THE UNIVERSITY OF CHICAGO

Braking Bad: The Effect of Decriminalizing Hard Drugs on Car Accident Rates

By

Sam Farnsworth

May 2025

A paper submitted in partial fulfillment of the requirements for the Master of Arts degree in the Social Sciences

Faculty Advisor: Jens Ludwig

Preceptor: Andrew Proctor

Abstract

Does decriminalizing hard drugs affect the risk of car accidents? The answer is hard to predict because the consequences of such a policy can run in different directions. To the extent that hard drugs impair coordination, judgment, and reaction time, legalization may increase accidents by making the drugs more easily available. But decriminalization might reduce accidents if people replace alcohol with substances that make them less likely to drive dangerously or drive at all. The net effect is an open question. This paper aims to answer it by looking at Oregon's Measure 110 as a natural experiment, using a synthetic difference-in-differences (DID) approach to compare accident trends before and after its implementation. The findings suggest that decriminalization had no detectable effect on the rate of car accidents.

1 Introduction

In 2020, Oregon passed the Drug Addiction Treatment and Recovery Act (Measure 110). The law decriminalized heroin, methamphetamine, LSD, and oxycodone. Rather than facing arrest, individuals found with these substances were issued a $100 fine; it could be waived if they completed a substance-use screening at an addiction-recovery center. The policy went into effect on February 1, 2021. It aimed to replace a punitive approach to drug abuse with a public-health model. The state also sought to use revenues from taxes on cannabis to fund programs that would provide treatment and recovery services to drug users.

Measure 110 was met with intense scrutiny and debate. Supporters argued that decriminalization would reduce incarceration rates, encourage treatment, and lower the harms associated with punitive drug laws, particularly for historically over-policed communities. Critics worried that decriminalization would lead to more drug use, public-safety problems, and a rise in accidents. In 2024, growing concerns about crime, overdose rates, and public disorder led Oregon to largely reverse the law.

In the time leading up to that reversal, public officials in Oregon claimed that changes in the law had caused, among other things, an increase in car accidents. In 2023, the Oregon Department of Transportation issued a statement claiming that Measure 110 "resulted in a predictable increase in drug-impaired driving crashes and related injuries and deaths." The state's Highway Safety Plan echoed that concern, asserting that "the implementation of State Ballot Measure 110 at the beginning of 2021 was likely responsible for a notable increase in drug-impaired and polysubstance impaired driving occurrences, and their related crashes."

Were those claims true? Oregon's temporary decriminalization offers a rare chance to test the real-world impact of such policies on accident rates. By comparing accident trends before and after Measure 110's implementation—and using data from neighboring states as a control—this study seeks to determine whether decriminalization led to an increase in traffic accidents or had neutral or even mitigating effects. The answer is hard to predict. If hard drugs impair motor skills, decision-making, or risk perception, then easier access to them could lead to more car crashes, workplace injuries, and other mishaps. On the other hand, legalization of some drugs might give the public an alternative to alcohol that makes them less likely to drive dangerously or drive at all. These competing possibilities produce a testable question: how does decriminalizing hard drugs affect accident rates?

2 Literature Review

Studies of the relationship between drug policy and accident rates have explored various channels through which decriminalization might influence public safety. Studies of alcohol and cannabis consumption, for example, suggest that both can impair reaction time and judgment, raising the likelihood of accidents (Anderson & Rees, 2014). Harder drugs such as opioids and stimulants have been linked to similar effects (Kaestner & Grossman, 1998). Kaestner and Grossman found the substances elevated the risk of workplace injuries for men. And studies of the effects of prescription opioids point in the same direction. Past research has shown that increased access to prescription opioids has been associated with rising workplace injuries and traffic fatalities (Currie, Jin, & Schnell, 2018). One might then reasonably expect decriminalizing opioids to be associated with an increase in car accidents as well.

Evidence from marijuana legalization also supports this concern. A systematic review of 29 peer-reviewed articles concluded that legalization of cannabis has generally had negative effects on traffic safety (González-Sala et al., 2023). While some studies found no clear relationship, the majority documented increases in traffic accidents following legalization, particularly among young male drivers. The review also highlighted substantial variation across studies; results differed somewhat by jurisdiction, research design, and patterns of co-use with alcohol. On the whole, however, the weight of the evidence suggests that marijuana legalization (at least in some contexts) has contributed to rising accident rates.

But other research suggests that the effects of drug legalization are more nuanced. A growing body of work explores potential substitution effects between substances. Some studies suggest that when cannabis is legalized, alcohol consumption may decline slightly (Dills et al., 2017), potentially reducing alcohol-related accidents. But both Dills et al. and a systematic review by Guttmannova et al. (2016) note that the evidence on substitution is mixed, with some studies showing complementary effects instead. A difference-in-differences study by Calvert and Erickson (2021) highlights the complexity of the relationship. Analyzing Nielsen Consumer Panel data, they found that cannabis legalization in Oregon was associated with a decline in purchases of alcohol, while Washington experienced an increase.

A final complicating factor is that hard drugs present a different set of risks compared to cannabis. Substances such as opioids, methamphetamines, and hallucinogens can more severely impair coordination, perception, and cognitive function. Opioids, in particular, have been associated with drowsiness and delayed reaction times, which can increase the likelihood of accidents (Currie, Jin, & Schnell, 2018). Those same risks, however, may make users less likely to get behind the wheel of a car. Stimulants such as methamphetamine, on the other hand, may lead to reckless driving by increasing impulsivity and overconfidence (Logan, 1996). These effects suggest that the relationship between hard drug decriminalization and accident rates could differ significantly from the patterns observed in marijuana-legalization studies. But studies examining the effect of hard drugs on accident rates are rare because natural experiments that shed light on the subject are uncommon.

The study that most directly addresses this question was published two months ago (Gunadi & Shi, 2025). (The idea for the present study was developed independently and prior to the publication of their article.) Their concurrent work, which focuses on fatal crashes, represents a parallel line of inquiry for the present analysis. Using a synthetic difference-in-differences approach, the authors estimated the impact of Oregon's Measure 110 on fatal traffic crashes, drawing on data from the Fatality Analysis Reporting System (FARS) between 2018 and 2021. Their findings revealed no statistically significant change in drug-related fatal crashes following the policy's implementation.

