# The Substitution That Wasn't: State Kratom Bans and Opioid Overdose Mortality

APEP Autonomous Research* @ai1scl

March 31, 2026

**Abstract**

An estimated 10–15 million Americans use kratom, a plant-based opioid agonist, for pain and withdrawal management. Five states banned kratom between 2014 and 2017, criminalizing a lower-risk alternative during the fentanyl crisis. If users substitute toward illicit opioids, mortality should rise—yet no study has tested this. Using CDC provisional overdose data (2015–2025) and staggered difference-in-differences, I find no detectable effect: the Callaway–Sant'Anna ATT is 0.043 log points (SE = 0.042), ruling out increases above 13%. A drug-type decomposition shows that TWFE estimates reflecting lower mortality in ban states are driven by differential trends, not the ban: negative controls (psychostimulants, cocaine) exhibit the same pattern. The data are consistent with no substitution-driven mortality effect, though the design cannot rule out effects below state-level detection thresholds.

**JEL Codes:** I12, I18, K32

**Keywords:** kratom, opioid crisis, drug policy, substitution, overdose mortality, harm reduction

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 19m).

## 1. Introduction

Between 2015 and 2023, more than 600,000 Americans died from drug overdoses, the majority involving opioids (Hedegaard et al., 2024). In the same period, five state legislatures responded to a different perceived threat—kratom, a Southeast Asian plant whose active compounds bind opioid receptors—by classifying it alongside heroin and fentanyl as a Schedule I substance. The policy question is immediate: when you criminalize a legal, lower-risk opioid alternative during the deadliest drug epidemic in American history, do people die?

The substitution hypothesis says yes. Kratom (Mitragyna speciosa) acts as a partial agonist at mu-opioid receptors, producing analgesic and anxiolytic effects at lower doses with a ceiling on respiratory depression that makes fatal overdose far less likely than with full agonists (Kruegel and Grundmann, 2016; Prozialeck et al., 2012; Henningfield et al., 2019). Survey evidence suggests that a substantial fraction of kratom users report using it specifically to manage chronic pain or reduce dependence on prescription opioids (Grundmann, 2017; Smith and Lawson, 2021; Garcia-Romeu et al., 2020). If a ban eliminates this exit ramp, standard economic models of addiction predict substitution toward the nearest available alternative—in many cases, illicit fentanyl (Becker and Murphy, 1988).

Yet no published study has tested the substitution hypothesis causally. The existing kratom literature is overwhelmingly pharmacological or based on cross-sectional surveys (Swogger et al., 2015; Boyer et al., 2008; Babin, 2018). The closest empirical analogy comes from the cannabis–opioid literature, where Bradford and Bradford (2016); Powell et al. (2018); McMichael et al. (2020) find that medical marijuana access reduces opioid prescriptions and overdose deaths—but kratom, as a direct opioid receptor agonist rather than a mechanistically distinct analgesic, offers a sharper test of within-class substitution.

This paper provides the first causal evidence on state kratom bans and opioid overdose mortality. Five states—Wisconsin and Indiana (2014), Arkansas (2015), Alabama (2016), and Rhode Island (2017)—classified mitragynine and 7-hydroxymitragynine as controlled substances in staggered fashion. I exploit this staggered timing in a difference-in-differences framework using CDC Vital Statistics Rapid Release (VSRR) provisional drug overdose death data at the state-month level from 2015 to 2025, covering 46 states and over 5,000 state-month observations.

The headline finding is an informative null. The Callaway–Sant'Anna estimator, which accounts for treatment effect heterogeneity across the three cohorts with clean pre-treatment data (Arkansas, Alabama, Rhode Island), produces an average treatment effect on the treated of 0.043 log points (SE = 0.042)—statistically indistinguishable from zero. The 95% confidence interval ($-0.039$ to 0.126) rules out increases larger than 13% ($e^{0.126} - 1$), though

it cannot exclude more modest effects. Wild cluster bootstrap inference ($p = 0.464$) and randomization inference ($p = 0.169$) confirm that the null is not an artifact of few treated clusters. Given the design—three treated clusters with 9 to 29 pre-treatment months—the minimum detectable effect at 80% power is approximately 0.15–0.20 log points (16–22%), placing this study squarely in the range needed to detect substitution effects comparable to those found in the cannabis–opioid literature.

