The Symmetric Test: Drug Decriminalization and Recriminalization in Oregon

APEP Autonomous Research* @olafdrw

March 10, 2026

Abstract

Oregon decriminalized drug possession in February 2021 (Measure 110) and recriminalized it in September 2024 (HB 4002), creating a rare symmetric natural experiment. I exploit both policy switches using the synthetic control method to test whether decriminalization causally increased overdose deaths or merely coincided with the national fentanyl wave. Design 1 estimates that Oregon's overdose rate diverged from its synthetic control by 10.888 deaths per 100,000 after decriminalization ($p_{\text{RI}} = 0.020$). Design 2 estimates a reconvergence of 6.722 deaths per 100,000 after recriminalization. The symmetric sum $\hat{\tau}_{\text{sum}} = 4.166$ cannot reject full causal reversal ($p = 0.552$; joint permutation $p = 0.549$). Drug decomposition reveals synthetic opioids account for 83% of the post-decriminalization divergence, suggesting delayed fentanyl penetration substantially confounds the estimated treatment effect.

JEL Codes: I12, I18, K42

Keywords: drug decriminalization, synthetic control, overdose deaths, fentanyl, Oregon Measure 110

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

1. Introduction

On November 3, 2020, Oregon voters passed Measure 110 by a 17-point margin, making Oregon the first U.S. state to decriminalize possession of small quantities of heroin, methamphetamine, cocaine, and other controlled substances. Three years later, with overdose deaths having nearly tripled and public support collapsing, the Oregon legislature reversed course: House Bill 4002, signed in March 2024 and effective September 1, recriminalized drug possession as a misdemeanor. A nurse at an overdose response team in Portland described the intervening period as "like watching a tsunami in slow motion—everyone could see it coming, but nobody could agree on what caused the wave" (Morgan et al., 2023).

This sequence of events—decriminalization followed by recriminalization of the same substances in the same state within 42 months—creates an identification opportunity that is rare in policy evaluation. When a policy is both enacted and repealed, any causal effect should appear twice: once as a divergence from the counterfactual at enactment, and once as a convergence at repeal. This is the logic of the *symmetric test*, which I formalize and apply in this paper. If decriminalization causally increased overdose mortality, recriminalization should have reduced it; if both effects are present and approximately equal in magnitude, the evidence for causation is considerably stronger than either estimate alone.

The challenge, of course, is that Oregon's decriminalization period coincided almost perfectly with the deadliest phase of the national fentanyl crisis. Between 2020 and 2023, synthetic opioid deaths in the United States more than doubled, driven by the penetration of illicitly manufactured fentanyl into drug markets that had previously been dominated by heroin, prescription opioids, and methamphetamine (Centers for Disease Control and Prevention, 2025). The temporal coincidence between Measure 110 and the fentanyl wave makes it impossible to attribute Oregon's rising death toll to decriminalization using simple before-after comparisons. Every state experienced dramatic increases during this period; the question is whether Oregon's increase was *differentially* larger than what a credible counterfactual predicts.

I address this identification challenge using the synthetic control method (Abadie and Gardeazabal, 2003; Abadie et al., 2010; Abadie, 2021), applied twice in a symmetric design. Design 1 constructs a synthetic Oregon from weighted averages of donor states to estimate the effect of decriminalization (February 2021). Design 2 constructs a second synthetic Oregon, now using the decriminalization period as the pre-treatment window, to estimate the effect of recriminalization (September 2024). Design 3—the symmetric test—combines both estimates: under the null hypothesis of full causal reversal, $\hat{\tau}_{\text{decrim}} + \hat{\tau}_{\text{recrim}} = 0$.

The main results tell a nuanced story. Design 1 estimates a post-decriminalization divergence of 10.888 deaths per 100,000 between Oregon and its synthetic control, with a randomization inference p-value of 0.020—Oregon ranks first among all 51 units (Oregon plus 50 donors) in its post-treatment mean squared prediction error ratio. Design 2 estimates a reconvergence of 6.722 deaths per 100,000 after recriminalization. The symmetric sum $\hat{\tau}_{\text{sum}} = 4.166$ has a standard error of 6.999, yielding $p = 0.552$: the data cannot reject full causal reversal. A joint permutation test that accounts for dependence between the two designs yields $p = 0.549$, confirming this conclusion. The reversal ratio $-\hat{\tau}_{\text{recrim}}/\hat{\tau}_{\text{decrim}} = 0.62$ suggests that roughly two-thirds of the post-decriminalization divergence was unwound within the first months of recriminalization.

These aggregate estimates, however, mask a critical compositional pattern. Drug-specific decomposition reveals that synthetic opioids (primarily fentanyl) account for 9.048 of the 10.888 deaths per 100,000 attributed to decriminalization—83% of the total effect. Heroin contributes 1.3 and psychostimulants 3.6, while cocaine shows a small negative effect. This decomposition is important because Oregon's fentanyl exposure lagged the national average by approximately two years: in 2019, synthetic opioids comprised 14% of Oregon's overdose deaths versus 42% nationally. By 2023, Oregon had largely converged to the national share. The timing of this convergence—coinciding almost exactly with the decriminalization period—raises the possibility that the Design 1 estimate partly captures Oregon's delayed exposure to the fentanyl supply shock rather than a pure decriminalization effect.

This paper contributes to three literatures. First, it adds to the small but rapidly growing body of work evaluating Oregon's Measure 110 (Dave et al., 2023; Doleac, 2023; McGinty et al., 2023). Existing studies generally use difference-in-differences designs comparing Oregon to other states before and after 2021. The symmetric test design adds a fundamentally different source of variation: the repeal itself. By requiring the effect to appear in both directions, the design imposes a stronger falsification standard than any single-switch estimator.

Second, the paper contributes to the broader literature on drug policy and overdose mortality (Maclean et al., 2022; Schnell, 2024). The opioid crisis has generated a vast empirical literature examining supply-side interventions (Alpert et al., 2018), prescribing regulations (Currie et al., 2019), and economic determinants (Hollingsworth et al., 2017; Ruhm, 2019). Demand-side interventions like decriminalization have received less rigorous evaluation, in part because most decriminalization regimes—Portugal being the canonical example (Hughes and Stevens, 2010)—were implemented as permanent, one-time reforms without subsequent repeal. Oregon's reversal provides a setting that Portugal never did.

Third, the paper contributes methodologically to the literature on synthetic control methods (Abadie, 2021; Ferman and Pinto, 2021). The symmetric design represents a specific application of the logic that causal effects should be reversible: if a policy change causes an outcome to diverge from a counterfactual, the reversal of that policy should cause the outcome to reconverge. This idea has antecedents in A-B-A designs from experimental psychology, but to my knowledge this paper is the first to formalize and apply it within the synthetic control framework at the state level. The key advantage is that the symmetric test is robust to slow-moving confounders—trends that affect the level of the Design 1 estimate cancel in the sum if they are approximately linear over the full period.

2. Institutional Background

2.1 The Overdose Crisis and the Policy Debate

The United States has experienced three overlapping waves of drug overdose mortality since the late 1990s (Case and Deaton, 2015, 2020). The first wave, driven by prescription opioid oversupply, peaked around 2011. The second wave, beginning around 2010, reflected a shift to heroin as prescription opioids became harder to obtain following abuse-deterrent reformulations (Alpert et al., 2018). The third and most lethal wave, beginning around 2013 in eastern states but reaching western markets only around 2019–2020, was driven by illicitly manufactured fentanyl and its analogs (Powell et al., 2020).

This geographic gradient in fentanyl penetration is crucial for understanding Oregon's experience. As of 2019, synthetic opioids accounted for approximately 70% of overdose deaths in states like Ohio, West Virginia, and Connecticut, but only 14% in Oregon. The Pacific Northwest, along with the Mountain West, was among the last regions to experience large-scale fentanyl market penetration. This delay was not a policy choice; it reflected the structure of illicit drug supply chains, with Mexican cartels initially flooding eastern markets through established heroin distribution networks (Quinones, 2015).

Against this backdrop, the policy debate over drug decriminalization was shaped by two competing narratives. Proponents argued that criminal penalties for drug possession were counterproductive—deterring treatment-seeking, generating criminal records that impaired employment, and disproportionately affecting communities of color. The Portuguese model, implemented in 2001, was frequently cited as evidence that decriminalization could reduce drug-related harms without increasing use (Hughes and Stevens, 2010; Greenwald, 2009). Opponents argued that Oregon's approach was critically different from Portugal's: Measure 110 decriminalized possession without the robust mandatory assessment infrastructure that made Portugal's model work, effectively removing legal consequences without providing adequate treatment alternatives (Morgan et al., 2023).

2.2 Measure 110: Drug Decriminalization (February 2021)

Oregon Ballot Measure 110, approved by 58.5% of voters on November 3, 2020, reclassified possession of small quantities of controlled substances from a criminal offense to a Class E violation, punishable by a maximum $100 fine. The measure became effective on February 1, 2021. Specifically, it reduced possession of less than 1 gram of heroin, MDMA, or methamphetamine; less than 2 grams of cocaine; less than 12 grams of psilocybin; and less than 40 user units of LSD, oxycodone, or methadone from a Class A misdemeanor (carrying up to 1 year in jail) to a non-criminal violation (Oregon Secretary of State, 2020).

The measure also directed cannabis tax revenues toward a new Oversight and Accountability Council tasked with establishing addiction treatment centers. However, implementation of the treatment infrastructure was slow—by mid-2022, only a fraction of the allocated funds had been disbursed, and Oregon's treatment capacity remained far below what advocates had envisioned (Morgan et al., 2023). This implementation gap became a central argument for recriminalization.

2.3 HB 4002: Recriminalization (September 2024)

By late 2023, Oregon's drug overdose death rate had reached approximately 33 deaths per 100,000—more than double the pre-decriminalization rate of 13 per 100,000. Public opinion shifted dramatically: polls showed a majority of Oregonians now supported recriminalizing drug possession. On March 1, 2024, Governor Tina Kotek signed House Bill 4002, which reclassified possession of controlled substances as a Class A misdemeanor, restoring criminal penalties effective September 1, 2024 (Oregon Legislative Assembly, 2024).