This paper goes beyond Gunadi and Shi's analysis in a few ways. First, while their study focused exclusively on fatal crashes, this paper considers all reported crashes. Fatal crashes are relatively rare, and many drug-impaired driving incidents do not result in death. Second, this study expands both the pre- and post-treatment periods. Third, the current study uses different variables to create a control group: monthly crash rates, vehicle miles traveled (VMT), and population density. This approach better captures the context relevant to accident risk, providing a more accurate comparison group for Oregon. Finally, this paper more directly investigates the claims made by Oregon officials about the effects of the state's decriminalization law. The claims made by those officials were not limited to fatalities. They were about accidents more broadly, and they may have contributed to the reversal of the law. Those claims deserve direct testing.

3 Methods and Data

3.1 Methods

This study uses a synthetic difference-in-differences (DID) design. DID designs allow us to estimate the effect of a policy change by comparing how outcomes change over time in a group that was affected by the policy and a group that was not. The DID design helps control for underlying trends that would have influenced outcomes even without the treatment.

DID designs start by measuring the outcome of interest (here, accident rates) before the policy goes into effect in both the treated group and a comparison group. They then measure the same outcome in both groups after the policy changes. The first “difference” is the change over time in the treated group: its after-minus-before accident rate. The second “difference” is the change over time in the control group: its after-minus-before accident rate. By subtracting the control-group change from the treated-group change, any background trends that would have affected both groups equally are removed. The remaining difference isolates the effect that can plausibly be attributed to the policy itself, under the assumption that without the intervention both groups would have followed similar paths.

More formally, DID designs identify a policy’s average treatment effect on the treated (ATT). The ATT is defined as:

$$\mathbb{E}[Y_i(1) - Y_i(0) \mid D_i = 1] \tag{1}$$

$Y_i(1)$ is the potential outcome for unit i under treatment, $Y_i(0)$ is the potential outcome under no treatment, $D_i = 1$ if unit i is treated, and $D_i = 0$ otherwise. Properly estimating the ATT requires an identification assumption. An identification assumption is a condition under which it’s fair to draw inferences from the data about the point under investigation. In this case, for example, it’s the assumption that the treated and untreated groups (here, Oregon and the otherwise similar states to which it’s being compared) would experience the same change over time. This is known more specifically as the parallel trends assumption, and is formally modeled as follows:

$$\mathbb{E}[Y_1(0) - Y_0(0) \mid D = 1, X = x] = \mathbb{E}[Y_1(0) - Y_0(0) \mid D = 0, X = x] \tag{2}$$

There is no way to directly test the empirical accuracy of the parallel trends assumption. The best we can do is examine whether the pre-treatment trends for Oregon and the control group move together over time. If trends in Oregon and the other states were very similar for years before Oregon changed its law, we assume they would’ve continued to be similar if Oregon hadn’t changed its law. If this assumption holds, then the average difference in the change in outcomes between the groups identifies the ATT.

The DID design in this study is synthetic. "Synthetic" means that instead of identifying a pre-existing control group, the control group is assembled for the purposes of the study in a way that aligns well with the treatment group (Arkhangelsky et al., 2021; Abadie, Diamond, & Hainmueller, 2015). This study created a "fake Oregon" from a mix of other states that, when blended together, closely match Oregon's trend before Measure 110 took effect. The weights assigned to each state in the blend were chosen by minimizing the squared differences between Oregon's observed accident rates and the weighted average of the donor states' accident rates during the pre-policy period.

With the synthetic control group constructed, two approaches were then used to measure the effect of Oregon's change in policy: a traditional difference-in-differences regression and a synthetic DID model. In the regression DID, Oregon was compared to all other states in the dataset that were not treated, while controlling for covariates. The regression approach models the outcome as:

$$Y_{it} = \alpha + \beta_1 \text{Treatment}_i + \beta_2 \text{Post}_t + \beta_3 (\text{Treatment}_i \times \text{Post}_t) + X_{it}\gamma + \varepsilon_{it} \tag{3}$$

Y_{it} represents the accident rate in region i at time t. Treatment_i is a binary indicator equal to 1 if the observation is from Oregon (treated unit) and 0 if it is from a control state. Post_t is a binary indicator equal to 1 for the period following the implementation of Measure 110 and 0 otherwise. The interaction term $\text{Treatment}_i \times \text{Post}_t$ captures the difference-in-differences estimate of the policy's effect. The term $X_{it}\gamma$ represents a vector of covariates such as traffic volume (VMT), population density, linear time trends, and a set of state fixed effects. In other words, we control for variables that might affect accident rates but that don't have to do with the change in law. ε_{it} is the error term.

Under the synthetic DID design, on the other hand, the ATT is calculated by comparing the post-policy change in Oregon's outcomes to the post-policy change in the synthetic control group, adjusting for any pre-existing differences between them. Using both methods (regression and synthetic) offers a robustness check: the regression controls for observable confounders, while the synthetic control method ensures close pre-treatment fit and helps guard against violations of the parallel trends assumption.

To evaluate whether any observed changes in Oregon's accident rates are significant, this study used placebo-based inferences. This approach provides a more realistic picture of uncertainty than traditional standard errors. In this case, the traditional approach would look at the data one month at a time, treating each of the months as independent and not considering the relationships between them. But in time-series data such as this, accident rates from one month to the next are related in ways that aren't random and need to be accounted for. This study therefore assigned a placebo treatment to each state in the donor pool, as if it had adopted Measure 110 at the same time as Oregon. Then a synthetic control was constructed from the remaining donor states, effectively running the DID design on every state that was part of the pool. The post-treatment gap between each state's actual outcomes and its synthetic control was then calculated, just as it was for Oregon. This produced a distribution of placebo ATT estimates for states that were not actually treated, representing the types of effects that could appear purely by chance. The standard error for Oregon's effect was computed as the standard deviation of these placebo estimates, and the p-value reflects the share of placebo effects at least as extreme as Oregon's. In other words, patterns of random error were generated using the control group and then compared to the results in Oregon.

