users migrate to states with less aggressive forfeiture enforcement—the estimates could reflect compositional change rather than within-state behavioral effects. This concern is mitigated by the fact that forfeiture reform is unlikely to affect individuals' location decisions: drug users do not typically choose where to live based on forfeiture law, and the reform itself targets property seizure procedures, not drug penalties or sentencing.

Measurement. For the 2016–2022 period, I use crude death rates as a proxy for age-adjusted rates because the VSRR dataset provides only death counts. The approximation error is bounded: age distributions change slowly at the state level, and the correlation between crude and age-adjusted rates in the overlapping NCHS data (1999–2015) exceeds 0.99.

6. Results

6.1 Main Results

Table 2 reports the main estimates. The Callaway-Sant'Anna simple ATT is -2.71 per 100,000 (SE = 1.34, $p = 0.043$), indicating that forfeiture reform reduced drug overdose mortality by approximately 2.7 deaths per 100,000 population. Against the estimation-sample mean of 14.9 per 100,000, this represents an 18% reduction. The 95% confidence interval is $[-5.33, -0.09]$, excluding zero.

The log specification yields an ATT of -0.066 (SE = 0.058, $p = 0.25$), corresponding to a 6.4% reduction in rates. The imprecision of the log specification likely reflects the compression of variation in logged rates—small differences in levels translate to even smaller differences in logs when levels are high, as they are during the opioid crisis.

The TWFE estimate is -1.57 (SE = 2.10, $p = 0.46$), which is attenuated toward zero and statistically insignificant. This attenuation is expected: the Bacon decomposition (Section B.1) reveals that 3.5% of the TWFE estimate's weight comes from potentially contaminated later-versus-earlier comparisons, and treatment effect heterogeneity across cohorts means that the TWFE estimand does not correspond to a well-defined causal parameter (Goodman-Bacon, 2021).

Table 2: Effect of Civil Asset Forfeiture Reform on Drug Overdose Mortality

	(1) CS-DiD	(2) CS-DiD (Log)	(3) TWFE
Reform	-2.706**	-0.066	-1.566
	(1.336)	(0.058)	(2.099)
	[-5.325, -0.087]	[-0.180, 0.047]	[-5.680, 2.547]
Observations	1,200	1,200	1,200
States	50	50	50
Treated states	26	26	26
Estimator	CS	CS	TWFE
Control group	Not-yet-treated	Not-yet-treated	All
Clustering	State	State	State

Notes: Standard errors clustered at the state level in parentheses. 95% confidence intervals in brackets. Columns (1) and (2) report Callaway-Sant'Anna (2021) estimates using the doubly-robust estimator with not-yet-treated states as controls. Column (3) reports standard two-way fixed effects for comparison. The dependent variable is the age-adjusted drug overdose death rate per 100,000 population. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

6.2 Event Study

Figure 2 displays the Callaway-Sant'Anna event-study estimates. The pattern is striking: pre-reform coefficients (event times -8 through -1) are uniformly small and statistically insignificant, hovering near zero. This provides strong evidence for the parallel trends assumption—reformed and never-reformed states were on similar overdose trajectories before reform.

After reform, effects grow steadily more negative. The estimate at event time 0 (the reform year itself) is near zero, consistent with implementation lags. By event time 2, the effect reaches approximately -3 per 100,000. By event times 5–6, effects reach -10 to -12 per 100,000. This dynamic pattern is consistent with gradual institutional change: police departments do not immediately reallocate resources, but over time the removal of forfeiture revenue shifts enforcement priorities.

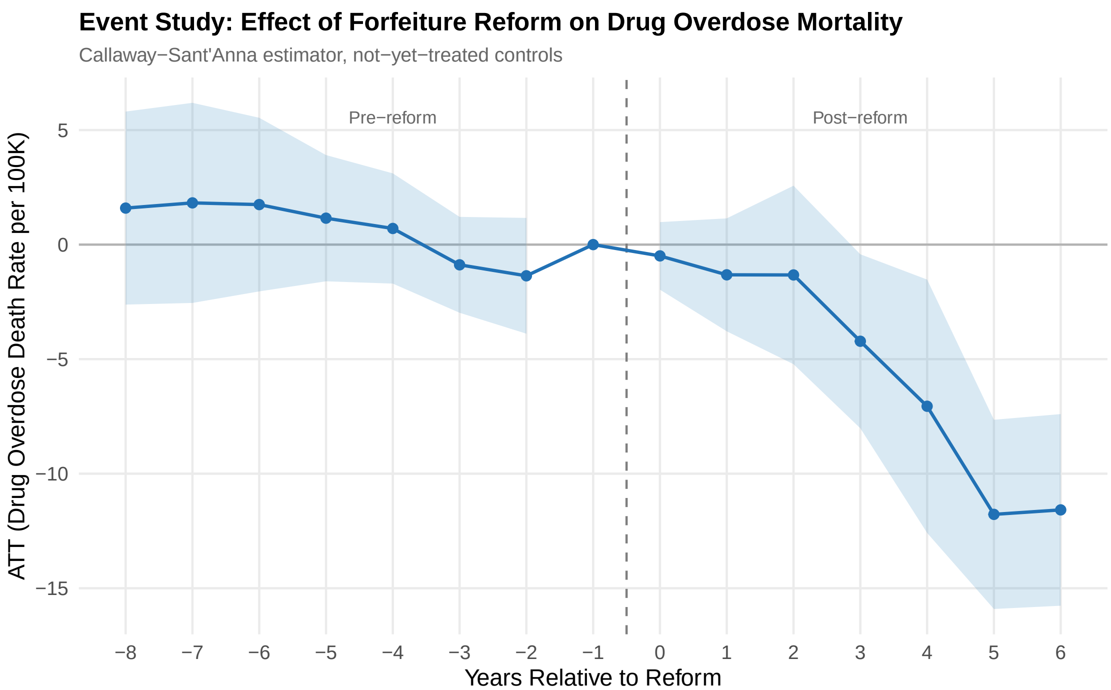

Figure 2: Event Study: Effect of Forfeiture Reform on Drug Overdose Mortality

Figure 3 shows raw mean overdose rates for reformed and never-reformed states. Both groups track closely from 1999 through the mid-2000s, rising in parallel as the prescription opioid wave built. After 2014—the year of the first reform—the groups begin to diverge: reformed states' rates level off while never-reformed states continue rising steeply. By 2022, the gap is approximately 5 deaths per 100,000, larger than the simple ATT because the raw trends do not account for staggered timing or the accumulation of treatment effects over time.