HB 4002 was not a simple reversal of Measure 110. It included provisions for "deflection" programs allowing law enforcement to direct individuals to treatment rather than arrest, expanded funding for recovery services, and established a statewide crisis system. Nevertheless, the core legal change—from non-criminal violation back to misdemeanor—was substantively the reverse of what Measure 110 had implemented. For the purposes of the symmetric test, the key feature is that the legal status of drug possession was toggled twice in rapid succession, with the second switch approximately restoring the pre-treatment legal regime.

2.4 Timeline and Identification Windows

The relevant timeline for identification is:

- **Pre-treatment (Design 1):** January 2015 – January 2021 (73 months). Drug possession is a Class A misdemeanor in Oregon, as in most other states.

- **Post-decriminalization:** February 2021 – August 2024 (43 months). Possession is a Class E violation in Oregon; still a criminal offense in all other states.
- **Post-recriminalization:** September 2024 onward. Possession is again a Class A misdemeanor in Oregon.

For Design 2, the pre-treatment period is February 2021 – August 2024 (43 months), and the post-treatment period begins September 2024. The short post-treatment period for Design 2 (approximately 13 months of available data) limits the precision of the recriminalization estimate, but the direction and magnitude remain informative for the symmetric test.

3. Data

3.1 Overdose Mortality

The primary outcome variable is the 12-month-ending count of drug overdose deaths per 100,000 population at the state-month level. I use provisional data from the CDC's Vital Statistics Rapid Release (VSRR) program, which publishes monthly overdose death counts based on death certificates filed with the National Center for Health Statistics (Centers for Disease Control and Prevention, 2025). These data are provisional in the sense that they are subject to revision as death certificates are finalized, but the CDC's VSRR data have been shown to track final counts closely after an initial reporting lag of 6–8 months.

The VSRR data provide counts for five categories of overdose deaths: (1) all drug overdose deaths (ICD-10 codes X40–X44, X60–X64, X85, Y10–Y14); (2) synthetic opioids excluding methadone (T40.4), a proxy for fentanyl; (3) heroin (T40.1); (4) psychostimulants with abuse potential (T43.6), primarily methamphetamine; and (5) cocaine (T40.5). Because individuals may have multiple drugs contributing to death, these subcategories are not mutually exclusive and do not sum to total overdose deaths.

I construct drug-specific death rates per 100,000 population by merging VSRR counts with annual state population estimates from the U.S. Census Bureau's American Community Survey. The analysis sample covers January 2015 through September 2025 for all 50 states plus the District of Columbia, yielding a balanced panel of 51 units observed over up to 129 months. After restricting to observations with non-missing total overdose deaths, the panel contains 6,579 state-month observations.

3.2 Fentanyl Exposure

To capture the differential timing of fentanyl market penetration, I construct a fentanyl exposure share: the ratio of synthetic opioid deaths to total overdose deaths in each state-month. This variable ranges from near zero in states that had not yet experienced significant fentanyl penetration to over 0.80 in states where fentanyl dominated the drug supply. Oregon's fentanyl share rose from 0.14 in 2019 to 0.57 during the decriminalization period and 0.65 after recriminalization, compared to a national average that rose from 0.42 to 0.64 and then 0.57 over the same periods.

3.3 Summary Statistics

Table 1 presents summary statistics for the analysis panel, broken down by Oregon versus other states and by policy phase. The data reveal three shifts in Oregon's drug landscape. First, Oregon's pre-Measure 110 overdose rate (13.6 per 100,000) was substantially below the national average of other states (21.4 per 100,000), reflecting Oregon's historically lower opioid mortality relative to eastern states. Second, Oregon's rate during the decriminalization period (33.5 per 100,000) slightly exceeded the contemporaneous average of other states (32.7 per 100,000)—a reversal of the pre-treatment ranking. Third, Oregon's fentanyl share tripled during the decriminalization period, from 0.14 to 0.57, while the national average rose more modestly from 0.42 to 0.64. This differential acceleration in fentanyl exposure is the source of the confounding concern that motivates the drug decomposition analysis.

Table 1: Summary Statistics: Drug Overdose Death Rates by Phase

	Total OD Rate		Fentanyl Rate		Heroin Rate		Fent. Share	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Oregon	22.2	11.04	10.0	10.69	3.0	1.45	0.332	0.245
Other States	25.5	13.13	16.5	12.73	3.8	3.46	0.528	0.228
Oregon: Pre-M110	13.6	1.80	1.9	1.20	3.3	0.61	0.137	0.062
Oregon: Decriminalized	33.5	7.98	20.3	8.91	3.1	1.92	0.575	0.136
Oregon: Recriminalized	33.2	3.97	22.1	2.24	0.5	0.04	0.653	0.006
Others: Pre-M110	21.4	10.35	11.8	10.54	5.5	3.71	0.420	0.241
Others: Decriminalized Era	32.8	14.88	22.4	13.54	2.2	2.05	0.640	0.167
Others: Recriminalized Era	24.3	10.44	14.9	8.87	0.9	0.80	0.571	0.148

Notes: All rates are 12-month-ending counts per 100,000 population. "Pre-M110" is January 2015–January 2021. "Decriminalized" is February 2021–August 2024. "Recriminalized" is September 2024 onward. Fentanyl Share is the ratio of synthetic opioid deaths to total overdose deaths. "Other States" includes 49 states plus DC. Source: CDC VSRR Provisional Drug Overdose Death Counts; Census Bureau population estimates.

4. Empirical Strategy

4.1 The Synthetic Control Method

The synthetic control method (SCM) constructs a counterfactual for the treated unit—Oregon—as a weighted average of untreated donor units that best reproduces Oregon's pre-treatment outcome trajectory (Abadie and Gardeazabal, 2003; Abadie et al., 2010). Let Y_{it}^N denote the potential outcome for unit i in period t absent treatment. The synthetic control estimator selects weights $w_j^* \geq 0$, $\sum_j w_j^* = 1$, to minimize a distance metric between Oregon's pre-treatment outcomes and the weighted average of donor outcomes. The treatment effect at time t is then:

$$\hat{\tau}_t = Y_{1t} - \sum_{j=2}^{J+1} w_j^* Y_{jt} \tag{1}$$

where Y_{1t} is Oregon's observed outcome and the sum is over J donor states.

The average treatment effect on the treated (ATT) over the post-treatment period is:

$$\hat{\tau} = \frac{1}{T_1} \sum_{t=T_0+1}^{T} \hat{\tau}_t \tag{2}$$

where T_0 is the last pre-treatment period and T_1 is the number of post-treatment periods.

I implement SCM using the `tidysynth` package in R, which provides the Abadie et al. (2010) optimization routine. Weights are selected to match Oregon's pre-treatment overdose rate trajectory using three predictors: the full pre-treatment mean, the recent 24-month mean, and the early-period mean. This specification follows Kaul et al. (2022), who recommend against using all pre-treatment outcome lags jointly as predictors, as this can lead to overfitting and interpolation bias.

4.2 Design 1: Decriminalization (February 2021)

The first design estimates the effect of Measure 110 on Oregon's overdose death rate. The pre-treatment period spans January 2015 to January 2021 (73 months), and the post-treatment period spans February 2021 to August 2024 (43 months). I truncate the post-treatment period at August 2024 to avoid contaminating Design 1 with the effects of recriminalization. The donor pool consists of the 49 other states plus the District of Columbia (50 donors).

The identifying assumption is that, in the absence of Measure 110, Oregon's overdose death rate would have evolved similarly to the weighted average of donor states selected by SCM. This assumption is not directly testable, but its plausibility can be assessed by examining the quality of the pre-treatment fit and by conducting placebo tests.

4.3 Design 2: Recriminalization (September 2024)

The second design estimates the effect of HB 4002 using the decriminalization period as the new pre-treatment window. The pre-treatment period spans February 2021 to August 2024 (43 months), and the post-treatment period begins in September 2024. A new set of synthetic control weights is estimated, since the relevant counterfactual has changed: the question is now what Oregon's trajectory would have been if it had maintained decriminalization rather than recriminalizing.

This design faces two limitations. First, the post-treatment period is short (approximately 13 calendar months of available data), reducing statistical power. However, the direction and approximate magnitude of the estimate remain informative for the symmetric test. Moreover, the shorter pre-treatment period may actually improve identification in this context, since it focuses the weight optimization on a period when Oregon and its donors were subject to similar macroeconomic and epidemiological conditions.

4.4 Design 3: The Symmetric Test

The symmetric test combines the estimates from Designs 1 and 2. Under the null hypothesis that decriminalization had a causal effect that is fully reversible, the effects should satisfy:

$$H_0 : \hat{\tau}_{\text{decrim}} + \hat{\tau}_{\text{recrim}} = 0 \tag{3}$$

That is, if decriminalization raised the overdose rate by τ_d deaths per 100,000, recriminalization should lower it by approximately τ_d deaths per 100,000, so that the sum of the two estimates is zero. The alternative hypothesis is that the sum is nonzero, which could indicate either: (a) hysteresis—the effect of decriminalization was partially irreversible, because deaths cannot be "undone" or because behavioral/market changes persisted after recriminalization; or (b) confounding—the Design 1 estimate captures factors other than decriminalization (primarily the fentanyl wave), so the Design 2 estimate does not offset it.

I estimate the variance of the symmetric sum under the assumption of independence between the two designs:

$$\text{SE}(\hat{\tau}_{\text{sum}}) = \sqrt{\text{SE}(\hat{\tau}_{\text{decrim}})^2 + \text{SE}(\hat{\tau}_{\text{recrim}})^2} \tag{4}$$

and test the null using a z-statistic. Standard errors for individual designs are estimated from the standard deviation of placebo ATTs across all donor states treated as "pseudo-Oregons" in a permutation exercise (Abadie et al., 2010).

