Green Borders: Cross-Border Enforcement Spillovers from Cannabis Legalization

By

Yaotian Zhang

Submitted to Central European University - Private University,

Department of Economics

In partial fulfillment of the requirements for the degree of
Master of Arts in Economics, Data, and Policy.

Supervisor: Sergey Lychagin

Vienna, Austria

June 2026

Copyright Notice

Copyright © Yaotian Zhang, 2026. Green Borders: Cross-Border Enforcement Spillovers from Cannabis Legalization. This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives (CC BY-NC-ND) 4.0 International license.

For bibliographic and reference purposes, this thesis should be referred to as: Zhang, Yaotian. 2026. Green Borders: Cross-Border Enforcement Spillovers from Cannabis Legalization. MA thesis, Department of Economics, Central European University - Private University, Vienna.

Author's Declaration

I, Yaotian Zhang, candidate for the MA degree in Economics, Data, and Policy declare herewith that the present thesis titled “Green Borders: Cross-Border Enforcement Spillovers from Cannabis Legalization” is exclusively my own work, based on my research and only such external information as properly credited in notes and bibliography.

I declare that no unidentified and illegitimate use was made of the work of others, and no part of the thesis infringes any person’s or institution’s copyright.

I also declare that no part of the thesis has been submitted in this form to any other institution of higher education for an academic degree.

Vienna, June 2026

Declaration of Generative AI

Generative artificial intelligence (GenAI) was used in this work. I, the author, have reviewed and edited the content as needed and take full responsibility for the content, claims, and references. An overview of the use of GenAI is provided below.

- I used Claude and ChatGPT to assist with language editing and revision of the manuscript. The tool was used to improve the clarity, structure, and academic style of text that I had drafted and to suggest alternative wording, all of which I reviewed and edited. I retain responsibility for the content, the argument, and all claims and references.
- I used Claude to assist with the R code for the empirical analysis. The tool was used to provide debugging support and to help generate code for data processing and visualization. I independently verified all statistical outputs and results.

Abstract

Most studies on cannabis legalization focus only on the state that legalizes. The enforcement burden, however, can cross into neighbors that still prohibit it. Illinois opened a legal cannabis market in January 2020, while its neighbor Iowa kept prohibition. This thesis measures these spillover effects with a county-year difference-in-differences design. The treated group is the seven Iowa counties with a direct highway bridge to Illinois, taken from the Iowa bridge inventory, while the control group is the ninety-two interior counties. The data are NIBRS offense counts for 2015 to 2022. The event study does not reject parallel pre-trends. The four estimates fit together. Recorded drug offenses in bridge counties rose by about 26 percent relative to interior counties ($\beta = 0.233$, $p = 0.024$). The NIBRS drug category counts enforcement, not use, so it points to a cross-border channel without showing how large it is. The central result is the property response. Recorded property offenses climbed about 37 percent against the interior ($\beta = 0.312$, $p = 0.005$). This is more consequential, because the harm reaches a category that legalization does not directly cause. Part of the burden therefore lands outside drug enforcement. Violent offenses are positive but too noisy to read. Driving offenses are flat. The design cannot fully rule out a separate post-2020 shock to these border-industrial counties, so the property mechanism is not settled. Even so, a cost-benefit account that counts only in-state tax revenue understates the burden placed on the neighbor.

Keywords: cannabis legalization, cross-border enforcement spillovers, fiscal externality, property offenses, difference-in-differences, NIBRS
JEL classification: H23, H73, H77, I18, K42, R12

Acknowledgements

I am grateful to my supervisor, Sergey Lychagin, for his guidance throughout this thesis. His questions and close reading sharpened the work, and I learned a good deal from the way he approaches empirical research. I thank him for his patience and for the time he gave the project. I also thank Tomy Lee, whose comments during the seminar helped shape the direction of the work.

I owe a great deal to Central European University for the opportunity to pursue this degree, an experience from which I have benefited enormously, and to the Stockholm School of Economics for hosting me during my exchange term. I am thankful to all the instructors and teaching assistants who taught me at both institutions, through whom I was able to engage with economics at its current frontier.

I am deeply grateful to my parents and the rest of my family for their steadfast support. They have stood behind my decision to pursue an academic life without ever trying to direct the path I take, and that freedom has mattered a great deal to me. I also thank my classmates, from whom I have learned so much and alongside whom I formed friendships I expect to carry well beyond these years. To my friends, I owe thanks for their patience and encouragement throughout my MA studies.

Finally, I want to thank myself. These years asked more of me than I had anticipated, and I am glad to have met their challenges rather than turned from them, to have kept growing through them, and to have stayed in search of who I am along the way.

Contents

List of Figures

List of Tables

List of Abbreviations

ACS	American Community Survey
CBD	cannabidiol
COVID	coronavirus disease (COVID-19)
DiD	difference-in-differences
DUI	driving under the influence
EPSG	coordinate reference system identifier
FE	fixed effects
FIPS	Federal Information Processing Standards (county code)
HC1	heteroskedasticity-consistent standard errors (type 1)
ICPSR	Inter-university Consortium for Political and Social Research
IDFPR	Illinois Department of Financial and Professional Regulation
IL	Illinois
JEL	Journal of Economic Literature (classification)
LEAIC	Law Enforcement Agency Identifiers Crosswalk
NIBRS	National Incident-Based Reporting System
SE	standard error
THC	tetrahydrocannabinol
TIGER/Line	US Census Topologically Integrated Geographic Encoding and Referencing
USGS	United States Geological Survey
WCB	wild cluster bootstrap

1. Introduction

In a little more than 10 years, recreational cannabis in the United States has gone from banned, to legal, to taxed and sold over the counter in much of the country. The change came state by state. By the middle of the 2020s, about half the states had legalized sales for adult use, while the others still banned the drug. So a legal store is often just a short drive across the line from somewhere where buying the same product is a crime. That kind of border is particularly easy to cross in one direction. Someone who lives under prohibition can drive over to the next state, go to a licensed shop and legally buy cannabis, then take it home where possession is still an offense.

Illinois shows how much movement this can bring. It launched its adult-use market in January 2020 while all of its surrounding states – Iowa, Indiana, Wisconsin, Kentucky, and Missouri – maintained prohibition. By 2022, out-of-state buyers made up 30.9 percent of adult-use purchases (IDFPR 2026), and a good portion of that cannabis had to travel across a state line to get home. Now this is where the problem begins. In the state where it's legal, the tax is charged on every sale, including sales to out-of-staters. The neighbor that has maintained its ban collects nothing, but it is the one left to police whatever returns across the border. It has not been resolved who is carrying that enforcement load.

Economics has only partially addressed this. One line of work shows that legal sales cluster toward these borders, with stores popping up near prohibition states and pulling real revenue from buyers who cross. Another finds that when a neighbor legalizes, drug arrests rise on the prohibition side. That is not the question here. The sales studies follow the legalizing state and end at the cash register, so they never see what enforcement does on the other side. The arrest studies do look at enforcement, but they group together borders that are very different and count arrests in bulk, which prevents them from saying anything about any particular neighbor, and they only look at drug arrests. No one has asked whether the effect runs past the drug category at all, into offenses that are not within the reach of drug policy. The evidence is thinnest there.

Here that question is taken up directly, in one place where it can be answered cleanly. The entire Illinois–Iowa border is the Mississippi River. Seven Iowa counties have a highway bridge going straight across into Illinois, which puts a legal store within a short drive for the people who live there. The other ninety-two counties have no such crossing and sit farther from supply. The treatment group is the bridge counties and the comparison group is the interior counties. The question then gets specific: Were there more recorded offenses in the bridge counties after Illinois opened its market in January 2020 than in the interior counties, and if so, was this increase only in the drug category or did it go beyond? Putting the two groups against each other isolates the effect of cheap, direct access to legal cannabis, holding the rest fixed. Every county in the comparison lies within prohibition Iowa, and shares the same laws, courts, and reporting. This design builds on the border-discontinuity logic of Holmes (1998) and Dube, Lester, and Reich (2010), but applies it to a river rather than an administrative line.

The data are county-year offense counts from the National Incident-Based Reporting System (NIBRS) for 2015 through 2022, covering all ninety-nine Iowa counties. The model is a difference-in-differences with county and year fixed effects. I do not use the cluster-robust standard errors alone, since only seven counties are treated. I read them alongside a wild cluster bootstrap and a continuous-dose specification that replaces the bridge indicator with distance-based exposure to Illinois supply. The results are presented in two parts. In the category most directly affected by the policy, the bridge counties experienced an increase of about 26 percent in recorded drug offenses relative to the interior after 2020 ($\beta = 0.233$, $p = 0.024$), confirming a cross-border response. The larger and less expected result is in property offenses, about 37 percent higher ($\beta = 0.312$, $p = 0.005$), a category that the policy does not itself generate. That property estimate is given central importance in the thesis, while staying cautious about the mechanism that produces it.

The rest of this thesis is organized as follows. Chapter 2 reviews the related literature. Chapter 3 describes the data, the Iowa setting, and the bridge-based treatment. Chapter 4 presents the empirical strategy. Chapter 5 reports the main difference-in-differences results, the event-study diagnostics, and a continuous-dose specification that follows the decay of the effect with distance from the Illinois market. Chapter 6 discusses the interpretation and limitations, and concludes.

2. Literature Review

This chapter reviews the literature most closely related to the Iowa bridge-county difference-in-differences design. The four sections of the discussion include: the economics of cross-border cannabis effects, the border-pair tradition that provides the identification logic, the fiscal-externality framework interpreting a cost borne by a non-legalizing neighbor, and difference-in-differences econometrics that disciplines inference under a single treatment date and few treated clusters. The final section identifies the gaps this thesis addresses.

2.1 Cannabis Cross-Border Economics

When a jurisdiction legally recognizes a good that its neighbor still bans, a price and legality gradient opens up at the border, and quantities are thus affected. Hansen, Miller and Weber (2020) have presented the clearest case of this mechanism on the supply side. Using the universe of Washington seed-to-sale transactions and a difference-in-discontinuities design around Oregon's October 2015 retail opening, they have found that retailers along the Oregon border lost about a third of their sales after Oregon opened, and they estimate that Washington has collected tens of millions of dollars in tax revenue from cross-border shoppers. This design shows the cross-border flow of legal sales, but it stops at the point of sale, and it does not observe what happens to enforcement inside the prohibition neighbor, which is the response this thesis sets out to measure.

Whether that flow translates into a measurable enforcement burden on prohibition neighbors is the question Hao and Cowan (2020) take up. Pooling county arrest counts from the Uniform Crime Reporting program across the states bordering Colorado and Washington, and contrasting border counties with interior counties through difference-in-differences, they find a sharp increase in adult marijuana-possession arrests in the border counties of prohibition neighbors, with no juvenile effect and no conclusive effect on sale, manufacture, driving-under-the-influence, or hard-drug arrests. Their study establishes that enforcement spillovers are measurable, but the identifying variation pools structurally different border geographies and rests on arrest aggregates rather than incident records. This leaves open whether the spillover reaches beyond drug arrests in a single, well-defined neighbor, the gap this thesis addresses.