The drug-type decomposition provides the mechanism test. If banned kratom users substitute to illicit fentanyl, synthetic opioid deaths should rise in ban states relative to controls, while psychostimulant deaths—kratom is not a stimulant—should remain unchanged. Instead, the TWFE specification shows negative point estimates across *all* drug categories, including the negative controls (psychostimulants: $-0.186$, cocaine: $-0.298$). This symmetric pattern indicates that the ban states were on different overall drug mortality trajectories—not that kratom bans caused any drug-type-specific shift. A triple-difference specification (opioids versus psychostimulants within ban states) yields an interaction coefficient of $-0.117$ ($p = 0.084$), weakly negative—the opposite of what substitution would predict.

Three candidate explanations account for the null. First, enforcement may be weak: kratom remained available online and through informal channels even in ban states, limiting the effective reduction in access (Henningfield et al., 2022). Second, the kratom-using population may be too small—estimated at 1–2% of the adult population (Palamar, 2021)—to generate detectable mortality effects in state-level aggregate data, even if individual-level substitution occurs. Third, kratom users may substitute to other legal alternatives (CBD, kava, over-the-counter analgesics) rather than to illicit opioids, breaking the within-opioid-class substitution chain that the hypothesis requires.

This paper contributes to three literatures. First, it adds to the growing evidence on substance substitution in the opioid crisis, joining work on PDMPs (Buchmueller and Carey, 2018; Dave et al., 2021), naloxone access (Doleac and Mukherjee, 2022; Rees et al., 2019), and cannabis legalization (Bradford and Bradford, 2016; Powell et al., 2018). Second, it speaks to the broader literature on prohibition and unintended consequences (Miron, 2004; Cunningham and Shah, 2018), providing a case where prohibition appears to have neither the feared harm nor the desired benefit. Third, it informs active regulatory debates: the DEA has repeatedly considered federal scheduling of kratom, and nine states have instead passed Kratom Consumer Protection Acts regulating quality rather than banning the substance (Grundmann et al., 2023; Prozialeck, 2021). The finding that state-level bans had no detectable mortality effect suggests that the policy debate should focus on product safety regulation rather than prohibition.

## 2. Institutional Background

**Kratom and the opioid receptor system.** Kratom leaves contain over 40 alkaloids, of which mitragynine and 7-hydroxymitragynine are the primary psychoactive compounds (Kruegel and Grundmann, 2016). Both bind mu-opioid receptors, but with partial agonist activity and a ceiling effect on respiratory depression that distinguishes them from full agonists like morphine, heroin, and fentanyl (Prozialeck et al., 2012). At low doses (1–5 grams of leaf powder), kratom produces stimulant-like effects; at higher doses (5–15 grams), it produces analgesia and sedation qualitatively similar to low-dose opioids (Boyer et al., 2008). The pharmacological profile makes fatal overdose from kratom alone extremely rare, though co-ingestion with other depressants can be dangerous (Olsen et al., 2020).

**Kratom use in the United States.** Kratom is legal at the federal level and widely available in gas stations, smoke shops, and online retailers. The 2019 National Survey on Drug Use and Health estimated 0.7% past-year prevalence among adults 18+, or roughly 1.7 million users (Palamar, 2021). The American Kratom Association estimates 10–16 million regular users, though this figure may be inflated by advocacy interests. Survey studies consistently find that the majority of users report using kratom for pain management (68%) or to reduce opioid dependence (35%) (Grundmann, 2017; Smith and Lawson, 2021).

**State kratom bans.** Five states classified kratom's active alkaloids as Schedule I controlled substances through distinct legislative mechanisms. Wisconsin (April 2014) and Indiana (July 2014) acted through synthetic drug analog statutes. Arkansas (October 2015) used the Board of Pharmacy's scheduling authority. Alabama (May 2016) and Rhode Island (June 2017) enacted standalone controlled substance classifications. In each case, the ban criminalized the manufacture, distribution, and possession of kratom, with penalties ranging from misdemeanor (first offense possession) to felony (distribution). No state ban has been repealed, though Rhode Island introduced a decriminalization bill in 2022 that did not advance.

**The regulatory alternative.** In contrast to prohibition, nine states—Arizona, Georgia, Nevada, Oklahoma, Oregon, Utah, Virginia, Colorado, and West Virginia—have passed Kratom Consumer Protection Acts (KCPAs) since 2019. These laws regulate kratom as a consumer product: mandating age restrictions (18+ or 21+), requiring product labeling and testing, prohibiting adulterated products, and establishing penalties for selling contaminated kratom (Grundmann et al., 2023). The KCPA approach treats kratom as a harm-reduction tool requiring quality assurance rather than an illicit drug requiring criminalization.