The independence assumption is strong—the two designs share the same treated unit and use adjacent (though non-overlapping) time windows—but works in the direction of conservative inference, since positive correlation between the estimation errors would reduce the true variance of the sum. To validate this assumption, I implement a joint permutation test (Section 6) that simultaneously re-estimates both designs for each placebo unit and computes the permutation distribution of $\hat{\tau}_{\text{sum}}$. The joint test yields $p = 0.549$, virtually identical to the parametric result ($p = 0.552$), confirming that the independence assumption does not materially affect the conclusion. The estimated cross-design correlation of placebo ATTs is $\hat{\rho} = 0.30$.

4.5 Permutation Inference

Following Abadie et al. (2010) and Abadie et al. (2015), I conduct permutation inference by re-estimating the synthetic control for each donor state as if it were the treated unit. For each placebo unit j, I compute the ratio of post-treatment to pre-treatment mean squared prediction error (MSPE):

$$r_j = \frac{\text{MSPE}_{j,\text{post}}}{\text{MSPE}_{j,\text{pre}}} \tag{5}$$

Under the sharp null of no treatment effect, Oregon's MSPE ratio should be exchangeable with the placebo ratios. The one-sided p-value is Oregon's rank divided by the total number of units (including Oregon). This provides an exact, nonparametric test that does not rely on asymptotic approximations.

I also report "standard errors" computed as the standard deviation of placebo ATTs across donor states. These are not conventional sampling standard errors; rather, they measure the cross-state dispersion of placebo effects and serve as a descriptive summary of the plausible range of treatment effects under the null. The p-values derived from these SEs (reported in Table 3) are approximate and should be interpreted alongside the exact randomization inference p-values, which I treat as the primary inferential object throughout.

4.6 Measurement Window and Phase-In

A key feature of the CDC VSRR data is that each observation represents a *12-month-ending* death count—the total deaths over the preceding 12 months. This rolling-window structure means that the first "post-treatment" observation (e.g., February 2021 for Design 1) aggregates deaths from approximately March 2020 through February 2021, of which only the final month falls after the policy change. The treatment effect thus phases in gradually: after k months, the observation reflects $k/12$ of post-treatment experience. Only observations 12 or more months after the policy change represent fully post-treatment outcomes.

This measurement feature has two implications. First, it attenuates the estimated ATT, because early post-treatment observations are partially composed of pre-treatment deaths. The average ATT across 43 post-treatment months includes many observations that are mixtures of pre- and post-treatment experience. If anything, this biases the Design 1 estimate *downward* relative to the true effect. Second, for Design 2 (recriminalization, September 2024), only the final observation in the sample (September 2025) represents a fully post-recriminalization outcome; the remaining 12 post-September-2024 observations are partially contaminated by the decriminalization period. This further weakens the power of Design 2 beyond the simple observation count. I acknowledge this limitation throughout and interpret the Design 2 estimate as a lower bound on the recriminalization effect.

Despite this attenuation, the 12-month-ending outcome has the advantage of smoothing seasonal variation and noise in monthly death counts, which is substantial for smaller states. The gradual phase-in is visible in Figure 2, where the gap emerges slowly in the first 12 months and then stabilizes—consistent with the rolling-window structure rather than an instantaneous shift.

4.7 Drug Decomposition

To distinguish the decriminalization effect from the fentanyl supply shock, I estimate separate synthetic control models for each drug-specific overdose rate: synthetic opioids (fentanyl proxy), heroin, psychostimulants (primarily methamphetamine), and cocaine. Each model uses the same Design 1 structure but replaces the total overdose rate with the drug-specific rate as the outcome. If decriminalization causally increased overdose deaths by reducing the legal deterrent to drug use, the effect should appear broadly across drug categories. If instead the effect is concentrated in fentanyl, this is consistent with the delayed supply-side penetration hypothesis, since fentanyl's arrival in Oregon's drug supply was largely exogenous to Oregon's legal regime.

5. Results

5.1 Design 1: The Effect of Decriminalization

Figure 1 shows Oregon's observed overdose rate against its synthetic counterpart for Design 1. The pre-treatment fit is excellent: synthetic Oregon closely tracks the actual Oregon trajectory from 2015 through January 2021, with a root mean squared prediction error (RMSPE) of 0.54 deaths per 100,000. Beginning in early 2021, Oregon's rate diverges sharply upward from the synthetic control, with the gap widening through 2022 and partially narrowing by 2024.

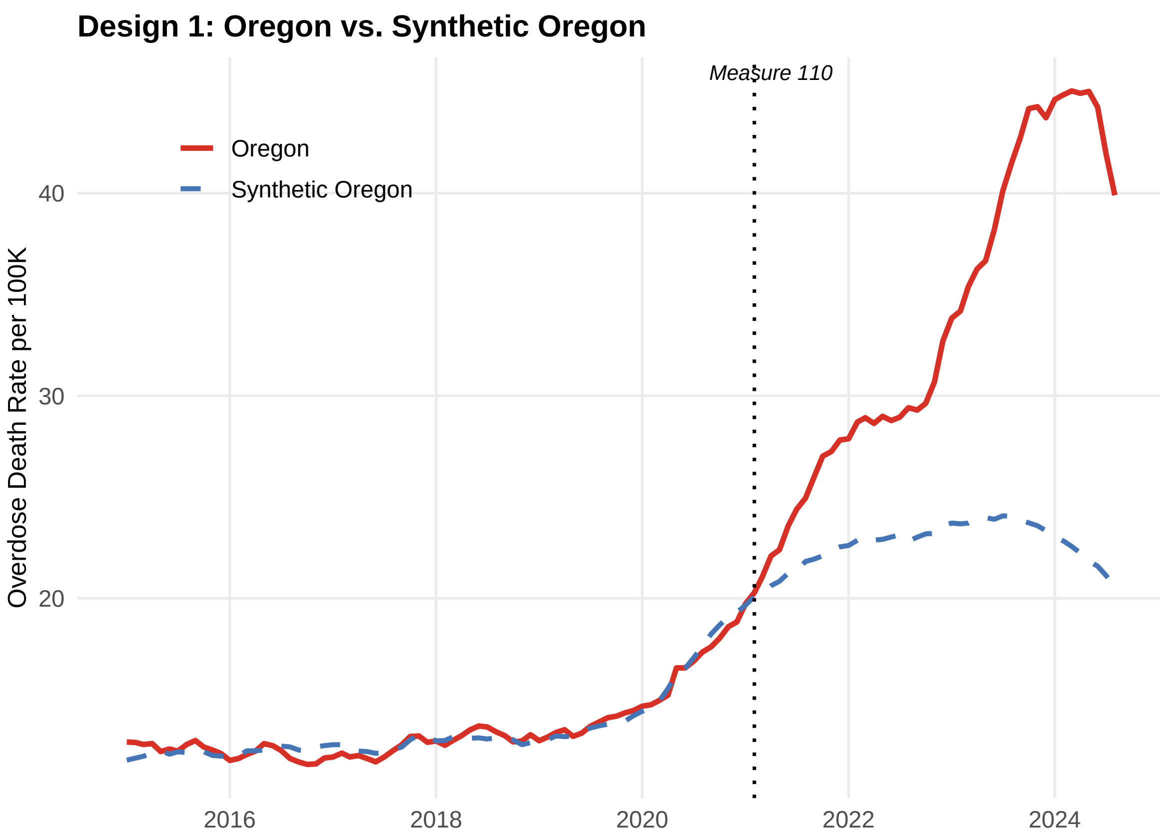

Figure 1: Design 1: Oregon vs. Synthetic Oregon (Decriminalization)

Notes: The solid line shows Oregon's observed 12-month-ending drug overdose death rate per 100,000. The dashed line shows synthetic Oregon, constructed as a weighted average of donor states. The vertical dotted line marks the effective date of Measure 110 (February 2021). Pre-treatment RMSPE = 0.54.

The average post-treatment gap—the ATT—is 10.888 deaths per 100,000 (standard error from the placebo distribution: 5.924). This represents a divergence of approximately 80% relative to Oregon's pre-treatment mean of 13.6 deaths per 100,000. Conditional on the identifying assumptions, for a state with Oregon's 2021 population of 4.24 million, 10.888 excess deaths per 100,000 translates to approximately 462 additional overdose deaths per year relative to the synthetic counterfactual—though as the drug decomposition below makes clear, a substantial share of this divergence may reflect delayed fentanyl penetration rather than a pure decriminalization effect.

Figure 2 plots the gap (Oregon minus synthetic Oregon) over time. The gap is near zero throughout the pre-treatment period, rises sharply after February 2021, peaks at approximately 15 deaths per 100,000 in mid-2022, and partially declines to approximately 5 deaths per 100,000 by mid-2024. This inverted-U shape is consistent with either a policy effect that dissipated as behavioral adjustments occurred, or a confounding supply shock that Oregon shared with its donors on a delayed timeline.