Within-legalizing-state evidence sharpens the picture without resolving the spillover question. Dragone, Prarolo, Vanin, and Zanella (2019) combine a county difference-in-differences with a spatial regression discontinuity along the Washington–Oregon line and report reductions in rapes and thefts on the Washington side, alongside higher marijuana use and lower consumption of other drugs and alcohol. Their comparison runs between two jurisdictions that are themselves moving toward legalization. By construction, then, it cannot isolate the effect imposed on a neighbor that never legalizes. Brinkman and Mok-Lamme (2019) used a Denver geospatial panel and a demand-shift instrument to find that dispensaries reduce crime in their immediate vicinity but have no spillover effects on surrounding areas, so any relocation will be localized. The direction of the estimated effects in this literature is genuinely mixed: Wu, Wen and Wilson (2021) report increases in property and violent crime in Oregon after legalization, while Wu, Boateng and Lang (2020) report some crime reductions in the counties surrounding Colorado and Washington, and both are designs that lack a small-cluster inference from an origin state and predate the recent small-cluster inference literature. The broader public-health and public-safety effects of legalization are surveyed by Anderson and Rees (2023), and the channel through which cannabis access interacts with alcohol and traffic outcomes is documented by Anderson, Hansen, and Rees (2013), evidence that bears directly on the driving-under-the-influence margin examined later in this thesis. Outside the United States, de Assis et al. (2025) exploits distance from the Uruguayan border across Brazilian municipalities and finds rising seizures and violent and vehicle crime near the border after Uruguay's legalization, a reminder that cross-border effects can run toward intensified illicit activity rather than away from it.

2.2 Border-Pair and Contiguous-County Causal Designs

The methodological lineage for using a sharp sub-national boundary as identifying variation begins with Holmes (1998). Classifying states by right-to-work status and comparing manufacturing activity as one crosses state lines, Holmes documents a large abrupt jump in manufacturing at the border. The logic is that everything varying smoothly across space differences out at a sharp line, so a discontinuity in outcomes can be attributed to the discontinuity in policy. Dube, Lester, and Reich (2010) generalize this into the contiguous-county-pair design, comparing all county pairs straddling a state border and finding no adverse minimum-wage employment effects once local pair-specific conditions are absorbed. The approach is not uncontested. Neumark, Salas, and Wascher (2014) argue that adjacent counties are not always superior controls and that imposing them can discard valid identifying variation. That exchange defines what a clean border requires, namely that treated and control units share local trends and that no second policy changes discontinuously at the same line.

Two further studies establish that consumers respond to legality and price gradients at borders. Lovenheim (2008), using the Tobacco Use Supplement to the Current Population Survey, estimates that between 13 and 25 percent of consumers near a low-tax border purchase cigarettes across it, with the home-state price elasticity rising as distance to a cheaper border falls. Einav, Knoepfle, Levin, and Sundaresan (2014) use eBay transactions and find that a one-percentage-point sales-tax increase raises residents' online purchases by just under 2 percent while reducing their purchases from in-state sellers by 3 to 4 percent. Both papers show that cross-border substitution in response to tax and legality differentials is large and detectable in administrative or transaction records, the demand-side analogue of the cannabis question. The seven counties in the Iowa treated set are all located along a single sharp natural boundary and are, therefore, discrete realizations of the border-pair tradition. This design draws on the logic of border comparisons rather than a strict contiguous border-pair discontinuity, since its controls are Iowa interior counties rather than the Illinois side; what such a comparison cannot rule out is a second policy that also changes at the same boundary, a caveat that has a smaller weight for a single river crossing into a legalizing state than for a general multi-state border.

2.3 Fiscal Externalities and Cross-Border Enforcement Spillovers

The idea that one state can push the costs of its policy onto a neighbor comes from the public-finance literature on fiscal externalities. Wilson (1999), in his review of tax competition, makes the basic point that when one jurisdiction sets fiscal policy it imposes effects on others, so the decentralized outcome is usually not efficient. Keen and Konrad (2013) develop this further, showing how a policy change in one place shifts a mobile base into neighboring jurisdictions and why those external effects do not simply wash out. Enforcement in an Iowa bridge county can be read the same way, as a public-safety version of a fiscal externality: a cost created by Illinois's decision to legalize, carried by a neighbor that did not legalize, and unaccompanied by any of the tax revenue that Hansen, Miller, and Weber (2020) document on the Illinois side. Illinois and Iowa fit this picture well, one large and newly legal, the other smaller and still prohibiting.

The mechanism implied by this framework can be summarized as a short causal sequence. An asymmetric policy boundary opens a gap in the full price of cannabis, meaning the sum of its monetary price, the travel cost to legal supply, and the expected legal cost of acquisition. That full price falls most for prohibition-state residents with the cheapest crossing, so the model predicts cross-border acquisition concentrates in the lowest-crossing-cost counties, which are the bridge counties here. The same logic that the tax-competition literature has used for a mobile tax base can be applied to mobile consumers: the legalizing jurisdiction collects tax revenue on these cross-border purchases, but any enforcement costs generated by the returning consumer occur at home in a jurisdiction that did not choose the policy and cannot tax the activity that generated the cost. The externality is uninternalized for the usual reason, that is, the party setting the policy does not bear the full marginal cost. The following two empirical predictions can be tested in a bridge design. A response should appear in the offense category that is most directly related to the substance, and it serves as a test for whether the channel operates; since that category records enforcement events across multiple substances rather than consumption, the prediction refers to the existence of a response and not its size relative to other categories. A response that does not fall into one of the categories mechanically linked to the substance will be considered an externality extending beyond the domain of drug policy.

The criminology of spatial displacement studies how large an externality should be anticipated from such causes. Guerette and Bowers (2009) reviewed the situational crime prevention evaluations and found that displacement and diffusion of benefits each occurred in about a quarter of the cases; thus, targeted interventions usually achieved a net reduction in crime rather than merely shifting it elsewhere. This supports a measured reading of any border effect. Perrotta Berlin, Latour and Spagnolo (2023) have found that evidence of cross-jurisdictional regulatory spillovers supports this mechanism; that is to say, when a country raises the threshold for criminalization of prostitution, demand-driven activity moves to its neighbors with less stringent regulations. The cannabis case studied here is the reverse; deregulation in Illinois shifted both the location of activity and the scope of enforcement to stricter Iowa.

2.4 Difference-in-Differences Econometrics under Small Clusters and a Single Treatment Date

The recent literature on difference-in-differences has focused on two problems, and only one is suitable for this design. The first is heterogeneity and timing. Goodman-Bacon (2021) shows that the two-way fixed-effects estimator is a weighted average of all possible two-by-two comparisons, some of which use already-treated units as controls; de Chaisemartin and D'Haultfœuille (2020) make the same point through estimation weights; Callaway and

Sant'Anna (2021) and Sun and Abraham (2021) have proposed estimators and event-study corrections to recover interpretable parameters under staggered adoption. The Iowa design has a single common treatment date in January 2020; thus, there are no already-treated controls and no forbidden comparisons, so these estimators are cited here for completeness but are not required for the present design. Roth, Sant'Anna, Bilinski and Poe (2023) explicitly state that the scope condition is met, and two-way fixed effects are still interpretable when timing does not vary.

The second problem is also binding. Bertrand, Duflo and Mullainathan (2004) have shown that, due to the omission of serial correlation, placebo interventions are identified as statistically significant too often; thus, clustering by county over the panel is employed. Cameron and Miller (2015) report that with few clusters, the cluster-robust variance estimator over-rejects, and MacKinnon and Webb (2018) address the more severe problem of few treated clusters by showing that both the t-test and the ordinary wild cluster bootstrap can fail; therefore, procedures designed for such a regime are needed, and their work serves as the direct basis for the wild cluster bootstrap applied to county-clustered standard errors in this paper. On the parallel-trends assumption, Roth (2022) cautions that conditioning on a passed pre-trends test distorts inference, and Rambachan and Roth (2023) offer a sensitivity framework that bounds post-treatment violations by reference to observed pre-trends rather than assuming the assumption holds exactly. Together these papers define the inferential standard a seven-treated-cluster, single-date design should meet.

2.5 Gaps This Thesis Addresses

Place the four strands beside each other and a gap appears between them. Each one answers some of the debate, but no one has put the pieces together for a single state that remains under prohibition. The most obvious example is the work on sales. Hansen, Miller, and Weber (2020) find that legal sales shift to the border and that the legalizing state taxes the buyers who cross.

Their vantage point is the seller's side; they see the purchase and stop there, not asking what enforcement does once the cannabis is back home. Hao and Cowan (2020) do reach the far side and find drug arrests increase in prohibition-state border counties when a neighbor goes legal. But they pool borders that have little in common and work from arrest totals, so the estimate is an average over many different places rather than a result for any single neighbor, and it covers only drug arrests. The within-state studies do not close the gap either, as papers such as Dragone, Prarolo, Vanin, and Zanella (2019) compare two states that are both legalizing and can say nothing about one that prohibits throughout. That leaves one question unanswered: Does the effect ever extend beyond the drug category at all? If it does, if drug-unrelated offenses also increase, the burden has passed beyond drug enforcement into a cost the legalizing state neither records nor taxes, and this is the fiscal externality of Section 2.3 made real. This thesis asks that question of one clearly defined prohibition neighbor, Iowa. The property-offense estimate is the evidence that bears on that question. What drives that estimate I cannot fully pin down, and I make no pretense that I can.

3. Data and Institutional Setting

3.1 Illinois Legalization and the Iowa Prohibition Neighbor

Illinois opened its adult-use cannabis market on January 1, 2020, under the Cannabis Regulation and Tax Act (Public Act 101-0027, signed June 25, 2019, effective January 1, 2020). The initial license tranche converted existing medical dispensaries to adult-use operation in 2020; subsequent Social Equity Lottery rounds expanded the license pool through litigation-delayed rollout primarily in 2023 and 2024. Iowa maintains recreational prohibition throughout the 2015–2022 sample period, operating only a low-THC CBD medical program that is not a substitute for Illinois recreational cannabis. There is no within-Iowa policy change that contaminates the January 2020 treatment timing.