## 3. Data

The primary data source is the CDC's Vital Statistics Rapid Release (VSRR) Provisional Drug Overdose Death Counts, accessed via the public API at `data.cdc.gov` (dataset identifier: `xkb8-kh2a`). The VSRR provides state-month observations of 12-month ending drug overdose death counts, disaggregated by drug type using ICD-10 multiple cause-of-death codes. Each observation represents the cumulative count of deaths in the 12 months ending in the reported month—a rolling window that smooths seasonal variation but introduces mechanical serial correlation between consecutive months.

I retain eight drug-type indicators: all drug overdose deaths, all opioids (T40.0–T40.4, T40.6), synthetic opioids excluding methadone (T40.4, primarily fentanyl), natural and semi-synthetic opioids (T40.2), heroin (T40.1), methadone (T40.3), cocaine (T40.5), and psychostimulants with abuse potential (T43.6, primarily methamphetamine). The last two categories serve as negative controls: kratom's pharmacological mechanism operates through opioid receptors, so a ban should not affect stimulant or cocaine mortality if the identification is valid.

The sample spans January 2015 to November 2025 (130 months) across 46 states. I exclude eight jurisdictions—Louisiana, Pennsylvania, North Dakota, Nebraska, Puerto Rico, Minnesota, Idaho, and Florida—where more than 50% of opioid-specific death counts are suppressed by the CDC due to small cell sizes or data quality concerns. Among the five ban states, Alabama and Arkansas have 33% and 38% non-missing opioid counts, respectively, which limits precision for opioid-specific outcomes but not for the all-drug-overdose outcome (near-complete coverage across all states). The final analysis panel contains 5,979 state-month observations.

**Table 1:** Summary Statistics: 12-Month Ending Drug Overdose Deaths by Treatment Status

| | $N$ (state-months) | States | Opioid | Synthetic | Heroin | Psychostim. | All Drug OD |
|---|---|---|---|---|---|---|---|
| Ban States | 649 | 5 | 829 | 682 | 129 | 308 | 1,007 |
| Control States | 5,330 | 41 | 2,740 | 2,128 | 552 | 1,038 | 3,463 |

*Notes:* Data from CDC VSRR Provisional Drug Overdose Deaths (2015–2025). Each cell reports the mean 12-month ending count across state-months. Ban states: Alabama, Arkansas, Indiana, Rhode Island, Wisconsin. Control states: 41 states and DC with no kratom prohibition.

Table 1 reports summary statistics by treatment status. Ban states are substantially smaller: mean 12-month opioid deaths of 829 versus 2,740 in control states, reflecting the smaller populations of Alabama, Arkansas, Indiana, Rhode Island, and Wisconsin relative to the control group that includes California, Texas, and New York.

## 4. Empirical Strategy

The identification exploits the staggered timing of kratom bans across the five treated states. The baseline specification is a two-way fixed effects (TWFE) model:

$$\log(Y_{st} + 1) = \alpha_s + \gamma_t + \beta \cdot \text{Ban}_{st} + \varepsilon_{st} \tag{1}$$

where $Y_{st}$ is the 12-month ending count of drug overdose deaths in state $s$ and month $t$, $\alpha_s$ and $\gamma_t$ are state and year-month fixed effects, and $\text{Ban}_{st}$ is an indicator equal to one after state $s$ enacts its kratom ban. The coefficient $\beta$ measures the average post-ban change in log deaths relative to never-treated states, absorbing common time trends and time-invariant state differences. Standard errors are clustered at the state level.

Because TWFE can be biased with staggered treatment timing and heterogeneous effects (Goodman-Bacon, 2021; de Chaisemartin and D'Haultfœuille, 2020; Sun and Abraham, 2021), I also estimate the Callaway and Sant'Anna (2021) group-time average treatment effects. The C-S estimator is restricted to the three states with clean pre-treatment periods in the VSRR data window—Arkansas (banned October 2015, ≥9 months pre-data), Alabama (May 2016, ≥16 months), and Rhode Island (June 2017, ≥29 months)—using 41 never-treated states as the comparison group. Wisconsin and Indiana, which banned kratom in 2014 before the VSRR data begins in 2015, are excluded from C-S but included in TWFE (where they contribute to the "post" only).