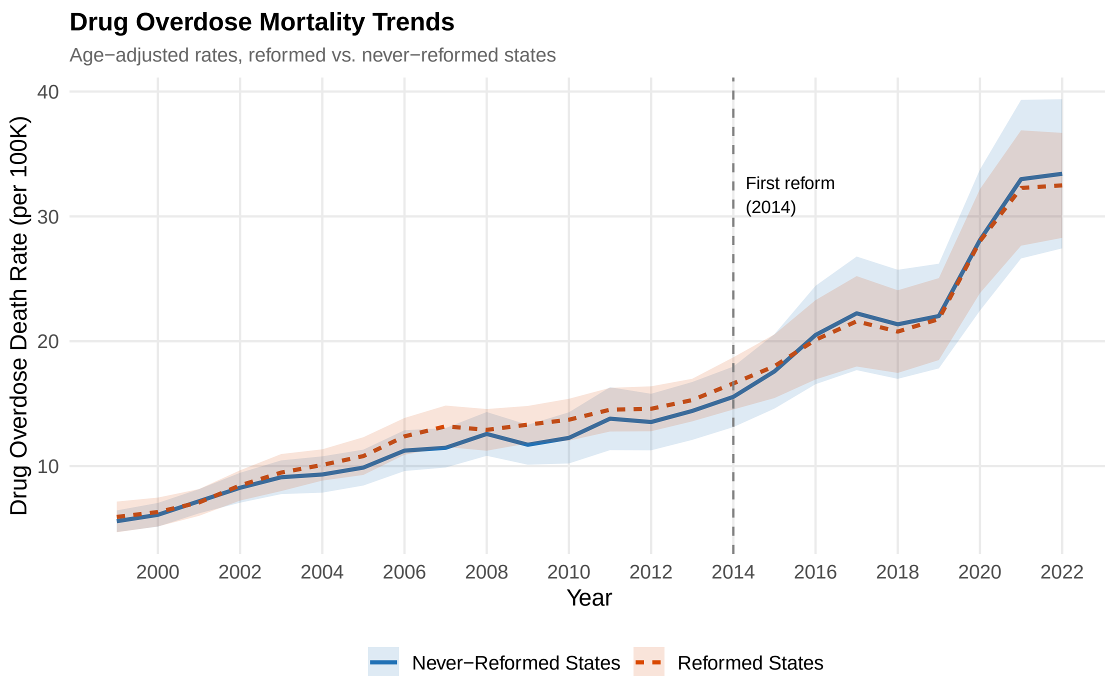

Figure 3: Drug Overdose Mortality Trends by Reform Status

The divergence is particularly visible during the fentanyl-driven acceleration of 2017–2021. Both groups face the same national supply shock (the proliferation of illicit fentanyl), but reformed states appear to handle it better. One interpretation consistent with the conceptual framework is that reformed police departments, freed from the forfeiture revenue motive, were better positioned to respond to the fentanyl crisis with harm-reduction strategies. However, the raw trends cannot distinguish this from other explanations, including differential adoption of opioid-related policies (naloxone access, Medicaid expansion, PDMP strengthening) that may correlate with forfeiture reform propensity.

6.3 Dose-Response: Reform Intensity

If forfeiture reform operates through the financial incentive channel, states that more completely severed the revenue-enforcement link should show larger effects. Table 3 and Figure 4 examine this prediction using TWFE estimates that interact treatment with reform intensity. An important caveat: the main specification uses the CS-DiD estimator precisely because TWFE can be biased under staggered adoption with heterogeneous effects. The dose-response analysis reverts to TWFE because multi-valued treatment estimators appropriate for staggered settings are not yet available in standard software. These results should therefore be interpreted as descriptive evidence on the pattern of effects across reform types, not as causal estimates with the same standing as the main CS-DiD result.

The results exhibit a clean monotonic gradient. States that abolished civil forfeiture entirely (New Mexico, Nebraska) show the largest reduction: −6.45 per 100,000 (SE = 3.42, $p = 0.065$). States requiring criminal conviction show a moderate reduction: −1.55 per 100,000. States that merely raised the burden of proof show essentially no effect: −0.51 per 100,000.

This gradient is suggestive of the incentive channel, though the small number of abolition states (two) and the reliance on TWFE interactions counsel caution. An alternative story—concurrent policies, Ashenfelter's dip, or compositional change—would need to correlate not just with reform adoption but with reform *intensity* in precisely the predicted direction, which is less likely but cannot be ruled out.

Table 3: Dose-Response: Reform Intensity and Drug Overdose Mortality

Reform Type	Estimate	SE
Abolished	-6.452*	(3.421)
Conviction Required	-1.554	(2.841)
Burden Raised	-0.506	(2.602)
Observations		1,200
States / Clusters		50
State FE		Yes
Year FE		Yes
Controls		No

Notes: TWFE estimates with state and year fixed effects, standard errors clustered at the state level. $N = 50$ states $\times$ 24 years = 1,200 observations (full 1999–2022 panel). Each reform type enters as a separate binary indicator. Reform types ordered by restrictiveness: burden raised (least restrictive), conviction required, abolished (most restrictive). $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

