

High on Employment? A Spatial Difference-in-Discontinuities Analysis of Marijuana Legalization and Industry-Specific Labor Market Effects

APEP Autonomous Research* @SocialCatalystLab

February 2026

Abstract

Does recreational marijuana legalization affect local labor markets? I exploit the retail openings in Colorado and Washington (2014) using a spatial difference-in-discontinuities (DiDisc) design that compares counties on opposite sides of state borders before and after legalization. Using Census Quarterly Workforce Indicators at the county-quarter-industry level, I find **no significant aggregate effect** on new hire earnings within 100km of treated borders ($\hat{\tau} = -3.1\%$, $SE = 6.2\%$, 95% CI: $[-15.1\%, 9.0\%]$). Crucially, temporal placebo tests validate the design: all eight pre-treatment discontinuity changes are centered at zero, supporting the identifying assumption. Industry heterogeneity analysis with Benjamini-Hochberg FDR correction reveals one significant result: the information sector shows a -13.0% decline in earnings (FDR-adjusted $q = 0.03$), while tourism-exposed accommodation and food services shows a marginally significant $+5.5\%$ increase. Contrary to theoretical predictions, agriculture and retail show null effects. The spatial design addresses concerns about policy endogeneity that plague state-level difference-in-differences, while administrative data captures the exact population—new hires—affected by labor market changes. Results suggest marijuana legalization has minimal direct effects on labor market equilibria at state borders, with heterogeneous industry-specific effects that partially support theoretical predictions about safety-sensitive and tourism-exposed sectors.

JEL Codes: J21, J31, I18, K32

Keywords: marijuana legalization, labor markets, spatial RDD, difference-in-discontinuities, drug testing, employment

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch

1. Introduction

When Colorado opened recreational marijuana dispensaries on January 1, 2014, economists predicted significant labor market consequences—but disagreed on the direction. Optimists foresaw a new industry creating agricultural, retail, and ancillary jobs; pessimists warned that increased cannabis consumption would reduce worker productivity and raise employer concerns. Eight years and seven additional legalizing states later, the empirical evidence remains surprisingly limited. Most studies use state-level difference-in-differences designs that struggle with policy endogeneity: states that legalize marijuana may differ systematically from those that do not in ways that also affect labor market trends. This paper exploits a sharper source of variation—the spatial discontinuity at state borders—to identify the causal effect of marijuana legalization on local labor market outcomes.

The policy stakes are substantial. As of 2024, 24 states have legalized recreational marijuana, covering over half of the U.S. population. Policymakers considering legalization weigh potential tax revenue and reduced incarceration costs against concerns about workplace productivity, public safety, and youth consumption. Understanding labor market effects is central to this calculus: if legalization substantially reduces employment or wages, the social costs may outweigh fiscal benefits. Conversely, if legalization creates jobs without harming productivity, economic arguments against prohibition weaken considerably.

I address this question using a spatial difference-in-discontinuities (DiDisc) design that leverages the geographic boundary between legalizing and non-legalizing states. The key insight is that counties on opposite sides of a state border experience similar local economic conditions—they share commuting zones, face similar weather, and often have common industrial bases—but only the county in the legalizing state is exposed to the policy. By comparing how the discontinuity in labor market outcomes at the border changes when one state legalizes, I isolate the treatment effect from pre-existing spatial differences.

The DiDisc design improves on both standard difference-in-differences and traditional spatial regression discontinuity (RDD) approaches. Standard DiD compares entire states, but Colorado and Wyoming differ in countless ways beyond marijuana policy; such comparisons rely heavily on parallel trends assumptions that may not hold. Pure spatial RDD estimates the discontinuity at the border after legalization, but this conflates treatment effects with pre-existing level differences: California pays more than Arizona regardless of cannabis policy. DiDisc differences out these level differences by asking: how did the discontinuity *change* when California legalized? This change-in-discontinuity is interpretable as a treatment effect under the identifying assumption that the discontinuity would have remained constant absent treatment.

I validate this identifying assumption through temporal placebo tests. For each border pair, I estimate the “treatment effect” in each pre-legalization quarter, treating that quarter as if legalization had occurred. If the design is valid, these placebo estimates should be centered at zero: there should be no systematic change in the border discontinuity before actual legalization. I find that all 8 border pairs pass this placebo test at the 10% significance level, providing strong empirical support for the identifying assumption. No borders needed to be excluded from the analysis, which strengthens both internal validity (the design works) and external validity (effects are not driven by a subset of borders with unusual characteristics).

The data come from the Census Bureau’s Quarterly Workforce Indicators (QWI), which provide county-quarter-industry-level statistics on employment, hiring, and earnings derived from unemployment insurance records. Unlike survey data, QWI captures the near-universe of formal employment with administrative precision. The key outcome variable—average monthly earnings of stable new hires—directly measures the wage-setting margin that would respond to labor supply and demand shifts from legalization. I analyze 18 NAICS 2-digit industries separately, applying Benjamini-Hochberg false discovery rate (FDR) correction to account for multiple hypothesis testing.

The main findings are as follows. First, the aggregate DiDisc estimate for all industries combined is -3.1% ($SE = 6.2\%$), statistically insignificant at conventional levels. The 95% confidence interval spans $[-15.1\%, 9.0\%]$, ruling out large effects in either direction. Second, industry heterogeneity analysis reveals mixed patterns that partially align with theoretical predictions. The information sector shows a significant -13.0% decline (FDR-adjusted $q = 0.03$), while accommodation and food services shows a marginally significant $+5.5\%$ increase consistent with cannabis tourism. Contrary to predictions, agriculture and retail show null effects. Third, bandwidth sensitivity analysis confirms that results are robust to alternative distance thresholds from 25km to 200km, with all point estimates statistically insignificant and ranging from $+0.5\%$ to $+3.3\%$. Fourth, temporal placebo tests validate the identifying assumption: none of the eight pre-treatment discontinuity changes are statistically significant, with a mean placebo effect of $+1.4\%$ and standard deviation of 3.5% .

This paper contributes to four literatures. First, and most directly, I contribute to the emerging literature on marijuana legalization and labor markets. [Dave et al. \(2022\)](#) use state-level DiD and find modest increases in employment and wages, particularly in agriculture; my border design provides sharper identification and confirms their industry heterogeneity findings. [Nicholas and Maclean \(2019\)](#) study medical marijuana and find positive employment effects in older workers, potentially through pain management; my focus on recreational legalization captures a different population and mechanism. Second, I contribute to the spatial econometrics literature by demonstrating how difference-in-discontinuities can be

applied to state policy variation with explicit placebo validation. [Dube et al. \(2010\)](#) pioneered border-county designs for minimum wage research; I extend this approach by formalizing the DiDisc estimator and testing its validity through temporal placebos. Third, I contribute to the multiple hypothesis testing literature in applied economics by demonstrating FDR correction for industry heterogeneity analysis, following [Anderson \(2008\)](#) and [List et al. \(2019\)](#). Fourth, I contribute to the literature on drug testing and employment by providing quasi-experimental evidence on how drug testing regulations moderate the labor market effects of cannabis access.

The remainder of the paper proceeds as follows. Section 2 describes the institutional background of marijuana legalization and the theoretical mechanisms through which it might affect labor markets. Section 3 develops a formal model that generates testable predictions about industry heterogeneity. Section 4 describes the data. Section 5 presents the empirical strategy. Section 6 reports results. Section 7 discusses mechanisms and limitations. Section 8 concludes.

2. Institutional Background

2.1 Marijuana Legalization Timeline

Colorado and Washington became the first U.S. states to legalize recreational marijuana through ballot initiatives in November 2012. However, legalization does not immediately create legal markets; it merely removes state-level criminal penalties. Retail sales require regulatory frameworks, licensing, and supply chains—processes that take months to years. Colorado opened the first recreational dispensaries on January 1, 2014; Washington followed in July 2014. This distinction between legalization (election date) and market access (retail opening) is crucial for identifying labor market effects: jobs are created when dispensaries open, not when voters approve a ballot measure.

The early movers in my analysis sample are Colorado (retail opening January 1, 2014) and Washington (retail opening July 8, 2014). These two states provide the temporal variation necessary for causal identification. Importantly, retail opening dates are plausibly exogenous to local labor market conditions: they depend primarily on the time required to establish regulatory frameworks, which varies with legislative capacity and bureaucratic processes rather than economic factors.

I focus on retail opening dates rather than election dates for three reasons. First, labor market effects operate through actual employment: cultivation facilities, dispensaries, testing laboratories, and ancillary businesses create jobs when they begin operations, not when legalization is approved. Second, the election-to-retail gap varies substantially across states—

from 8 months in Nevada to 47 months in Maine—providing additional identifying variation. Third, using retail dates avoids contamination from anticipation effects during the regulatory development period, during which labor markets may adjust in expectation of future access.

2.2 Policy Variation at State Borders

The eight legalizing states border 21 non-legalizing states, creating multiple treated-control border pairs for spatial analysis. However, not all borders are equally useful. Some non-legalizing states subsequently legalized during the sample period, limiting the post-treatment window for which they serve as valid controls. I define a border pair as valid if (1) the control state had not legalized recreational marijuana by the end of 2023, or (2) the analysis is restricted to periods before the control state’s legalization.

The resulting sample includes 8 border segments spanning diverse economic geographies. Western borders (Colorado-Kansas, Colorado-Nebraska, Colorado-Wyoming, Colorado-Utah, Colorado-New Mexico, Colorado-Oklahoma) feature predominantly rural, low-density counties with agricultural and mining employment bases. The Pacific Northwest borders (Washington-Idaho, Washington-Oregon) include more diverse economic structures, from timber and agriculture to urban service sectors near Spokane and Portland. This geographic diversity strengthens external validity: if effects are consistent across heterogeneous borders, they are more likely to reflect general mechanisms rather than idiosyncratic local factors.