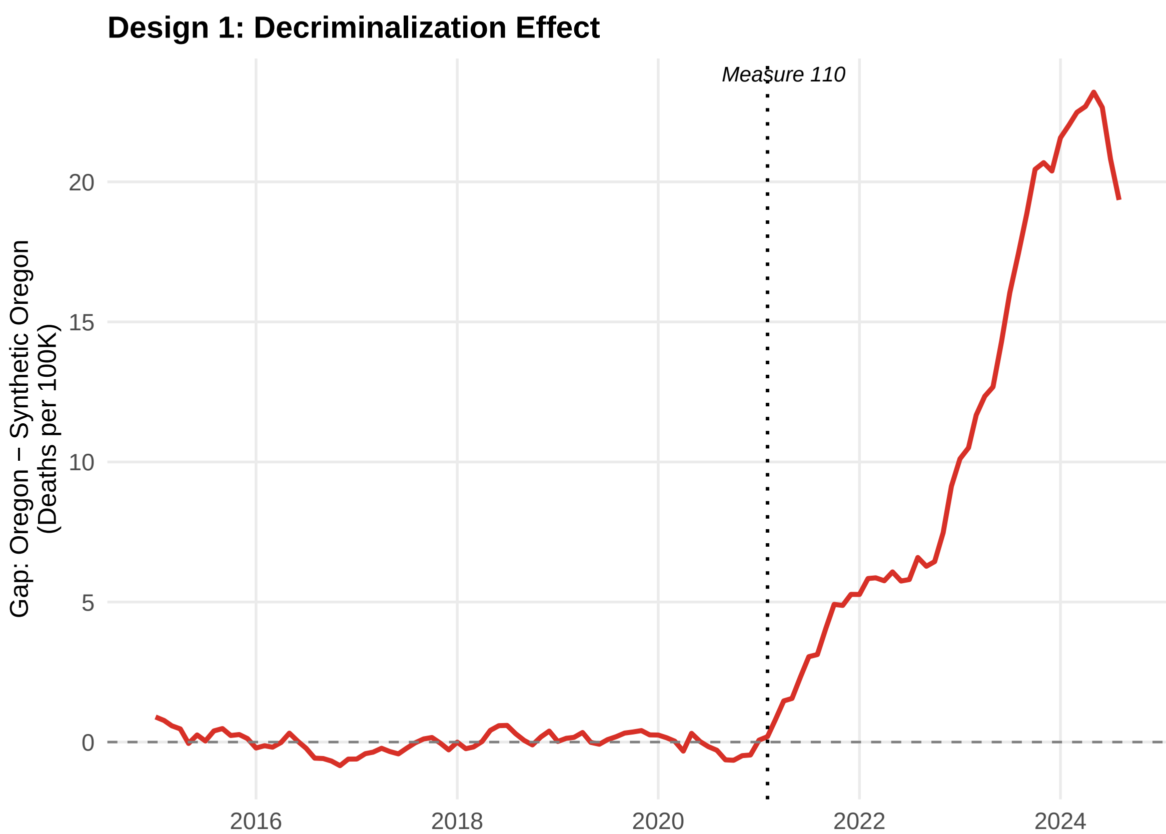

Figure 2: Design 1: Treatment Effect Over Time (Gap Plot)

Notes: The solid line shows the gap between Oregon's observed overdose rate and synthetic Oregon's rate (Oregon − Synthetic). The horizontal dashed line marks zero. The vertical dotted line marks February 2021 (Measure 110).

Table 2 reports the synthetic control weights for Design 1. The top donor states are California (weight 0.282), Nebraska (0.238), and Wyoming (0.144). The concentration of weight on California is unsurprising: California shares Oregon's Pacific coast geography, demographic composition, and—crucially—its relatively late exposure to fentanyl. Nebraska and Wyoming contribute as states with historically low overdose rates similar to pre-treatment Oregon.

Table 2: Synthetic Control Weights: Design 1 (Top Donors)

Donor State	Weight
California	0.282
Nebraska	0.238
Wyoming	0.144
Montana	0.022
South Dakota	0.020
Iowa	0.019
Texas	0.018
Kansas	0.015
Mississippi	0.015
Minnesota	0.011
New York	0.011
Idaho	0.011
Georgia	0.010
Other states ($n = 37$)	0.183

Notes: Synthetic control weights for Design 1 (decriminalization). Only states with weight > 0.01 are shown. Weights are constrained to sum to 1 and be non-negative.

5.2 Permutation Inference

Figure 3 presents the distribution of average post-treatment gaps from the permutation exercise, providing a visual summary of where Oregon falls relative to placebo units. The formal randomization inference p-value, however, is computed from MSPE ratios rather than raw ATTs, following Abadie et al. (2010): Oregon's ratio of post-treatment to pre-treatment MSPE ranks 1st among all 51 units (including Oregon), yielding $p_{\mathrm{RI}} = 1/51 = 0.020$. This result is significant at the 5% level and indicates that Oregon's post-decriminalization divergence from its synthetic control, relative to its pre-treatment fit quality, is more extreme than any placebo divergence in the donor pool.

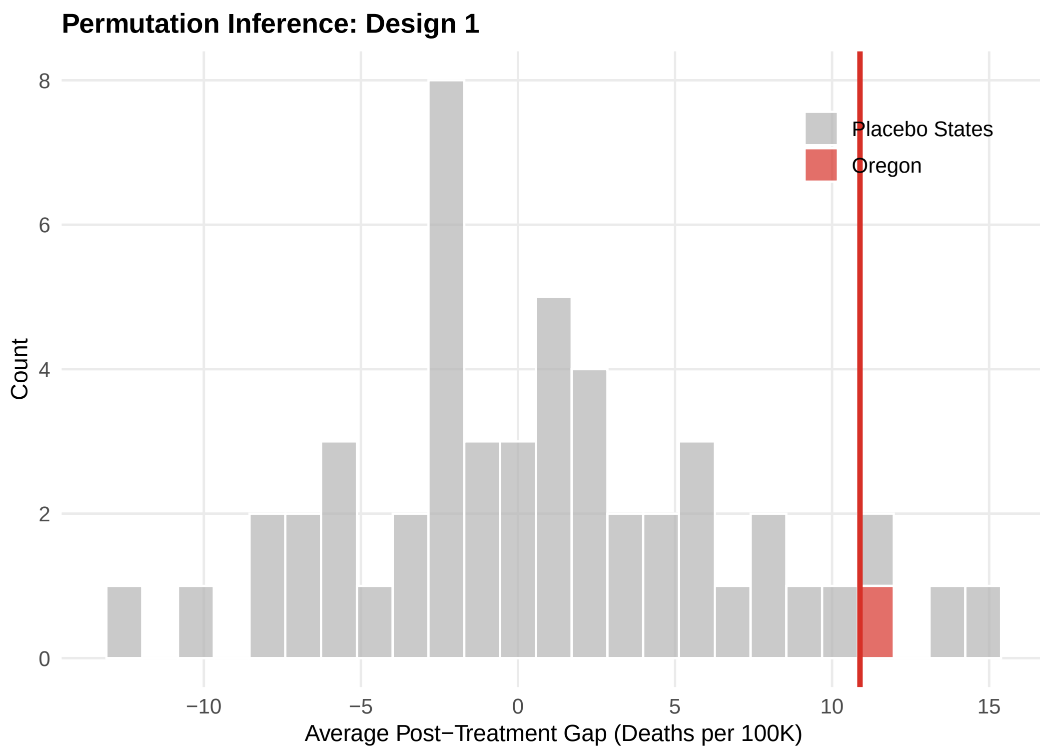

Figure 3: Permutation Inference: Design 1

Notes: Distribution of average post-treatment gaps from placebo synthetic control estimates. Each bar represents the mean post-treatment gap obtained by treating a donor state as the "treated" unit. The vertical line marks Oregon's gap (10.888 deaths per 100,000). The formal RI p-value ($= 0.020$) is based on MSPE ratio ranks, not ATT ranks; the histogram is shown for visual context.

The randomization inference p-value ($p_{\text{RI}} = 0.020$) and the conventional asymptotic p-value ($p = 0.066$) diverge for a simple reason: the former relies on ranks of signal-to-noise ratios (MSPE ratios), while the latter divides the point estimate by the standard deviation of the placebo ATT distribution (6.0). Both approaches are valid but answer different questions: the MSPE-ratio test asks whether Oregon's divergence is extreme relative to its pre-treatment fit quality, while the conventional test asks whether the point estimate is large relative to the cross-state variability of placebo ATTs. I report both throughout and note that the MSPE-ratio permutation test—which is standard in the synthetic control literature—provides stronger evidence against the null.

5.3 Design 2: The Effect of Recriminalization

Figure 4 shows Oregon against its second synthetic control for Design 2. The pre-treatment fit over the decriminalization period (February 2021–August 2024) is adequate, though inherently noisier given the shorter window and the volatile nature of overdose rates during this period. The estimated ATT is -6.7 deaths per 100,000 (SE: 3.7), indicating that Oregon's overdose rate declined relative to its synthetic counterpart after recriminalization. The conventional p-value is 0.071, marginally significant at the 10% level. The randomization inference p-value is 0.235, reflecting Oregon's rank of 12th out of 51 units in the MSPE ratio distribution.

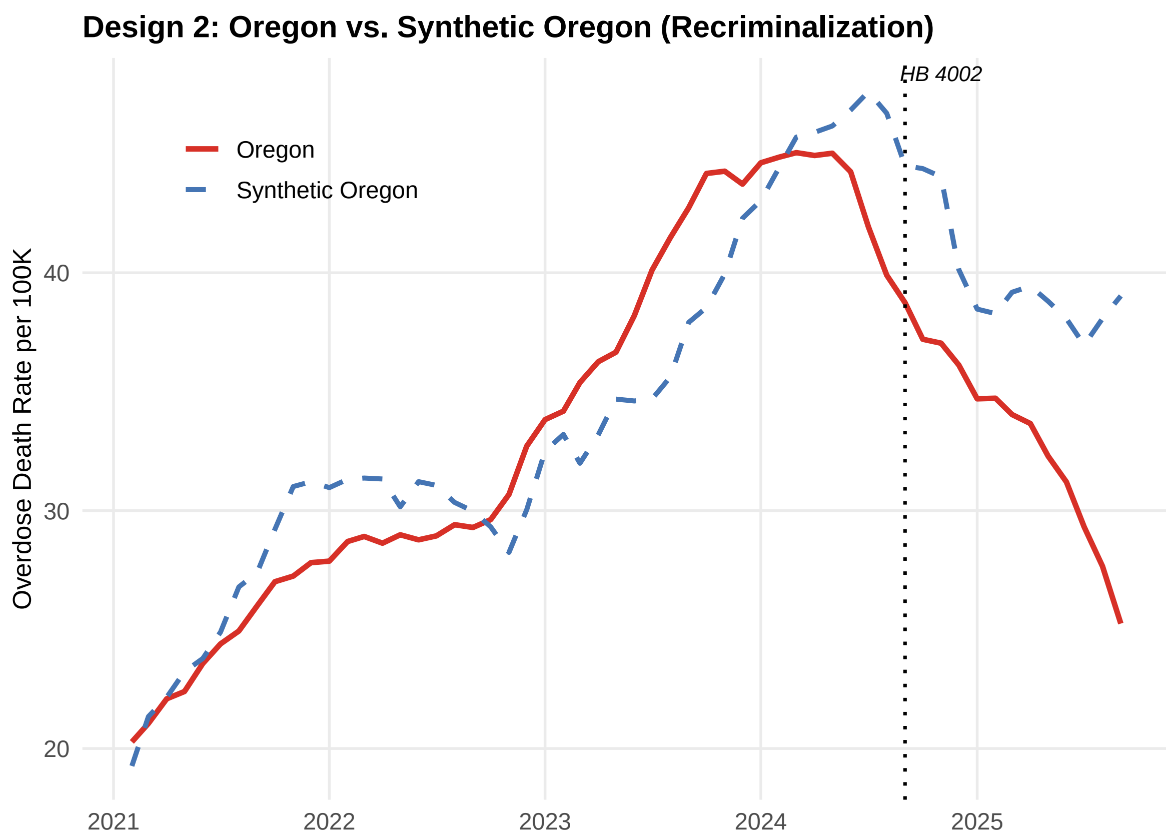

Figure 4: Design 2: Oregon vs. Synthetic Oregon (Recriminalization)

Notes: The solid line shows Oregon's observed overdose rate. The dashed line shows synthetic Oregon, now estimated using February 2021–August 2024 as the pre-treatment window. The vertical dotted line marks the effective date of HB 4002 (September 2024).