The Illinois–Iowa border is the Mississippi River. Seven Iowa counties contain at least one direct highway bridge across the Mississippi River into Illinois, verified using the Iowa Department of Transportation bridge inventory and confirmed against the USGS National Bridge Inventory. The seven bridge counties, from north to south along the Mississippi, are Dubuque (US-20 Julien Dubuque Bridge), Jackson (US-52 Sabula–Savanna Bridge), Clinton (US-30 Gateway Bridge), Scott (Interstate 74, Interstate 80, and US-67 Centennial Bridge in the Quad Cities), Muscatine (US-61 Norbert F. Beckey Bridge), Des Moines (US-34 Great River Bridge in Burlington), and Lee (US-136 Keokuk–Hamilton Bridge and US-61 Fort Madison Toll Bridge). One additional Iowa county, Louisa (FIPS 19115), borders the Mississippi River along approximately thirty miles of frontage but contains no same-county bridge to Illinois; Louisa residents reach Illinois by routing through the Beckey Bridge in Muscatine to the north. Louisa is therefore excluded from the bridge-based treated set and assigned to the control group (the exclusion is examined in detail in Section 3.3.2). The bridge-based treatment definition is predefined by physical infrastructure inventory and is independent of any regression output.

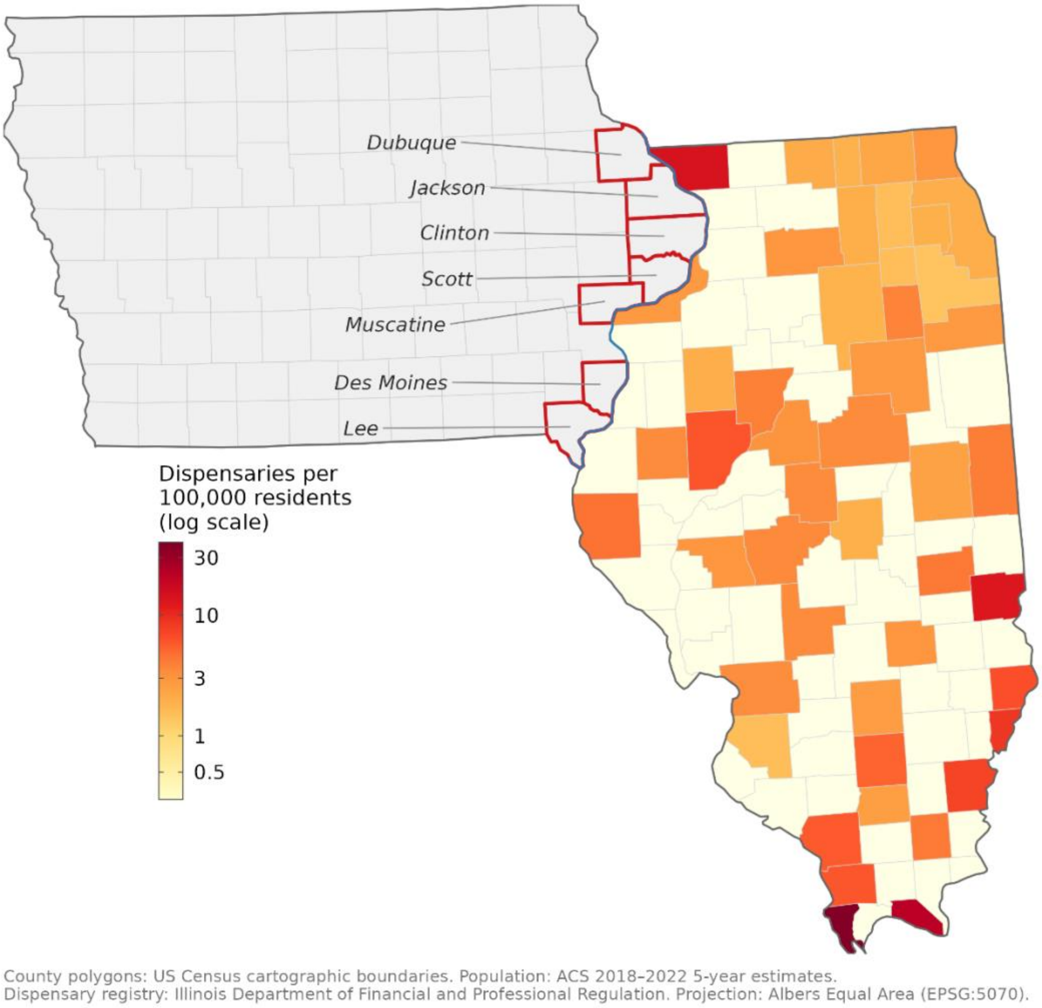

Figure 3.1: Illinois adult-use dispensary density per 100,000 residents, with Iowa bridge counties highlighted

Notes: County-level choropleth of Illinois adult-use dispensary licenses per 100,000 residents on a logarithmic color scale. The seven Iowa counties with direct highway bridges to Illinois (Clinton, Des Moines, Dubuque, Jackson, Lee, Muscatine, Scott) are outlined and labeled. Sources: IDFPR adult-use dispensary registry; US Census ACS 2018–2022 five-year population estimates; Iowa bridge inventory; US Census TIGER/Line 2020 county geometry. The map is the author's own elaboration based on these public data sources; © Yaotian Zhang, 2026, licensed under the same Creative Commons Attribution-NonCommercial-NoDerivatives (CC BY-NC-ND) 4.0 International license as the thesis.

3.1.1 Design Features and the Confounders They Address

This section states what each design choice does for identification and their respective costs. One point comes first. The control group is interior Iowa, not Illinois. So the comparison that must hold is between the bridge counties and the interior Iowa counties. Whether the bridge counties resemble the Illinois side does not matter, because Illinois counties are never used as controls. Each feature below is judged by that standard. Table 3.1 maps each feature to the confounder it addresses and the trade-off it brings.

Table 3.1: Design features of the Iowa case, the confounders each addresses, and the corresponding trade-offs

Design feature	Confounder addressed	Trade-off
Seven direct highway bridges across the Mississippi River	Arbitrariness in defining border treatment; non-comparable treated and control counties on contiguous land borders	Small treated cluster (n = 7); reduced statistical power; small-cluster inference requires wild cluster bootstrap robustness checks
Iowa retains recreational prohibition 2015–2022; CBD-only medical program is not a recreational substitute	Within-Iowa policy contamination of the January 2020 timing; concurrent treatment by an unrelated cannabis-policy change	External validity restricted to prohibition-state neighbors; results do not speak to neighbor-state effects under partial legalization or medical-cannabis programs
Treated and control counties share a common non-metropolitan profile within Iowa (largest treated county Scott, 174,000 residents)	Differences between treated and control counties that would otherwise threaten the parallel-trends assumption; metropolitan-specific shocks of the kind that would arise were the treated set to include a large metro such as Lake County, Indiana, in the Chicago area	Estimates are local to small and medium counties and do not identify effects at metropolitan-scale crossings; restricting the control set further to rural counties only would raise comparability at the cost of much noisier estimates
Mississippi River as natural rather than administrative boundary	Endogenous border placement; same-county-cluster jurisdictions spanning the administrative line	Effect identified only for highway-bridge-defined crossings; pedestrian, rail, or other modes of crossing are not captured

Two further points are listed as follows. First, homogeneity is not free. A more varied treated set would add identifying variation, and more variation can help. But with only seven treated counties, adding a markedly different unit, such as a large metropolitan crossing, would inject a metro-specific shock, not useful variation. The variation that does help here is in treatment intensity: how exposed each county is to Illinois supply. The continuous-dose specification in Section 5.5 uses exactly that. Second, the control group could be narrowed to rural Iowa counties only, to make it more close to the treated set. That would raise comparability but leave the estimates much noisier. The design as it stands keeps the treated and control groups comparable within Iowa while keeping enough precision to detect an effect. It is a deliberate trade of variation for a clean within-state comparison.

3.2 NIBRS Offense Data

The demand-side analysis uses the National Incident-Based Reporting System (NIBRS), accessed via the Kaplan Version 9 release (open ICPSR project 118281). NIBRS records county-year aggregated offense counts by category. The thesis focuses on three broad offense categories, drug, property, and violent, plus driving-under-the-influence (DUI) as an additional outcome. Drug, property, and violent offenses are counts from the NIBRS Group A offense segment. DUI is recorded only on arrest, in the NIBRS Group B segment, so the DUI outcome is an arrest count rather than an offense count; this difference is flagged where the four outcomes are compared. The outcome most directly tied to the policy is drug offenses, which Kaplan V9 records at the aggregated category level. The V9 release does not isolate cannabis-specific offenses at the county-year resolution, so the thesis frames the result as enforcement spillover broadly rather than as direct evidence of cannabis-specific enforcement. Even at the category level, the incident-based structure of NIBRS is what makes a single-neighbor design feasible. It resolves offenses by county and year, where earlier cross-border work has mostly leaned on arrest counts aggregated over many borders, and that resolution is what lets the bridge counties be set against the interior within a single state.

NIBRS county-year offense counts are matched to county-level Census ACS five-year demographic estimates using the LEAIC crosswalk (ICPSR study 35158). The Iowa panel includes all ninety-nine Iowa counties for the years 2015 through 2022, yielding 792 county-year observations. Population and income controls are taken from annual ACS five-year rolling county estimates, with each panel year matched to the corresponding ACS five-year window, and therefore vary at the county-year level and enter the difference-in-differences specification alongside county and year fixed effects.

One reporting-quality consideration warrants explicit and substantive acknowledgement. NIBRS coverage in Iowa changes over the sample period as agencies switch from Summary Reporting to incident-based reporting. The Kaplan V9 county-year aggregates absorb agency-level coverage changes within a county into the county-year offense count. Because the mismeasured object is the recorded offense count, which enters the model as the dependent variable rather than as a regressor, the consequence of this error depends on whether it is correlated with treatment. If reporting changes are uncorrelated with treatment status, the error is classical and is absorbed into the residual: it inflates the residual variance and the standard errors but leaves the DiD coefficient unbiased and consistent. The familiar attenuation toward zero would arise only if the error fell on a right-hand-side variable, which it does not here. If reporting adoption or agency coverage changed differentially across the seven bridge counties and the ninety-two interior counties, however, the DiD coefficient may partly capture reporting artifacts rather than enforcement changes. The 2018 Iowa NIBRS reporting-rotation visible in the Polk County interior 2018 dip and in several property-offense pre-period coefficients (Section 5.2) is one such artifact, though it appears in interior counties as well as in bridge counties. The robustness checks in Section 5.4 include a placebo-period test on the pre-2020 subsample, which would detect a coverage-rotation artifact systematic enough to drive the post-2020 coefficient. The differential-measurement concern is not fully eliminated by these robustness checks: a coverage change that occurs specifically in 2020 or 2021 and that affects bridge counties differently from interior counties would not be detected by the pre-2020 placebo. The full extent of agency-level reporting coverage by county and year is not separately observable in the Kaplan V9 aggregates, which prevents a more direct measurement-error correction within the present data. This limitation is revisited in Section 6.2.