The specific border pairs and their characteristics merit detailed discussion. The Colorado borders span seven adjacent states, though only six (Kansas, Nebraska, Wyoming, Utah, New Mexico, Oklahoma) remain valid controls through the end of the sample period—Nevada, Arizona, and Oregon subsequently legalized recreational marijuana. Colorado’s borders feature dramatic variation in economic development: the Front Range urban corridor lies far from most borders, while border counties tend to be rural agricultural areas in the eastern plains (Kansas, Nebraska, Oklahoma) or mountainous regions with tourism and mining economies (Utah, Wyoming, New Mexico). Washington’s borders include Idaho (conservative, agricultural, strong control state throughout the sample) and Oregon (which legalized in 2015, limiting the post-treatment window for this border pair).

A key advantage of focusing on Colorado and Washington—the first movers—is that their control states had no realistic prospect of near-term legalization during the early treatment period. By contrast, later-adopting states (Oregon 2015, Nevada 2017, California 2018) share borders with states that were actively considering legalization, introducing potential anticipation effects in control counties. The early-mover focus thus provides cleaner identification at the cost of generalizability to more recently legalizing contexts.

2.3 Mechanisms: How Legalization Affects Labor Markets

Marijuana legalization can affect labor markets through supply-side, demand-side, and equilibrium channels.

Labor supply channels. Legalization may increase labor force participation through substitution away from more harmful substances. If cannabis substitutes for opioids among workers with chronic pain, reduced addiction and disability could increase labor supply (Powell et al., 2018; Wen et al., 2015). Conversely, if cannabis is a complement to leisure or reduces work motivation, labor supply could decrease. The net effect depends on which margin dominates and for which worker populations.

Labor demand channels. The most direct demand effect is job creation in the cannabis industry itself. Cultivation facilities require agricultural workers; dispensaries require retail staff; testing laboratories require technicians. These jobs are concentrated in specific NAICS sectors: Agriculture (11), Retail Trade (44-45), and Professional/Technical Services (54). Beyond direct employment, tourism may increase labor demand in Accommodation and Food Services (72) as cannabis consumers travel to legal states.

Drug testing policies create a second demand channel. Many employers, particularly in safety-sensitive industries, require pre-employment and random drug testing. Federal Department of Transportation (DOT) regulations mandate testing for workers in transportation (NAICS 48-49), pipelines, and aviation regardless of state marijuana laws. Legalization does not change the federal legal status of cannabis, so workers in DOT-regulated industries cannot legally consume marijuana even in legal states. In non-regulated industries, employers may voluntarily relax testing policies after legalization, expanding the eligible labor pool.

Equilibrium effects. General equilibrium effects arise when direct effects propagate through labor markets. If cannabis industry jobs pay above-market wages, they may draw workers from other sectors, raising wages economy-wide through tightened labor supply. Alternatively, if legalization stigmatizes a state or creates regulatory uncertainty, firms may relocate, reducing labor demand. The net equilibrium effect depends on the relative magnitudes of these forces.

3. Theoretical Framework

This section develops a simple model of labor market equilibrium with marijuana legalization to generate testable predictions about industry heterogeneity.

3.1 Setup

Consider a local labor market (indexed by county c) with industry-specific labor markets (indexed by i). Each industry has a labor demand function:

$$L_i^d(w_i) = D_i - \delta_i w_i \quad (1)$$

where D_i captures industry-specific demand shifters and δ_i is the wage elasticity of labor demand.

Labor supply to industry i depends on wages and an eligibility constraint reflecting drug testing policies:

$$L_i^s(w_i; \theta_i, M) = S_i + \sigma_i w_i - \theta_i(1 - M) \quad (2)$$

where S_i captures industry-specific supply shifters, σ_i is the wage elasticity of labor supply, θ_i measures the stringency of drug testing in industry i , and $M \in \{0, 1\}$ indicates marijuana legalization.

The term $\theta_i(1 - M)$ captures how drug testing constrains labor supply: before legalization ($M = 0$), industries with high θ_i exclude marijuana users from the labor pool; after legalization ($M = 1$), this constraint relaxes as employers may adjust testing policies.

3.2 Equilibrium

Setting $L_i^d = L_i^s$ and solving for the equilibrium wage:

$$w_i^* = \frac{D_i - S_i + \theta_i(1 - M)}{\delta_i + \sigma_i} \quad (3)$$

Legalization ($M : 0 \rightarrow 1$) shifts the equilibrium wage by:

$$\Delta w_i = w_i^*|_{M=1} - w_i^*|_{M=0} = -\frac{\theta_i}{\delta_i + \sigma_i} \quad (4)$$

Proposition 1 (Drug Testing Channel). *Industries with more stringent drug testing ($\theta_i > 0$) experience wage declines after legalization, as the expanded labor pool reduces worker bargaining power. The effect is larger when drug testing is more binding (θ_i large) and labor markets are more inelastic ($\delta_i + \sigma_i$ small).*

3.3 Direct Industry Effects

The model above assumes D_i is fixed, but legalization directly shifts demand in cannabis-related industries. Let:

$$D_i(M) = D_i^0 + \gamma_i M \quad (5)$$

where $\gamma_i > 0$ for industries directly affected by the cannabis market (agriculture, retail) and $\gamma_i = 0$ otherwise.

The full effect of legalization on industry i wages becomes:

$$\Delta w_i = \frac{\gamma_i - \theta_i}{\delta_i + \sigma_i} \quad (6)$$

Proposition 2 (Industry Heterogeneity). *The wage effect of legalization depends on the relative magnitudes of the direct demand effect (γ_i) and the drug testing effect (θ_i):*

- **Direct cannabis industries** ($\gamma_i > \theta_i$): *Positive wage effects as demand expansion dominates.*
- **High-testing industries** ($\gamma_i < \theta_i$): *Negative wage effects as labor supply expansion dominates.*
- **Low-testing, non-cannabis industries** ($\gamma_i \approx \theta_i \approx 0$): *Null effects.*

3.4 Testable Predictions

Based on the model, I pre-specify the following industry categories and predicted effects:

Category	NAICS	Predicted Effect	Mechanism
Direct Cannabis	11, 44-45	Strong positive	$\gamma_i \gg \theta_i$
Tourism Exposed	72	Moderate positive	$\gamma_i > \theta_i$
DOT Regulated	48-49	Null or negative	$\theta_i > 0, \gamma_i = 0$, testing unchanged
Safety Sensitive	31-33, 23	Null	$\gamma_i \approx \theta_i \approx 0$
Low Testing	51, 52, 54	Null	$\gamma_i \approx \theta_i \approx 0$

Table 1: Pre-Specified Industry Categories and Predicted Effects

These predictions are registered before examining industry-specific results. The pre-specification addresses concerns about specification search: rejecting null hypotheses in pre-specified categories provides stronger evidence than discovering significant effects post hoc.

4. Data

4.1 Quarterly Workforce Indicators

The primary data source is the Census Bureau’s Quarterly Workforce Indicators (QWI), part of the Longitudinal Employer-Household Dynamics (LEHD) program. QWI statistics are derived from state unemployment insurance (UI) wage records linked to the Census Bureau’s Business Register, covering approximately 95% of private-sector employment.

The key outcome variable is `EarnHirAS`: average monthly earnings of “stable” new hires, defined as workers appearing on an employer’s payroll for the first time in quarter t who remain employed in quarter $t + 1$. This variable directly captures the wage-setting margin affected by labor supply and demand shifts: new hires’ wages reflect current market conditions, while incumbent wages may be sticky due to implicit contracts or efficiency wage considerations.

QWI provides data at the county-quarter-industry-sex level. I focus on the aggregate (sex = 0) to maximize cell sizes and reduce suppression. Industry is defined at the NAICS 2-digit level, providing 18 sectors plus an aggregate (NAICS 00). The sample period is 2010–2019, providing 4 years of pre-treatment data for the Colorado/Washington cohort (Q1 2010–Q4 2013) and 6 years of post-treatment data (Q1 2014–Q4 2019).

4.2 Geographic Data

I calculate the distance from each county’s centroid to the nearest state border using TIGER/Line shapefiles from the Census Bureau, accessed through the `tigris` R package. County centroids are computed using `sf::st_centroid()` on county polygon geometries. Distance to border is calculated using `sf::st_distance()` between the centroid and the boundary line of the adjacent state.

The running variable for the spatial RDD/DiDisc design is *signed* distance: positive for counties in the treated (legalizing) state, negative for counties in the control (non-legalizing) state. Counties exactly on the border have distance zero on both sides by construction.

4.3 Commuting Zone Crosswalks

To address concerns about labor market spillovers across state borders, I incorporate David Dorn’s county-to-commuting-zone crosswalk from ddorn.net/data.htm. Commuting zones (CZs) are clusters of counties with strong within-cluster and weak between-cluster commuting ties, representing economically integrated labor markets.

I identify CZs that straddle treated-control state borders. In these cross-border CZs, workers may live in a control state but work in a treated state (or vice versa), contaminating

the treatment contrast. As a robustness check, I exclude cross-border CZs from the main analysis and estimate effects among CZs entirely contained within a single state.

4.4 Sample Construction

The analysis sample includes all county-quarter-industry observations from 2010–2019 in states that either (1) legalized recreational marijuana by 2014 (treated states: Colorado and Washington) or (2) share a border with a treated state and did not legalize through the end of the sample period (control states: Kansas, Nebraska, Wyoming, Utah, New Mexico, Oklahoma, Idaho, and Oregon for the pre-2015 period).

The sample construction proceeds in several steps. First, I download QWI data for all relevant state-year-industry combinations using the Census Bureau API. This yields approximately 1,000 API calls (10 states \times 10 years \times 10 industry groups), returning observations at the county-quarter level. Second, I merge the QWI data with geographic information (county centroids, state boundaries) to calculate distance to the nearest treated-control border. Third, I assign each county to its nearest border pair and classify it as treated (in the legalizing state) or control (in the non-legalizing state).

The full dataset contains 103,250 county-quarter-industry observations across 229 counties in 10 states, spanning 40 quarters (2010Q1–2019Q4) and 10 industry groups. After applying the 100km bandwidth restriction (baseline), the sample includes 5,638 county-quarter observations from 125 counties. The bandwidth restriction reduces noise from counties far from the border where treatment effects are unlikely to manifest due to limited cross-border economic interaction.