The placebo-based p-values used in this study are relatively coarse. This method asks, in effect, how extreme an outcome might be produced in the control states by randomness. The relevant result amounts to a fraction showing what share of the donor states had outcomes that were as extreme as Oregon's but were produced by random chance. Because there were ten donor states, the answer to this inquiry can only be a discrete value such as 1/10, 2/10, and so on. A larger donor pool would allow for more precise measurements, but the placebo approach nevertheless offers a useful way to benchmark how surprising Oregon's observed effect would be if the law had never changed.

3.2 Data

The primary data source for the treatment group in this analysis is the Oregon Department of Transportation (ODOT) Crash Database, which provides detailed records of traffic accidents, including whether the crash involved drug and alcohol impairment (but without specifying the character of it). This study focuses on crash reports from 2013 to 2023, allowing for a sufficient pre- and post-policy period to observe meaningful trends. Variables of interest include the total number of accidents, monthly vehicle miles traveled, and population density. Taken together, the latter two variables provide a proxy for the amount of traffic and congestion within a state, and the corresponding number of opportunities for accidents.

The source of the data comes with limitations. By law, Oregon relies on drivers to make reports of their own accidents, and police are not required to respond to or document them. While officers attend serious crashes, minor ones often go unreported. Only 50% of crash reports in the state include police documentation (ODOT, 2020). The completeness of the state database depends on drivers reporting their own crashes, so certain types of incidents are likely to be underreported. Most notably for this study, people involved in crashes while under the influence of drugs may be less likely to report the incident or seek police involvement. These measurement concerns suggest that the crash data in Oregon might understate the true numbers of drug-related accidents. One consideration that cuts against this concern, however, is that accidents usually involve more than one car. And the non-intoxicated party, if there was one, has the same incentive to report as ever.

The synthetic control group was constructed using weighted averages of ten donor states: Washington, California, Wyoming, Idaho, Colorado, Utah, Arizona, Wisconsin, North Carolina, and South Dakota. These states were selected for their similar geographies, traffic patterns, and policy environments to Oregon. Table 1 displays the weights assigned to each state in the construction of the synthetic control group. The zeros are displayed for transparency; the states were included in the dataset but not given any weight as donors. They remain relevant for their use in constructing donor subsets, as will be discussed later.

Table 1: Synthetic Control Weights

Donor State	**Weight**
Washington	0.4805
California	0.2593
Wyoming	0.1263
Idaho	0.0660
Colorado	0.0471
Utah	0.0208
Arizona	0.0000
Wisconsin	0.0000
North Carolina	0.0000
South Dakota	0.0000

Before proceeding with the weighting process, this study considered whether to remove Washington from the donor pool because of its policy environment. The Washington Supreme Court struck down the state's felony drug possession law in February 2021, at which point legislators passed a temporary measure reclassifying possession of small amounts of hard drugs as a misdemeanor. The law took effect in July 2021, around the same time as Oregon's Measure 110, and ran until July 1, 2023 (Gunadi & Shi, 2025). Washington was initially included in the pool despite this policy change because of the differences between the two states' approaches: Oregon fully decriminalized hard drugs, while Washington continued to impose criminal penalties, albeit reduced. Regardless, to assess whether Washington's legal shift influenced the estimated treatment effect, a sensitivity analysis later in this paper re-estimates the synthetic DID model while excluding Washington from the donor pool.

This study also considered another approach entirely: not looking at entire states, but instead comparing the effects in counties that might be considered comparable. That strategy might seem to permit, for example, appealing comparisons between counties in Oregon and counties just on the other side of the state line in Washington, California, and Idaho. Such an approach would have the advantage of comparing counties that might well be similar in demographic and other respects—except for the change in law that applies in one of them but not the other. But these comparisons would be made distinctly problematic by the character of the inquiry in this case, which involves driving. Drivers can and do go back and forth between states freely; there is no barrier at the border, merely a "Welcome" sign. For example, Portland, Oregon lies only about fifteen minutes from the Washington state line. Many Washington residents commute there every day for work and play. Other counties would produce different patterns of this kind, which could only be accounted for through elaborate controls that would be challenging to design. Granted, three of the states in the synthetic pool do border Oregon, which might seem to recreate the difficulties just explained. But of course most of the counties in those states are not adjacent to Oregon, and indeed many of them are hundreds of miles away. This greatly reduces the problems of bleed-over effect that can occur along state lines.

Accident data for each state was pulled directly from its own Department of Transportation. Monthly vehicle miles traveled and population statistics were obtained from federal registries. Because data was limited in some cases, this study imputed plausible numbers at some points to allow for consistent time-series comparisons across states. For example, Colorado provided only quarterly accident data; this study distributed that data evenly over each month in the quarter. California's 2023 data was available only at the annual level and was similarly divided across months. For Wisconsin, four years of crash data were only available at the yearly level; these annual totals were evenly distributed across months to approximate a monthly series. While these imputations were necessary to align the temporal structure of the data, they introduce some measurement uncertainty that should be acknowledged when interpreting the results. The weights were chosen to make the combined crash trend from those states look as similar as possible to Oregon's trend before Measure 110 took effect. The result of the procedure was that states whose accident patterns were close to Oregon's before the policy were given more influence in the synthetic control group, while states that looked less like Oregon were given zero weight.

The weighting process led to three states receiving substantial weights in the final model: Washington, California, and Wyoming. To assess whether a weighted combination of these states serves as a reasonable counterfactual, it is useful to visually inspect their pre-policy crash trends relative to Oregon's.

Figure 1: Monthly crash rates in Oregon vs. selected donor states.