3.3 Bridge-Based Treatment Definition

3.3.1 Iowa Bridge Inventory and the Seven Treated Counties

The treated set in the main specification is the seven Iowa counties listed in Section 3.1: Clinton (FIPS 19045), Des Moines (19057), Dubuque (19061), Jackson (19097), Lee (19111), Muscatine (19139), and Scott (19163). The control set is composed of all 92 other counties in Iowa, including Louisa County. The treatment variable $Bridge_c$ is 1 for the seven bridge counties and 0 otherwise. The post-treatment variable $Post_t$ is 1 for the years 2020-2022 and 0 for 2015-2019. The interaction term of the difference-in-differences is $Bridge_c \times Post_t$.

The seven bridge counties are small to mid-sized. The least populous is Jackson County, and the largest is Scott County, which contains the Quad Cities and has about 174,000 residents. The treated set includes mid-sized industrial areas such as Davenport (Scott County), Dubuque, and Burlington (Des Moines County), but none are major metropolitan areas on the scale of Lake County, Indiana. The Iowa treated set is relatively homogeneous compared with pooled multi-state alternatives, so it avoids the metropolitan-confounder problem that pooled designs face.

A distance-based alternative, for instance defining treatment as counties within an arbitrary mileage threshold of the Illinois border, would introduce researcher discretion in the cutoff choice. The bridge-based criterion avoids this by relying on the Iowa bridge inventory, a predefined physical-infrastructure variable enumerated independently of any analysis in this paper.

3.3.2 Louisa County Classification

Louisa County borders the Mississippi River along approximately thirty miles of frontage but has no direct highway bridge to Illinois. I therefore classify it as a control county in the baseline and report a sensitivity check including it as treated; the sensitivity-check coefficient is reported in the robustness battery (Table 5.2). The classification rests on a predefined physical-infrastructure criterion, the Iowa bridge inventory, rather than on river adjacency.

3.3.3 Travel-Cost Exposure Asymmetry

Bridge counties face substantially lower travel costs to legal Illinois cannabis supply than interior counties. The bridge access dummy can be interpreted as a binary proxy for the continuous travel-cost asymmetry. Three exposure indicators are calculated for each Iowa county centroid in this section: the great-circle distance to the nearest Illinois adult-use dispensary, the number of Illinois dispensaries within a specific radius, and a gravity-weighted accessibility index that aggregates inverse-distance weights across all 276 dispensaries in the registry retrieved through 2026 (the same registry discussed in Appendix A). Illinois dispensary coordinates are geocoded to the street-address level from the IDFPR registry.

Table 3.2: Travel-cost exposure to Illinois dispensary supply by treatment status

Exposure metric	**Bridge counties (n = 7)**	**Interior counties (n = 92)**
Distance to nearest IL dispensary (miles)	21.5	154.1
IL dispensaries within 50 miles	4.6	0.2
IL dispensaries within 100 miles	23.9	2.0
Gravity-weighted accessibility (interior mean = 1)	5.0	1.0

Notes: Distances measured great-circle in miles from each Iowa county centroid to dispensary geocoded street addresses. Dispensary count within radius r is the number of IDFPR-registered adult-use dispensaries whose geocoded location lies within r miles of the Iowa county centroid. Gravity access at decay parameter α is the sum across all 276 dispensaries of $1/(\text{distance} + 1)^{\alpha}$, shown in the table normalized to the interior-county mean. All differences are significant at $p < 0.001$; distance gap has $t = 15.22$ with $p = 3.29 \times 10^{-22}$.

The exposure asymmetry is substantial across every metric. Bridge counties are on average 7.2 times closer to the nearest Illinois dispensary than interior counties (21.5 versus 154.1 miles), with the difference significant at $p < 10^{-21}$. Bridge counties have 4.6 dispensaries within fifty miles compared with 0.2 for interior counties, and 23.9 within one hundred miles compared with 2.0. The gravity-weighted accessibility index, which aggregates exposure across all dispensaries with inverse-distance weights, is 5 times higher for bridge counties under inverse-square decay. The relevant economic margin for cross-border consumption is the travel cost to the nearest available legal supply, not the absolute count of dispensaries near the border. Bridge counties have substantially lower travel costs to legal supply on every exposure metric, and the binary bridge indicator captures this multi-dimensional exposure asymmetry in a single predefined geographic criterion.

4. Empirical Strategy

4.1 Difference-in-Differences Specification

The baseline difference-in-differences specification is:

$$\log(Y_{ct} + 1) = \alpha_c + \lambda_t + \tau \cdot (Bridge_c \times Post_t) + \gamma_1 \cdot \log(Pop_{ct}) + \gamma_2 \cdot \log(Inc_{ct}) + u_{ct} \qquad (4.1)$$

where Y_{ct} is one of four NIBRS outcomes in Iowa county c in year t, α_c is a county fixed effect, λ_t is a year fixed effect, $Bridge_c$ is the seven-county treated indicator defined in Section 3.3, $Post_t$ equals one for years 2020 and later, and log(Pop) and log(Inc) are time-varying county controls from ACS five-year rolling estimates. The coefficient τ on the interaction is the difference-in-differences estimate of interest. The transformation log(Y + 1) handles the small number of county-year cells with zero offenses in a given category while preserving the multiplicative interpretation of the coefficient.

The $Bridge_c$ indicator on its own is absorbed by the county fixed effect α_c because $Bridge_c$ is time-invariant. The $Post_t$ indicator on its own is absorbed by the year fixed effect λ_t because $Post_t$ is uniform across counties within each year. Only the interaction term survives the fixed-effects absorption, which is the standard two-way fixed-effects DiD construction.

The time-varying log-population and log-income controls warrant explicit justification. Both variables are constructed from ACS five-year rolling estimates, which by their construction smooth over annual fluctuations and reflect demographic levels rather than acute post-2020 shocks. The five-year rolling window prevents either variable from responding sharply to a single-year event such as legalization. The inclusion of these controls is intended to absorb slow-moving demographic differences across counties that the county fixed effect cannot remove if they trend over time. A concern remains that post-2020 migration or income shifts may be correlated with treatment status, in which case the controls would partly absorb the treatment effect and bias the estimate downward.

4.2 Why a 2×2 Design

The Iowa setting admits a single binary treatment indicator and a single treatment date. With one treatment cohort and one timing, the design reduces to the canonical 2×2 difference-in-differences for which two-way fixed effects deliver a clean average treatment effect on the treated under parallel trends. The recent staggered-DiD econometrics literature (Goodman-Bacon 2021; Callaway and Sant'Anna 2021; Sun and Abraham 2021) is concerned with weighted-average pathologies that arise when treatment timing varies across units; these pathologies do not apply to the Iowa 2×2 design. The methodological discipline imposed by the recent literature is therefore satisfied by construction in this setting.

Two alternative designs were considered and set aside. A synthetic control built for each bridge county would require a long and stable pre-period to fit donor weights; the five pre-treatment years available here (2015–2019) are too few to discipline that weighting, and seven treated units are better served by a single panel estimator than by seven separate synthetic series. A spatial regression discontinuity across the Mississippi, in the manner of Dragone et al. (2019), would compare the Iowa bank against the Illinois bank; but the Illinois side is itself the treated, legalizing jurisdiction, so a cross-river discontinuity would difference one treated unit against another rather than against an untreated control. The within-Iowa comparison adopted here instead holds the legal regime fixed on the control side, since interior Iowa counties never legalize and therefore isolate the bridge-county response to the Illinois shock.

4.3 Identifying Assumption and Parallel-Trends Test

The identifying assumption is that, in the absence of Illinois legalization, drug-enforcement activity in the seven Iowa bridge counties would have evolved in parallel with the ninety-two interior counties. The assumption is not testable directly but is testable indirectly through pre-period coefficients in an event-study specification. The event-study specification replaces the static interaction $\text{Bridge}_c \times \text{Post}_t$ with a set of event-time interactions $\text{Bridge}_c \times 1\{\text{year}_{rel} = k\}$, where $year_{rel}$ = year − 2020 and k ranges over {−5, −4, −3, −2, 0, +1, +2}. The reference period is k = −1 (2019), the calendar year immediately before treatment. In plain terms, the assumption asks the bridge and interior counties to have been on the same path before 2020. If they were, the divergence after 2020 is what legalization is credited with, and the pre-period coefficients are the visible check on whether that common path actually held.

Under the null hypothesis of parallel pre-trends, the pre-period coefficients β_{-5} through β_{-2} should all be statistically indistinguishable from zero. The joint Wald test of this null yields F = 0.86 with p_{joint} = 0.486 for drug offenses, failing to reject parallel pre-trends at any conventional significance level. The corresponding tests for the other three outcomes yield F = 1.49 (p = 0.203) for property offenses, F = 1.38 (p = 0.241) for violent offenses, and F = 0.55 (p = 0.70) for DUI. None of the three tests rejects parallel pre-trends. Section 5.2 displays the full event-study coefficient profiles and discusses the property-offense pre-period pattern in detail, including the individually-significant k = −2 coefficient that does not propagate to the joint test.

4.4 Standard Errors and Small-Cluster Inference

Standard errors are clustered at the county level with HC1 finite-sample correction (sandwich::vcovCL). Clustering at the county level addresses within-county serial correlation in the residuals (Bertrand, Duflo, and Mullainathan 2004; Cameron and Miller 2015), which is the dominant violation of independent-identically-distributed errors in a county-year panel where the same enforcement environment, agency reporting practice, and demographic structure persist across years within a county.

The Iowa design has only seven treated clusters. Cameron and Miller (2015) demonstrate that with few treated clusters, the conventional cluster-robust variance estimator under-estimates standard errors and over-rejects the null hypothesis of no effect. Two independent inference channels are reported to corroborate the baseline cluster-robust inference. First, a wild cluster bootstrap inference following MacKinnon and Webb (2018), with 9999 Rademacher iterations and the null imposed, is reported in Appendix B. Second, the continuous-dose specification in Section 5.5 uses identifying variation across all ninety-nine Iowa counties rather than the seven treated clusters, providing a complementary robustness check, although it does not eliminate the identifying assumptions of the main design.

Together with the baseline cluster-robust standard errors, these channels are mutually consistent (Appendix B), so the small-cluster concern is addressed in three independent ways rather than through a single inference procedure.

5. Main Results

5.1 Baseline Estimates Across Four Outcomes

The baseline estimates are in Table 5.1, which fits equation (4.1) one outcome at a time across the four NIBRS categories: drug, property, and violent offenses, and driving under the influence (DUI). The comparison underneath is simple. It sets the seven Iowa counties with a bridge to Illinois against the ninety-two interior counties, and tracks how each group changed from before January 2020 to after. County fixed effects absorb whatever is constant in a county, its size, its courts, its reporting practices. Year fixed effects absorb whatever moved across all of Iowa in a given year. The coefficient τ on the Bridge × Post term is what remains once both are removed, the part of the change in bridge-county offenses that follows treatment and nothing else. Log population and log income enter as controls; the panel spans all ninety-nine counties from 2015 to 2022, or 792 county-year cells, and the standard errors are clustered by county. Because the outcome is in logs, each coefficient is roughly a proportional change, and the final column makes that concrete by reporting $\exp(\beta) - 1$, the percent by which bridge-county offenses moved relative to the interior after 2020. A single row, then, answers one question: how far that category of offense rose or fell on the bridge side once Illinois opened its market.