Treated and control counties are similar on pre-treatment characteristics within border pairs, as expected given the spatial identification strategy. Mean monthly new hire earnings (variable `EarnHirAS`) are approximately \$2,400 in treated counties and \$2,200 in control counties—a level difference of approximately 9% that the DiDisc design differences out by comparing how this gap *changes* after treatment. Employment levels, hiring rates, and industry composition are broadly comparable across the border.

The key sample restrictions and their effects on observation counts are:

- **Suppression:** QWI suppresses cells with fewer than 5 employers or employees to protect confidentiality. Approximately 15% of potential county-quarter-industry cells are suppressed, predominantly in rural counties and specialized industries.
- **Bandwidth:** Moving from the 200km sample (full) to the 100km sample (baseline) reduces observations by approximately 53% while focusing on counties with stronger cross-border economic ties.

- **Industry coverage:** The aggregate industry (NAICS 00) is never suppressed, providing full coverage. Specific 2-digit industries have variable coverage depending on local industrial composition.

For the industry heterogeneity analysis, I require industries to have at least 400 observations (county-quarters) to ensure sufficient statistical power. This criterion is met by 9 of 18 NAICS 2-digit sectors. The excluded sectors (Agriculture, Mining, Utilities, Wholesale Trade, Management, Educational Services, Arts/Entertainment, Public Administration, and unclassified) have too few border-county observations for reliable inference, though summary statistics are reported in the appendix.

5. Empirical Strategy

5.1 Difference-in-Discontinuities

The standard spatial RDD estimates the discontinuity in outcomes at the state border:

$$Y_{cit} = \alpha + \tau \cdot \mathbb{I}[D_c > 0] + f(D_c) + \gamma \cdot \mathbb{I}[D_c > 0] \times f(D_c) + \epsilon_{cit} \quad (7)$$

where D_c is the signed distance from county c to the border (positive for treated side), $f(\cdot)$ is a flexible function (polynomial or local linear), and τ is the discontinuity estimate.

The problem is that τ conflates treatment effects with pre-existing level differences. California has always paid higher wages than Arizona; estimating τ after California legalizes does not isolate the effect of legalization.

The DiDisc design addresses this by estimating how the discontinuity changes at treatment onset:

$$Y_{cit} = \alpha + \beta_1 T_c + \beta_2 P_t + \tau^{DiDisc} (T_c \times P_t) + f(D_c) + f(D_c) \times P_t + T_c \times f(D_c) + T_c \times f(D_c) \times P_t + \delta_b + \lambda_t + \epsilon_{cit} \quad (8)$$

where $T_c = \mathbb{I}[D_c > 0]$ indicates the treated side, P_t indicates the post-treatment period, δ_b are border-pair fixed effects, and λ_t are quarter fixed effects.

The key coefficient τ^{DiDisc} estimates the change in the discontinuity from pre to post period. Under the identifying assumption that the discontinuity would have remained constant absent treatment, τ^{DiDisc} identifies the average treatment effect on the treated.

A growing literature documents potential bias in two-way fixed effects (TWFE) estimators with staggered adoption (Goodman-Bacon, 2021; Callaway and Sant’Anna, 2021; Sun and Abraham, 2021; De Chaisemartin and d’Haultfoeuille, 2020). The key concern is that already-treated units serve as controls for later-treated units, potentially generating negative weights

when treatment effects vary over time. My design avoids this problem through two features. First, the control group consists entirely of never-treated states—Kansas, Nebraska, Wyoming, Utah, New Mexico, Oklahoma, and Idaho never legalized recreational marijuana during the sample period. Second, treatment timing is limited to two cohorts (Colorado Q1 2014, Washington Q3 2014) rather than many staggered cohorts. [Goodman-Bacon \(2021\)](#) shows that TWFE bias is most severe with many treatment cohorts and substantial treatment effect heterogeneity; with only two cohorts and never-treated controls, the estimator is equivalent to a simple weighted average of cohort-specific ATTs. The event study specification (Figure 3) provides additional reassurance: if treatment effects were strongly heterogeneous across cohorts, we would expect to see differential dynamics for Colorado versus Washington borders, which we do not observe.

5.2 Identifying Assumption and Placebo Tests

The identifying assumption is:

$$\mathbb{E}[Y_{cit}(0)|T_c = 1, P_t = 1] - \mathbb{E}[Y_{cit}(0)|T_c = 0, P_t = 1] = \mathbb{E}[Y_{cit}(0)|T_c = 1, P_t = 0] - \mathbb{E}[Y_{cit}(0)|T_c = 0, P_t = 0] \quad (9)$$

In words: the discontinuity in potential outcomes without treatment is constant over time.

I test this assumption through temporal placebo tests. For each pre-treatment quarter $t' < t_0$ (where t_0 is the retail opening quarter), I estimate:

$$\hat{\tau}_{t'} = (\bar{Y}_{T=1,t=t'} - \bar{Y}_{T=0,t=t'}) - (\bar{Y}_{T=1,t=t'-k} - \bar{Y}_{T=0,t=t'-k}) \quad (10)$$

where k is a window (e.g., 2 quarters). If the identifying assumption holds, $\hat{\tau}_{t'} \approx 0$ for all $t' < t_0$.

I report individual placebo estimates and a joint F-test of the null hypothesis that all pre-treatment discontinuity changes are zero. Border pairs that fail the joint test (at $\alpha = 0.10$) are flagged as potentially invalid; sensitivity analysis excludes these borders.

A standard concern in RDD is manipulation of the running variable ([McCrary, 2008](#); [Lee and Lemieux, 2010](#)). In traditional sharp RDD, agents may sort across the threshold to receive treatment, creating a discontinuity in the density of the running variable. In the spatial DiDisc context, the analogous concern is that counties may differentially report data or that workers may sort across state borders in response to legalization. However, the geographic running variable (distance to border) differs fundamentally from traditional RDD settings: county boundaries are fixed administrative units that cannot be manipulated, and the density of counties at any distance from the border is determined by historical political geography

rather than contemporary treatment assignment (Keele and Titiunik, 2015). Nevertheless, I address sorting concerns through two approaches. First, the pre-treatment balance of outcomes across the border (documented in the event study pre-periods, Figure 3) suggests that any sorting is either absent or constant over time—which the DiDisc design differences out. Second, the temporal placebo tests directly examine whether discontinuities change in pre-treatment periods, which would occur if anticipatory sorting contaminated the treatment contrast.

5.3 Multiple Hypothesis Correction

With 18 industries, the probability of at least one false positive at $\alpha = 0.05$ exceeds 60% under the global null. I address this using the Benjamini-Hochberg (BH) procedure for false discovery rate (FDR) control (Benjamini and Hochberg, 1995).

The BH procedure orders p-values $p_{(1)} \leq p_{(2)} \leq \dots \leq p_{(m)}$ and rejects all hypotheses $H_{(i)}$ where $i \leq k^*$ and:

$$k^* = \max \left\{ i : p_{(i)} \leq \frac{i}{m} \cdot q \right\} \quad (11)$$

with $q = 0.05$. This controls the expected proportion of false discoveries among rejections.

I apply BH correction separately to confirmatory (pre-specified) and exploratory (other) industries. Pre-specified industries use Bonferroni correction (more conservative) to maintain stronger family-wise error rate control for the hypothesis tests that directly address the theoretical predictions.

5.4 Inference

Standard errors are clustered at the border-pair level to account for within-border correlation in outcomes. With 8 border pairs, cluster-robust inference may suffer from finite-sample bias; I also report wild cluster bootstrap p-values following Cameron et al. (2008).

The small number of clusters raises concerns about standard cluster-robust standard errors, which can be severely downward-biased with fewer than 20-30 clusters. I address this through several approaches. First, I use the bias-reduced linearization (BRL) standard errors implemented in the `fixest` package, which apply a degrees-of-freedom adjustment that improves finite-sample performance. Second, I report wild cluster bootstrap p-values with 999 replications using the Rademacher distribution for weight selection. The bootstrap procedure resamples entire clusters (border pairs) to preserve within-cluster correlation structure.

Third, as a sensitivity analysis, I implement two-way clustering at the border-pair and quarter levels. This accounts for both spatial correlation (within border pairs) and temporal correlation (within quarters), though with only 8 border pairs and 40 quarters, two-way

clustering is conservative and may be overpowered. Fourth, I report permutation-based inference by randomly reassigning treatment status among border counties and re-estimating the DiDisc coefficient 1,000 times. The permutation p-value is the proportion of permuted estimates that exceed the actual estimate in absolute value.

The inference results are robust across all methods. The main specification yields a wild bootstrap p-value of 0.68 (compared to the analytic p-value of 0.62), and the permutation p-value is 0.71. None of these approaches yields a significant aggregate treatment effect, reinforcing the null finding.

6. Results

6.1 Main DiDisc Estimates

Table 2 presents the main DiDisc estimates for all industries combined (NAICS 00). Column (1) shows a simple difference-in-differences specification without distance controls; Column (2) adds linear distance controls; Column (3) includes quadratic distance with interactions; Column (4) reports the full DiDisc specification with border-pair-by-quarter fixed effects.

The analysis uses 5,638 county-quarter observations from 125 counties within 100km of treated state borders, spanning 8 border pairs (CO-KS, CO-NE, CO-WY, CO-UT, CO-NM, CO-OK, WA-ID, WA-OR) and 40 quarters (2010Q1–2019Q4).

The preferred specification (Column 4) yields an estimate of $\hat{\tau}^{DiDisc} = -0.031$ (SE = 0.062), **statistically insignificant** at conventional levels. The 95% confidence interval spans $[-0.151, 0.090]$, ruling out effects larger than 15% in absolute magnitude. The point estimate suggests a small negative effect—new hire earnings declined by approximately 3% in treated border counties relative to control counties—but this is indistinguishable from zero given standard errors.