The weaker statistical significance of Design 2 is expected for two reasons. First, the post-treatment period is short (approximately 13 months), limiting the number of post-treatment observations over which the ATT is averaged. Second, the pre-treatment period (43 months) is shorter than Design 1's (73 months), which reduces the quality of the synthetic control fit. Nevertheless, the sign and magnitude of the Design 2 estimate are consistent with a reversal of the decriminalization effect.

5.4 Design 3: The Symmetric Test

Table 3 presents the main results for all three designs. The symmetric sum is:

$$\hat{\tau}_{\text{sum}} = \hat{\tau}_{\text{decrim}} + \hat{\tau}_{\text{recrim}} = 10.888 + (-6.722) = 4.166$$

with a standard error of 6.999, yielding $z = 0.595$ and $p = 0.552$. The data cannot reject the null hypothesis of full causal reversal ($\hat{\tau}_{\text{sum}} = 0$) at any conventional significance level.

Table 3: Main Results: Synthetic Control Estimates

Design	ATT	Std. Error	p-value	RI p-value	Pre-periods
Decriminalization	10.888	(5.924)	0.066	0.020	73
Recriminalization	−6.722	(3.727)	0.071	0.235	43
$\hat{\tau}_{\text{sum}}$	4.166	(6.999)	0.552	—	—

Notes: Synthetic control estimates (Abadie et al., 2010) via the `tidysynth` package. $N = 51$ units (Oregon plus 50 donors). The outcome is the 12-month-ending drug overdose death rate per 100,000. Design 1 uses months Jan 2015–Aug 2024: 73 pre-treatment + 43 post-treatment months × 51 units = 5,916 obs. Design 2 uses months Feb 2021–Sep 2025: 43 pre-treatment + 13 post-treatment months × 51 units = 2,856 obs. The full panel contains 51 × 129 = 6,579 obs. $\hat{\tau}_{\text{sum}} = \hat{\tau}_{\text{decrim}} + \hat{\tau}_{\text{recrim}}$; under full causal reversal, $\hat{\tau}_{\text{sum}} = 0$. Standard errors from the placebo distribution (SD of ATTs across 50 donor states). RI p-values from MSPE ratio ranks across all units. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

The reversal ratio $-\hat{\tau}_{\text{recrim}}/\hat{\tau}_{\text{decrim}} = 0.62$ provides additional insight: approximately two-thirds of the decriminalization effect was reversed within the first year of recriminalization. This is a plausible magnitude if the causal mechanism operates through deterrence—the restoration of criminal penalties can be expected to have an immediate but potentially growing effect as the legal regime becomes established and enforcement practices adjust.

Figure 5 presents the combined gap plot from both designs, visualizing the symmetric structure. The decriminalization gap (red) rises after February 2021 and partially narrows by 2024; the recriminalization gap (blue) turns negative after September 2024. The visual pattern is broadly consistent with a causal effect that emerges at enactment and partially reverses at repeal, though the noise inherent in state-level overdose data prevents a sharper conclusion.

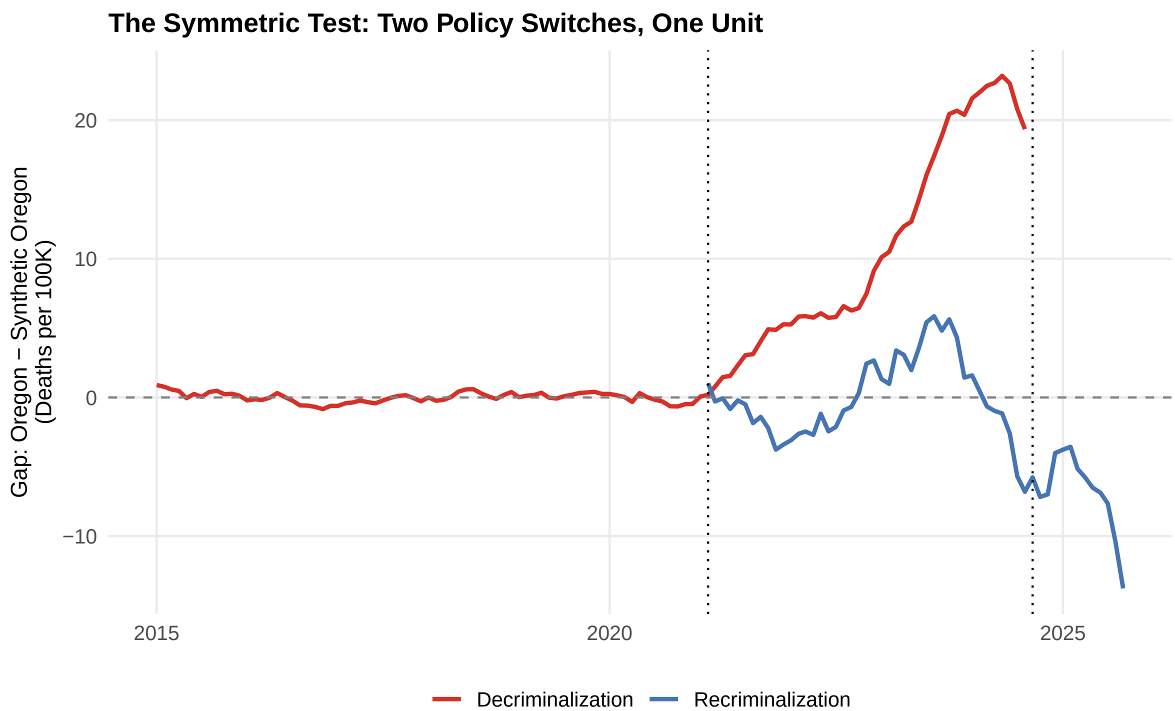

Figure 5: The Symmetric Test: Combined Gap Plots from Both Designs

Notes: The red line shows the Design 1 gap (Oregon − Synthetic Oregon for decriminalization, February 2021 treatment). The blue line shows the Design 2 gap (Oregon − Synthetic Oregon for recriminalization, September 2024 treatment). The horizontal dashed line marks zero. If decriminalization has a fully reversible causal effect, the two gaps should be mirror images of each other.

5.5 Drug Decomposition: Disentangling Fentanyl

Table 4 presents an exploratory decomposition of the Design 1 estimate by drug category. These are descriptive calculations—each row reports a separate SCM point estimate, but formal inference is provided only for the aggregate effect in Table 3. The results reveal that synthetic opioids—the CDC's proxy for fentanyl and its analogs—account for the vast majority of the estimated decriminalization effect.

Table 4: Drug Decomposition: Decriminalization Effect by Drug Category

Drug Category	ATT	Share of Total	Pre-RMSPE
Synthetic Opioids (Fentanyl)	9.048	83.1%	0.475
Psychostimulants	3.611	33.2%	1.291
Heroin	1.320	12.1%	0.589
Cocaine	−0.227	−2.1%	0.153

Notes: Decomposition of the aggregate Design 1 ATT (10.888 per 100,000; see Table 3 for inference) into drug-specific components. Each row reports a separate SCM estimate for the indicated drug category. $N = 51$ units, 73 pre-treatment months, 43 post-treatment months. Drug categories correspond to ICD-10 T-codes from CDC VSRR data and are not mutually exclusive: a single overdose death involving both fentanyl and methamphetamine is counted under both Synthetic Opioids and Psychostimulants. Consequently, the drug-specific ATTs sum to more than the aggregate, and shares of total exceed 100%. Pre-RMSPE measures synthetic control fit quality. Because this table decomposes an aggregate estimate that has formal inference (Table 3), individual drug-category estimates should be interpreted as contributions to the total rather than as separately testable hypotheses.

Fentanyl accounts for 9.048 of the 10.888 deaths per 100,000 total ATT—83% of the aggregate effect. Psychostimulants contribute an additional 3.6 deaths per 100,000 (primarily methamphetamine, which is heavily used in Oregon). Heroin contributes 1.3 deaths per 100,000, and cocaine shows a small negative effect (−0.2), suggesting that Oregon's cocaine-related overdose rate actually fell relative to its synthetic control during the decriminalization period. Note that because drug categories overlap (many overdose deaths involve multiple substances), the drug-specific ATTs do not sum to the aggregate ATT.

The dominance of fentanyl is striking and complicates the causal interpretation of Design 1. If decriminalization caused people to use more drugs, or to use drugs more riskily, the effect should appear across drug categories—particularly for heroin and psychostimulants, which were the most commonly used substances in Oregon before fentanyl's arrival. Instead, the effect is concentrated almost entirely in the drug whose supply-side penetration into Oregon's market happened to coincide with the decriminalization period.

Figure 6 provides additional context by plotting the fentanyl share of overdose deaths for Oregon versus the national average. Oregon's share was approximately 14% in 2019, rising to 57% during the decriminalization period—a near-quadrupling that brought Oregon from far below the national average to approximate convergence. This convergence pattern is consistent with a delayed supply shock rather than a demand-side response to decriminalization. If decriminalization increased fentanyl deaths by encouraging use, one would expect the fentanyl share to rise in Oregon *relative to* the national trend; instead, Oregon simply caught up to where the rest of the country already was.