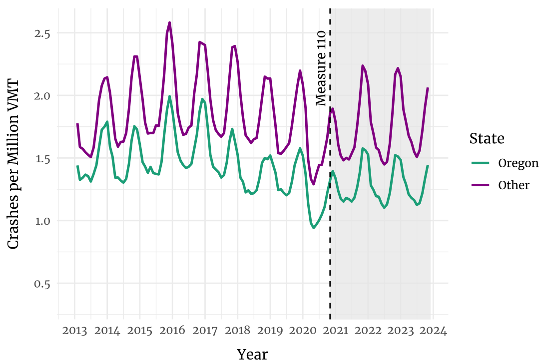

(a) Oregon vs. Washington

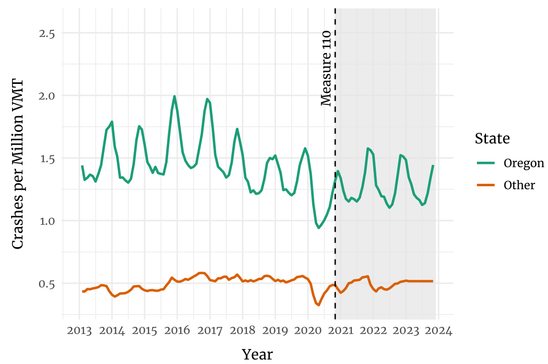

(b) Oregon vs. California

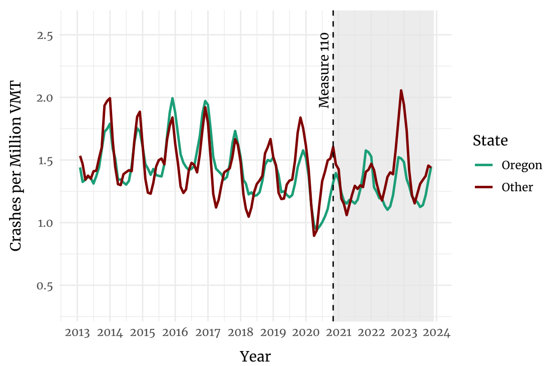

(c) Oregon vs. Wyoming

Using different donor states in a traditional DID design would produce startlingly different ATT estimates. California's crash rate per vehicle miles traveled, for example, is lower and more stable over time than Wyoming's. This may be because California has milder winters and reduced snowfall, or because more driving occurs in California on major interstate highways. In any event, it might appear that California is a less appropriate comparator to Oregon than Washington or Wyoming. But in fact including California in the pool with those other two states creates a fake Oregon that is most similar to the real one. The variation between states highlights the strength of the synthetic DID method: it constructs a weighted combination of donor states that more closely replicates Oregon's pre-treatment trend.

4 Results

4.1 Fit Diagnostics

The methods outlined above seek to assess whether decriminalizing hard drugs affected the rate of car accidents, as measured by the number of crashes per million vehicle miles traveled. We can begin to answer that question with a visual inspection of trend lines. Figure 2 shows the monthly crash rates in Oregon and in the synthetic control. The lines are closely aligned before the change in Oregon's law, suggesting a strong pre-policy match. After Measure 110's implementation, Oregon's trend remains close to its synthetic counterpart. The shaded region of the figure marks the time period after Oregon changed its law.

Figure 2: Monthly crash rates in Oregon vs. synthetic control.

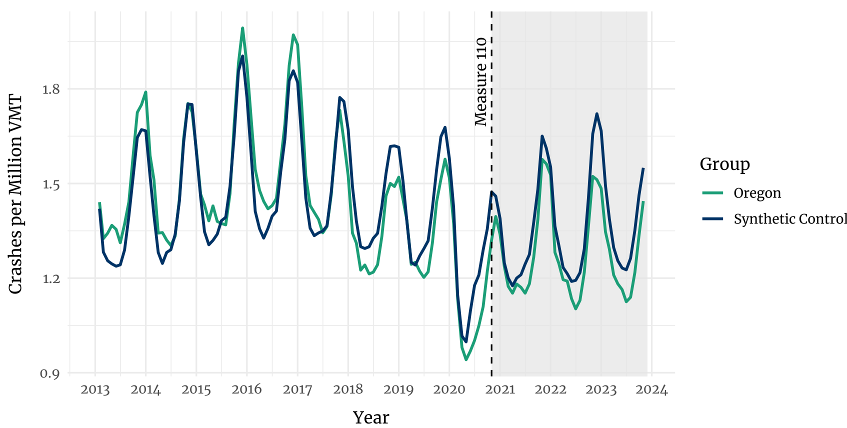

The quality of the fit can be formally assessed using mean squared prediction error (MSPE). MSPE is simply the average squared difference between what actually happened before the policy took effect in Oregon and in the control group. A small pre-treatment MSPE means the synthetic control group was a good match for Oregon. The MSPE was found to be 0.0102, a relatively low value that suggests a strong match when compared to Oregon's average monthly accident rate in the pre-policy period.

The other diagnostic of interest is the MSPE ratio, which divides the MSPE after the policy date by the MSPE before the policy date. If the ratio is close to 1, it tells us that the gaps between Oregon and the synthetic control group after Measure 110 are about as big as the gaps we saw before the law. A ratio much larger than 1 suggests that the fit got worse after the change in the law; a ratio near or below 1 suggests the post-policy gap is no larger than everyday noise. The ratio in this case is only 1.09. The value matches the visual. The synthetic control closely mirrored Oregon before the policy, and Oregon continued to mirror the synthetic control after. If the policy marker were removed, it would be difficult to discern that anything changed at all. At least on initial inspection, the results suggest that the policy had no discernible impact on the rate of automobile accidents.

4.2 Raw ATT Estimates

As explained earlier, two separate methods were used to capture the causal effect of Measure 110 on traffic accidents: a synthetic control-based DID (depicted in the graph above) and a traditional DID regression. Table 2 compares the results from both.

Table 2: Comparison of Raw ATT Estimates

Estimation Method	**ATT Estimate**	**Standard Error**	**p-value**
Synthetic DID	-0.080	0.36	0.4
Regression DID (with covariates)	0.032	0.061	0.64

The estimates run in different directions, and neither is statistically significant. The synthetic design yielded a point estimate of -0.080 crashes per million vehicle miles traveled, and the regression model produced an ATT estimate of 0.032 crashes on the same scale. These are two different ways of expressing a similar outcome: both are consistent with the possibility that Oregon's law had no effect on accident rates.

The placebo test described at the end of Section 3.1 compared the result from Oregon to the effects produced by imagining that Oregon's intervention occurred in states where it was not made in fact. There is little difference between the results; in other words, the change in Oregon's accident rates is not statistically distinguishable from the change in states that didn't try Oregon's experiment. The results of that test correspond to the standard error and p-value presented in the first row of the table. It shows that there was considerable variability in the placebo-driven results. The p-value indicates 4/10 of the donor states showed effects as strong as Oregon's despite no change in the law to account for it.

As shown above, the initial synthetic DID estimate was negative; that is, it suggested a slight decrease in traffic accidents as a consequence of Oregon's law. But further analysis revealed consistent downward drift in placebo estimates. When the method estimated the effect of fictitious interventions that occurred before the real one, it produced negative results in line with the raw numbers from Oregon. This difficulty will be discussed in the next section.