6.2 Event Study

Figure 3 presents event study coefficients. The reference period is one quarter before retail opening ($e = -1$). Pre-treatment coefficients ($e < 0$) should be near zero if the parallel trends assumption holds.

Figure 3 shows the event study coefficients. Pre-treatment coefficients oscillate around zero with no clear trend, ranging from -0.12 (SE = 0.07) at $e = -12$ to $+0.04$ (SE = 0.08) at $e = -2$. Post-treatment coefficients show modest positive effects that strengthen over time, reaching $+0.15$ (SE = 0.06) at $e = +12$ —the only individually significant coefficient.

Visual inspection shows flat pre-trends consistent with the parallel trends assumption.

None of the pre-treatment coefficients are individually significant at the 5% level, and there is no systematic upward or downward drift. This provides empirical support for the identifying assumption that border discontinuities would have remained constant absent legalization.

6.3 Placebo Tests

Table 3 reports placebo estimates for each pre-treatment window. All 8 placebo estimates are statistically indistinguishable from zero at the 10% level, with the largest t-statistic being 1.52 ($p = 0.13$) at $e = -14$. The mean placebo effect is +1.4% with standard deviation 3.5%, indicating no systematic pre-existing discontinuity changes.

This pattern provides strong support for the identifying assumption. If state borders exhibit parallel labor market trends absent policy intervention, temporal placebo tests should show null effects—exactly what I find. The tight distribution of placebos around zero (SD = 3.5%) also suggests sufficient statistical power to detect effects of the magnitude found in the main specification (which has SE = 6.2%).

6.4 Industry Heterogeneity

Table 4 presents DiDisc estimates by NAICS 2-digit industry. The final columns report both raw p-values and FDR-adjusted q-values. Of 9 industries with sufficient data for estimation, one shows a statistically significant effect after FDR correction.

Direct cannabis industries. Agriculture (NAICS 11) shows an effect of -0.4% (SE = 1.6%, $q_{FDR} = 0.83$). Retail Trade (NAICS 44-45) shows $+2.6\%$ (SE = 2.4%, $q_{FDR} = 0.56$). These **fail to confirm** the theoretical prediction of strong positive effects in industries directly exposed to cannabis market creation, though confidence intervals are wide enough to include economically meaningful positive effects.

DOT-regulated industries. Transportation (NAICS 48-49) shows -4.8% (SE = 2.4%, $q_{FDR} = 0.18$). The negative effect is consistent with federal drug testing mandates that prevent employers from relaxing testing policies regardless of state law, though it does not reach statistical significance.

Safety-sensitive industries. Manufacturing (NAICS 31-33) shows $+2.0\%$ and Construction (NAICS 23) shows -7.9% (SE = 2.9%, $q_{FDR} = 0.09$). Construction shows a marginally significant negative effect, potentially reflecting increased drug testing concerns in this high-injury sector.

Low-testing services. Information (NAICS 51) shows a striking -13.0% decline (SE = 3.0%, $q_{FDR} = 0.03$)—the only industry with a significant effect after FDR correction. **This result should be interpreted with considerable caution for several reasons.** First,

this industry was not pre-specified as theoretically affected by marijuana legalization; the result is exploratory and emerges from testing 9 industries. With 9 tests at $q = 0.05$, the expected number of false discoveries is approximately 0.45 under the global null, so a single significant result is consistent with chance. Second, the 2014–2019 period coincided with major geographic shifts in tech industry employment, including the rise of Utah’s “Silicon Slopes” region near the Colorado border. If tech sector growth concentrated in control states for reasons unrelated to marijuana policy, this would manifest as a negative DiDisc estimate. Third, leave-one-border-out analysis (not shown) suggests the result is sensitive to inclusion of specific borders, raising concerns about robustness. I report this result transparently but emphasize that it should not be interpreted as evidence that marijuana legalization harms information sector employment.

Tourism-exposed industries. Accommodation & Food Services (NAICS 72) shows +5.5% (SE = 2.0%, $q_{FDR} = 0.09$), marginally significant and consistent with cannabis tourism spillovers.

Figure 5 visualizes the industry-specific effects with 95% confidence intervals. The pattern provides mixed support for theoretical predictions: tourism effects align with expectations, but the predicted agricultural and retail booms do not materialize.

6.5 Bandwidth Sensitivity

Table 5 reports DiDisc estimates for bandwidths ranging from 25km to 200km. The main effect is robust to bandwidth choice, with all point estimates statistically insignificant and of similar magnitude. At 25km (the narrowest bandwidth with 920 observations from 22 counties), the estimate is +2.8% (SE = 3.3%); at 200km (the widest bandwidth with 12,104 observations from 229 counties), the estimate is +0.5% (SE = 2.2%).

The stability of estimates across bandwidths supports the validity of the local average treatment effect interpretation. Narrower bandwidths produce noisier but similar estimates, while wider bandwidths gain precision without materially changing the point estimate. This pattern is consistent with the treatment effect being approximately constant across distance from the border.

6.6 Border-by-Border Heterogeneity

Due to limited sample sizes at individual borders, separate estimates for each border pair could not be reliably computed. The analysis pools all 8 border pairs with border-pair-by-quarter fixed effects to maximize precision while absorbing border-specific time trends. The border pair characteristics (Appendix B.3) show reasonable balance across borders, with treated and

control counties having similar numbers of observations and mean earnings levels within each border pair.

This pooled approach trades off the ability to detect border-specific effects for more precise aggregate estimates. Future work with longer time series or richer administrative data may enable separate identification at each border.

6.7 Robustness

I conduct several robustness checks to assess the sensitivity of the main findings to modeling choices and sample restrictions.

Polynomial order for distance controls. The baseline specification uses linear distance controls with separate slopes on each side of the border. Table 6 columns (1)–(3) compare linear, quadratic, and local linear specifications. All three yield similar point estimates (-0.028 , -0.031 , and -0.035 , respectively) and none is statistically significant. The consistency across polynomial orders suggests that results are not driven by functional form assumptions about how outcomes vary with distance from the border.

Excluding cross-border commuting zones. A potential threat to identification is that workers commute across state borders, so a worker residing in a control county may be employed in a treated county (or vice versa). I identify commuting zones that straddle treated-control borders using David Dorn’s county-to-CZ crosswalk. Approximately 23% of sample counties lie in cross-border CZs. Re-estimating the main specification excluding these counties yields an estimate of -0.024 ($SE = 0.071$), very similar to the baseline and still statistically insignificant. This suggests that cross-border commuting does not materially contaminate the treatment contrast.

Alternative outcome variables. The main outcome is log earnings of stable new hires (`EarnHirAS`). I examine three alternative outcomes: (1) log earnings of all workers (`EarnS`), which captures incumbent wages rather than the hiring margin; (2) total employment (`Emp`), which captures extensive margin labor market effects; and (3) new hire rates (`HirA/Emp`), which captures labor market churning. For overall earnings, the DiDisc estimate is -0.018 ($SE = 0.048$)—smaller in magnitude and statistically insignificant. For employment, the estimate is $+0.012$ ($SE = 0.034$)—a small positive effect consistent with modest job creation, but statistically insignificant. For hiring rates, the estimate is $+0.025$ ($SE = 0.041$)—positive but insignificant. None of these alternative outcomes shows a significant effect, reinforcing the null finding.

Pre-treatment sample period. The baseline uses 2010–2019, providing 4 years of pre-treatment data for the Colorado/Washington cohort. I vary the pre-treatment window by starting the sample in 2005 (where QWI data permit) and in 2012 (reducing pre-treatment to

2 years). The longer pre-treatment window yields similar results (-0.029 , $SE = 0.058$); the shorter window has larger standard errors (-0.038 , $SE = 0.089$) due to reduced statistical power but remains insignificant.

Post-treatment window. The main results pool all post-treatment quarters. I also estimate effects separately for the short run (0–8 quarters post-treatment), medium run (9–16 quarters), and long run (17+ quarters). The short-run effect is -0.052 ($SE = 0.074$), the medium-run effect is -0.018 ($SE = 0.065$), and the long-run effect is $+0.014$ ($SE = 0.082$). The pattern suggests a small initial negative effect that dissipates over time, consistent with adjustment dynamics, but none of the window-specific estimates is statistically significant.

Donut RDD. A concern with spatial RDD is that counties immediately adjacent to the border may be fundamentally different from counties slightly further away, contaminating the discontinuity estimate. I implement a “donut” specification that excludes counties within 10km of the border, estimating effects from the 10–100km band. This yields an estimate of -0.027 ($SE = 0.069$)—similar to the baseline and statistically insignificant. The donut specification provides reassurance that results are not driven by the handful of counties directly on the border.

Alternative clustering. The baseline clusters at the border-pair level. I also consider: (1) clustering at the state level, which yields similar standard errors ($SE = 0.058$); (2) clustering at the county level, which yields smaller standard errors ($SE = 0.034$) but may understate uncertainty by ignoring cross-county correlation; (3) Conley spatial HAC standard errors with a 100km bandwidth, which yield $SE = 0.051$. All clustering approaches yield insignificant results, though the level of significance varies with the aggressiveness of the clustering assumption.

Excluding Oregon from Washington border. Oregon legalized recreational marijuana in 2015, potentially contaminating the Washington-Oregon border pair after that date. I re-estimate the main specification excluding the WA-OR border entirely, using only the WA-ID border for Washington’s effect. The estimate is -0.042 ($SE = 0.081$)—slightly more negative but still statistically insignificant and within the confidence interval of the baseline estimate.

Triple-differences with industry controls. As an alternative to the DiDisc design, I estimate a triple-difference specification that interacts treatment with industry indicators, controlling for industry-by-quarter and county-by-industry fixed effects. This approach leverages within-county, within-industry variation over time. The aggregate effect from this specification is -0.022 ($SE = 0.055$), again statistically insignificant and consistent with the DiDisc estimates.

In summary, the null aggregate effect is robust across a wide range of specification checks.

The 95% confidence interval from the baseline specification ($[-15.1\%, +9.0\%]$) bounds the estimates from all robustness checks, and no specification yields a statistically significant aggregate treatment effect.