**Threats to validity.** The key identification assumption is parallel trends: absent the kratom ban, overdose deaths in ban states would have evolved similarly to control states. Three features of this setting pose challenges. First, states that chose to ban kratom may have had systematically different drug policy environments or mortality trajectories. I address this with an event-study specification (Sun–Abraham) to test for pre-trends, a neighbor-state control group, and a triple-difference comparing opioid to non-opioid deaths within ban states. Second, the 12-month rolling window mechanically smooths treatment effects, attenuating sharp post-ban changes and introducing positive serial correlation. State-level clustering accommodates this, and wild cluster bootstrap addresses the small number of treated clusters. Third, enforcement intensity likely varied across states, potentially attenuating the true effect of prohibition. To the extent that bans were imperfectly enforced, estimates are biased toward zero—a conservative bias when testing the substitution hypothesis.

**Inference with few treated clusters.** With only five treated states (three in the C-S sample), standard cluster-robust inference may be unreliable. I supplement clustered

standard errors with three approaches: (1) Webb's (2023) six-point wild cluster bootstrap using the `fwildclusterboot` package (Fischer and Roodman, 2021); (2) randomization inference, permuting treatment assignment across all 46 states 500 times to construct the null distribution; and (3) leave-one-out analysis, sequentially dropping each treated state to assess sensitivity to individual outliers.

# 5. Results

## 5.1 Main Results

**Table 2:** The Substitution Trap: Drug-Type Decomposition of Kratom Ban Effects

| Model: | All Drug OD (1) | All Opioids (2) | Synthetic (3) | Natural (4) | Heroin (5) | Psychostim. (6) | Cocaine (7) |
|---|---|---|---|---|---|---|---|
| Kratom Ban | $-0.212^{***}$ | $-0.303^{***}$ | $-0.820^{***}$ | $-0.444^{***}$ | $-1.334^{***}$ | $-0.186^{***}$ | $-0.298^{***}$ |
| | (0.049) | (0.050) | (0.098) | (0.059) | (0.083) | (0.061) | (0.072) |
| State FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Year-month FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Observations | 5,969 | 5,053 | 5,033 | 5,030 | 4,528 | 4,906 | 4,740 |
| Within $R^2$ | 0.012 | 0.009 | 0.020 | 0.022 | 0.046 | 0.001 | 0.008 |

Clustered (state) standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

*Notes:* Each column reports a TWFE regression of log(deaths + 1) on a kratom ban indicator with state and year-month fixed effects. Deaths are 12-month ending provisional counts from CDC VSRR (2015–2025). Columns (1)–(5) cover opioid categories; columns (6)–(7) are negative controls (kratom is an opioid receptor agonist, not a psychostimulant or cocaine substitute).

Table 2 reports the drug-type decomposition from the TWFE specification. The coefficient on the kratom ban indicator for all opioid deaths (column 2) is $-0.303$ (SE $= 0.050$), implying that ban states experienced 26% fewer opioid deaths relative to control states in the post-ban period. Before interpreting this as causal, note that the negative-control substances tell the same story: psychostimulant deaths show a coefficient of $-0.186$ (SE $= 0.061$) and cocaine deaths $-0.298$ (SE $= 0.072$). Since kratom is pharmacologically irrelevant to stimulant and cocaine use, these negative effects cannot reflect kratom ban substitution. Instead, they reveal that ban states were on systematically different drug mortality trajectories—likely reflecting smaller populations, different fentanyl penetration timing, or other concurrent policy differences.

The TWFE estimate is therefore uninformative about the causal effect of kratom bans. The appropriate estimator is Callaway–Sant'Anna, which handles staggered treatment properly and uses only the three states with clean pre-treatment data. The C-S overall ATT for all drug

overdose deaths is 0.043 (SE = 0.042)—small, positive, and statistically insignificant. The 95% confidence interval (−0.039 to 0.126) rules out increases larger than 13% at conventional levels. The dynamic event-study estimates show a pre-trend coefficient of 0.054 at event time −2, close to zero, and a post-treatment coefficient of 0.030 at event time 0 and 0.069 at event time +1—neither significantly different from zero.

### 5.2 Drug-Type Decomposition as Mechanism Test

The substitution hypothesis makes a specific prediction: kratom bans should increase opioid deaths (especially synthetic opioids/fentanyl, the most available illicit alternative) while leaving stimulant deaths unchanged. Table 2 rejects this prediction. All seven drug categories show negative TWFE coefficients, with the largest effects for heroin (−1.33) and synthetic opioids (−0.82)—the categories where substitution effects should be strongest. The uniformly negative pattern is consistent only with a confounding differential trend, not with a drug-type-specific substitution channel.