4.3 Placebo Tests and Bias Correction

To fairly judge the accuracy of the method used in this paper, it is prudent to examine whether it's prone to giving false positives—i.e., to suggesting the existence of effects when none were present. If noise can cause the method to show an effect when there wasn't one, that would call into question the reliability of the approach in general. This study tests for that problem using two methods. The first, already mentioned, asks how the data react if we imagine that Oregon's intervention occurred in other states where in fact it didn't. The second, discussed here, asks how the data react if we imagine that Oregon's change in law occurred earlier than it actually did.

This last method tests for a specific danger: that the synthetic DID estimates might result from "overfitting" in the pre-policy period. Overfitting occurs when the model captures not just the underlying trend but also random changes or noise in the data. As a consequence, random drops or spikes after the date of implementation can falsely appear as meaningful policy effects. What looks like a sharp divergence after treatment could simply reflect the model's tendency to overreact to noise, rather than capturing the true ATT.

The risk of overfitting can be tested through use of a pseudo-treatment test. This involves pretending the policy change occurred earlier than it actually did and checking for a perceived effect. If the model consistently detects significant effects during these placebo periods, it suggests that the model is overfitting the pre-policy data and might be falsely identifying randomness as an effect of the change in law.

The method starts by removing all post-policy data and then selecting various "pseudo" treatment dates from the pre-policy period. For each pseudo date, the synthetic DID method is applied as if the change in policy had occurred then. By repeating this process across many dates, it becomes possible to see how often the method detects an "effect" when no real intervention occurred. Under a well-specified model, few of these placebo effects should be statistically significant. Standard errors for the estimates were constructed using the same process outlined in Section 3.1 of this paper. If the actual result in Oregon resembles the results under placebo conditions, that would suggest that Oregon's change in policy did not have a significant effect on accident rates.

Figure 3 (below) displays the results of the fictitious earlier interventions with 95% confidence intervals (the red dot reflects the result of the real change in law). As the chart shows, the effects of the fake interventions are small but consistently negative. In other words, under placebo conditions, the result appears to be a reduction in accident rates. The bias could come from a concentration of weights on only a few donor states. If those heavily weighted states experience a slight downward drift in crash rates, the synthetic control will produce a small negative gap even though no treatment occurred. This bias requires rectification. The average ATT across all placebo periods is –0.134, exceeding the magnitude of the original estimate. The presence of this downward tilt suggests the model may overstate negative effects even in the absence of a real policy change.

Figure 3: Raw ATT Estimates By Pseudo-Treatment Date

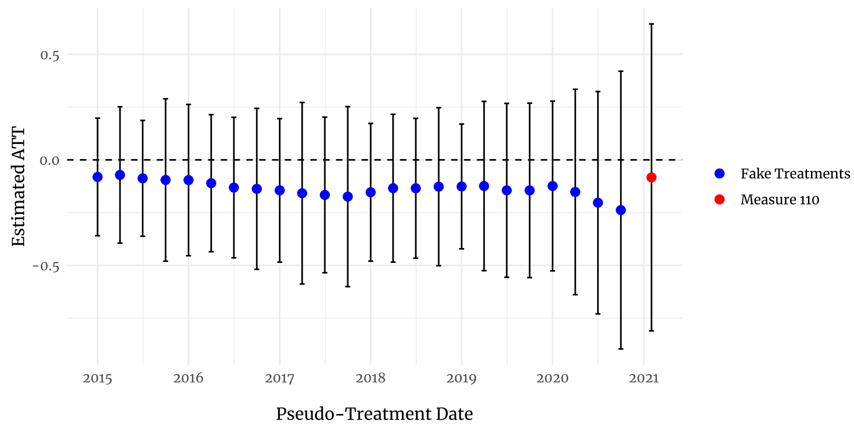

This calls for correction of our investigatory method. We begin by determining what adjustment is needed to account for the bias found in the placebo condition. Then we apply that same adjustment to our actual conditions. The change from Figure 3 to Figure 4 reflects the adjustment.

Figure 4: Bias-Corrected ATT Estimates By Pseudo-Treatment Date

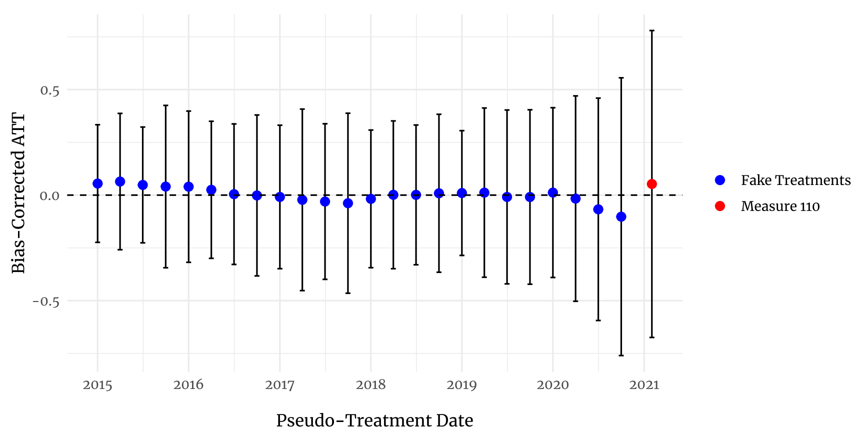

After the correction, the true ATT switches sign and becomes +0.043. Put more simply, the adjustment changes our interpretation of the effect of Oregon's law. It now appears to have produced a slight increase in accident rates, in line with the regression estimate.

The placebo estimates also provide a way to assess how striking the Oregon result should be considered. Figure 5 (below) displays this. It shows the results produced by falsely assuming, over many trials, that Oregon changed its law before it actually did. Randomness, at times, causes those false interventions to appear to produce an effect in the world. We can compare Oregon's result to the fake results to see how unusual it really is.