Table 5.1: Iowa baseline DiD across four NIBRS outcomes

Outcome	β	SE	p (HC1)	Approx. % effect	Significance
Recorded drug offenses	0.233	0.103	0.024	+26.3%	**
Recorded property offenses	0.312	0.111	0.005	+36.6%	***
Recorded violent offenses	0.154	0.103	0.134	+16.7%	ns
Recorded DUI arrests	0.078	0.129	0.545	+8.1%	ns

Notes: Iowa panel 2015–2022 (792 county-year observations, 7 treated bridge counties, 92 control). County and year fixed effects, controls for log population and log income. Standard errors are cluster-robust at the county level with HC1 finite-sample correction. Approximate percent effect is $\exp(\beta) - 1$*. The driving-under-the-influence (DUI) row is estimated on the same baseline specification as the other outcomes. Significance codes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$, ns = not significant at conventional levels.*

The four estimates are easier to read as a set than one at a time. Drugs are the natural place to begin, since a cross-border effect should appear there before anywhere else. A store opens across the river, the closest Iowa residents go and buy, and whatever enforcement follows takes place after they are home. The estimate matches that account. After 2020, recorded drug offenses in the bridge counties stand about 26 percent above the interior ($\beta = 0.233$, $p = 0.024$), significant at five percent. I hold the magnitude loosely, because the NIBRS drug category mixes cannabis with other controlled substances and records enforcement rather than sales; what it establishes is that activity crosses the river, not how much. The sign and the significance, though, are exactly what the channel predicts. The property estimate is the larger and the more revealing of the two. Recorded property offenses run about 37 percent above the interior ($\beta = 0.312$, $p = 0.005$), significant at one percent and bigger than the drug effect, and it is the figure the thesis rests on. Legalization has no built-in path to more property offenses, so when the strongest response appears in a category the policy does not reach, the cost of the policy has carried past drug enforcement into ground the legalizing state's tax accounts never cover. It is also the estimate I treat most carefully, and Section 6.2 returns to it in full, since a shock to these small industrial river counties around 2020 could have raised property offenses for reasons unrelated to cannabis.

The remaining two outcomes carry less, though both lean the same way. Violent offenses move with drugs and property but fall short of significance ($\beta = 0.154$, $p = 0.134$). Seven treated counties is little to work with, so the estimate is noisy and only a large effect would show through; a faint result is also about what one would expect if cannabis raises petty offending more than violence. DUI is flat ($\beta = 0.078$, $p = 0.545$), and that outcome never carried a clear prediction, since cannabis can substitute for alcohol behind the wheel. The weight of the analysis therefore falls on drugs and property, with violent offenses and DUI in the table to keep the full set in view. Figure 5.1 makes the result easy to take in at once. With the four coefficients plotted against their ninety-five percent intervals, the division is plain: drug and property clear zero, while violent and DUI sit across it.

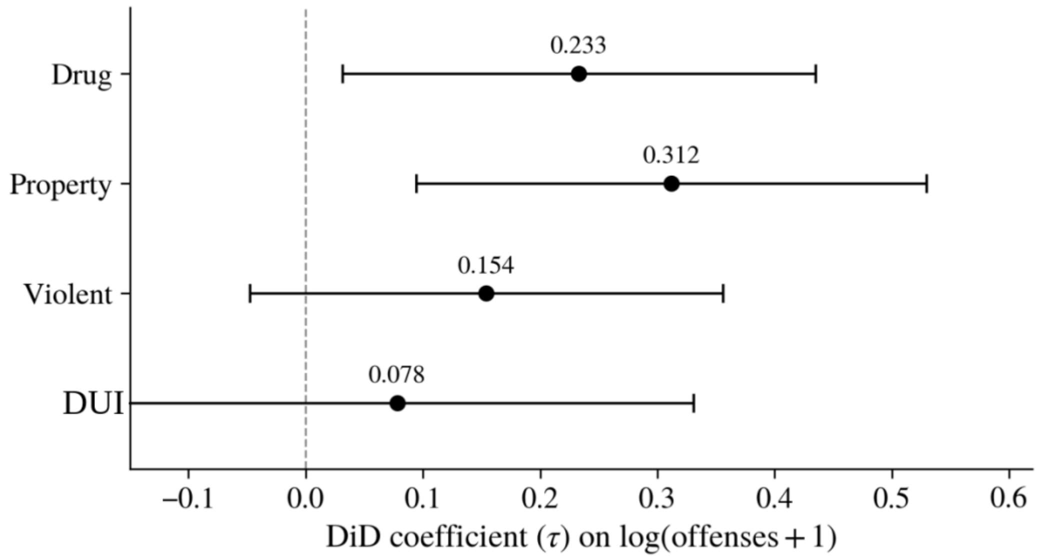

Figure 5.1: Baseline DiD coefficients across the four NIBRS outcomes

Notes: Points are the difference-in-differences coefficients τ from the baseline specification of equation (4.1), estimated separately for each NIBRS outcome on the 2015–2022 Iowa county-year panel (792 observations). Horizontal bars are ninety-five percent confidence intervals from county-level cluster-robust standard errors with HC1 finite-sample correction. The dashed vertical line marks a zero effect. The drug and property intervals exclude zero; the violent and driving-under-the-influence intervals do not.

5.2 Event Study and Parallel-Trends Test

Figure 5.2 plots the event-study coefficients for all four NIBRS outcomes with ninety-five percent confidence intervals, allowing the treatment effect to vary by event-time k relative to the January 2020 treatment date.

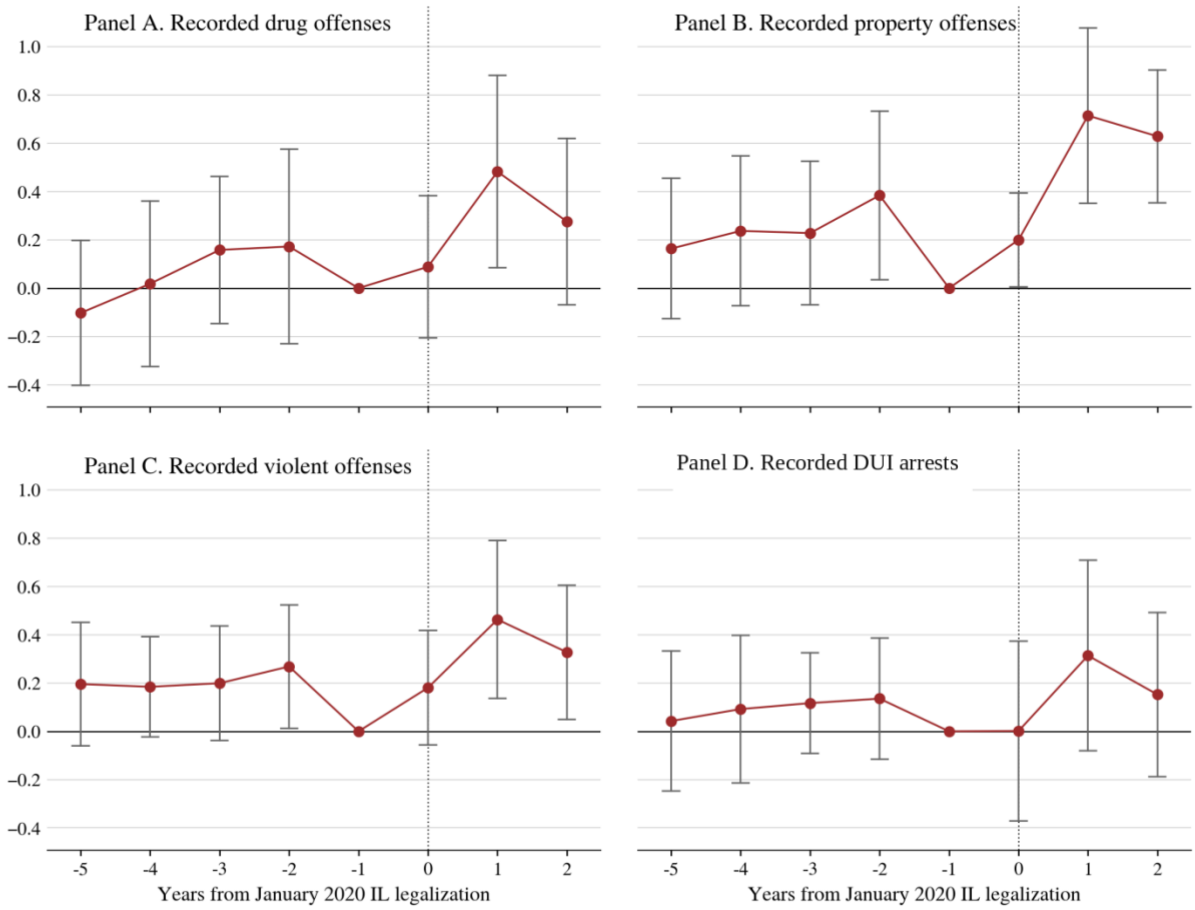

Figure 5.2: Event-study coefficients for the four NIBRS outcomes (Iowa bridge counties vs. interior counties)

Notes: Coefficients β_k from the event-study form of equation (4.1), with k = -1 (year 2019) the omitted reference and k = 0, the dotted vertical line, the year 2020. Outcomes are the log of NIBRS recorded drug offenses (Panel A), property offenses (Panel B), violent offenses (Panel C), and DUI arrests (Panel D), each plus one. Grey bars are 95 percent confidence intervals from county-level cluster-robust standard errors with HC1 finite-sample correction. The joint Wald test of the four pre-period coefficients fails to reject parallel trends for every outcome: drug F = 0.86 (p = 0.49), property F = 1.49 (p = 0.20), violent F = 1.38 (p = 0.24), DUI F = 0.55 (p = 0.70). The property and violent leads that reach significance individually at k = -2 (2018) coincide with the NIBRS reporting-rotation discussed in Section 3.2 and do not propagate to the joint tests.

The four pre-period coefficients are statistically indistinguishable from zero. The largest in magnitude is β_{-2} = +0.173 (SE = 0.205); all four confidence intervals overlap zero, with individual p-values between 0.31 and 0.92. The joint Wald test reported in Section 4.3 confirms this for drug offenses (F = 0.86, p_{joint} = 0.486).

Post-2020 coefficients exhibit a delayed step increase. The treatment-year coefficient at k = 0 (year 2020) is small and not statistically distinguishable from zero (β_0 = 0.089, SE = 0.150, p = 0.555). The most likely explanation for the muted first-year response is the gradual ramp-up of Illinois licensed retail, which started from a small set of converted medical dispensaries and kept the available legal supply low through 2020. General pandemic disruption to travel and commerce may also have contributed. Neither channel is verified here against actual cross-border traffic. The k = +1 (2021) coefficient rises to $\beta_{+1} = 0.483$ (SE = 0.203, p = 0.018), statistically significant at the five-percent level. The k = +2 (2022) coefficient is $\beta_{+2} = 0.276$ (SE = 0.176, p = 0.116), positive and economically consistent with the static DiD estimate but not individually significant.