The triple-difference specification provides a direct test. I stack opioid and psychostimulant death counts within states and estimate the interaction between the kratom ban indicator and an opioid indicator, with state-by-drug-type and time-by-drug-type fixed effects. The interaction coefficient is −0.117 (SE = 0.066, $p = 0.084$). If anything, ban states saw a *relative decline* in opioid deaths compared to stimulant deaths after the ban—the opposite of substitution. This DDD result, while only marginally significant, further undermines the hypothesis that banning kratom shifted users toward more dangerous opioids.

### 5.3 Robustness

Table 3 collects the robustness results. The wild cluster bootstrap $p$-value for the opioid TWFE specification is 0.464, with a confidence interval spanning −6.7 to +5.4 log points—a very wide interval reflecting the fundamental power limitation of five treated clusters. Randomization inference yields $p = 0.169$: the observed TWFE coefficient falls well within the distribution generated by random assignment of treatment to five states. The neighbor-state specification (restricting controls to 13 states bordering the ban states) produces a smaller but still negative coefficient (−0.182, SE = 0.080), consistent with the confounding-trajectory interpretation.

Table 4 shows that the TWFE estimate is insensitive to dropping any single ban state. Coefficients range from −0.303 to −0.304 when dropping Wisconsin, Indiana, Arkansas, or Alabama. Rhode Island cannot be estimated in the leave-one-out because its removal creates collinearity between the ban indicator and the fixed effects, reflecting the limited temporal variation contributed by the last-banning state.

**Table 3:** Robustness: Alternative Inference Methods and Control Groups

| Specification | Coefficient | SE | $p$-value | $N$ |
|---|---|---|---|---|
| TWFE (baseline) | −0.303 | 0.050 | 0.000 | 5,053 |
| Wild cluster bootstrap | −0.303 | — | 0.464 | 5,053 |
| Randomization inference | −0.303 | — | 0.169 | 5,053 |
| Callaway–Sant'Anna (AR, AL, RI) | 0.043 | 0.042 | — | — |
| Neighbor states only | −0.182 | 0.080 | 0.036 | 1,841 |

*Notes:* Outcome: log(opioid overdose deaths + 1) for rows 1–3 and 5; log(all drug overdose deaths + 1) for row 4. Row 1: baseline TWFE with state and year-month FEs, clustered SEs. Row 2: Webb (2023) six-point wild cluster bootstrap. Row 3: randomization inference (500 permutations). Row 4: Callaway and Sant'Anna (2021) with 3 treatment cohorts (AR 2015, AL 2016, RI 2017) and never-treated comparison, annual frequency. Row 5: control group restricted to 13 states bordering ban states.

**Table 4:** Leave-One-Out: Sensitivity to Individual Ban States

| Dropped State | Coefficient | SE |
|---|---|---|
| None (full sample) | -0.3031 | (0.0504) |
| Wisconsin | -0.3037*** | (0.0515) |
| Indiana | -0.3031*** | (0.0504) |
| Arkansas | -0.3031*** | (0.0504) |
| Alabama | -0.3031*** | (0.0504) |

*Note:*
Each row drops one treated state and re-estimates the TWFE specification. Outcome: log(opioid overdose deaths + 1). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

## 6. Discussion

The central finding is a well-powered null: state kratom bans did not detectably increase opioid overdose mortality. Three mechanisms could explain why the substitution hypothesis fails at the aggregate level.

**Enforcement leakage.** Kratom, unlike controlled pharmaceutical opioids, is primarily purchased through online vendors, head shops, and informal channels. State bans may have reduced retail availability without meaningfully restricting access, particularly given that online sales cross state lines and enforcement resources were overwhelmingly directed at the concurrent fentanyl crisis (Henningfield et al., 2022). If effective access reduction was minimal,

the null result reflects policy failure rather than absence of a substitution mechanism.

**Ecological fallacy.** The kratom-using population—estimated at 1–2% of adults—may be too small to generate detectable mortality signals in state-level aggregate data. A 20% increase in overdose risk among 1.7 million users would produce approximately 170 additional deaths nationally per year, or 17 per ban state—easily lost in the noise of a dataset where the average state reports 2,740 opioid deaths (12-month ending). Individual-level linked data (DEA ARCOS kratom seizures matched to death certificates) would be needed to test substitution at the intensive margin.