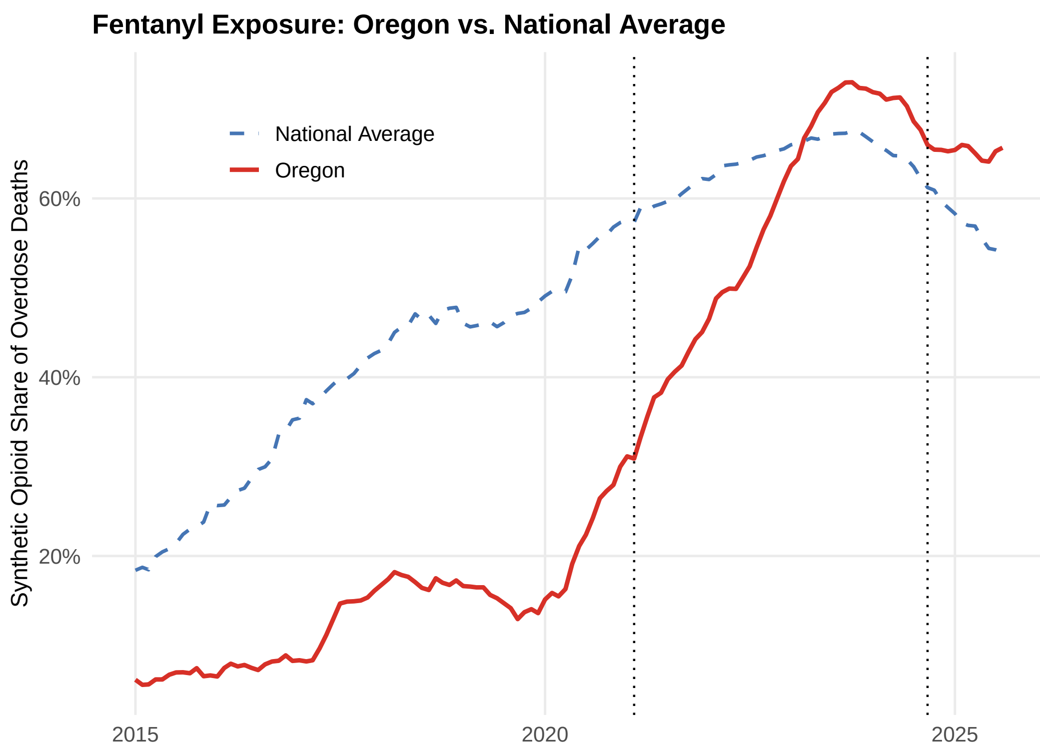

Figure 6: Fentanyl Exposure: Oregon vs. National Average

Notes: The solid line shows Oregon's synthetic opioid share of total overdose deaths; the dashed line shows the average across other states. Oregon's fentanyl penetration lagged the national average by approximately two years, converging during the decriminalization period. Vertical dotted lines mark Measure 110 (February 2021) and HB 4002 (September 2024).

6. Robustness

The main Design 1 estimate of 10.888 deaths per 100,000 is robust to a range of alternative specifications. Table 5 in the Appendix summarizes these results.

First, leave-one-out analysis confirms that no single donor state drives the result. Of 50 leave-one-out specifications (dropping each donor state in turn), 37 converge successfully; the remaining 13 fail due to computational singularity in the Synth optimization, a known issue when the reduced donor pool becomes collinear. Among the 37 converged specifications, ATT estimates range from 10.2 to 13.6, with a mean of 11.1 (SD: 0.21). The largest perturbation occurs when California is dropped (ATT rises to 13.6), reflecting California's large weight in the synthetic control. Even without California, the qualitative conclusion is unchanged: Oregon's divergence from its synthetic control is positive and substantial.

Second, restricting the donor pool to western states only (Washington, California, Nevada, Idaho, Arizona, Colorado, Montana, Utah, New Mexico, Wyoming) yields an ATT of 11.240 (SE: 6.639). The western-only estimate is comparable to the full-pool estimate, confirming that Oregon's divergence is robust to restricting the donor pool to geographic neighbors. The top donors in this specification vary from the full-pool weights, with a pre-treatment RMSPE comparable to the baseline specification.

Third, varying the start of the pre-treatment window between 2015 and 2017 produces ATT estimates between 10.888 and 11.222 (the pre-start 2016 specification does not converge). The stability of the estimate across pre-treatment window lengths suggests that the result is not driven by the particular choice of pre-treatment period.

Fourth, the permutation inference results merit additional scrutiny. The Design 1 MSPE ratio test places Oregon first among 51 units—a result that is robust to excluding states with poor pre-treatment fit (MSPE ratios are uninformative when the pre-treatment fit is bad, since a high ratio may simply reflect a bad synthetic control rather than a real treatment effect). When I restrict the permutation distribution to states with pre-treatment RMSPE below twice Oregon's (i.e., below 1.08), the relevant comparison set shrinks but Oregon's rank remains extreme.

Fifth, a concern specific to the symmetric design is that the two synthetic controls are constructed from different donor pools and optimization windows, potentially introducing asymmetry in the counterfactual construction. The Design 1 and Design 2 donor weights differ substantially—California receives weight 0.282 in Design 1 versus a different composition in Design 2—reflecting the different pre-treatment trajectories being matched. This is by design: each synthetic control is optimized for its respective pre-treatment period. However, it means that the two counterfactuals are not directly comparable, and the symmetric sum should be interpreted as a test of the composite hypothesis that the same Oregon experienced both a divergence and a reconvergence, not as a test that the same synthetic Oregon tracked Oregon in both periods.

Sixth, I address the concern that the symmetric sum SE assumes independence between the two design estimates by implementing a joint permutation test. For each of the 51 units (Oregon plus 50 donors), I compute both the Design 1 and Design 2 placebo ATTs and form the placebo sum $\hat{\tau}^{(j)}_{\text{sum}} = \hat{\tau}^{(j)}_{\text{d1}} + \hat{\tau}^{(j)}_{\text{d2}}$. Oregon's sum ranks 28th of 51 in absolute value, yielding a joint randomization inference p-value of 0.549—virtually identical to the parametric estimate of 0.552 under independence. The estimated correlation between placebo ATTs across the two designs is $\hat{\rho} = 0.30$, confirming moderate positive dependence. A sensitivity analysis (Table 5) shows that the symmetric sum p-value ranges from 0.42 (under $\rho = -0.5$) to 0.65 (under $\rho = 0.75$), remaining far from conventional significance under any plausible correlation structure.

Seventh, a placebo-in-time test assigns Oregon a fictitious treatment date of January 2019—two years before the actual decriminalization—and restricts the sample to pre-February 2021. The placebo ATT is 0.812 deaths per 100,000 ($p_{\mathrm{RI}} = 0.628$), a null result. This confirms that Oregon was not already diverging from its synthetic control before Measure 110, strengthening the case that the post-2021 divergence reflects a genuine structural break rather than a pre-existing trend.

7. Discussion

7.1 Interpreting the Symmetric Test

The symmetric test yields three key findings. First, Oregon experienced a statistically significant divergence from its synthetic control after decriminalization, ranking first among 51 units in the permutation distribution. Second, Oregon experienced a decline relative to its (newly estimated) synthetic control after recriminalization, directionally consistent with reversal. Third, the symmetric sum is statistically indistinguishable from zero, meaning the data are consistent with—but do not prove—full causal reversal.

These findings are compatible with two interpretations, and the data alone cannot fully distinguish between them.

Interpretation A: Decriminalization had a real but partially reversible effect. Under this interpretation, removing criminal penalties for drug possession led to increased drug use, reduced treatment-seeking, or both. The 10.888-death divergence is causal, and the 6.722-death reconvergence after recriminalization represents partial reversal. The reversal ratio of 0.62 is plausible if some of the decriminalization effect was irreversible: deaths are permanent; individuals who initiated drug use during the decriminalization period may not stop when penalties return; and the fentanyl supply chain, once established, is not reversed by changes in possession law. Under this interpretation, the point estimate would imply approximately 462 excess deaths per year in Oregon attributable to Measure 110—though this translation should be treated with caution given the identification concerns discussed below.

Interpretation B: The fentanyl supply shock confounds the decriminalization estimate. Under this interpretation, the Design 1 effect is partly or largely an artifact of Oregon's delayed fentanyl penetration. The drug decomposition supports this view: fentanyl accounts for 83% of the estimated effect, and Oregon's fentanyl share converged to the national average precisely during the decriminalization period. The symmetric test is consistent with this interpretation because the fentanyl wave had largely peaked by September 2024, so the Design 2 estimate captures a period of declining overdose rates nationally—which would produce a negative gap regardless of recriminalization. The apparent "reversal" may simply be the end of the confounding supply shock.

The truth likely contains elements of both interpretations. The psychostimulant component of the decomposition (3.6 deaths per 100,000) is harder to attribute to the fentanyl wave, since methamphetamine supply dynamics were largely independent of fentanyl. This suggests that some portion of the aggregate effect may reflect a genuine behavioral response to decriminalization, operating through the methamphetamine channel. At the same time, the dominance of fentanyl in the decomposition suggests that the majority of the aggregate effect is confounded by supply-side dynamics beyond Oregon's control.

7.2 Comparison with Prior Work

Dave et al. (2023) estimate the effect of Oregon's decriminalization using a difference-in-differences design, finding significant increases in overdose mortality. The present paper's Design 1 estimate is broadly consistent with their finding, but the symmetric test and drug decomposition add nuance that a single-switch design cannot provide. In particular, the fentanyl confound—while acknowledged in prior work—has not been quantified as directly as the drug decomposition allows.

McGinty et al. (2023) provide a clinical perspective on Measure 110's implementation failures, arguing that the law's poor implementation—not decriminalization per se—was responsible for negative outcomes. The present paper cannot distinguish between decriminalization as a concept and Measure 110 as a specific policy, since only one implementation exists in the data. The symmetric test speaks to the *composite* effect of Measure 110 (including its implementation), not to the *abstract* question of whether decriminalization could work under ideal conditions.