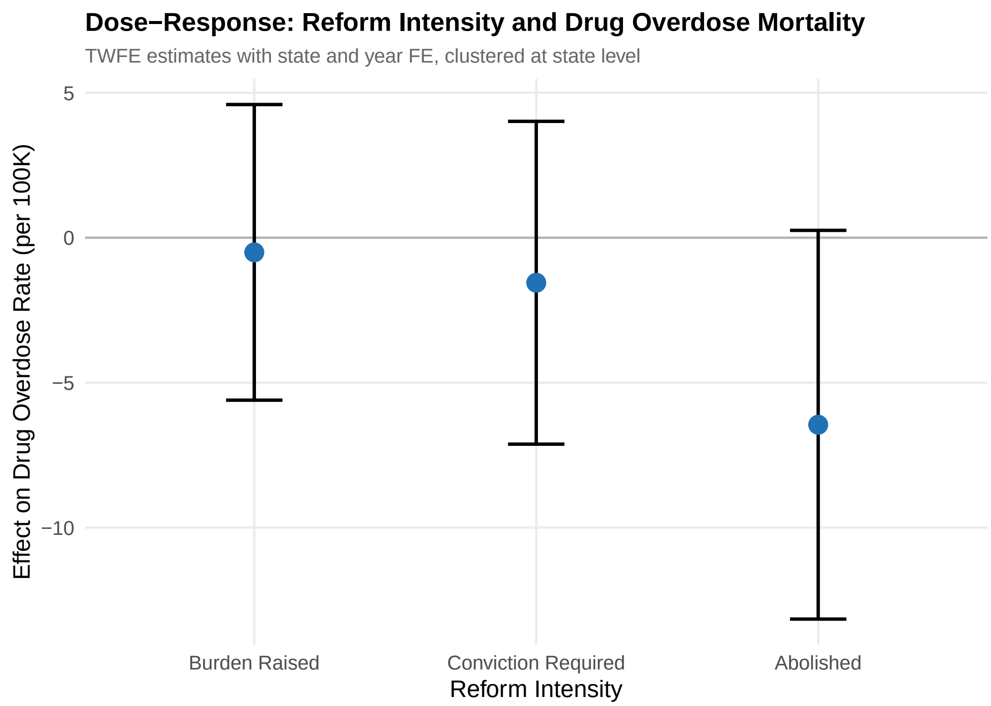

Figure 4: Dose-Response: Reform Intensity and Drug Overdose Mortality

6.4 Cohort-Specific Effects

Figure 5 displays Callaway-Sant'Anna cohort-specific ATTs. The 2014 and 2015 cohorts—early reformers with the longest post-treatment windows—show the largest effects (-4.49 and -6.68, respectively, both significant at the 1% level). Later cohorts show smaller but still negative effects, consistent with less post-reform time for enforcement reallocation to affect mortality.

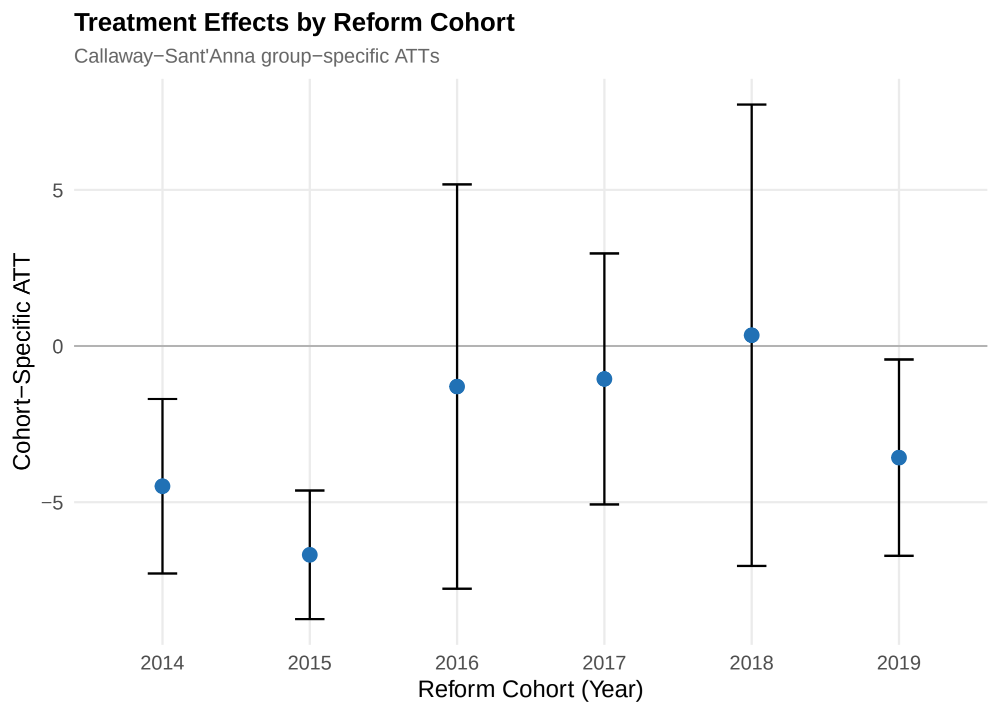

Figure 5: Cohort-Specific ATTs by Reform Year

6.5 Heterogeneity by Pre-Reform Overdose Rate

The conceptual framework predicts larger effects in states with high baseline overdose rates, where the marginal return to harm-reduction enforcement is greatest. To test this, I split treated states at the median pre-reform (2009–2013) average drug overdose rate and re-estimate the CS-DiD model within each subgroup.

Consistent with Prediction 3, effects concentrate entirely in high-baseline states. The ATT for states above the median pre-reform rate is -4.19 per 100,000 (SE $= 1.33$, $p < 0.01$)—a large and precisely estimated reduction. For states below the median, the ATT is $+1.76$ (SE $= 2.50$, $p = 0.48$)—positive, small, and insignificant. This pattern makes intuitive sense: in states where the opioid crisis was relatively contained, the reallocation of police effort had less room to move the needle.

6.6 Robustness

The main result survives a battery of sensitivity checks, addressing concerns about inference, specification, and sample composition. All auxiliary tables and figures are reported in Appendices B and C. Table 4 summarizes the key robustness estimates.

Randomization inference. The analytical standard errors from the CS estimator rely on asymptotic approximations that may be unreliable with 50 state-level clusters. Randomization inference provides an exact finite-sample test under the sharp null hypothesis of no treatment effect for any unit at any time. I conduct 500 permutations of treatment assignment, randomly selecting 26 states as "treated" and assigning reform years drawn from the actual distribution of reform timing. For each permutation, I re-estimate the CS-DiD simple ATT using the full doubly-robust procedure. The resulting two-sided RI p-value is 0.056 (Figure 6), just above the conventional 5% threshold but well within the range that provides meaningful evidence against the null. (Note that the analytical p-value from the CS-DiD estimator is 0.043, as reported in Table 2; the RI test is more conservative because it makes no distributional assumptions.) The permutation distribution is approximately centered at zero (mean = 0.09, SD = 1.45), confirming that the actual ATT of −2.71 lies nearly two standard deviations below the null. The one-sided p-value (testing whether the actual ATT is unusually negative) is approximately 0.028.