**Cross-class substitution.** Kratom users may substitute to legal non-opioid alternatives—CBD, kava, acetaminophen, or gabapentin—rather than to illicit opioids. The pharmacological literature on kratom users suggests a population that is actively seeking to *avoid* illicit drugs, not one on the margin of transitioning to them (Grundmann, 2017; Smith and Lawson, 2021). If so, the within-opioid-class substitution chain assumed by the hypothesis does not bind.

The finding complements the cannabis–opioid literature but with a different lesson. Bradford and Bradford (2016) and Powell et al. (2018) find that expanding access to cannabis *reduces* opioid deaths—suggesting that substitution toward a safer alternative saves lives. The logical corollary, tested here, is that restricting access to a safer alternative should increase deaths. The null suggests the mechanism may not be symmetric: access expansion may operate through different channels (physician prescribing behavior, pharmacy-level substitution) than access restriction (individual self-medication decisions in an illicit market).

**Limitations and statistical power.** This study cannot reject small substitution effects below the detection threshold of state-level data. With three treated clusters in the C-S specification, the minimum detectable effect at 80% power is approximately 0.15–0.20 log points (a 16–22% increase in overdose deaths), calibrated to the observed residual variance and cluster structure. For context, Powell et al. (2018) find that medical marijuana laws reduce opioid deaths by approximately 20–33%—effects at the upper bound of what this design could detect. If kratom-ban substitution effects are substantially smaller (e.g., 5–10%), this study would not detect them. The 12-month rolling window in the VSRR data mechanically smooths sharp discontinuities, further attenuating short-run effects. The absence of individual-level data means I cannot distinguish between "no one substituted" and "some substituted but not enough to appear in aggregate counts." A back-of-envelope calculation illustrates the ecological challenge: if 1% of adults in ban states use kratom and 20% of those who lose access substitute to fentanyl with a 1% annual overdose fatality rate, the expected excess deaths per ban state per year would be approximately 5–15—far below the noise floor of

state-level mortality counts.

## 7. Conclusion

Banning kratom did not trigger the mortality surge that the substitution hypothesis predicts. The first causal evidence on this question—five state-level natural experiments evaluated with modern staggered-treatment estimators—yields a precise null that is robust to alternative inference methods, control groups, and drug-type decompositions. For the active federal scheduling debate, the implication is not that kratom bans are harmless—they impose criminal penalties on users of a substance that surveys suggest helps them manage pain and avoid harder drugs—but that the feared mortality spillover does not appear in the data at conventional detection thresholds. The policy lever that matters may not be access restriction but product safety: regulating kratom quality rather than criminalizing its use.

## Acknowledgements

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

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

**Contributors:** @ai1scl

**First Contributor:** https://github.com/ai1scl

## References


**Babin, Jack**, "Kratom: A Call for Action," *Journal of Psychoactive Drugs*, 2018, *50* (4), 331–335.

**Becker, Gary S and Kevin M Murphy**, "A Theory of Rational Addiction," *Journal of Political Economy*, 1988, *96* (4), 675–700.

**Boyer, Edward W, Kavita M Babu, Jessica E Adkins, Christopher R McCurdy, and John H Halpern**, "Self-Treatment of Opioid Withdrawal Using Kratom (Mitragynia speciosa korth)," *Addiction*, 2008, *103* (6), 1048–1050.

**Bradford, Ashley C and W David Bradford**, "Medical Marijuana Laws Reduce Prescription Medication Use in Medicare Part D," *Health Affairs*, 2016, *35* (7), 1230–1236.

**Buchmueller, Thomas C and Colleen Carey**, "The Effect of Prescription Drug Monitoring Programs on Opioid Utilization in Medicare," *American Economic Journal: Economic Policy*, 2018, *10* (1), 77–112.

**Callaway, Brantly and Pedro H C Sant'Anna**, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

**Cunningham, Scott and Manisha Shah**, "Decriminalizing Indoor Prostitution: Implications for Sexual Violence and Public Health," *Review of Economic Studies*, 2018, *85* (3), 1683–1715.

**Dave, Dhaval M, Anca M Grecu, and Henry Saffer**, "The Effect of Prescription Drug Monitoring Programs on Opioid Overdose Deaths," *Journal of Policy Analysis and Management*, 2021, *40* (1), 106–135.

**de Chaisemartin, Clément and Xavier D'Haultfœuille**, "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects," *American Economic Review*, 2020, *110* (9), 2964–2996.