Figure 5: Distribution of Bias-Corrected Placebo Treatment Effects

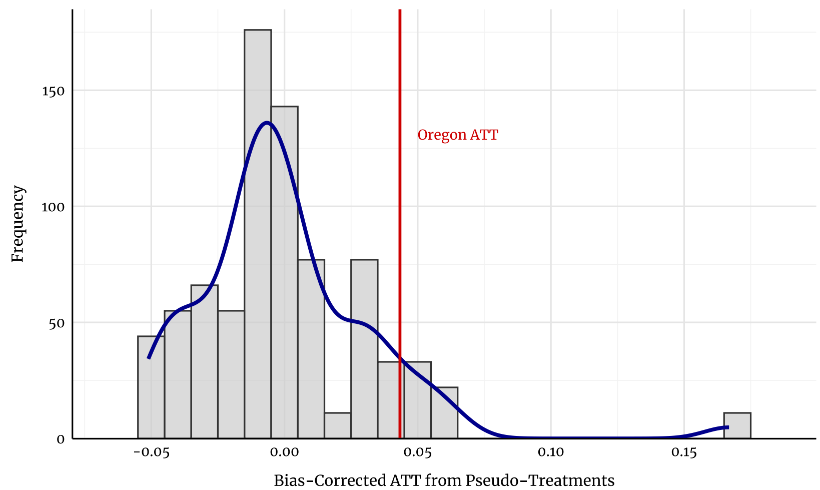

The heights of the gray bars represent the number of placebo treatments that produce a given result. Their overall shape is very roughly that of a bell curve. The tallest bar corresponds to results near zero. The shorter bars to the left and right of the peak represent placebo trials with small positive and negative effects, and the heights of the bars decrease as the illusory effects become greater. Oregon's actual experience is reflected by the red line and would fall between the third and fourth bin from the right. Oregon's result, then, does not stand outside the results driven by placebo interventions. To the contrary, the result in Oregon was weaker than a number of the results produced by placebo trials; to be specific, 11% of the placebo interventions produced stronger effects than Oregon experienced in fact. If the effect in Oregon had been substantial, we would expect it to be further in the tail of the distribution or outside of the placebo effects altogether.

This section has discussed a placebo method that differs from the method used earlier to create standard errors. There, we imagined that Oregon's change in law occurred in states *where* it did not; here, we imagine that it occurred at times *when* it did not. The first method—in which we imagine that the intervention occurred in other states—is better for evaluating the significance of the findings here. Such an approach is commonly used to generate p-values and confidence intervals in studies seeking to understand the impact of policies within a state (Abadie et al., 2010; Arkhangelsky et al., 2021). The other placebo test—in which the intervention is imagined to occur in the same state at different times—is better understood as a diagnostic tool. It tests whether the method is detecting false effects in the absence of treatment by looking at the same context in which the real treatment did occur and may have had effects at a different time (Abadie et al., 2015; Chen & Yan, 2023). Shifting the policy date re-uses nearly the same data from month to month. The in-time placebos are therefore highly correlated; in other words, they share much in common even when they're treated as independent trials. That means the results produced by such a test (and presented in this section) can show us if the test is working and help adjust for bias, but they aren't as reliable a source for further inferences.

4.4 Regression Estimates

To provide another perspective, we can assess the same question using a conventional method of regression. The regression approach has the advantage of controlling for potential confounders. It compares the result in Oregon with the result from all donor states averaged together, regardless of exactly how well each of them matches Oregon. Of course some of the states are more similar to Oregon than others. The synthetic approach weights them to account for that variation; the regression approach does not. But the regression remains a useful way to double-check the results of the synthetic approach, and perhaps to confirm or cast doubt on them.

Table 3: Regression Estimates from Difference-in-Differences Model

Variable	**ATT Estimate**	**t value**	**p-value**
Treated × Post	0.032 (0.061)	0.53	0.60
Vehicle Miles Traveled (VMT)	-0.000064 (0.000054)	-1.19	0.24
Population Density	0.00018 (0.0084)	0.021	0.98

Note: Standard errors in parentheses. $^{*}p < 0.05$

Table 3 presents a full set of regression results, including the treatment effect and the effects of two covariates. Those covariates are population density and vehicle miles traveled. The dependent variable in Table 3 is the accident rate, which was constructed by dividing the monthly number of crashes by the corresponding number of vehicle miles traveled.

The results from the regression above are muted in a way that is consistent with the debiased results from the synthetic analysis shown earlier. The rate of accidents in Oregon after the change in law does not differ significantly from the rate of accidents in other states—once we control for the covariates in the regression (miles traveled and population density). The effect of the policy reflected in the first row of the table above is positive but statistically insignificant, suggesting that the policy's effect on accident rates cannot meaningfully be distinguished from zero.

5 Sensitivity Analysis

It is worth taking a moment to consider how different decisions in this study could have affected its results. A critical decision in a difference-in-differences analysis is the choice of a comparison group. As we've seen, a synthetic control design has the advantage of constructing a weighted comparison group that closely matches Oregon's pre-treatment outcomes. This makes it more robust than a classical DID estimate based on preexisting populations. Still, the accuracy of this approach depends on the weights assigned to donor states. Shifting weight from California to Utah, for example, would change the results produced by the comparison group and affect the study's conclusions. Because different weighting schemes can produce different results, a thorough analysis must include sensitivity checks that test whether the main findings hold under different ways of creating the donor pool and assigning weights to its members.

Two alternative donor pools are examined here. The first limits the pool to Oregon's neighbors: Washington, Idaho, and California. This subset offers a geographically intuitive comparison group, but also constrains the quality of the synthetic match. Since the goal of the method is to create a pool that best matches Oregon, Washington receives the majority of the weight in this specification. This suggests that Washington's pre-treatment trends most closely resemble Oregon's. But a smaller donor set increases the risk that unobserved and idiosyncratic variables may affect the comparisons between those states and Oregon. For example, suppose Washington seemed similar to Oregon before the change in law because they had similar policies with respect to reporting accidents or collecting data. A change in Washington's policies on those scores might create a misleading impression when comparing the results after Oregon made changes of other kinds. There is no way to control for all such possibilities; but putting more states in the pool helps to smooth them out.

The second specification removes Washington from the donor pool. Although Washington's policy shift differed in scope from Oregon's, both states relaxed penalties for hard drug possession around the same time. As noted earlier, possession of hard drugs remained criminal in Washington but the punishments were reduced. This overlap in timing raises the concern that concurrent legal changes could obscure the comparison. If both states' accident rates went up at the same time, for example, they might appear stationary when compared to one another. Excluding Washington from the donor pool tests this possibility.