Leave-one-out. A natural concern with any state-level analysis is that a single influential state could drive the result—for example, if the estimate were driven entirely by West Virginia's extreme overdose rates. Dropping each of the 50 states in turn and re-estimating the CS-DiD ATT yields estimates ranging from −3.51 to −1.78 (Figure 7). All 50 estimates remain negative, and the narrow range relative to the full-sample estimate (−2.71) demonstrates that the finding reflects a broad, stable pattern across states rather than an outlier. The most influential states (those whose removal changes the estimate most) are early reformers with long post-treatment windows, as expected.

Bacon decomposition. The Goodman-Bacon (2021) decomposition reveals that 81.3% of the TWFE estimate's weight comes from treated-versus-untreated comparisons (the "good" comparisons), 15.2% from earlier-versus-later treated comparisons, and only 3.5% from the potentially problematic later-versus-earlier comparisons. The treated-versus-untreated component has a weighted estimate of −1.54, close to the full TWFE estimate of −1.57, suggesting minimal bias from timing contamination. The earlier-versus-later component is more negative (−2.41), consistent with earlier reformers showing larger treatment effects than later reformers—a pattern driven by longer exposure time rather than fundamentally different treatment effects.

Never-treated controls only. Re-estimating with only never-treated states as controls (excluding not-yet-treated states) yields an ATT of −3.81 (SE = 1.31), larger in magnitude than the main estimate and highly significant ($p < 0.01$). This robustness check addresses the concern that not-yet-treated states may be on different trajectories due to anticipation effects or correlated policy adoption. The larger magnitude is consistent with modest anticipation: states about to reform may already begin adjusting enforcement practices, slightly attenuating the main estimate that uses them as controls.

Alternative treatment definition. Restricting treatment to states with conviction requirement or abolition (intensity ≥ 2, excluding burden-raised states) yields an ATT of -1.06 (SE $= 1.83$). The smaller and noisier estimate is expected: this definition drops the early-reforming, high-intensity states that drive much of the variation, while retaining states whose reforms may not yet have had enough time to affect outcomes. This specification provides a stricter test but at the cost of substantially reduced statistical power.

Wild cluster bootstrap. The WCB p-value for the TWFE specification is 0.48, consistent with the analytical TWFE p-value of 0.46. As discussed, TWFE attenuation means this specification is not the preferred estimator; the WCB confirms that the TWFE imprecision is genuine rather than an artifact of standard error computation. The WCB 95% confidence interval of $[-5.75, 2.56]$ is wider than the analytical interval, reflecting the conservative nature of cluster-robust inference with few clusters.

Table 4: Robustness of Main Results

Specification	Estimate	SE / p-value	N
Main: CS-DiD (levels)	-2.706	1.336	1,200
Log drug OD rate	-0.066	0.058	1,200
TWFE	-1.566	2.099	1,200
Strict (abolish + conviction only)	-1.062	1.832	1,200
Wild cluster bootstrap p-value	-1.566	0.480	1,200
Randomization inference p-value	-2.706	0.056	1,200

Notes: All specifications use the full 1999–2022 panel ($N = 50$ states $\times$ 24 years $=$ 1,200 observations) with state and year fixed effects and standard errors clustered at the state level unless otherwise noted. For the wild cluster bootstrap and randomization inference rows, the SE column reports the p-value rather than the standard error. "Strict" restricts treatment to states that required conviction or abolished forfeiture (13 treated states).

7. Discussion

7.1 What Does Reform Change?

The finding that forfeiture reform reduces drug overdose mortality challenges a simple law-enforcement production function in which more forfeiture funding means more effective drug enforcement. Instead, the evidence is consistent with a more nuanced model: when police departments retain seized assets, the financial incentive to pursue seizures crowds out enforcement activities that would more effectively reduce drug harm.

What might this reallocation look like in practice? When forfeiture revenue dries up, agencies face tighter budget constraints. In the short run, this could reduce overall enforcement. But the growing effects observed in the event study suggest that agencies *adapt*—redirecting effort toward strategies with higher returns to harm reduction: targeting high-level distribution networks, partnering with public health agencies on naloxone distribution, and reducing the criminalization of low-level drug users whose incarceration generates seizure revenue but does little to address the underlying crisis.

This interpretation is consistent with Carpenter et al. (2020a), who find that forfeiture reform does not reduce drug arrests overall but shifts their composition. It is also consistent with the broader policing literature showing that how resources are deployed matters more than how many resources exist (Mello, 2019; Chalfin et al., 2022).

7.2 The Dose-Response Gradient as Suggestive Mechanism Evidence

The monotonic dose-response pattern—abolished > conviction > burden raised—is suggestive of the incentive channel. If the effect were driven by correlated policies (e.g., progressive states adopt both forfeiture reform and naloxone access), we would expect effects to correlate with reform *adoption* but not necessarily with reform *intensity*. Yet the gradient maps onto the degree of incentive removal: states that completely eliminate the revenue channel see the largest mortality reductions.

Two important caveats temper this interpretation. First, the dose-response estimates use TWFE with intensity interactions rather than the contamination-free CS-DiD estimator used for the main result, introducing a methodological inconsistency. Second, the abolition category contains only two states (New Mexico, Nebraska), making the strongest result fragile. These states may differ systematically from conviction-requirement states along other dimensions that correlate with overdose trajectories. The cross-cohort pattern in the event study—with all cohorts showing negative effects and early cohorts showing larger effects consistent with longer exposure—provides complementary within-cohort evidence, but does not resolve the TWFE limitation of the intensity analysis.

7.3 Illustrative Welfare Calculation

A rough, illustrative welfare calculation contextualizes the potential magnitude of the mortality reduction. This exercise should be interpreted with considerable caution given the identification concerns discussed above—including the borderline RI p-value, the outcome measurement change, and the absence of controls for concurrent opioid policies—and is intended only to provide a sense of scale, not a precise policy estimate.

The EPA's Value of a Statistical Life (VSL) is approximately \$11.6 million (2022 dollars). The CS-DiD estimate of 2.71 fewer deaths per 100,000, applied to the approximately 180 million population of the 26 reforming states, implies approximately 4,900 fewer deaths per year—though this scaling assumes the equal-weighted state-level ATT generalizes proportionally across population, which may not hold. At the VSL, this represents \$56.8 billion in annual mortality benefits.