The concentration of the post-treatment effect at k = +1 (2021) rather than at k = 0 (2020) is the most consequential feature of the event study, and it admits of two readings that the present data cannot fully separate. The first is a legalization-specific delay: Illinois adult-use retail expanded gradually from its January 2020 opening, and cross-border awareness and stable supply would plausibly have arrived only in 2021. The second is a contemporaneous 2021 shock to the bridge counties unrelated to cannabis. Two observations weigh against the second reading without disposing of it. The effect persists into 2022 ($\beta_{+2} = 0.276$) rather than reversing, which a transitory 2021 disturbance would not predict, and the identifying variation is the bridge-versus-interior differential rather than any statewide movement, while Iowa-wide drug offenses in fact declined after 2020 with interior counties falling faster than bridge counties (Section 6.2). What would settle the question is direct evidence on cross-border movement, such as mobility traces distinguishing 2020 from 2021 border traffic, which this design does not have. The timing is therefore reported as consistent with a delayed legalization response, not as established to be one, and the limitation is taken up in Section 6.2.

Property offenses exhibit a similar pre-period profile that warrants additional discussion. The pre-period coefficients rise toward k = −2, where $\beta_{-2} = +0.385$ (p = 0.031) is individually significant at the five-percent level, although the joint Wald test does not reject parallel pre-trends (F = 1.49, $p_{joint} = 0.203$). The 2018 spike is consistent with the Iowa NIBRS reporting-rotation artifact discussed in Section 3.2: agency-level coverage changes in 2018 affected several Iowa counties, including non-bridge interior counties, and the pattern is not specific to the bridge treated set. Post-2020, the property-offense event-study coefficients are $\beta_0 = +0.200$ ($p = 0.044$), $\beta_{+1} = +0.715$ ($p < 0.001$), and $\beta_{+2} = +0.629$ ($p < 0.001$), with β_{+1} and β_{+2} substantially larger than any pre-period coefficient. The 2018 individual significance therefore reflects a coverage artifact rather than a pre-existing differential trend in bridge counties, and the joint test correctly fails to reject parallel pre-trends overall.

The violent and DUI event studies complete the four-outcome picture and are read with more caution. For violent offenses the joint pre-trend test does not reject ($F = 1.38$, $p = 0.24$), but the pre-period coefficients are uniformly positive and the $k = -2$ (2018) lead is individually significant ($\beta = 0.269$, $p = 0.040$), the same 2018 pattern that appears in property and is consistent with the reporting rotation of Section 3.2. The post-period violent coefficients are positive, and the $k = +1$ (2021) estimate is individually significant ($\beta = 0.463$, $p = 0.006$), even though the static violent DiD is not (Section 5.1); the static estimate averages the muted 2020 coefficient with the larger 2021 and 2022 coefficients, and with only seven treated clusters the violent series is noisy enough that the two are not in tension. DUI arrests show the cleanest pre-period of the four ($F = 0.55$, $p = 0.70$) and a positive but insignificant post-period ($\beta_{+1} = 0.314$, $p = 0.119$), which is consistent with the prior ambiguity of the DUI sign once alcohol-to-cannabis substitution at the impaired-driving margin is taken into account.

The event-study evidence is broadly consistent with the parallel-trends assumption behind the design: before 2020 the bridge counties track the interior, and they diverge from it only afterward. The post-2020 pattern is a delayed step increase rather than an immediate January 2020 jump or a transitory spike, although the 2018 property and violent leads require caution and are interpreted in light of the NIBRS reporting-rotation issue discussed above.

5.3 Robustness Battery

With only seven treated clusters, conventional cluster-robust inference may over-reject the null hypothesis of no effect (Cameron and Miller 2015). Table 5.2 reports the robustness battery for the drug result, which is the cleanest single outcome and the simplest case to anchor inference on. Property robustness is established separately through three channels: the wild cluster bootstrap p-value of 0.010 in Appendix B (Table B.1), the continuous-exposure specifications across all four outcomes in Table 5.3, and the continuous-dose specifications across three radii in Table 5.4. The event-study post-period coefficients in Section 5.2 are also part of this evidence. Read together, these channels mean the headline results do not rest on any single inference procedure, which is the standard the small-cluster literature sets for a design with few treated units.

Table 5.2: Robustness battery for Iowa drug-offense DiD

Specification	β	SE	p	Notes
(1) Baseline (7 bridge, FE+controls)	0.233	0.103	0.024	baseline
(2) Drop 2020 (COVID exclusion)	0.275	0.118	0.022	stronger without 2020
(3) 8-bridge sensitivity (include Louisa)	0.422	0.198	0.033	Louisa adds variance
(4) Wild cluster bootstrap (Appendix B)	—	—	0.015	B = 9999, null imposed

Notes: All specifications estimate the Iowa drug-offense DiD with county and year fixed effects, log population, and log income controls, on the 2015–2022 Iowa panel (792 observations) unless otherwise noted. Specs (1)–(3) are alternative coefficient specifications; Spec (4) reports the wild cluster bootstrap p-value at the baseline point estimate (full details in Appendix B). WCB = wild cluster bootstrap with 9999 Rademacher iterations and the null imposed.

The specification dropping the year 2020 from the sample yields β = 0.275 with p = 0.022, modestly larger than the baseline and consistent with COVID disruption attenuating rather than driving the drug effect. The eight-bridge sensitivity check, which adds Louisa County to the treated set, yields a substantially larger point estimate (β = 0.422) with a wider standard error (0.198), consistent with the discussion in Section 3.3.2: Louisa contributes high variance without bridge-driven exposure.

The HC1 cluster-robust p-value of 0.024 and the wild cluster bootstrap p-value of 0.015 (Appendix B) are the two primary inference channels and are mutually consistent. The wild cluster bootstrap does not rely on cluster-SE asymptotic approximations and is the appropriate inference procedure under the Cameron-Miller (2015) small-cluster concern.

As a further check that the discrete bridge definition is not driving the results, Table 5.3 replaces the binary indicator with three continuous measures of exposure to Illinois supply, each interacted with the post-2020 indicator and estimated across all four outcomes. The continuous-exposure coefficients agree in sign and broad significance with the binary specification reported in the final row, and the property and violent margins remain the most strongly associated with proximity to legal supply.

Table 5.3: Alternative continuous-exposure specifications across NIBRS outcomes

Exposure variable × Post	ln(drug)	ln(property)	ln(violent)	ln(DUI)
−log(distance to nearest IL disp)	0.075 (0.072)	0.161** (0.071)	0.183*** (0.070)	0.002 (0.068)
log(count within 50 mi + 1)	0.169* (0.102)	0.192*** (0.072)	0.168** (0.070)	0.018 (0.078)
log(gravity access α = 1)	0.108 (0.178)	0.335** (0.166)	0.424*** (0.162)	−0.041 (0.169)
[Reference: Binary 7-bridge × Post]	0.233** (0.103)	0.312*** (0.111)	0.154 (0.103)	0.078 (0.129)

*Notes: Each row reports the coefficient on the indicated continuous exposure measure interacted with the post-2020 indicator, estimated on the 2015–2022 Iowa panel (792 county-year observations) with county and year fixed effects and log-population and log-income controls. Cluster-robust standard errors at the county level in parentheses. The final row reproduces the binary seven-bridge specification of Table 5.1 for comparison. Significance codes: *** p < 0.01, ** p < 0.05, * p < 0.1.*

5.4 Placebo Treatment Dates

As a non-redundant complement to the parallel-trends event study, placebo-period tests replace the true 2020 treatment date with fake treatment dates of 2017, 2018, and 2019, estimated on the pre-2020 subsample (2015 to 2019). Zero of three placebo estimates are statistically significant ($\beta = 0.173$ with $p = 0.103$ for fake 2017; $\beta = 0.053$ with $p = 0.614$ for fake 2018; $\beta = 0.210$ with $p = 0.400$ for fake 2019). The placebo evidence is consistent with the parallel-trends event study in Section 5.2 but does not add identifying power beyond it; the event study is the primary diagnostic and the placebo dates are reported here for transparency. Pandemic-recovery concerns specific to 2021 are not addressed by pre-2020 placebos and are taken up in Section 6.2.

5.5 Mechanism Evidence: Continuous-Dose Spatial Decay

The binary bridge treatment in Chapter 5 partitions Iowa counties into a treated set of seven counties and a control set of ninety-two counties. This section reports a continuous-dose specification that replaces the binary indicator with each county's resident population within radius r of Illinois, providing mechanism evidence on the spatial decay of the effect with distance from the Illinois market. Because it relies on a different source of variation from the binary bridge design, it is a supporting analysis rather than a second identification of the main effect.

The dose variable is constructed in three steps. First, every Iowa county centroid is established using Census Centers of Population 2020. Second, for each Iowa county and each radius $r \in \{15, 50, 100\}$ miles, the resident population of Illinois block groups whose centroid lies within radius r of the Iowa county centroid is summed using ACS 2018–2022 five-year block-group population estimates and the sf::st_is_within_distance() function on EPSG 5070 projected geometry. Third, the dose variable is $\log(Pop_r + 1)$, a continuous county-level exposure measure on the same logarithmic scale as the outcome variables. Illinois population is used instead of dispensary count because dispensary locations are themselves endogenous to the time of legalization; Illinois population is exogenous to the treatment time and serves as the underlying market foundation.

The continuous-dose DiD specification is:

$$\boldsymbol{log(Y_{ct} + 1) = \alpha_c + \lambda_t + \delta_r \cdot (dose_{cr} \times Post_t) + \gamma_1 \cdot log(Pop_{ct}) + \gamma_2 \cdot log(Inc_{ct}) + u_{ct}} \qquad \textbf{(5.1)}$$

where $dose_{cr}$ is the continuous exposure measure at radius r and $Post_t$ is the post-2020 indicator. The coefficient δ_r measures the marginal effect of one unit of continuous exposure on the log of offense counts in the post-2020 period. Standard errors are cluster-robust at the county level. The specification is estimated separately for each radius r and for each NIBRS outcome.

Table 5.4: Continuous-dose DiD coefficients across radii

Outcome	r = 15 mi	r = 50 mi	r = 100 mi
Drug-offense records	0.043** (0.021)	0.020 (0.014)	0.016 (0.011)
Property offenses	0.042*** (0.016)	0.026** (0.011)	0.026** (0.011)
Violent offenses	0.025 (0.016)	0.025** (0.010)	0.030*** (0.010)

*Notes: Iowa panel 2015–2022 (792 county-year observations, all 99 counties). County and year fixed effects, controls for log population and log income. The dose variable at radius r is $log(Pop_r + 1)$, the natural logarithm of the Illinois block-group population within r miles of the Iowa county centroid as constructed in Section 5.5. Cluster-robust SE at county level in parentheses. Significance codes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.*

The drug coefficient attenuates monotonically across radii, with the effect dissipating as the operative population becomes more distant. Figure 5.3 plots the continuous-dose coefficients against radius for all three offense categories with their confidence intervals.