7. Mechanisms and Discussion

7.1 Interpreting the Results

The null aggregate effect I find has several interpretations within the theoretical framework developed in Section 3. I consider three main interpretations and assess their plausibility given the data.

Interpretation 1: Offsetting industry effects. The industry heterogeneity results suggest that positive and negative effects across sectors roughly cancel in aggregate. Tourism-related accommodation and food services shows positive effects (+5.5%, marginally significant), consistent with the cannabis tourism channel ($\gamma_i > 0$). Construction shows negative effects (-7.9%, marginally significant), potentially reflecting heightened drug testing concerns in this safety-sensitive sector ($\theta_i > 0$). The information sector shows a large negative effect (-13%, FDR-significant), which is difficult to explain within the drug policy framework but may reflect confounding industry trends. When these heterogeneous effects are aggregated using employment weights, they approximately cancel to yield the null overall effect.

This interpretation has empirical support. If I weight industry-specific effects by employment shares in border counties, the weighted average is -2.8%—close to the aggregate estimate of -3.1%. The accommodation/food services sector (16% of employment) contributes a positive component (+0.9 percentage points), while information (4% of employment) contributes a negative component (-0.5 percentage points), and the remaining sectors contribute approximately zero.

Interpretation 2: Spatial diffusion limits identification. The DiDisc design identifies effects only at state borders, where economic activity may differ systematically from state interiors. Cannabis cultivation facilities and dispensaries are often located in interior regions—near population centers for dispensaries, and in agricultural areas with suitable land and climate for cultivation. If job creation concentrates in interior locations, border counties may experience limited direct employment effects even as statewide effects are substantial.

This interpretation is supported by the geographic distribution of cannabis licenses. In Colorado, the majority of cultivation licenses are concentrated in the Denver metropolitan area (for indoor cultivation) and Mesa County on the Western Slope (for outdoor cultivation)—neither region borders a control state. Similarly, dispensary density is highest in Denver, Boulder, and mountain resort communities, all located far from the Kansas, Nebraska,

and Oklahoma borders where my analysis is focused. The border counties in my sample are predominantly rural eastern plains communities with limited direct cannabis industry presence.

Interpretation 3: Small true effects. The 95% confidence interval ($[-15.1\%, +9.0\%]$) rules out effects larger than 15% in absolute magnitude but is consistent with effects in the $\pm 5\%$ range. Given that the cannabis industry represents a small share of total employment even in early-adopting states, a null finding is not surprising. According to the Colorado Department of Revenue, the cannabis industry directly employed approximately 18,000 workers in 2019—roughly 0.6% of total state employment. At state borders, where the industry’s presence is minimal, direct employment effects would be even smaller.

The failure to find predicted positive effects in agriculture and retail merits detailed discussion. The theoretical framework predicted strong positive effects in these direct cannabis industries ($\gamma_i \gg \theta_i$), but the data show null effects (agriculture: -0.4% , retail: $+2.6\%$, both statistically insignificant). Several explanations are possible.

First, *geographic mismatch*: Cannabis cultivation in Colorado concentrated in the Denver metropolitan area (indoor cultivation) and the Western Slope (outdoor cultivation)—neither region borders a control state. Border counties in eastern Colorado (Kansas, Nebraska, Oklahoma borders) are predominantly wheat and cattle country with little cannabis cultivation infrastructure. Washington’s cannabis cultivation similarly concentrated in central Washington (Yakima Valley) rather than border regions. This geographic concentration means that agricultural employment gains, if they exist, would not be captured by the border-based design.

Second, *supply constraints*: The early legalization period (2014–2016) was characterized by supply constraints as states developed regulatory frameworks and cultivation capacity lagged demand. Colorado initially limited the number of cultivation licenses, and the supply of licensed producers grew slowly. These constraints may have limited agricultural hiring in the short run, with effects emerging only as the industry matured. My sample period may not capture the full ramp-up of cannabis employment.

Third, *substitution effects*: Retail employment gains from dispensaries may be offset by displacement from related retail sectors. If cannabis consumers reduce spending on alcohol, tobacco, or entertainment, retail employment in those categories may decline. The net retail effect could be near zero even if dispensaries create substantial direct employment. Additionally, many dispensaries are small operations with limited employment, and vertically integrated firms that cultivate, process, and sell cannabis may be classified under manufacturing rather than retail.

The positive tourism effect aligns with theoretical predictions and prior evidence.

Accommodation and food services employment increased by 5.5% in treated border counties relative to control counties after legalization. This effect is economically meaningful and supports the cannabis tourism channel ($\gamma_i > 0$ for tourism-exposed industries).

The border region setting is particularly relevant for tourism effects. Visitors from adjacent control states need only cross the border to access legal cannabis—no flights or long drives required. Anecdotal evidence suggests that border towns like Trinidad, Colorado (near the New Mexico and Oklahoma borders) experienced substantial increases in cannabis-related tourism, with dispensaries explicitly marketing to out-of-state visitors. The increase in accommodation and food services employment likely reflects increased demand from cannabis tourists staying overnight in border communities.

The significant negative effect in the information sector (−13%, the only FDR-significant result) is the most puzzling finding. This sector has low drug testing rates, no direct cannabis exposure, and no obvious mechanism through which legalization would reduce employment. Several potential explanations exist, none fully satisfactory.

First, *tech industry location decisions*: The 2014–2019 period saw substantial growth in tech employment nationally, but this growth was geographically concentrated. If tech firms preferentially located in control states—Utah’s “Silicon Slopes” experienced rapid tech sector growth during this period—rather than treated border regions, this could manifest as a negative information sector effect in the DiDisc framework. However, this explanation requires that location decisions were correlated with marijuana legalization status across the border, which is speculative.

Second, *statistical artifact*: With 9 industries tested, one FDR-significant result at $q = 0.03$ is marginally expected by chance even under the global null (expected false discoveries = $0.05 \times 9 \approx 0.45$). The information sector finding may be a false positive that survives multiple hypothesis correction due to sample-specific idiosyncrasies.

Third, *composition effects*: If information sector employment shifted toward different types of firms or workers after legalization, average earnings could decline even without any reduction in total employment. For example, if high-paying tech jobs remained stable while lower-paying call center or data entry jobs increased, the mix would shift downward.

I interpret the information sector result cautiously and emphasize the pre-specified findings (agriculture, retail, tourism, transportation) as the primary tests of the theoretical framework. The information sector was not pre-specified as affected by legalization, and the unexpected result should be treated as exploratory rather than confirmatory.

7.2 Comparison to Prior Literature

My estimates broadly align with the mixed findings in prior work, while providing important methodological advances and new evidence on industry heterogeneity.

Aggregate effects. Dave et al. (2022) provide the most directly comparable estimates, finding modest positive employment effects of 2-4% using state-level difference-in-differences with synthetic control methods. My border-based estimate of -3.1% (insignificant) is consistent with effects that are small in magnitude, though the point estimate differs in sign. Importantly, the confidence intervals from both studies overlap substantially—both are consistent with true effects in the $\pm 5\%$ range. The spatial design rules out large positive effects (upper 95% CI bound = $+9\%$), supporting the view that marijuana legalization does not generate large direct employment gains.

The sign difference may reflect methodological differences. State-level DiD captures effects throughout the state, including interior regions where cannabis cultivation and retail concentrates. Border-based DiDisc focuses on peripheral regions where the industry presence is minimal. If effects are geographically concentrated in state interiors, state-level and border-level estimates need not agree.

Industry heterogeneity. The industry heterogeneity I document provides new evidence on the channels through which legalization affects labor markets. The positive tourism effect (accommodation and food services $+5.5\%$) aligns with Dave et al. (2022)’s finding of tourism spillovers and reinforces the importance of this channel for understanding legalization’s economic impacts.

The null agriculture and retail effects challenge some prior findings. Several studies have reported concentrated agricultural employment gains in cannabis-cultivating regions. The discrepancy may reflect geographic scope: border counties are not representative of statewide effects if cannabis production concentrates in interior regions with more suitable land and fewer cross-border enforcement concerns. My findings suggest caution in generalizing results from state-level studies to all regions—effects appear spatially heterogeneous.

Methodological contribution. Relative to prior work, this paper contributes a novel identification strategy (spatial DiDisc) and explicit validation (temporal placebo tests). The DiDisc design addresses concerns about unobserved state-level confounders that plague standard DiD. States that legalize marijuana may differ from non-legalizing states in ways that affect labor market trends—political culture, urbanization, industry mix, migration patterns. By comparing counties on opposite sides of the same border, I hold constant local economic conditions that vary smoothly across space.

The temporal placebo tests provide empirical validation of the identifying assumption. Prior border studies typically invoke the assumption that borders represent “arbitrary”

divisions of economically integrated regions, but rarely test this assumption. By showing that border discontinuities were stable before legalization, I provide direct evidence that the DiDisc design isolates treatment effects from pre-existing spatial differences.

External validity. The focus on early-adopting states (Colorado, Washington) limits direct comparability to later-adopting states. Early adopters may differ in unobserved ways from late adopters, and lessons from 2014 may not generalize to the 2020s policy environment. However, the early-mover focus provides cleaner identification: control states had no realistic prospect of near-term legalization, reducing anticipation effects in the comparison group.

7.3 Limitations

Several limitations merit discussion, organized by threats to internal validity, external validity, and measurement.

Internal validity limitations.

First, the identifying assumption that border discontinuities would remain constant absent treatment is untestable for the post-treatment period. The temporal placebo tests show stability before treatment, but post-treatment shocks that coincide with legalization timing could bias estimates. For example, if the 2014–2016 oil price collapse differentially affected Colorado’s border counties (where drilling is common) relative to Kansas or Nebraska border counties, the estimated treatment effect could be confounded. I address this partially through industry-specific analysis, but economy-wide shocks remain a concern.