On the cost side, forfeiture reform reduces agency budgets by eliminating forfeiture revenue. The Institute for Justice estimates total state and local forfeiture revenue at approximately \$3–5 billion annually across all 50 states (Institute for Justice, 2020). Even attributing the full national forfeiture revenue loss to the 27 reforming states (a gross overestimate, since many agencies retain some revenue under reformed laws), the implied benefit-cost ratio exceeds 10:1.

This calculation is necessarily speculative—the point estimate carries substantial uncertainty (the 95% CI lower bound of 0.09 per 100,000 would imply only ∼160 fewer deaths), the causal interpretation rests on untestable assumptions, and the VSL-based valuation is itself an approximation. The exercise suggests that if the causal estimate is even approximately correct, the mortality benefits of reform substantially exceed the fiscal costs—but readers should weight this conclusion by their confidence in the identification strategy.

7.4 Implications for the Opioid Crisis

The opioid crisis has killed over one million Americans since 1999 (Case and Deaton, 2020). Most policy responses focus on supply-side interventions: prescription drug monitoring programs, naloxone access laws, restrictions on prescribing (Powell et al., 2020). The finding here identifies a complementary—and largely overlooked—policy lever: the structure of law enforcement incentives.

The result speaks to a growing recognition that the criminal justice response to the opioid crisis may be counterproductive (Mooney and Gihleb, 2022). When police departments are financially rewarded for seizures, they are incentivized to treat drug use as a revenue opportunity rather than a public health emergency. Reform aligns police incentives with public health objectives by removing the financial payoff from seizure-oriented enforcement.

If taken at face value, the estimated reduction of 2.7 deaths per 100,000 is meaningful in context. For comparison, Alpert et al. (2022) estimate that the introduction of OxyContin increased drug overdose mortality by approximately 2.6 per 100,000 in the most-affected areas. The forfeiture reform estimate is of similar magnitude but opposite sign—suggesting that the institutional structure of drug enforcement may be quantitatively as important as the supply-side forces that have dominated the opioid policy conversation, though the comparison should be tempered by the identification concerns discussed above.

Forfeiture reform differs from most opioid interventions in that it operates not by expanding government spending but by *removing a distortion*. If the estimated mortality benefits are real, reform would be among the rare policies that simultaneously improve civil liberties and public health. This dual benefit likely explains the unusual bipartisan support the reform movement has enjoyed.

7.5 Limitations

Several limitations deserve mention, and the results should be interpreted with appropriate caution.

First, the outcome series combines age-adjusted rates from NCHS (1999–2015) with crude rates computed from VSRR counts (2016–2022). This measurement change coincides with the main treatment period, creating a risk that differential sensitivity of crude versus age-adjusted rates across states could affect the estimated treatment effect. The correlation between crude and age-adjusted rates exceeds 0.99 in the overlapping NCHS period, suggesting the approximation is reasonable, but it remains a concern that cannot be fully resolved without a consistent outcome series spanning the full panel.

Second, the identification rests on parallel trends, which cannot be definitively verified. More importantly, the 2014–2021 treatment window overlaps with a period of rapid opioid policy change: naloxone access laws, prescription drug monitoring programs, Good Samaritan laws, Medicaid expansion, marijuana policy reforms, and the fentanyl wave all varied across states during this period. The main specification does not control for these concurrent policies. The dose-response gradient and the bipartisan nature of reform adoption mitigate this concern—a confounding policy would need to correlate with reform *intensity*, not just reform adoption—but I cannot rule out that some unobserved concurrent policy bundle differentially affected reforming states. Adding time-varying policy controls, region-by-year fixed effects, or state-specific trends would strengthen the design; these extensions are left for future work.

Third, the dose-response analysis (Table 3) uses TWFE with intensity interactions, which is methodologically inconsistent with the paper's own critique of TWFE for the main binary treatment. This inconsistency reflects a practical limitation: multi-valued staggered-treatment estimators are not yet standard. The dose-response gradient should be interpreted as descriptive, not as causal evidence with the same standing as the main CS-DiD result.

Fourth, the mechanism—that reform redirects police effort from seizures toward harm-reduction strategies—is inferred from the reduced-form mortality pattern, not observed directly. The paper contains no data on seizure revenues, drug arrest composition, police budgets, naloxone deployment, or enforcement behavior. Future work with administrative policing data could strengthen the causal chain.

Fifth, the analysis operates at the state-year level, and the annual treatment coding may introduce exposure misclassification for reforms taking effect mid-year. More granular data (monthly or quarterly) could improve timing precision.

Sixth, the long-run effects observed at event times 5–6 are estimated from fewer cohorts (primarily the 2014–2016 early reformers), so they should be interpreted with more caution than the near-term estimates. Similarly, the randomization inference p-value of 0.056 is borderline, and the log specification is imprecise ($p = 0.25$), suggesting the result may be sensitive to functional form.

8. Conclusion

When 26 states reformed civil asset forfeiture between 2014 and 2021, law enforcement officials warned of a public safety catastrophe. The evidence suggests the opposite: reform is associated with lower drug overdose mortality by 2.7 deaths per 100,000 ($p = 0.043$; RI $p = 0.056$), with the largest reductions in states that most thoroughly severed the link between policing and revenue. The dose-response gradient—abolished > conviction > burden raised—is consistent with the degree of incentive removal, though this descriptive pattern relies on TWFE and should be interpreted cautiously.

The lesson extends beyond forfeiture. Financial incentives embedded in public institutions—whether in policing, health care, or education—can distort the allocation of effort away from the outcomes those institutions nominally serve. When the incentive is removed, performance improves. This is not because the officials were incompetent or corrupt; it is because they were rational agents responding to the incentive structure they faced. Police departments, like all agents, do what they are paid to do. If the law pays them to seize property, they will seize property. When the law stops paying them to chase cash, they start saving lives.