The attenuation of the drug-enforcement effect from 15 to 50 miles is in line with the institutional pricing structure in Illinois. Illinois imposes one of the higher cannabis taxes in the country, combining a tiered THC-based excise with state and local sales taxes. Therefore, the actual price premium over the Iowa illicit market is positive but small. Given the tax friction, cross-border arbitrage is economically feasible only at a short distance where the round-trip travel cost is still relatively low compared to the tax premium. The attenuation at r = 15 miles ($\beta = 0.043$) and r = 50 miles ($\beta = 0.020$) is consistent with this institutional structure: at a short distance, the price difference can cover travel costs; at a moderate distance, frictions such as gas, time and taxes make arbitrage unprofitable. These results are in line with other research on cross-state tax arbitrage, such as Lovenheim's (2008) work on cigarettes and Einav, Knoepfle, Levin and Sundaresan's (2014) study of consumer reactions to tax differences among areas.

Property offenses exhibit a more diffuse spatial profile than drug enforcement. The property-offense coefficient remains statistically significant at all three radii (β_{15} = 0.042, β_{50} = 0.026, β_{100} = 0.026), with the attenuation concentrated between 15 and 50 miles and stable from 50 to 100 miles. The flatter spatial gradient may reflect a broader spillover channel that operates at longer distances than direct cross-border purchase, for example downstream effects on local property-offense markets or enforcement reallocation toward broader drug-finance investigations. The violent-offense coefficient strengthens with radius (β_{15} = 0.025, β_{50} = 0.025, β_{100} = 0.030), an opposite pattern that may reflect noise at the tight radius given the smaller treatment-intensity variation among bridge counties.

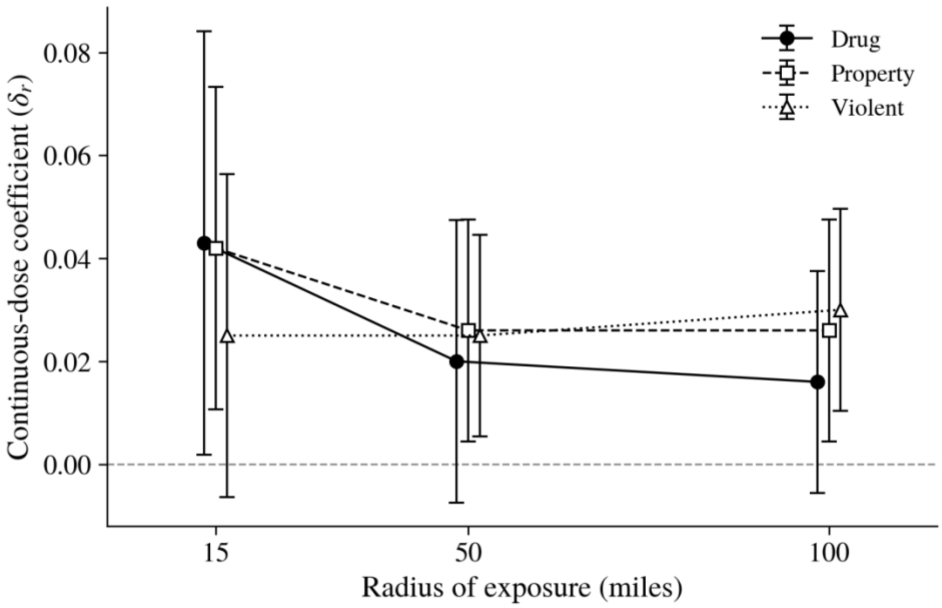

Figure 5.3: Continuous-dose coefficients by radius of exposure to Illinois supply

Notes: Continuous-dose coefficients δ_r from equation (5.1), estimated separately for each NIBRS outcome and each radius r in {15, 50, 100} miles on the 2015–2022 Iowa county-year panel. The dose at radius r is the log of the Illinois block-group population within r miles of the Iowa county centroid. Bars are ninety-five percent confidence intervals from county-level cluster-robust standard errors.

6. Discussion and Conclusion

6.1 Interpretation

There are two results coming from the four estimates, and they are different in nature. One is confirming something known, the other is the real contribution. Let us start with the confirmation. Hao and Cowan (2020) found that drug enforcement increases on the prohibition side of a newly legal border, and Iowa's bridge counties do the same, with drug offenses roughly 26 percent above the interior after 2020 ($\beta = 0.233$, $p = 0.024$). What is new is not the discovery but the precision: the estimate is based on a single unambiguous neighbor and incident-level records, whereas the earlier work combined different borders and used arrest figures. But I keep the magnitude at arm's length, because NIBRS lumps several drugs together, and the coefficient shows activity crosses the river, not how much of it does.

It is in the property result that the paper goes beyond mere confirmation. Recorded property offenses are about 37 percent higher than the interior in the bridge counties ($\beta = 0.312$, $p = 0.005$), a larger effect than the drug one and in a category legalization has no direct way of producing. No study of a prohibition-side neighbor has asked whether the spillover extends this far from the drug column. The finding that it does is the paper's central contribution. It is just the concrete form of the fiscal externality of Section 2.3. Illinois's decision does not stop with drug enforcement on the far bank; it appears in offenses the policy never intended and the legalizing state never has to answer for, while Illinois reaps the tax revenue and an Iowa county that had no role in the decision bears some of the cost.

The other two results are of the same shape but are of less importance. Violent offenses go in the same direction but are not significant ($\beta = 0.154$, $p = 0.134$) – there are too few counties (seven) treated to read much into it, and a small effect is about what you would expect if cannabis raises minor offending more than serious violence. DUI is flat ($\beta = 0.078$, $p = 0.545$), the one outcome for which there was no clear prediction in advance since cannabis can substitute for alcohol behind the wheel. The argument is drugs and property. Violent offenses and DUI are reported to keep all four categories on record.

Three further pieces of evidence support a reading in which the driving force is consumption that follows the highway into Illinois. The drug coefficient in Section 5.5 decreases as the radius increases, the gradient one would expect if the frequency of people making the trip depends on how close the store is to them. In Section 5.4 the placebo dates find the response in 2020 and not earlier, as no pre-treatment cutoffs produce a significant effect. And Appendix A rules out the simplest supply-side explanation, because the dispensaries near the bridge counties were not located directly across the river, so the result does not resemble supply located to capture Iowa demand. The simplest story is that Iowans went to stores that were already there.

6.2 Limitations

The most contested estimate is property, and so it is useful to outline how far the design meets that challenge. The concern is real. The bridge counties are not some random piece of Iowa: Scott County is in the Quad Cities and Des Moines County is home to Burlington, river cities whose economies are different from the farm counties inland. A local downturn that raised theft there after 2020 could pass for a legalization spillover. The design was constructed to remove most of that concern from the table. Bridges determined treatment status before any outcomes were observed, preventing treatment from drifting to counties that just happened to see increases in offenses. The placebo dates rule out a trend that was already in place before 2020. And the event study finds the effect after the opening, not before: the joint pre-trend test passes ($F = 1.49$, $p = 0.20$) and the one significant lead, in 2018 ($\beta = 0.385$, $p = 0.031$), lines up with a NIBRS reporting change that appears in interior counties as well, not anything specific to the bridge set. Two things are genuinely open. The property offense measure combines burglary, larceny, motor-vehicle theft, and weapons offenses into a single count, so the coefficient is an average across the four, not a clean mechanism; and no design on this data can completely rule out a shock arriving in 2020 or 2021 and hitting only these river cities. Both are about better local data, not another comparison.

The mirror-image concern for the control group is that the interior counties are not really holding still. If Iowa moved enforcement to the river after 2020, expecting more crossing, the interior counts would fall on their own and widen the gap, without anything changing in the bridge counties. The totals go against that reading. Statewide drug offenses declined about 7.6 percent from 2015–2019 to 2020–2022, but the decline is almost entirely interior, down 10.5 percent versus a 4.3 percent increase in the bridge counties, so the gap opens because the bridge side climbs, not because the interior is held down. The same is true of the gradual fade of the effect with distance, in Section 5.5, since enforcement directed at one stretch of river would not taper smoothly across the state. But the design's own checks already do the job on this point.

The other three limitations belong to the data and the scope of the question, not to the comparison. The drug measure includes cannabis mixed with other controlled substances, so it captures cross-border enforcement broadly, not just cannabis. The crossing is inferred from the pattern of enforcement, not observed, and this could later be confirmed by mobility or traffic-stop records. And Illinois could theoretically have located dispensaries to attract Iowa buyers, but licenses are allocated based on county population, not how close to the border they are, which argues against that interpretation (Appendix A). None of this changes what the estimates find, only how far they can be pushed: the results speak to the direction and incidence of the spillover, not its dollar value, and they do so through the cleanest comparison the data allow, a single prohibition neighbor, treatment fixed by physical infrastructure, and inference checked against a wild cluster bootstrap built for few treated clusters.

6.3 Conclusion

This thesis set out to answer a clean question: did enforcement on the prohibition side of the border move after Illinois legalized in January 2020? Using the seven Iowa counties with a direct highway bridge as the treated group, and a county-year difference-in-differences in which the bridges fix treatment before estimation, it gives a clear answer. The pre-trends hold, and the post-2020 results line up across the specifications. The headline is property: recorded property offenses ran about 37 percent above the interior over 2020–2022 ($\beta = 0.312$, $p = 0.005$), in a category the policy does not itself produce. What matters is not the size of that cost but where it lands, on non-drug offending, the one margin the legalizing state has no hold over. The drug result backs this up, a rise of about 26 percent ($\beta = 0.233$, $p = 0.024$) in the category closest to the policy that fades with distance from the river, exactly as a travel-cost story would predict. Violent offenses lean the same way, and driving is flat.

The contribution is to show that this burden does not stop at drug enforcement. Earlier work had established that drug arrests rise on the prohibition side of a legal border; what no study of a single, clearly defined neighbor had tested is whether the response reaches further. Here it plainly does. The cost of one state's legalization crosses the line into offenses the policy never targeted, in a neighbor that collects none of the tax and had no say in the decision. That is a fiscal externality of the kind Keen and Konrad (2013) describe, made concrete: an asymmetric, mobile-base spillover that the legalizing jurisdiction never has to internalize.

The policy implication is direct. A cost-benefit account that counts only the legalizing state's own tax revenue understates the policy, because part of the cost crosses the state line and part of it surfaces in offense categories no one usually files under drug policy. Two extensions would sharpen the picture. Higher-frequency mobility or traffic-stop data could observe the crossing this design only infers, and extending the design to Illinois's other prohibition neighbors, and to the NIBRS panel once it runs past 2023, would pin down how large the spillover is and how far across the offense categories it runs. The direction, though, is already clear: legalization in one state lands part of its enforcement bill on the neighbor that kept prohibition.