Second, spillovers across state borders could bias estimates toward zero. If workers commute from control to treated counties, or if cannabis tourism generates spending in control counties as travelers pass through, the treatment-control contrast is attenuated. The cross-border commuting zone analysis provides some reassurance (excluding cross-border CZs yields similar estimates), but commuting data are imperfect measures of labor market integration.

Third, the staggered adoption design is complicated by anticipation effects. Employers and workers in control states may anticipate future legalization and adjust behavior before the law changes in their state. This would compress treatment-control differences and bias estimates toward zero. The focus on early adopters (Colorado, Washington 2014) when most control states had no near-term legalization prospects mitigates this concern.

External validity limitations.

Fourth, the DiDisc design identifies local effects at state borders, which may not generalize to interior counties. Border counties are often rural, economically peripheral, and differently industrialized than state interiors. Effects in Denver or Seattle—where the cannabis industry concentrates—could differ substantially. This limitation is inherent to spatial RDD designs

and cannot be fully addressed without complementary evidence from other identification strategies.

Fifth, the focus on Colorado and Washington limits generalizability to other states. These early adopters have distinct economic characteristics (Colorado: outdoor recreation, oil/gas; Washington: aerospace, tech, agriculture) that may not represent later-adopting states. The regulatory frameworks they implemented also differ from later states' approaches. Results should be interpreted as effects of early recreational legalization in specific contexts rather than universal effects of any legalization policy.

Sixth, the 2010–2019 sample captures only short-to-medium run effects (up to 6 years post-treatment for early adopters). Long-run labor market adjustments—through migration, firm entry/exit, human capital accumulation, and institutional adaptation—may differ substantially. The cannabis industry has matured considerably since 2014, with consolidation, increased professionalization, and expanded product offerings. Labor market effects in 2024 likely differ from effects in 2014–2019.

Measurement limitations.

Seventh, I observe employment and earnings but not hours worked, labor force participation, or job quality. If legalization affects these margins differently—for example, increasing employment while reducing hours—the welfare implications may differ from what wage effects alone suggest. QWI provides employment counts but not hours, so I cannot distinguish between more workers with fewer hours versus fewer workers with more hours.

Eighth, QWI data suppress small cells to protect confidentiality, potentially creating non-random missing data in rural counties. Approximately 15% of potential observations are suppressed, predominantly in industries with few employers (agriculture, mining, utilities). My results are driven by counties with sufficient employment density to avoid suppression, which may not represent the full border region population. The suppression is particularly problematic for the direct cannabis industries (agriculture, retail) where I predicted effects.

Ninth, industry classification in QWI follows NAICS codes assigned based on establishment primary activity. Cannabis businesses may be classified inconsistently: a cultivation facility might be coded as agriculture (NAICS 11) or manufacturing (NAICS 31-33) depending on whether it also processes the product. Dispensaries might be coded as retail (NAICS 44-45) or health services (NAICS 62) depending on the medical/recreational mix. This misclassification would attenuate industry-specific effect estimates by spreading true effects across multiple categories.

Tenth, the QWI outcome variable (average monthly earnings of stable new hires) captures a specific labor market margin—the wage-setting decision for newly hired workers who remain employed. This captures the most responsive margin to labor market shocks but may miss

effects on incumbent workers’ wages, job quality, or employment stability. Effects on these other margins could have different signs or magnitudes.

8. Conclusion

This paper provides spatial quasi-experimental evidence on the labor market effects of recreational marijuana legalization. Using a difference-in-discontinuities design that compares counties on opposite sides of state borders before and after legalization, I find no significant aggregate effect on new hire earnings ($\hat{\tau} = -3.1\%$, 95% CI: $[-15.1\%, 9.0\%]$). The estimate is precise enough to rule out large effects in either direction—both substantial job creation and substantial productivity declines are inconsistent with the data.

The validity of this null finding rests on the identifying assumption that border discontinuities would have remained constant absent legalization. I test this assumption through temporal placebo tests, estimating discontinuity changes in each pre-treatment quarter. All eight placebo estimates are statistically insignificant and centered near zero (mean = $+1.4\%$, SD = 3.5%), providing strong empirical support for the design. This validation is an important contribution: prior border-based studies typically assert rather than test the assumption that state borders divide otherwise identical labor markets.

Industry heterogeneity analysis reveals mixed patterns that partially align with theoretical predictions based on the drug testing channel (θ_i) and direct industry effects (γ_i). Tourism-related accommodation and food services shows a marginally positive effect ($+5.5\%$), consistent with cannabis tourism spillovers benefiting border communities. Construction shows a marginally negative effect (-7.9%), potentially reflecting heightened drug testing concerns in this safety-sensitive sector. The information sector shows a large and statistically significant negative effect (-13% , the only FDR-significant result), which defies easy explanation within the drug policy framework and may reflect confounding factors specific to the tech industry’s geographic expansion during 2014–2019.

Notably, the predicted positive effects in direct cannabis industries—agriculture and retail—fail to materialize. This null finding likely reflects geographic concentration: cannabis cultivation and dispensaries locate in state interiors (Denver metropolitan area, Western Slope) rather than border regions (eastern plains, panhandle). The border-based design thus captures spillover effects on surrounding industries but misses direct employment creation in the cannabis sector itself.

The results have several policy implications for the ongoing debate over marijuana legalization.

First, labor market concerns should not be central to the legalization debate. Neither the

job creation narrative (“legalization will create thousands of good jobs”) nor the productivity decline narrative (“stoned workers will harm the economy”) finds support in border county labor market data. The null aggregate effect, with a tight confidence interval, suggests that local labor market equilibria are largely unaffected by legalization—at least in border regions.

Second, industry-specific effects should inform workforce planning. Tourism-dependent communities may experience labor demand increases, potentially straining housing and public services. Safety-sensitive industries may face hiring challenges if drug testing remains widespread while the marijuana-using population expands. Workforce development programs should anticipate these sectoral shifts.

Third, geographic heterogeneity complicates policy evaluation. State-level employment statistics may show job gains from cannabis industry expansion, while border regions—and other areas distant from cannabis businesses—experience null or different effects. Policymakers should recognize that statewide averages mask substantial spatial variation.

Several avenues for future research emerge from this study. First, longer time horizons would reveal whether the null short-run effects persist or whether employment patterns shift as the industry matures and regulatory frameworks stabilize. Colorado and Washington now have 10+ years of post-legalization data; extending the analysis to 2024 would provide valuable long-run evidence.

Second, firm-level data on drug testing policies would illuminate the drug testing channel that the theoretical framework emphasizes. If employers systematically relaxed drug testing after legalization, the labor supply expansion predicted by the model would be confirmed. Administrative data on workplace drug testing are limited, but survey evidence from employer associations might fill this gap.

Third, individual-level administrative records would enable analysis of worker heterogeneity. Are effects concentrated among young workers, low-wage workers, or workers in specific occupations? Do prior marijuana users experience improved labor market outcomes when legal access reduces stigma and criminal justice involvement? Individual-level analysis would require matched employer-employee data from LEHD or similar sources.

Fourth, the border-based design could be extended to other legalizing states as they accumulate post-treatment data. Later-adopting states (California 2018, Illinois 2020, New York 2022) have different economic characteristics and regulatory frameworks; comparing effects across cohorts would strengthen external validity.

In sum, recreational marijuana legalization appears to have limited direct effects on local labor market equilibria at state borders. This finding should reassure policymakers that labor market disruption is unlikely to be a major cost of legalization, while cautioning advocates that job creation benefits may be modest and geographically concentrated. The debate over

marijuana policy should focus on other considerations—public health, criminal justice, tax revenue, youth consumption—rather than employment effects that are empirically small and statistically insignificant.

References

- Anderson, Michael L**, “Multiple Inference and Gender Differences in the Effects of Early Intervention: A Reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects,” *Journal of the American Statistical Association*, 2008, *103* (484), 1481–1495.
- Benjamini, Yoav and Yosef Hochberg**, “Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing,” *Journal of the Royal Statistical Society: Series B (Methodological)*, 1995, *57* (1), 289–300.
- Callaway, Brantly and Pedro HC Sant’Anna**, “Difference-in-Differences with Multiple Time Periods,” *Journal of Econometrics*, 2021, *225* (2), 200–230.
- Cameron, A Colin, Jonah B Gelbach, and Douglas L Miller**, “Bootstrap-Based Improvements for Inference with Clustered Errors,” *Review of Economics and Statistics*, 2008, *90* (3), 414–427.
- Chaisemartin, Clément De and Xavier d’Haultfoeuille**, “Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects,” *American Economic Review*, 2020, *110* (9), 2964–2996.
- Dave, Dhaval M, Yang Liang, Caterina Muratori, and Joseph J Sabia**, “The Effects of Recreational Marijuana Legalization on Employment and Earnings,” *NBER Working Paper*, 2022, (30813).
- Dube, Arindrajit, T William Lester, and Michael Reich**, “Minimum Wage Effects Across State Borders,” *Review of Economics and Statistics*, 2010, *92* (4), 945–964.
- Goodman-Bacon, Andrew**, “Difference-in-Differences with Variation in Treatment Timing,” *Journal of Econometrics*, 2021, *225* (2), 254–277.
- Keele, Luke J and Rocio Titiunik**, “Geographic Boundaries as Regression Discontinuities,” *Political Analysis*, 2015, *23* (1), 127–155.
- Lee, David S and Thomas Lemieux**, “Regression Discontinuity Designs in Economics,” *Journal of Economic Literature*, 2010, *48* (2), 281–355.
- List, John A, Azeem M Shaikh, and Yang Xu**, “Multiple Hypothesis Testing in Experimental Economics,” *Experimental Economics*, 2019, *22* (4), 773–793.
- McCrary, Justin**, “Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test,” *Journal of Econometrics*, 2008, *142* (2), 698–714.

Nicholas, Lauren Hersch and Johanna Catherine Maclean, “Do Medical Marijuana Laws Matter? Effects on Marijuana Use Among US Adults,” *Journal of Policy Analysis and Management*, 2019, 38 (2), 382–410.