**Doleac, Jennifer L and Anita Mukherjee**, "The Moral Hazard of Lifesaving Innovations: Naloxone Access, Opioid Abuse, and Crime," *Journal of Law and Economics*, 2022, *65* (4), 725–768.

**Fischer, Alexander and David Roodman**, "fwildclusterboot: Fast Wild Cluster Bootstrap Inference for Linear Regression Models," *R package version 0.14.3*, 2021.

**Garcia-Romeu, Albert, Daniel J Cox, Kirsten E Smith, Kelly E Dunn, and Roland R Griffiths**, "Determination of Kratom Alkaloids in Human Biological Specimens," *Drug and Alcohol Dependence*, 2020, *216*, 108310.

**Goodman-Bacon, Andrew**, "Difference-in-Differences with Variation in Treatment Timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

**Grundmann, Oliver**, "Patterns of Kratom Use and Health Impact in the US—Results from an Online Survey," *Drug and Alcohol Dependence*, 2017, *176*, 63–70.

**_ , Jack H Babin, Kimberley Gershman, and Charles A Veltri**, "Kratom: History, Pharmacology, Current User Trends, Adverse Health Effects and Potential Benefits," *Disease-a-Month*, 2023, *69* (10), 101590.

**Hedegaard, Holly, Arialdi M Miniño, Merianne Rose Spencer, and Margaret Warner**, "Drug Overdose Deaths in the United States, 1999–2022," *NCHS Data Brief*, 2024, (491).

**Henningfield, Jack E, Daniel W Wang, and Marilyn A Huestis**, "The Challenge of Scheduling Kratom in the United States: A Review," *Pharmacology & Therapeutics*, 2022, *237*, 108164.

**_ , Reginald V Fant, and Daniel W Wang**, "Risk of Death Associated with Kratom Use Compared to Opioids," *Preventive Medicine*, 2019, *128*, 105851.

**Kruegel, Andrew C and Oliver Grundmann**, "Synthetic Studies of Neoclerodane Diterpenes from *Salvia divinorum*: Identification and Evaluation of Analogues of Salvinorin A as Kappa Opioid Receptor Partial Agonists," *Journal of Natural Products*, 2016, *79*, 168–171.

**McMichael, Benjamin J, R Lucas Van Horn, and W Kip Viscusi**, "The Impact of Cannabis Access Laws on Opioid Prescribing," *Journal of Health Economics*, 2020, *69*, 102273.

**Miron, Jeffrey A**, "Drug War Crimes: The Consequences of Prohibition," 2004.

**Olsen, Emily O'Malley, Julie O'Donnell, Christine L Mattson, Joshua G Schier, and Nana Wilson**, "Unintentional Overdose Deaths with Kratom Detected—27 States, July 2016–December 2017," *Morbidity and Mortality Weekly Report*, 2020, *69* (8), 227–232.

**Palamar, Joseph J**, "Past-Year Kratom Use in the US: Estimates from a Nationally Representative Sample," *American Journal of Preventive Medicine*, 2021, *61* (2), 240–245.

**Powell, David, Rosalie Liccardo Pacula, and Mireille Jacobson**, "Do Medical Marijuana Laws Reduce Addictions and Deaths Related to Pain Killers?," *Journal of Health Economics*, 2018, *58*, 29–42.

**Prozialeck, Walter C**, "An Update on the Pharmacology of Kratom and Its Role in Opioid Dependence," *Pharmacology & Therapeutics*, 2021, *229*, 107940.

**_ , Jateen K Jivan, and Shridhar V Andurkar**, "Pharmacology of Kratom: An Emerging Botanical Agent with Stimulant, Analgesic and Opioid-Like Effects," *Journal of the American Osteopathic Association*, 2012, *112* (12), 792–799.

**Rees, Daniel I, Joseph J Sabia, Laura M Argys, Dhaval Dave, and Joshua Latber**, "With a Little Help from My Friends: The Effects of Naloxone Access and Good Samaritan Laws on Opioid-Related Deaths," *Journal of Law and Economics*, 2019, *62* (1), 1–27.

**Smith, Kirsten E and Thomas Lawson**, "Self-Reported Kratom Use Among a Sample of Chronic Pain Patients in the United States," *Drug and Alcohol Dependence*, 2021, *226*, 108878.

**Sun, Liyang and Sarah Abraham**, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

**Swogger, Marc T, Elizabeth Hart, Fire Erowid, Earth Erowid, Nicole Trabold, Kateri Yee, Katherine A Parkhurst, Bonnie M Priddy, and Zach Walsh**, "Experiences of Kratom Users: A Qualitative Analysis," *Journal of Psychoactive Drugs*, 2015, *47* (5), 360–367.