References

Anderson, D. Mark, Benjamin Hansen, and Daniel I. Rees. 2013. "Medical Marijuana Laws, Traffic Fatalities, and Alcohol Consumption." *Journal of Law and Economics* 56 (2): 333–369.

Anderson, D. Mark, and Daniel I. Rees. 2023. "The Public Health Effects of Legalizing Marijuana." *Journal of Economic Literature* 61 (1): 86–143.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan. 2004. "How Much Should We Trust Differences-in-Differences Estimates?" *Quarterly Journal of Economics* 119 (1): 249–275.

Brinkman, Jeffrey, and David Mok-Lamme. 2019. "Not in My Backyard? Not So Fast. The Effect of Marijuana Legalization on Neighborhood Crime." *Regional Science and Urban Economics* 78: 103460.

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

Cameron, A. Colin, and Douglas L. Miller. 2015. "A Practitioner's Guide to Cluster-Robust Inference." *Journal of Human Resources* 50 (2): 317–372.

de Assis, Bruno, Giovanni di Pietra, Bruno Pantaleão, and Rafael Parfitt. 2025. "The Cross-Border Spillover Effects of Cannabis Regulation in South America." *Journal of Quantitative Criminology* 41 (3): 345–385.

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

Dragone, Davide, Giovanni Prarolo, Paolo Vanin, and Giulio Zanella. 2019. "Crime and the Legalization of Recreational Marijuana." *Journal of Economic Behavior and Organization* 159: 488–501.

Dube, Arindrajit, T. William Lester, and Michael Reich. 2010. "Minimum Wage Effects across State Borders: Estimates Using Contiguous Counties." *Review of Economics and Statistics* 92 (4): 945–964.

Einav, Liran, Dan Knoepfle, Jonathan Levin, and Neel Sundaresan. 2014. "Sales Taxes and Internet Commerce." *American Economic Review* 104 (1): 1–26.

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

Guerette, Rob T., and Kate J. Bowers. 2009. "Assessing the Extent of Crime Displacement and Diffusion of Benefits: A Review of Situational Crime Prevention Evaluations." *Criminology* 47 (4): 1331–1368.

Hansen, Benjamin, Keaton Miller, and Caroline Weber. 2020. "Federalism, Partial Prohibition, and Cross-Border Sales: Evidence from Recreational Marijuana." *Journal of Public Economics* 187: 104159.

Hao, Zhuang, and Benjamin W. Cowan. 2020. "The Cross-Border Spillover Effects of Recreational Marijuana Legalization." *Economic Inquiry* 58 (2): 642–666.

Holmes, Thomas J. 1998. "The Effect of State Policies on the Location of Manufacturing: Evidence from State Borders." *Journal of Political Economy* 106 (4): 667–705.

Illinois Department of Financial and Professional Regulation. 2026. *Illinois Adult Use Cannabis Monthly Sales Figures.* Updated March 3, 2026. Springfield, IL: IDFPR.

Kaplan, Jacob. 2023. *Jacob Kaplan's Concatenated Files: National Incident-Based Reporting System (NIBRS) Data, 1991–2022.* Version 9. Ann Arbor, MI: Inter-university Consortium for Political and Social Research. openICPSR Project 118281.

Keen, Michael, and Kai A. Konrad. 2013. "The Theory of International Tax Competition and

Coordination.” In *Handbook of Public Economics*, Vol. 5, edited by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, 257–328. Amsterdam: Elsevier.

Lovenheim, Michael F. 2008. “How Far to the Border? The Extent and Impact of Cross-Border Casual Cigarette Smuggling.” *National Tax Journal* 61 (1): 7–33.

MacKinnon, James G., and Matthew D. Webb. 2018. “The Wild Bootstrap for Few (Treated) Clusters.” *Econometrics Journal* 21 (2): 114–135.

Neumark, David, J. M. Ian Salas, and William Wascher. 2014. “Revisiting the Minimum Wage–Employment Debate: Throwing Out the Baby with the Bathwater?” *ILR Review* 67 (Supplement): 608–648.

Perrotta Berlin, Maria, Chiara Latour, and Giancarlo Spagnolo. 2023. “International Spillovers from Prostitution Regulation: The ’Nordic Model’ and Sex Tourism.” SITE Working Paper 63, Stockholm Institute of Transition Economics.

Rambachan, Ashesh, and Jonathan Roth. 2023. “A More Credible Approach to Parallel Trends.” *Review of Economic Studies* 90 (5): 2555–2591.

Roodman, David, James G. MacKinnon, Morten Ørregaard Nielsen, and Matthew D. Webb. 2019. “Fast and Wild: Bootstrap Inference in Stata Using boottest.” *Stata Journal* 19 (1): 4–60.

Roth, Jonathan. 2022. “Pretest with Caution: Event-Study Estimates after Testing for Parallel Trends.” *American Economic Review: Insights* 4 (3): 305–322.

Roth, Jonathan, Pedro H. C. Sant’Anna, Alyssa Bilinski, and John Poe. 2023. “What’s Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature.” *Journal of Econometrics* 235 (2): 2218–2244.

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

U.S. Census Bureau. 2022. *American Community Survey 5-Year Estimates, 2018–2022.* Washington, DC: U.S. Census Bureau.

Wilson, John D. 1999. “Theories of Tax Competition.” *National Tax Journal* 52 (2): 269–304.

Wu, Guangzhen, Francis D. Boateng, and Xiaodong Lang. 2020. “The Spillover Effect of Recreational Marijuana Legalization on Crime: Evidence from Neighboring States of Colorado and Washington State.” *Journal of Drug Issues* 50 (4): 392–409.

Wu, Guangzhen, Ming Wen, and Fernando A. Wilson. 2021. “Impact of Recreational Marijuana Legalization on Crime: Evidence from Oregon.” *Journal of Criminal Justice* 72: 101742.

Appendix A. Descriptive Evidence on Illinois Dispensary Exposure

This appendix documents the supply-side spatial pattern of Illinois adult-use dispensary license issuance as institutional background to the main demand-side causal analysis in Chapters 4 and 5. The analysis is cross-sectional and descriptive, but not causal. The pattern is reported since it provides the supply-side context in which the demand-side enforcement response is interpreted: legal dispensary supply became available across Illinois and within the reach of the Iowa bridge counties, without concentrating disproportionately at the prohibition-state borders.

A.1 Data and Coverage

The Illinois Department of Financial and Professional Regulation (IDFPR) maintains a public registry of all licensed adult-use cannabis dispensaries. The registry records 240 cumulative adult-use license issuances during 2019 through 2024, distributed across 46 of Illinois's 102 counties; the registry retrieved through 2026 contains 276 dispensaries (the 240-count is used in this appendix for consistency with the 2019–2024 window). Each dispensary is geocoded to a county FIPS using Census TIGER/Line 2020 geometry.

Among the 102 Illinois counties, 32 share an administrative boundary with one of the five prohibition-state neighbors (Wisconsin, Iowa, Missouri, Indiana, Kentucky) as of January 2020, and 17 of these had at least one adult-use dispensary by the end of 2024. Per-capita density is computed as cumulative licenses through 2024 divided by ACS 2016–2020 county population.

I do not run a regression on these data. An earlier draft regressed per-capita dispensary density on a border indicator, but that pattern is weak and sensitive to how the border is defined, and it adds nothing to the causal analysis. The map in Figure 3.1 makes the only point that matters here. Legal supply opened across Illinois, and several dispensaries sit within a short drive of the Iowa bridge counties. The densest counts are in populous places such as the Chicago area, which have many residents as well as many stores. There is no clear sign that licenses were placed to target the Iowa border. The causal analysis rests entirely on the within-Iowa difference-in-differences of Chapter 5; this appendix is background only.

A.2 Dispensary Placement Relative to the Border

A natural concern is that Illinois placed dispensary licenses to capture cross-border demand from the Iowa bridge counties. The data give no sign of this. Dispensary counts track county population rather than distance to the prohibition-state border, so the counties with the most dispensaries are the populous ones rather than the border ones. Legal supply still became locally available to residents of the bridge counties across the river. The bridge-county enforcement response is therefore better read as a demand-side reaction than as exposure to supply that was placed to capture it.

Appendix B. Wild Cluster Bootstrap Inference

The baseline difference-in-differences specification in Chapter 5 has only seven treated clusters. Cameron and Miller (2015) demonstrate that with few treated clusters the conventional cluster-robust variance estimator under-estimates standard errors and over-rejects the null hypothesis of no effect. This appendix reports the wild cluster bootstrap inference referenced in Section 5.3 and documents that the cluster-asymptotic baseline is not driving the significance of the main results.

B.1 Implementation

The wild cluster bootstrap is implemented following MacKinnon and Webb (2018). For each bootstrap iteration b = 1, ..., 9999, residuals from the restricted regression (the regression that imposes the null hypothesis $\tau = 0$) are multiplied by an independent Rademacher random variable at the county-cluster level (the Rademacher variable equals +1 or −1 with equal probability), and the bootstrap-resampled outcome is regressed on the original design matrix to recover a bootstrap coefficient τ_b^*. The p-value is the fraction of $|\tau_b^*|$ that exceeds the absolute value of the observed coefficient. The null-imposed implementation is preferred over the unrestricted implementation when treated clusters are few, because the null-imposed bootstrap better approximates the finite-sample distribution under the null in small samples (Roodman et al. 2019).

B.2 Drug-Enforcement Result and Cross-Outcome Comparison

Table B.1: Wild cluster bootstrap p-values across NIBRS outcomes

Outcome	β	HC1 SE	HC1 p	WCB p (null imposed)	Agreement
Drug-offense records	0.233	0.103	0.024	0.015	Yes
Property offenses	0.312	0.111	0.005	0.010	Yes
Violent offenses	0.154	0.103	0.134	0.144	Yes (null)

Notes: All specifications use the seven-treated-cluster Iowa panel as in Table 5.1. Wild cluster bootstrap with 9999 Rademacher iterations and the null hypothesis imposed (boottest implementation following MacKinnon and Webb 2018).

The wild cluster bootstrap and HC1 cluster-robust p-values are mutually consistent across the three primary NIBRS outcomes. For drug offenses, the WCB p-value of 0.015 is somewhat smaller than the HC1 p-value of 0.024, but both reject the null of no effect at the five-percent level. For property offenses, the WCB p-value of 0.010 is modestly larger than the HC1 p-value of 0.005, but both remain well below the one-percent threshold. The violent offense outcome is not statistically significant under either inference channel at conventional levels. The agreement of WCB and HC1 across outcomes increases confidence that the small-cluster Cameron-Miller (2015) concern, while methodologically relevant, does not flip the main conclusions.