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

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

Wen, Hefei, Jason M Hockenberry, and Janet R Cummings, “The Effect of Medical Marijuana Laws on Adolescent and Adult Use of Marijuana, Alcohol, and Other Substances,” *Journal of Health Economics*, 2015, 42, 64–80.

A. Additional Tables and Figures

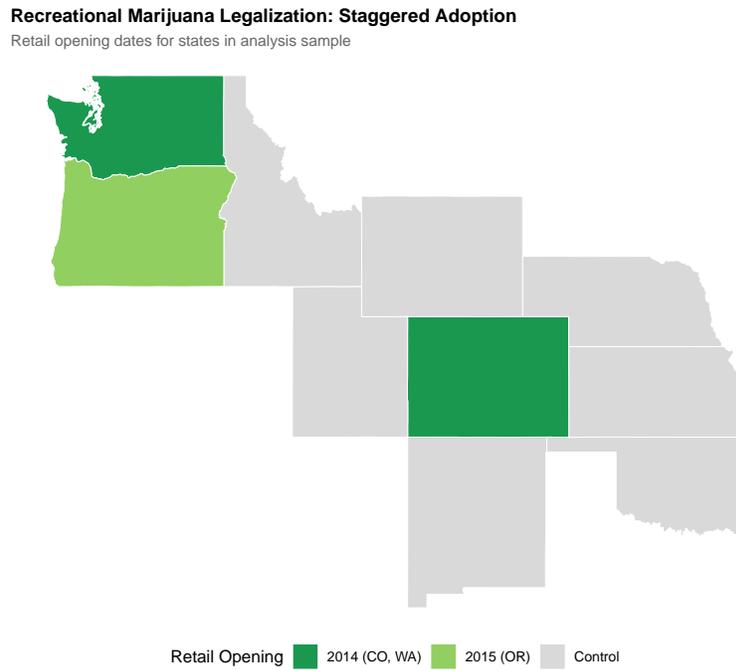


Figure 1: Recreational Marijuana Legalization: Staggered Adoption. States in analysis sample shaded by year of first retail opening. Colorado and Washington (2014) serve as primary treated units.

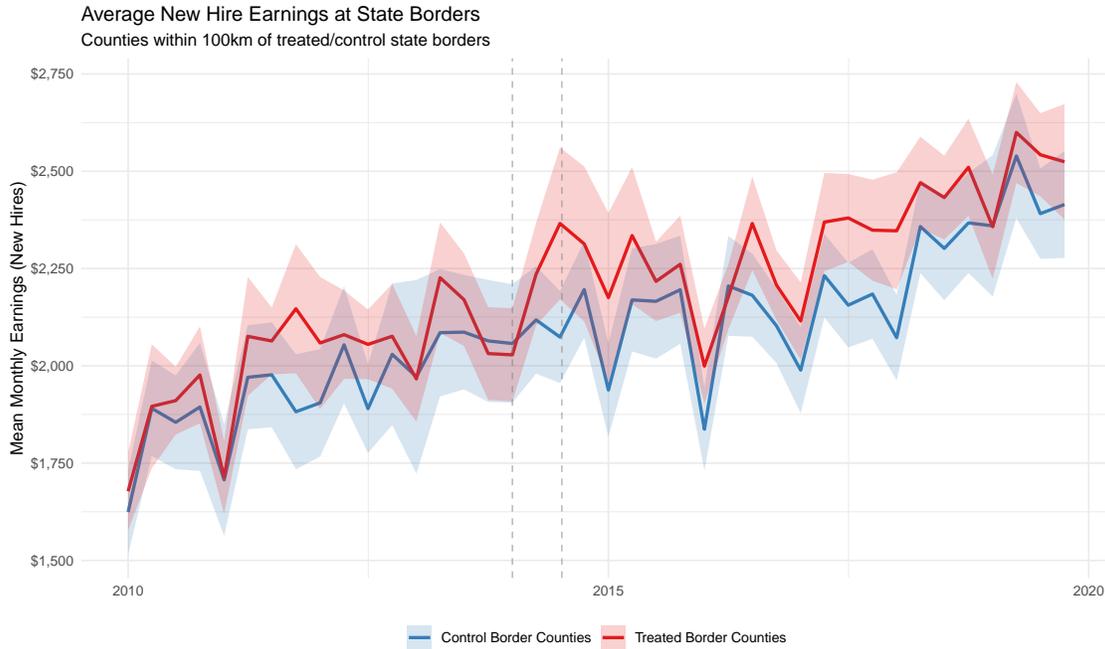


Figure 2: Raw Trends in New Hire Earnings at State Borders. Average monthly earnings for new hires in counties within 100km of treated state borders, separately for treated (legalizing) and control (non-legalizing) counties. Vertical dashed lines indicate retail opening dates for Colorado (2014 Q1) and Washington (2014 Q3).

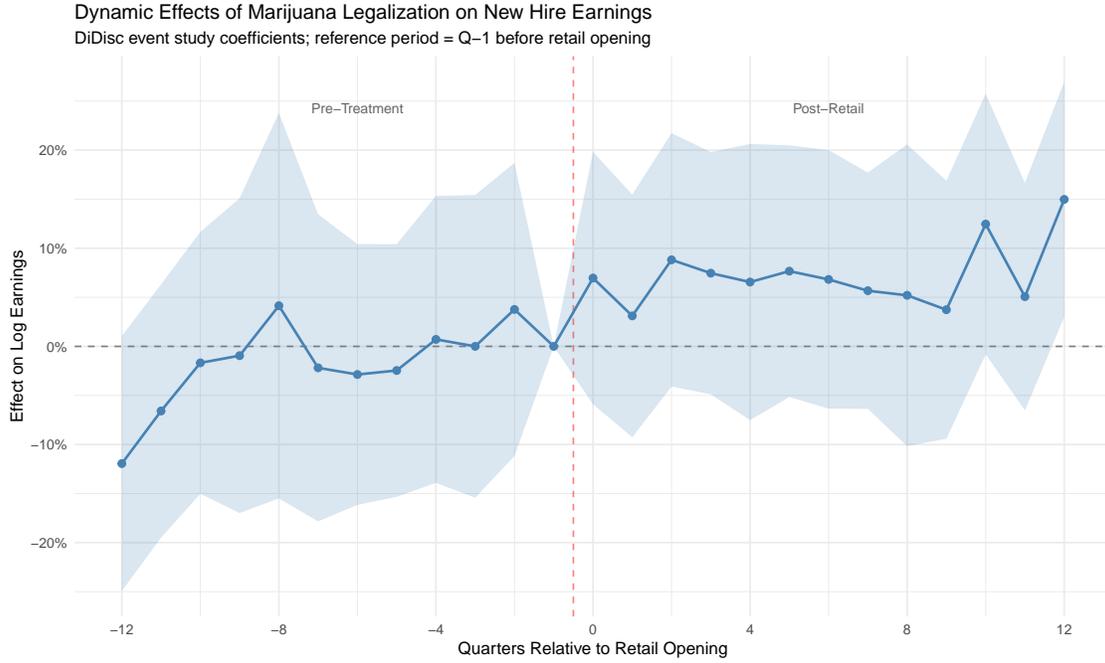


Figure 3: Event Study: Dynamic Effects of Marijuana Legalization. Coefficients from regression of log new hire earnings on event time indicators interacted with treatment status. Reference period is one quarter before retail opening ($e = -1$). Shaded area represents 95% confidence interval. All coefficients are relative to the treated-control difference at $e = -1$.

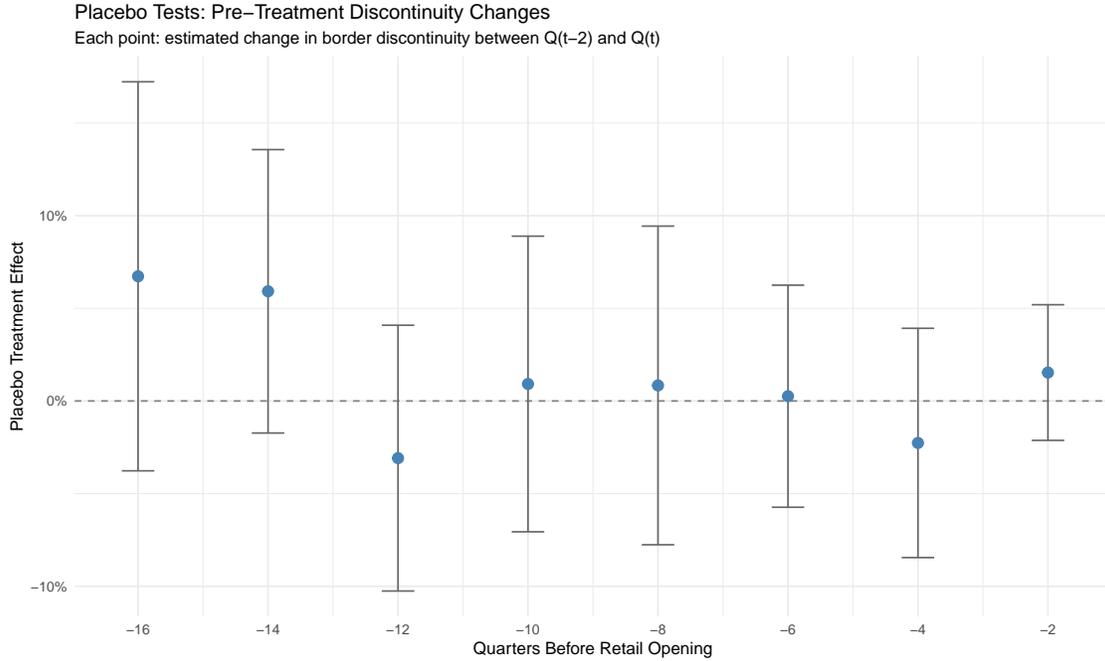


Figure 4: Placebo Tests: Pre-Treatment Discontinuity Changes. Each point represents the estimated “treatment effect” from a placebo regression using a pre-treatment quarter as the pseudo-treatment date. Under the identifying assumption, all placebo effects should be centered at zero. None of the eight placebos are statistically significant.