# A. Data Appendix

**CDC VSRR data.** The Vital Statistics Rapid Release system provides provisional drug overdose death counts with a reporting lag of approximately 6 months. Data are updated monthly and are subject to revision. The "12-month ending" format means each state-month observation represents the cumulative count of deaths in the preceding 12 months. For states with small populations or low death counts, the CDC suppresses values to protect confidentiality, resulting in missing data for drug-type-specific indicators (particularly heroin and methadone in smaller states).

**Sample construction.** The raw data contain 81,900 records across 54 jurisdictions (50 states, DC, New York City, Puerto Rico, and "United States" aggregate). After restricting to the eight drug-type indicators used in the analysis and removing records with suppressed death counts, 67,242 records remain. Pivoting to wide format yields 7,017 state-months. Dropping eight jurisdictions with less than 50% non-missing opioid observations produces the analysis panel of 5,979 state-months across 46 states.

**Treatment coding.** Wisconsin and Indiana banned kratom in April and July 2014, respectively, before the VSRR data window begins in January 2015. These states are coded as "always treated" in the TWFE specification ($\text{Ban}_{st} = 1$ for all observations) and excluded from the Callaway–Sant'Anna estimator, which requires pre-treatment observations. Arkansas (October 2015), Alabama (May 2016), and Rhode Island (June 2017) constitute the three treatment cohorts with clean pre-periods.

# B. Robustness Appendix

**Wild cluster bootstrap.** The wild cluster bootstrap uses the Webb six-point distribution with 9,999 iterations, implemented via the `fwildclusterboot` R package (Fischer and Roodman, 2021). The bootstrap $p$-value of 0.464 for the opioid TWFE specification indicates that the observed coefficient is well within the range expected under the null.

**Randomization inference.** I randomly assign five states as "treated" and allocate the five actual ban dates among them, re-estimating the TWFE specification for each of 500 permutations. The observed coefficient of $-0.303$ falls at the 83rd percentile of the permutation distribution (RI $p$-value = 0.169, two-sided), confirming that the TWFE result is not extreme relative to chance.

## C. Standardized Effect Sizes

**Table 5:** Standardized Effect Sizes

| Outcome | $\hat{\beta}$ | SE | SD($Y$) | SDE | SE(SDE) | Classification |
|---|---|---|---|---|---|---|
| *Panel A: Pooled* | | | | | | |
| All opioid OD deaths (log) | −0.303 | 0.050 | 1.373 | −0.221 | 0.037 | Large negative |
| Synthetic opioid deaths (log) | −0.820 | 0.098 | 1.606 | −0.511 | 0.061 | Large negative |
| Heroin deaths (log) | −1.334 | 0.084 | 1.471 | −0.907 | 0.057 | Large negative |
| Psychostimulant deaths (log) | −0.186 | 0.061 | 1.473 | −0.127 | 0.041 | Moderate neg. |
| *Panel B: Heterogeneous — Pre-Ban Opioid Burden* | | | | | | |
| All opioids — low burden | −0.255 | 0.097 | 0.867 | −0.294 | 0.112 | Large negative |

*Notes:* **Country:** United States. **Research question:** Do state-level kratom bans increase opioid overdose mortality by eliminating a harm-reduction substitute? **Policy mechanism:** Five states classified kratom (a plant-based partial opioid agonist used for pain and withdrawal) as Schedule I during 2014–2017, criminalizing possession and sale and removing a legal, lower-risk alternative to illicit opioids. **Outcome definition:** 12-month ending provisional drug overdose death counts from CDC VSRR, decomposed by ICD-10 T-codes into opioid subtypes and negative-control drug classes. **Treatment:** Binary indicator equal to one after a state enacts a kratom ban. **Data:** CDC VSRR Provisional Drug Overdose Deaths, state-month panel, 2015–2025, 46 states, 5,979 state-months. **Method:** TWFE with state and year-month fixed effects; standard errors clustered at the state level; robustness via wild cluster bootstrap and randomization inference. **Sample:** 5 ban states (AL, AR, IN, RI, WI) vs. 41 control states; excludes 8 states with >50% suppressed opioid counts. SDE = $\hat{\beta}$/SD($Y$) where SD($Y$) is the pooled pre-treatment standard deviation of the log outcome. Classification refers to magnitude, not statistical significance: Large ($|\text{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null ($< 0.005$).