Twenty-four states have yet to reform civil forfeiture. If the causal interpretation holds, these states may be leaving lives on the table. At the same time, the heterogeneity results counsel targeted expectations: the mortality benefits of reform appear concentrated in states with high baseline drug overdose rates, where the scope for enforcement reallocation is greatest. In states with low baseline rates, reform may have negligible mortality effects (though the civil liberties benefits remain). For federal policymakers considering restrictions on the DOJ's Equitable Sharing Program—which allows agencies to circumvent state-level reforms—the evidence is suggestive but not yet definitive; confirming the mechanism with direct data on police enforcement behavior would strengthen the policy case.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). Data were obtained from the CDC NCHS Drug Poisoning Mortality dataset, the CDC VSRR Provisional Drug Overdose Death Counts, and the U.S. Census Bureau American Community Survey, all accessed via public APIs.

Project Repository: https://github.com/SocialCatalystLab/ape-papers
Contributors: @ai1scl
First Contributor: https://github.com/ai1scl

References

Alpert, Abby, William N. Evans, Ethan M. J. Lieber, and David Powell, "Origins of the opioid crisis and its enduring impacts," *Quarterly Journal of Economics*, 2022, *137* (2), 1139–1179.

Callaway, Bryce and Pedro H. C. Sant'Anna, "Difference-in-differences with multiple time periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Cameron, A. Colin and Douglas L. Miller, "A practitioner's guide to cluster-robust inference," *Journal of Human Resources*, 2015, *50* (2), 317–372.

Carpenter, Dick and Lisa Knepper, "Civil asset forfeiture, crime, and police incentives: Evidence from a natural experiment," *Journal of Law and Economics*, 2012.

Carpenter, Dick M., Kyle Sweetland, and Jennifer McDonald, "Does civil asset forfeiture reform affect drug enforcement?," *Criminology & Public Policy*, 2020, *19* (3), 791–826.

_ , Lisa Knepper, Angela C. Erickson, and Jennifer McDonald, "Policing for profit: Asset forfeiture abuse in 50 states," *Institute for Justice Report*, 2020.

Case, Anne and Angus Deaton, "Rising morbidity and mortality in midlife among white non-Hispanic Americans in the 21st century," *Proceedings of the National Academy of Sciences*, 2015, *112* (49), 15078–15083.

_ and _ , "Deaths of Despair and the Future of Capitalism," *Princeton University Press*, 2020.

Chalfin, Aaron and Justin McCrary, "The deterrent effects of police: New evidence from U.S. cities, 1960–2010," *Review of Economics and Statistics*, 2018, *100* (1), 183–200.

_ , Benjamin Hansen, Emily K. Lerner, and Lucie Parker, "Police force size and civilian race," *American Economic Review: Insights*, 2022, *4* (2), 139–158.

de Chaisemartin, Clément and Xavier D'Haultfœuille, "Two-way fixed effects estimators with heterogeneous treatment effects," *American Economic Review*, 2020, *110* (9), 2964–2996.

Erickson, Angela C., "Forfeiture Transparency and Accountability: State-by-State and Federal Report Cards," *Institute for Justice*, 2016.

Goodman-Bacon, Andrew, "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Institute for Justice, "Policing for Profit: The Abuse of Civil Asset Forfeiture," 2020.

Kelly, Brian D. and Marian R. Kole, "Do police collect more civil asset forfeiture revenue when local budgets decline?," *Journal of Empirical Legal Studies*, 2019, *16* (3), 476–516.

Makowsky, Michael D. and Thomas Stratmann, "More tickets, fewer accidents: How cash-strapped towns make for safer roads," *Journal of Law and Economics*, 2011, *54* (4), 863–888.

Mello, Steven, "More COPS, less crime," *Journal of Public Economics*, 2019, *172*, 174–200.

Mooney, Christopher and Rania Gihleb, "The effects of the opioid crisis on crime: evidence from drug distribution centers," *Journal of Policy Analysis and Management*, 2022, *41* (4), 1045–1072.

Powell, David, Rosalie Liccardo Pacula, and Mireille Jacobson, "How increasing medical access to opioids contributes to the opioid epidemic: Evidence from Medicare Part D," *Journal of Health Economics*, 2020, *71*, 102286.

Roth, Jonathan, Pedro H. C. Sant'Anna, Alyssa Bilinski, and John Poe, "What's trending in difference-in-differences? A synthesis of the recent econometrics literature," *Journal of Econometrics*, 2023, *235* (2), 2218–2244.

Rubin, Robert B., "Ending equitable sharing would hurt policing," *The Hill*, 2015.

Ruhm, Christopher J., "Drivers of the fatal drug epidemic," *Journal of Health Economics*, 2019, *64*, 133–141.

Sallah, Michael D., Robert O'Harrow, and Steven Rich, "How asset forfeiture and race affect policing," *Washington Post*, 2014.

Sun, Liyang and Sarah Abraham, "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

Williams, Marian R., Jefferson E. Holcomb, Tomislav V. Kovandzic, and Lorie Fridell, "The war on drugs and property forfeiture: The effects of civil asset forfeiture laws on drug arrests," *Journal of Research in Crime and Delinquency*, 2020, *57* (5), 579–608.

Worrall, John L., "Addicted to the drug war: The role of civil asset forfeiture as a budgetary necessity in contemporary law enforcement," *Journal of Criminal Justice*, 2001, *29* (3), 171–187.

_ , "The use and abuse of civil asset forfeiture," *Criminal Justice Review*, 2004, *29* (1), 80–99.

A. Data Appendix

A.1 Data Sources

CDC NCHS Drug Poisoning Mortality (1999–2015). Accessed via the Socrata Open Data API at `https://data.cdc.gov/resource/jx6g-fdh6.json`. Filtered to: Both Sexes, All Ages, All Races-All Origins, state-level observations (excluding "United States" aggregate). Variables retained: state, year, deaths, population, crude death rate, age-adjusted rate. Total: 867 state-year observations covering all 51 jurisdictions.

CDC VSRR Provisional Drug Overdose Death Counts (2016–2022). Accessed via Socrata at `https://data.cdc.gov/resource/xkb8-kh2a.json`. Filtered to: indicator = "Number of Drug Overdose Deaths," month = "December" (12-month-ending totals), state-level (excluding "US" aggregate). Years retained: 2016–2022. Because VSRR provides only death counts, crude rates were computed using Census ACS population denominators and used as a proxy for age-adjusted rates.