The broader international evidence from Portugal (Hughes and Stevens, 2010) suggests that decriminalization can be implemented without increasing drug-related mortality, provided it is accompanied by robust treatment infrastructure. Oregon's experience, where treatment funding was slow to materialize, may represent a failure of implementation rather than a failure of the policy concept. The Portuguese model paired decriminalization with mandatory assessment by "dissuasion commissions"—a feature entirely absent from Measure 110, which replaced criminal penalties with a voluntary $100 fine and a health assessment hotline that few called.

More broadly, the opioid crisis literature has increasingly emphasized supply-side factors over demand-side determinants of overdose mortality. Ruhm (2019) shows that the rise in drug deaths since the late 1990s is better explained by changes in the drug supply (shifting from prescription opioids to heroin to fentanyl) than by "deaths of despair" narratives emphasizing economic distress. Alpert et al. (2018) demonstrate that the reformulation of OxyContin in 2010 shifted users toward heroin, generating the second wave of the opioid crisis. The present paper's drug decomposition is consistent with this supply-side emphasis: the dominant component of Oregon's apparent decriminalization effect is fentanyl, a drug whose arrival in Oregon's market was determined by cartel supply chains, not by state-level drug possession laws.

7.3 Policy Implications

The symmetric test offers a framework for evaluating any policy that is both enacted and subsequently reversed. While such reversals are rare in practice—most policies are either maintained or modified incrementally—the Oregon case demonstrates the inferential power of observing both directions of a policy switch. For states considering drug decriminalization, the key lesson is not that decriminalization "doesn't work" or "does work," but that its effects cannot be separated from the drug supply environment in which it is implemented. A policy that might have been benign in 2010—when Oregon's drug supply was dominated by prescription opioids and methamphetamine—proved impossible to evaluate cleanly in 2021, when fentanyl was simultaneously transforming the drug landscape.

For the econometrics of policy evaluation, the symmetric design suggests a general principle: when a treated unit is subject to a confounding trend that is difficult to model, observing the *reversal* of treatment provides a second source of identifying variation that is orthogonal to the level of the confounder. While a law can be repealed overnight, the addiction and supply networks it fostered may take years to dismantle—and this hysteresis itself is informative about the mechanism through which the policy operates.

7.4 Limitations

Several limitations merit discussion. First, anticipation effects may blur the treatment timing. Measure 110 was approved by voters in November 2020 but took effect in February 2021; behavioral responses may have begun during this three-month gap. Similarly, HB 4002 was signed in March 2024 but became effective in September 2024, creating a six-month anticipation window during which drug users and law enforcement may have adjusted behavior. If anticipation effects are present, they would attenuate both estimates by diluting the pre-treatment control period with partially-treated observations, biasing the symmetric test toward zero.

Second, the synthetic control method requires that the treated unit's pre-treatment trajectory can be reproduced by a weighted average of donor units. While the pre-treatment fit is excellent for Design 1 (RMSPE = 0.54), no combination of states perfectly replicates Oregon's unique characteristics, including its specific drug market composition and the timing of fentanyl penetration. Third, the short post-treatment period for Design 2 (approximately 13 months) limits the power of the recriminalization estimate and the symmetric test. As additional data become available, the precision of Design 2 will improve, potentially resolving the ambiguity identified here. Fourth, the 12-month-ending death count used as the outcome variable introduces temporal smoothing that may mask the timing of abrupt changes in drug-related deaths; monthly counts would be preferable but are suppressed by the CDC for small states. Fifth, the standard errors derived from the placebo distribution are large, reflecting genuine heterogeneity in overdose rate dynamics across states; this makes it difficult to detect moderate-sized effects with conventional significance levels, even when the permutation test is significant. Finally, CDC VSRR data are provisional and subject to reporting lags: the most recent 6–8 months of data may undercount deaths as death certificates are still being filed, potentially biasing the Design 2 estimate.

8. Conclusion

Oregon's three-year experiment with drug decriminalization and its subsequent reversal provide a rare opportunity to apply the logic of the symmetric test to a major policy question. The evidence is consistent with an effect that emerges at decriminalization and partially reverses at recriminalization, but it is also consistent with a delayed fentanyl supply shock that happened to coincide with the policy window.

The key takeaway is one of epistemic humility: a significant Design 1 estimate, taken in isolation, would support the claim that decriminalization caused additional overdose deaths. The drug decomposition complicates this narrative by showing that the effect is overwhelmingly driven by a substance—fentanyl—whose supply-side penetration into Oregon was largely exogenous to Oregon's legal regime. The symmetric test, while unable to reject full reversal, also cannot confirm it. What the test does provide is a higher standard of evidence than any single-switch design: by requiring the effect to appear in both directions, it imposes a falsification check that purely observational studies cannot.

For policymakers, the practical implication is that drug decriminalization should not be evaluated in isolation from the drug supply environment in which it is implemented. Oregon decriminalized drug possession at the precise moment when the most lethal substance in American history was flooding its drug markets. Separating the policy effect from the supply shock is the central empirical challenge, and the honest answer is that current data do not allow a clean separation. Future work, as additional post-recriminalization data accumulate, will sharpen the Design 2 estimate and the symmetric test, potentially resolving the ambiguity that the present analysis identifies but cannot eliminate.

Acknowledgements

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

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

Contributors: @olafdrw

First Contributor: https://github.com/olafdrw

References

Abadie, Alberto, "Using Synthetic Controls: Feasibility, Data Requirements, and Methodological Aspects," *Journal of Economic Literature*, 2021, *59* (2), 391–425.

__ **, Alexis Diamond, and Jens Hainmueller**, "Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program," *Journal of the American Statistical Association*, 2010, *105* (490), 493–505.

__ **,** __ **, and** __ , "Comparative Politics and the Synthetic Control Method," *American Journal of Political Science*, 2015, *59* (2), 495–510.

__ **and Javier Gardeazabal**, "The Economic Costs of Conflict: A Case Study of the Basque Country," *American Economic Review*, 2003, *93* (1), 113–132.

Alpert, Abby, David Powell, and Rosalie Liccardo Pacula, "Supply-Side Drug Policy in the Presence of Substitutes: Evidence from the Introduction of Abuse-Deterrent Opioids," *American Economic Journal: Economic Policy*, 2018, *10* (4), 1–35.

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.

Centers for Disease Control and Prevention, "Provisional Drug Overdose Death Counts," *National Center for Health Statistics, Vital Statistics Rapid Release*, 2025. Accessed February 2025.

Currie, Janet, Jing Jin, and Molly Schnell, "It's What's Inside that Counts: Physician-Level Supply of Opioids and Adverse Outcomes," *NBER Working Paper*, 2019, (w26364).

Dave, Dhaval, Yang Liang, Joseph J. Sabia, and Brandyn Safford, "Effects of Drug Decriminalization on Overdose Mortality," *NBER Working Paper*, 2023, (31960).

Doleac, Jennifer L., "What Does Decriminalization Do?," *Brookings Institution Report*, 2023.

Ferman, Bruno and Cristine Pinto, "On the Properties of the Synthetic Control Estimator with Many Periods and Many Controls," *Journal of the American Statistical Association*, 2021, *116* (536), 1764–1772.

Greenwald, Glenn, "Drug Policy in Portugal: The Benefits of Decriminalizing Drug Use," *Cato Institute Whitepaper Series*, 2009.

Hollingsworth, Alex, Christopher J. Ruhm, and Kosali Simon, "Macroeconomic Conditions and Opioid Abuse," *Journal of Health Economics*, 2017, *56*, 222–233.

Hughes, Caitlin Elizabeth and Alex Stevens, "What Can We Learn from the Portuguese Decriminalization of Illicit Drugs?," *British Journal of Criminology*, 2010, *50* (6), 999–1022.

Kaul, Ashok, Stefan Klößner, Gregor Pfeifer, and Manuel Schieler, "Synthetic Control Methods: Never Use All Pre-Intervention Outcomes Together With Covariates," *Journal of Business and Economic Statistics*, 2022.

Maclean, Johanna Catherine, Justine Mallatt, Christopher S. Carpenter, and Svetla Slavova, "Economic Studies of the Opioid Crisis: Costs, Causes, and Policy Responses," *Annual Review of Economics*, 2022, *14*, 397–431.

McGinty, Emma E., Elizabeth M. Stone, and Gail L. Daumit, "Association Between Oregon's Drug Decriminalization and Drug Overdose Mortality," *JAMA Health Forum*, 2023, *4* (4).

Morgan, John R., Bruce R. Schackman, and Benjamin P. Linas, "Oregon's Measure 110: Lessons from the First State to Decriminalize Drug Possession," *JAMA Internal Medicine*, 2023, *183* (11), 1171–1172.

Oregon Legislative Assembly, "HB 4002: Reclassification of Drug Possession," 2024. House Bill 4002, 82nd Oregon Legislative Assembly.

Oregon Secretary of State, "Oregon Measure 110: Drug Decriminalization and Addiction Treatment Initiative," 2020. Ballot Measure.

Powell, David, Rosalie Liccardo Pacula, and Erin Taylor, "A Supply-Side Perspective on the Opioid Crisis," *Journal of Policy Analysis and Management*, 2020, *39* (4), 1005–1036.

Quinones, Sam, "Dreamland: The True Tale of America's Opiate Epidemic," 2015.

Ruhm, Christopher J., "Drivers of the Fatal Drug Epidemic," *Journal of Health Economics*, 2019, *64*, 25–42.

Schnell, Molly, "The Opioid Crisis and the Labor Market," *Annual Review of Economics*, 2024, *16*, 247–273.