The results from these two ways of creating a comparison pool are shown in Table 4. The ATTs for each subgroup were debiased using the same procedure outlined above.

Table 4: Estimates From Different Specifications

Subset	**Debiased ATT Estimate**	**Standard Error**
Neighbors	0.097	0.68
Exclude Washington	0.012	0.38

The two variations suggest that Oregon's change in law produced a slight increase in the rate of car accidents. But neither result is statistically significant. The direction and magnitude of the results are in line with those produced using the full set of donors. All estimates are positive, fall below 0.1, and are accompanied by relatively large standard errors. This consistency suggests that the estimated effect of Measure 110 is not sensitive to the specific composition of the comparison group. These pools, like the complete pool used as the basis for this study, do not produce evidence that the change in law affected accident rates.

6 Discussion

This paper has sought to assess how decriminalizing hard drugs affects accident rates, using Oregon's experience as a case study. Oregon's numbers require careful interpretation. After Oregon decriminalized hard drugs, the rate of car accidents in the state (corrected for the amount of driving done) rose slightly. But that number was also rising slightly in places where such a change in law had not occurred. Oregon's relative increase—that is, its increase compared to what happened in otherwise comparable states—was small enough to lack statistical significance; it's consistent with no actual effect (or possibly a negative one).

Over the eleven years studied overall, Oregon averaged around three billion vehicle miles traveled each month. If the SDID model used in this study were perfectly accurate, it would suggest that Measure 110 was responsible for around 130 more car crashes each month. But the result is noisy; the point estimate is associated with a confidence interval of [-0.676, 0.779] crashes per million VMT. In other words, given the number of variables involved, the raw numbers we see would be consistent with a modest negative or positive effect on the rate of car accidents. If the true effect lay at the lower end of this interval, Measure 110 would be associated with roughly 2,000 fewer crashes each month. At the upper end, it would imply about 2,300 more crashes per month. Because the interval spans zero (and indeed includes both reductions and increases), it confirms that the estimated effect is not statistically distinguishable from no effect.

There are a number of possible reasons why Measure 110 may have contributed to a decline in crash rates, or to a smaller increase that some predicted or perceived. One possibility is a substitution effect between controlled substances. As discussed earlier, research on cannabis legalization has found evidence (although mixed) that increased access to one substance may reduce consumption of another (such as alcohol) that is more strongly associated with impaired driving and traffic accidents (Dills et al., 2017; Guttmannova et al., 2016). If similar substitution occurred in response to the decriminalization of hard drugs, with some individuals shifting from alcohol to substances less likely to be consumed in social or mobile settings, this could plausibly reduce alcohol-related accidents. Unlike alcohol, many hard drugs are more often used in stationary environments such as private homes, decreasing the likelihood of impaired driving immediately following use. Opioids and certain other hard drugs tend to induce drowsiness, sedation, or general passivity as well. Those effects can reduce the likelihood of driving. By contrast, alcohol can increase impulsivity and overconfidence, traits associated with riskier driving behavior (Logan, 1996).

Not all hard drugs have the same physical effects, of course. Substituting a stimulant such as methamphetamine for alcohol has very different safety implications than substituting an opioid such as heroin. Because Oregon's Measure 110 decriminalized a range of hard drugs simultaneously, state-level data does not allow us to isolate the crash-rate impact of one drug class versus another. If drug users in Oregon were replacing alcohol with different kinds of hard drugs (or using combinations of these substances together), the effects of the policy on car accidents could be muted.

Another possible explanation is that the results reflect limitations in the accident data rather than real behavioral change. As noted earlier, Oregon's traffic database partially relies on individuals to report their own accidents, particularly for less severe incidents. If someone is using drugs when they crash, they might be less likely to call the police. But again, an accident can be reported by any party to it, not just by an impaired driver. But it's nevertheless possible that some accidents were more likely to go unreported after decriminalization, making it appear that crash rates went down even if they did not.

Beyond measurement and behavioral explanations, it's also worth considering potential confounders. Measure 110 was not just a standalone decriminalization law. It redirected substantial public funding toward addiction recovery services, treatment programs, and harm reduction initiatives. These concurrent policy components may independently influence accident trends. If such changes occurred around the same time as Measure 110's implementation, they may have confounded the estimated treatment effect and limit our ability to isolate the impact of the drug policy alone.

Even so, the findings here offer worthwhile insights. Small increases in car-crash rates can produce significant social and economic harms. Car accidents impose substantial costs, from medical expenses and lost productivity to property damage and emergency response. According to the U.S. Department of Transportation, 14.2 million motor car crashes in 2019 led to a total economic cost of $339.8 billion. This implies that the average crash cost approximately $24,000. If Measure 110 led to roughly 130 more crashes per month (as the synthetic DID point estimate might suggest), that would translate to $3.1 million in monthly costs or nearly $75 million over the two years after implementation. These back-of-the-envelope calculations illustrate that even modest changes in crash rates can have large downstream effects on public resources. At the same time, the large standard error around the estimate means the policy could just about as plausibly have generated net savings of similar magnitude. The uncertainty underscores the need for further research and more precise estimates to guide decisions about drug reform and road safety.

Perhaps most striking, however, is what this study does not find. Contrary to some public claims noted at the outset of this paper, Measure 110 does not appear to have caused a dramatic increase in accident rates in Oregon relative to other states. After Oregon changed its law, its accident rates did go up. But accident rates went up similarly in a collection of other states that have long had accident trends much like Oregon's. If the change in Oregon's law caused an increase in accidents, we would expect to see not just a raw increase in those accidents but a *greater* increase than is seen in similar states that made no such change. That greater increase is not found.

7 Conclusion

The findings of this study have significant practical implications. The war on drugs has been expensive; decriminalization has costs of its own. An accurate comparison of those costs is important for the creation of rational policies going forward. A higher rate of automobile accidents may not be the first form of cost that policymakers consider when they think about decriminalization, but it ought to be part of the calculus—and it certainly was in Oregon. But fears that decriminalization would trigger a surge in drug-impaired collisions were not borne out by the data, nor does the policy appear to have markedly enhanced traffic safety. These results thus challenge common assumptions about the risks of decriminalization. More broadly, this study illustrates the danger that can arise when a state focuses on its own experience without making careful comparisons to other states that are similar in relevant ways. Data seen in isolation can create alarm that dissolves when the numbers are viewed more rigorously and with appropriate context.