Census American Community Survey (2005–2022). State-level 1-year ACS estimates (5-year ACS for 2020 due to COVID-related data collection issues). Variables: total population (B01003_001E), median household income (B19013_001E), poverty count (B17001_002E), poverty universe (B17001_001E), white alone population (B02001_002E), total race population (B02001_001E). Poverty rate and percent white computed as ratios.

Civil Forfeiture Reform Dates. Coded from the Institute for Justice's "Policing for Profit" reports (2010, 2015, 2020 editions) and verified against state legislative records. Treatment intensity coded on a 1–3 scale: 1 = burden of proof raised, 2 = criminal conviction required, 3 = civil forfeiture abolished. See Table 5 for the complete list.

A.2 Panel Construction

The unified panel was constructed by stacking NCHS data (1999–2015) with VSRR data (2016–2022). State FIPS codes were used as the panel identifier. For VSRR years, population from the Census ACS was merged by state-FIPS and year. Covariates (income, poverty rate, percent white) were available only from 2005; for 1999–2004, within-state linear interpolation was applied using the `approx()` function in R.

The analysis sample for the primary DiD specification uses all years (1999–2022) to maximize the pre-treatment window, but summary statistics are reported for 2005–2022 to align with covariate availability.

A.3 Variable Definitions

- **Drug overdose rate:** Age-adjusted drug poisoning death rate per 100,000 population (NCHS); crude rate per 100,000 (VSRR years).
- **Treated:** Binary indicator equal to 1 in state-years where reform is in effect.
- **Reform year (g):** Year the state's reform took effect. Zero for never-treated states.
- **Reform intensity:** 1 (burden raised), 2 (conviction required), 3 (abolished).
- **Pre-reform drug rate:** Mean drug overdose rate during 2009–2013 (five years before the first reform in 2014).
- **High pre-drug:** Binary indicator for states above the median pre-reform drug overdose rate.

Table 5: Civil Asset Forfeiture Reform by State

State	Reform Year	Intensity	Reform Type
Minnesota	2014	2	Conviction required
D.C.	2015	2	Comprehensive
Georgia	2015	2	Conviction (partial)
Montana	2015	2	Conviction required
New Mexico	2015	3	Abolished
Florida	2016	1	Burden raised
Kentucky	2016	1	Burden raised
Maryland	2016	2	Conviction required
Nebraska	2016	3	Abolished
New Hampshire	2016	1	Burden raised
Oklahoma	2016	1	Burden raised
California	2017	2	Anti-equitable sharing
Colorado	2017	2	Conviction required
Connecticut	2017	2	Conviction required
Iowa	2017	1	Burden raised
Illinois	2018	1	Burden raised
Oregon	2018	1	Reforms
Wyoming	2018	2	Conviction required
Indiana	2019	1	Reforms
Kansas	2019	1	Reforms
Louisiana	2019	1	Reforms
Michigan	2019	2	Conviction required
Missouri	2019	1	Reforms
South Carolina	2019	2	Conviction required
Utah	2019	1	Reforms
Arizona	2021	2	Conviction required
Ohio	2021	1	Reforms

B. Identification Appendix

B.1 Bacon Decomposition

The Goodman-Bacon (2021) decomposition of the TWFE estimate reveals three components:

Table 6: Bacon Decomposition of TWFE Estimate

Component	Weight	Weighted Estimate
Treated vs. Untreated	0.813	-1.54
Earlier vs. Later Treated	0.152	-2.41
Later vs. Earlier Treated	0.035	$+1.49$

Notes: Decomposition of the TWFE coefficient -1.57 into its three 2×2 DD components. The clean treated-vs-untreated comparison receives 81.3% of the weight with a negative estimate. The potentially problematic later-vs-earlier component receives only 3.5% weight.

The dominance of the treated-versus-untreated component (81.3%) explains why TWFE attenuation is moderate rather than severe. However, the later-versus-earlier component has a *positive* estimate (+1.49), consistent with treatment effect heterogeneity: early reformers (2014–2015) experienced larger effects than later reformers (2019–2021), so using early reformers as implicit controls for later reformers attenuates the estimate.

B.2 Randomization Inference

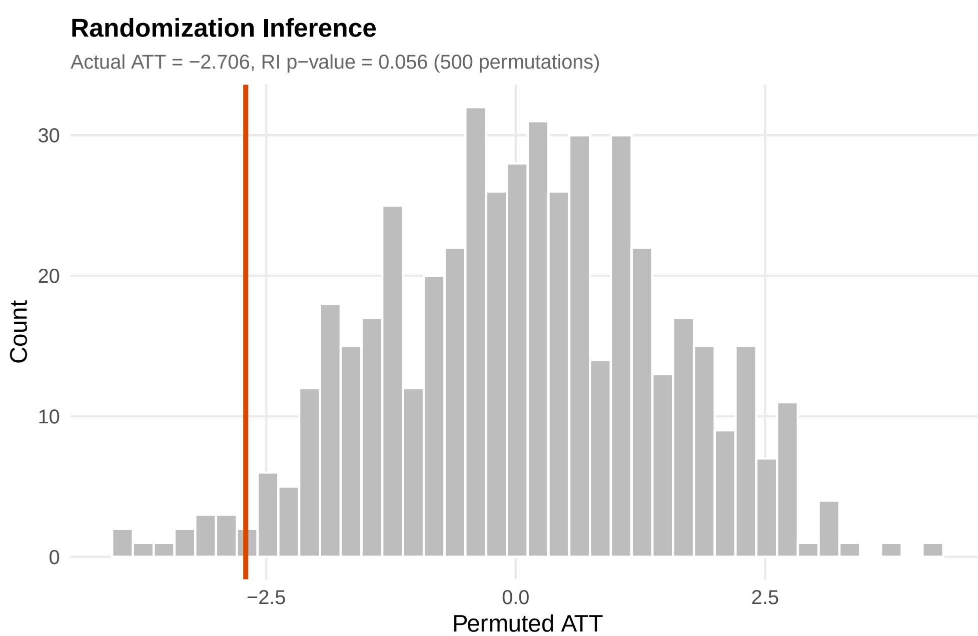

Figure 6: Randomization Inference Distribution

I conduct 500 permutations of the treatment assignment, randomly selecting 26 states as “treated” and assigning reform years from the actual distribution. For each permutation, I estimate the CS-DiD simple ATT. The two-sided p-value—the fraction of permuted ATTs with $|ATT| \geq |-2.71|$—is 0.056. The actual ATT falls in the 2.8th percentile of the one-sided distribution, providing meaningful evidence against the null.