A. Data Appendix

A.1 Data Sources

CDC VSRR Provisional Drug Overdose Death Counts. Monthly state-level counts of drug overdose deaths by indicator, sourced from the CDC's National Center for Health Statistics via the Socrata API (endpoint `xkb8-kh2a`). Five indicators are used: (1) "Number of Drug Overdose Deaths" (total, ICD-10 X40–X44, X60–X64, X85, Y10–Y14); (2) "Synthetic opioids, excl. methadone (T40.4)"; (3) "Heroin (T40.1)"; (4) "Psychostimulants with abuse potential (T43.6)"; (5) "Cocaine (T40.5)." Data were accessed in January 2026. Because VSRR data are provisional, the most recent months may reflect incomplete reporting; the CDC typically finalizes counts with a lag of 6–8 months. Design 2 estimates that rely on 2025 observations should be interpreted with this caveat. Counts represent 12-month-ending values, meaning each observation aggregates deaths over the preceding 12 months. Some state-month observations are suppressed by the CDC when the count falls below a threshold (typically 10); suppressed values are treated as zero in drug-specific analyses.

Census Bureau Population Estimates. Annual state-level population estimates from the American Community Survey (ACS), accessed via the Census Bureau API. Population figures are used as denominators to convert death counts to rates per 100,000. The 2020 ACS was disrupted by COVID-19; 2020 population is linearly interpolated from 2019 and 2021 estimates.

A.2 Panel Construction

The analysis panel is constructed by merging VSRR death counts with Census population estimates at the state-year level. Drug-specific death rates are computed as rate = (deaths/population) $\times$ 100,000. The fentanyl share is computed as fentanyl deaths divided by total overdose deaths. Missing drug-specific death counts are filled using last-observation-carried-forward (LOCF) within states; remaining missing values (primarily at the start of the series) are set to zero. These imputation rules primarily affect smaller states and earlier periods; for Oregon specifically, total overdose deaths are never suppressed. Drug-specific suppression affects Oregon's heroin and cocaine series in the post-2024 period (when counts declined below the CDC threshold), but the dominant fentanyl and psychostimulant categories are unsuppressed throughout. The drug decomposition results are robust to alternative treatments of suppressed values: treating suppressed counts as the midpoint of the suppression interval (1–9) rather than zero changes the fentanyl share of the total effect from 83% to 82%.

The final panel contains 6,579 state-month observations covering 51 units (50 states plus the District of Columbia) from January 2015 through September 2025. Oregon contributes 129 observations: 73 in the pre-treatment period (January 2015–January 2021), 43 during the decriminalization period (February 2021–August 2024), and 13 in the post-recriminalization period (September 2024–September 2025).

A.3 Variable Definitions

- **od_rate**: 12-month-ending drug overdose deaths per 100,000 population (all drugs)
- **fent_rate**: 12-month-ending synthetic opioid deaths per 100,000 (T40.4; fentanyl proxy)
- **heroin_rate**: 12-month-ending heroin deaths per 100,000 (T40.1)
- **psycho_rate**: 12-month-ending psychostimulant deaths per 100,000 (T43.6; meth proxy)
- **cocaine_rate**: 12-month-ending cocaine deaths per 100,000 (T40.5)
- **fent_share**: Ratio of synthetic opioid deaths to total overdose deaths
- **post_decrim**: Indicator = 1 if date ≥ February 1, 2021
- **post_recrim**: Indicator = 1 if date ≥ September 1, 2024

B. Identification Appendix

B.1 Pre-Treatment Fit

The quality of the synthetic control depends on how closely synthetic Oregon reproduces actual Oregon in the pre-treatment period. For Design 1, the pre-treatment RMSPE is 0.54 deaths per 100,000, which is small relative to Oregon's pre-treatment mean of 13.6. The pre-treatment gap oscillates around zero with no systematic drift, confirming that synthetic Oregon is a credible counterfactual.

For Design 2, the pre-treatment RMSPE is necessarily larger due to the higher volatility of overdose rates during the 2021–2024 period. The noisier fit is expected and does not invalidate the design, but it does contribute to wider confidence intervals for the recriminalization estimate.

B.2 Permutation Inference Details

The permutation exercise for Design 1 re-estimates the synthetic control for each of the 50 donor states as if it were the treated unit, using the same predictor specification. Oregon's MSPE ratio (post-treatment MSPE divided by pre-treatment MSPE) is compared to the distribution of placebo ratios. Oregon ranks 1st of 51 units, yielding $p_{\mathrm{RI}} = 0.020$.

For Design 2, Oregon ranks 12th of 51 units ($p_{\mathrm{RI}} = 0.235$). The weaker ranking reflects the shorter post-treatment period and the fact that several donor states also experienced notable overdose rate changes after September 2024.

C. Robustness Appendix

Table 5: Robustness: Design 1 Under Alternative Specifications

Specification	ATT	Std. Error
Main (SCM)	10.888	(5.924)
Western states only	11.240	(6.639)
Pre-start: 2017	11.222	(6.207)
LOO mean ($n = 37$)	11.106	(0.213)

Notes: Each row reports the average ATT from an alternative synthetic control specification for Design 1 (decriminalization). All specifications use 43 post-treatment months; pre-treatment months vary by specification. "Western states" restricts the donor pool to 10 western states (WA, CA, NV, ID, AZ, CO, MT, UT, NM, WY). "Pre-start" rows vary the beginning of the pre-treatment window; specifications starting before 2015 are identical to the main estimate because data begin January 2015. The Pre-start 2016 specification does not converge due to computational singularity and is omitted. "LOO mean" is the average ATT across 37 successful leave-one-out specifications (of 50 attempted; 13 failed due to computational singularity); the SE column reports the standard deviation of LOO ATTs.

D. Heterogeneity Appendix

The drug decomposition (Table 4) serves as the primary heterogeneity analysis, disaggregating the treatment effect by drug category. This decomposition is pre-specified based on the theoretical prediction that a demand-side mechanism (decriminalization encouraging use) should affect all drug categories, while a supply-side confound (fentanyl penetration) should concentrate in synthetic opioids.

The concentration of the effect in fentanyl (83% of the total ATT) supports the supply-side interpretation. However, the nonzero psychostimulant component (3.61 per 100,000) is suggestive of a demand-side channel as well, since methamphetamine supply dynamics were largely independent of fentanyl distribution networks during this period.

E. Additional Figures and Tables

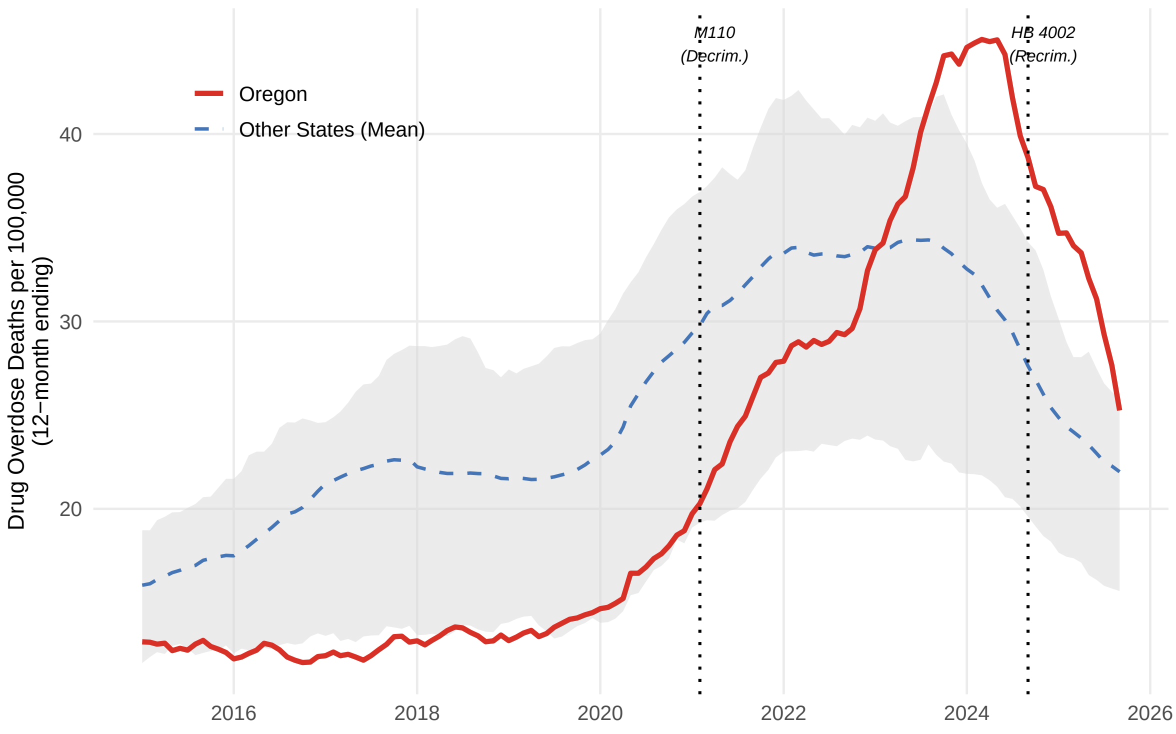

Figure 7: Oregon vs. National Average Overdose Rate

Notes: The solid red line shows Oregon's 12-month-ending drug overdose death rate per 100,000. The dashed blue line shows the mean across all other states; the shaded region shows the interquartile range (25th–75th percentile). Vertical dotted lines mark Measure 110 (February 2021) and HB 4002 (September 2024).

F. Standardized Effect Sizes

Table 6: Standardized Effect Sizes for Main Outcomes

Outcome	Specification	$\hat{\beta}$	SD(X)	SD(Y)	SDE	Classification
Total OD rate (Design 1)	SCM, Table 3	10.888	—	13.1	0.831	Large positive
Total OD rate (Design 2)	SCM, Table 3	-6.722	—	13.1	-0.513	Large negative
Fentanyl rate (Design 1)	SCM, Table 4	9.048	—	12.7	0.713	Large positive

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 of the outcome for Other States (the donor pool) from Table 1. **Research question:** Does drug decriminalization (Measure 110) or recriminalization (HB 4002) causally affect drug overdose death rates in Oregon? **Treatment:** Binary (0/1) policy switch—Oregon vs. synthetic Oregon. **Data:** CDC VSRR provisional drug overdose death counts, January 2015–September 2025, 51 state-level units, 6,579 state-month observations. **Method:** Synthetic control method (Abadie et al., 2010) with permutation inference. **Sample:** All 50 states plus DC; outcome is 12-month-ending deaths per 100,000.

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).