The costs and benefits of Measure 110 may have justified the decision to reverse the law in Oregon entirely apart from the inquiry made here. This paper takes no position on that question. But since public officials regarded the effect of the law on accident rates as relevant, it should be understood accurately. Lessons may be drawn from the Oregon experience for the sake of other kinds of laws in further settings. Data, not fear, should shape the road ahead.

Data and Code Availability Statement

The dataset used in this study was compiled from publicly available sources. Monthly traffic crash data and vehicle miles traveled (VMT) were obtained from the Oregon Department of Transportation and comparable state agencies, as well as the Bureau of Transportation Statistics. The data and code needed to replicate the analyses presented in this paper can be found at the following repository: https://github.com/samfarnsworth/Measure110-CrashRates.

References

[1] Abadie, A., Diamond, A., & Hainmueller, J. (2015). Comparative politics and the synthetic control method. *American Journal of Political Science*, *59*(2), 495–510. https://doi.org/10.1111/ajps.12116

[2] Abadie, A., Diamond, A., & Hainmueller, J. (2010). Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program. *Journal of the American Statistical Association*, *105*(490), 493–505.

[3] Anderson, D. M., & Rees, D. I. (2014). The legalization of recreational marijuana: How likely is the worst-case scenario? *Journal of Policy Analysis and Management*, 33(1), 221–232.

[4] Arkhangelsky, D., Athey, S., Hirshberg, D. A., Imbens, G. W., & Wager, S. (2021). Synthetic difference in differences. *National Bureau of Economic Research Working Paper No. 25532*. https://www.nber.org/system/files/working_papers/w25532/w25532.pdf

[5] Blincoe, L. J., Miller, T. R., Zaloshnja, E., & Lawrence, B. A. (2015). *The economic and societal impact of motor vehicle crashes, 2010 (Revised)*. National Highway Traffic Safety Administration. https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/812013

[6] Calvert, C. M., & Erickson, D. J. (2021). Recreational cannabis legalization and alcohol purchasing: A difference-in-differences analysis. *Journal of Cannabis Research*, 3(27). https://doi.org/10.1186/s42238-021-00085-x

[7] Chen, Q., & Yan, G. (2023). A mixed placebo test for synthetic control method. *Economics Letters*, *224*, 111004. https://doi.org/10.1016/j.econlet.2023.111004

[8] Currie, J., Jin, J. Y., & Schnell, M. (2018). U.S. employment and opioids: Is there a connection? *Brookings Papers on Economic Activity*, Spring, 1–44.

[9] Dills, A. K., Goffard, S., & Miron, J. (2017). The effects of marijuana liberalizations: Evidence from Monitoring the Future. *National Bureau of Economic Research Working Paper No. 23779*.

[10] González-Sala, F., Tortosa-Pérez, M., Peñaranda-Ortega, M., & Tortosa, F. (2023). Effects of cannabis legalization on road safety: A literature review. *International Journal of Environmental Research and Public Health*, 20(4655). https://doi.org/10.3390/ijerph20054655

[11] Gunadi, C., & Shi, Y. (2025). *Drug decriminalization and fatal traffic crashes: Evidence from BM110 in Oregon. Health Economics*, 34(5), 815–820. https://doi.org/10.1002/hec.4944

[12] Guttmannova, K., Fleming, C. B., Rhew, I. C., Kosterman, R., Lee, C. M., Kilmer, J. R., & Larimer, M. E. (2015). Impacts of changing marijuana policies on alcohol use in the United States. *Alcoholism: Clinical and Experimental Research*, 40(1), 33–46. https://doi.org/10.1111/acer.12942

[13] Hughes, C., & Stevens, A. (2010). What can we learn from the Portuguese decriminalization of illicit drugs? *British Journal of Criminology*, 50(6), 999–1022.

[14] Kaestner, R., & Grossman, M. (1998). The effect of drug use on workplace accidents. *Labour Economics*, 5(3), 267–294.

[15] Logan, B. K. (1996). Methamphetamine and driving impairment. *Journal of Forensic Sciences*, 41(3), 457–464.

[16] National Highway Traffic Safety Administration. (2023). *The economic and societal impact of motor vehicle crashes, 2019 (Report No. DOT HS 813 403)*. https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/813403.pdf

[17] National Safety Council. (2022). Costs of motor-vehicle crashes. *Injury Facts*. https://injuryfacts.nsc.org/all-injuries/costs/guide-to-calculating-costs/data-details/

[18] Oregon Department of Transportation. (2020). *ODOT crash data system (CDS) "Read Me"*. October.

[19] Oregon Department of Transportation. (2022). *Oregon FY 2023 1300 NHTSA Grant Application*. https://www.oregon.gov/odot/Safety/Documents/Oregon_FY_2023_1300_NHTSA_Grant_Application_08-11-2022.pdf

[20] Oregon Department of Transportation. (2023). *Oregon Highway Safety Performance Plan: Fiscal Year 2024 Planning Workshop*. https://www.oregon.gov/odot/Safety/Documents/Oregon_HSP_2024_Planning_Workshop_Book_FINAL_opt.pdf

[21] National Highway Traffic Safety Administration. (2023). *Oregon FY24–26 Highway Safety Plan*. https://www.nhtsa.gov/sites/nhtsa.gov/files/2024-01/OR_FY24-26HSP-tag.pdf

[22] Oregon State Legislature. (2024). *Oregonians win as the Oregon Senate recriminalizes hard drugs through historic vote in repealing/reforming Measure 110*. https://www.oregonlegislature.gov/smithd/Documents/Measure%20110%20repeal%20reform.pdf

[23] Rosenbaum, P. R. (1999). Choice as an alternative to control in observational studies. *Statistical Science*, 14(3), 259–278. http://www.jstor.org/stable/2676761

[24] Selsky, A. (2020, November 4). Oregon leads the way in decriminalizing hard drugs. *Associated Press.* https://apnews.com/article/oregon-first-decriminalizing-hard-drugs-01edca37c776c9ea8bfd4afdd7a7a33e