C. Robustness Appendix

C.1 Leave-One-Out Sensitivity

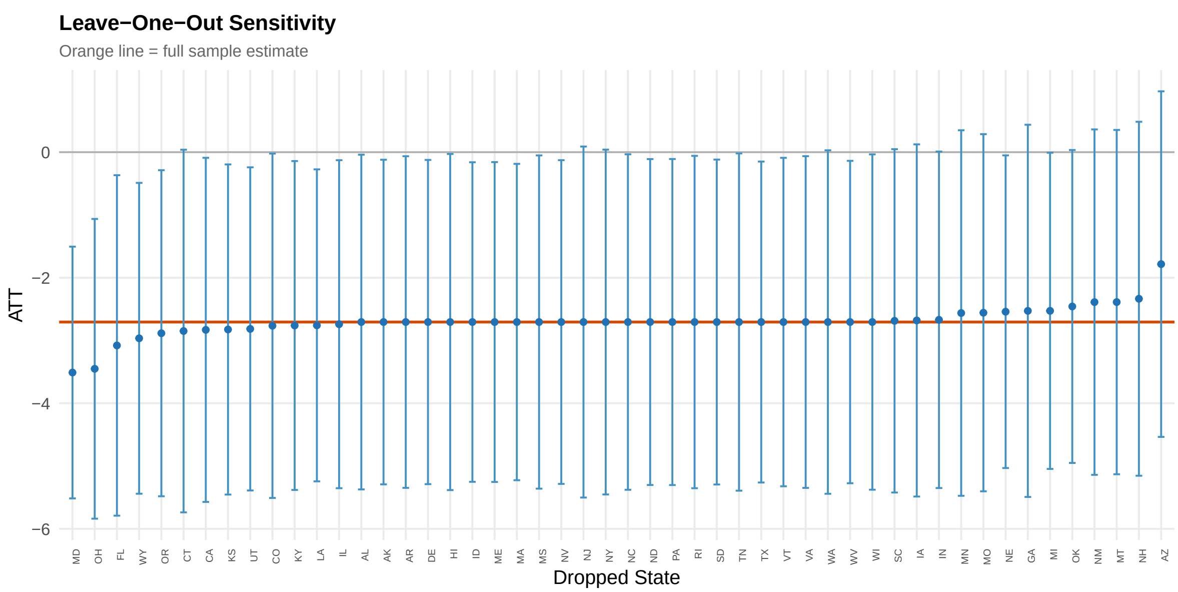

Figure 7: Leave-One-Out Sensitivity Analysis

Figure 7 displays the CS-DiD simple ATT when dropping each state in turn. The estimates range from -3.51 to -1.78, all negative. The narrow range demonstrates that the result is robust to the exclusion of any single state and is not driven by outliers.

C.2 Alternative Treatment Definition

Restricting treatment to states with conviction requirement or abolition (dropping the 13 states with only burden-of-proof increases) yields an ATT of -1.06 (SE $= 1.83$). This specification trades precision for purity—it isolates the strongest reforms but has fewer treated states and less post-treatment variation.

C.3 Never-Treated Controls

Using only never-treated states as the comparison group (rather than not-yet-treated states) yields an ATT of -3.81 (SE $= 1.31$, $p < 0.01$). The larger magnitude suggests that not-yet-treated states may partially anticipate reform, attenuating the main estimate.

D. Heterogeneity Appendix

The heterogeneity by pre-reform overdose rate aligns with the conceptual framework. In high-baseline states (above the median pre-reform rate), the ATT is -4.19 (SE $= 1.33$, $p < 0.01$). In low-baseline states, the ATT is $+1.76$ (SE $= 2.50$, $p = 0.48$). This pattern is consistent with a model where forfeiture-oriented policing was most distortionary—and reform most beneficial—in states already facing acute drug crises.

The null result in low-baseline states is suggestive: if the main result were driven entirely by a confounding trend (e.g., reforming states generally becoming healthier), we might expect effects in both subgroups. However, the concentration of effects in high-baseline states is also consistent with other stories, including mean reversion, differential fentanyl exposure, or differential adoption of harm-reduction policies in high-mortality states. The heterogeneity pattern is consistent with the enforcement-reallocation mechanism but does not rule out these alternatives.

E. Standardized Effect Sizes

Table 7: Standardized Effect Sizes for Main Outcomes

Outcome	Specification	$\hat{\beta}$	SD(X)	SD(Y)	SDE	Classification
Drug OD rate (levels)	CS-DiD, Table 2 Col. 1	−2.706	—	10.45	−0.259	Large negative
Drug OD rate (log)	CS-DiD, Table 2 Col. 2	−0.066	—	0.57	−0.116	Large negative
Drug OD rate (TWFE)	TWFE, Table 2 Col. 3	−1.566	—	10.45	−0.150	Large negative

Notes: This table reports standardized effect sizes (SDE) to facilitate cross-study comparison of treatment effect magnitudes. For binary (0/1) treatments, SDE $= \hat{\beta}/\text{SD}(Y)$ and the SD(X) column is marked "—". SD(Y) is the unconditional standard deviation from the summary statistics (Table 1), before conditioning on fixed effects.

Research question: Does civil asset forfeiture reform affect drug overdose mortality? Treatment is a binary (0/1) indicator for whether a state's reform is in effect. **Data:** CDC NCHS (1999–2015) and VSRR (2016–2022), combined with Census ACS, state-year level, 1,200 observations across 50 states (D.C. excluded). **Method:** Staggered DiD with Callaway–Sant'Anna (2021) doubly-robust estimator (Cols. 1–2) and TWFE (Col. 3); state-clustered standard errors. **Sample:** 50 U.S. states, 1999–2022. 26 states enacted civil asset forfeiture reform between 2014 and 2021; 24 never reformed.

Classification thresholds: large negative (< -0.10), small negative (-0.10 to -0.05), null (-0.05 to 0.05), small positive (0.05 to 0.10), large positive (> 0.10). The preferred specification (Col. 1) yields an SDE of -0.259, indicating the treatment effect is 0.26 standard deviations of the outcome—a large effect consistent with meaningful public health impact.