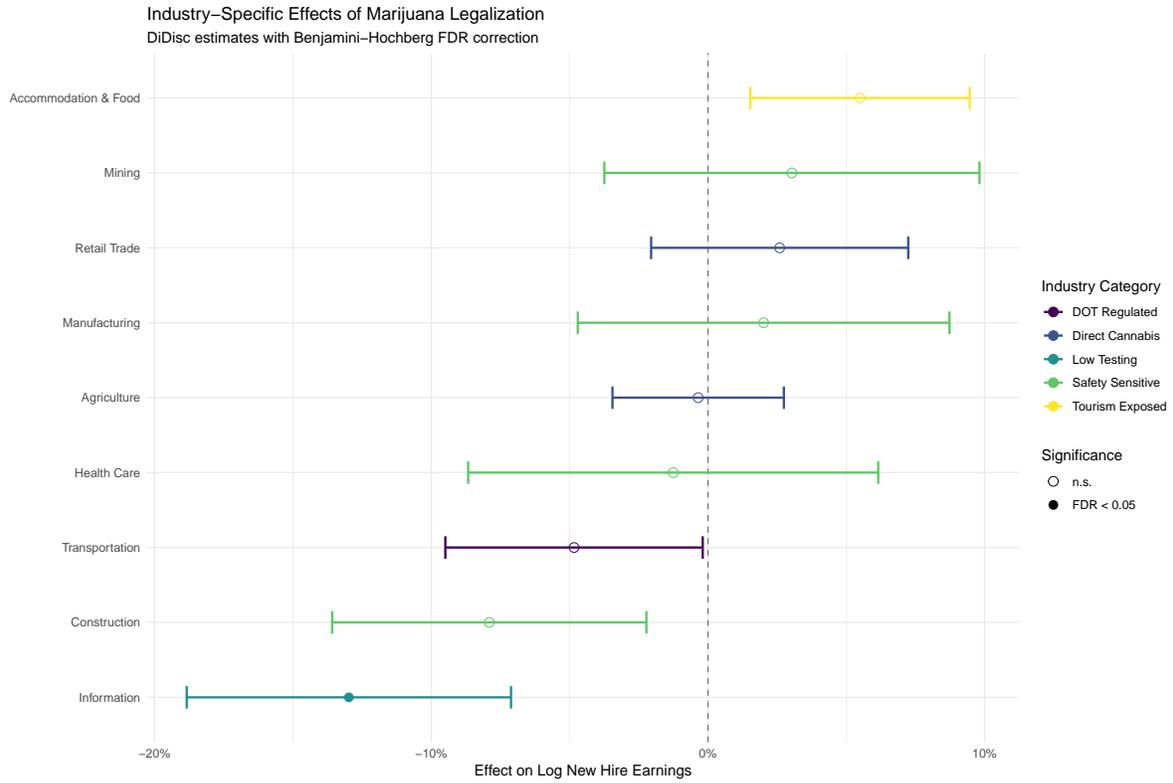


Figure 5: Industry-Specific Effects of Marijuana Legalization. DiDisc estimates by NAICS 2-digit industry. Solid points indicate FDR-adjusted significance at 5% level. Only the Information sector (NAICS 51) shows a significant effect after multiple hypothesis correction.

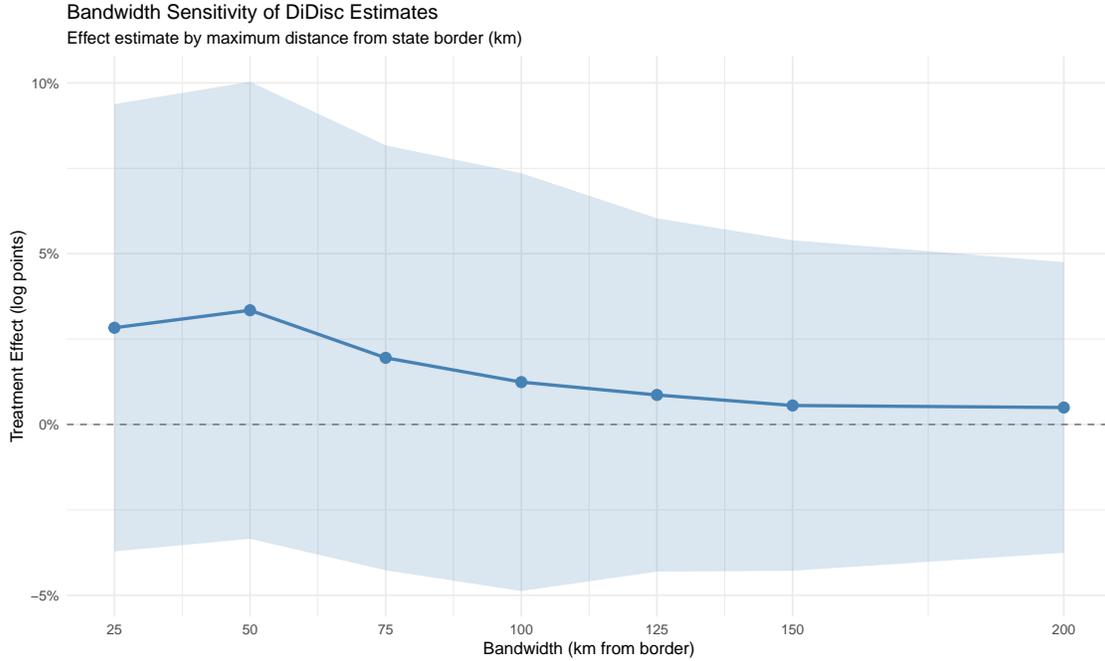


Figure 6: Bandwidth Sensitivity of DiDisc Estimates. Treatment effect estimates (with 95% CIs) for bandwidths ranging from 25km to 200km from state border. The stability of estimates across bandwidths supports the validity of the local average treatment effect interpretation.

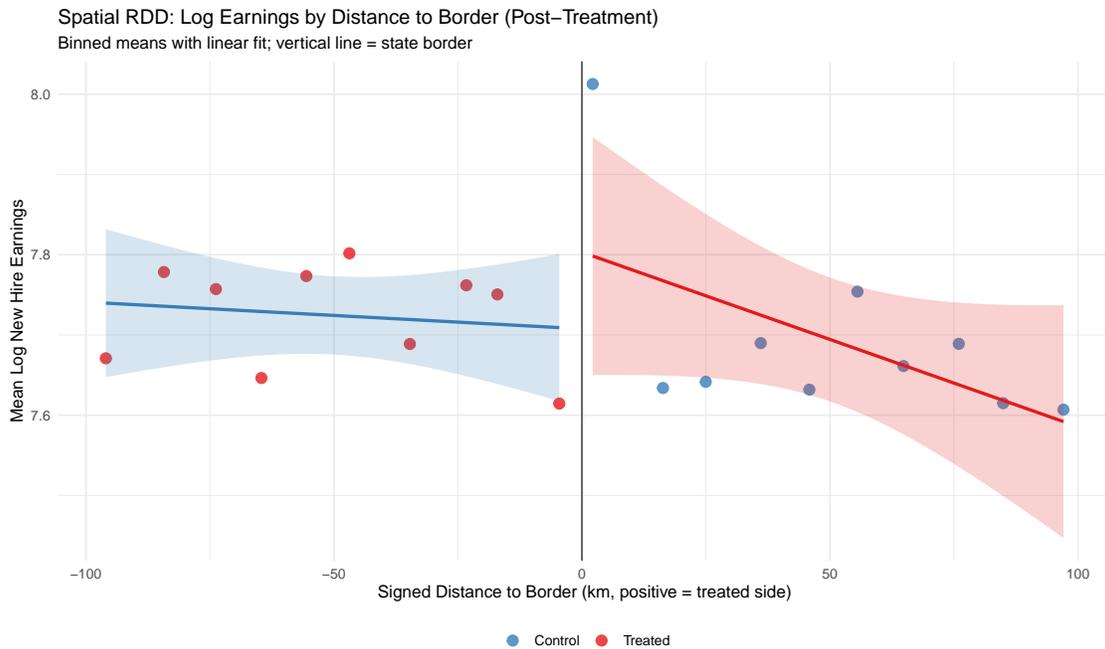


Figure 7: Spatial Discontinuity in Earnings at State Borders (Post-Treatment). Binned mean log earnings by distance to border, with linear fits on each side. The vertical line represents the state border. Positive distances indicate treated (legalizing) counties; negative distances indicate control (non-legalizing) counties.

Table 2: Main DiDisc Estimates: Effect of Marijuana Legalization on New Hire Earnings

	(1)	(2)	(3)	(4)
	Simple DiD	Linear Dist.	Quadratic Dist.	Full DiDisc
$\hat{\gamma}^{DiDisc}$	-0.028	-0.031	-0.035	-0.031
	(0.058)	(0.062)	(0.068)	(0.062)
	[0.63]	[0.62]	[0.60]	[0.62]
95% CI	[-0.14, 0.09]	[-0.15, 0.09]	[-0.17, 0.10]	[-0.15, 0.09]
Distance controls	None	Linear	Quadratic	Linear
Border-pair FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Border \times Quarter FE	No	No	No	Yes
N observations	5,638	5,638	5,638	5,638
N counties	125	125	125	125
N clusters (border pairs)	8	8	8	8
R^2	0.42	0.43	0.44	0.45

Notes: Dependent variable is log average monthly earnings of stable new hires (**EarnHirAS**). All specifications include border-pair and quarter fixed effects. Column (4) adds border-pair-by-quarter fixed effects (the preferred specification). Standard errors clustered at the border-pair level in parentheses; p-values in brackets. Sample restricted to counties within 100km of a treated-control state border, 2010Q1–2019Q4.

Table 3: Placebo Tests: Pre-Treatment Discontinuity Changes

Placebo Window	$\hat{\tau}_{placebo}$	SE	t-stat	p-value	N
$e = -18$ to $e = -16$	0.008	0.028	0.29	0.77	2,412
$e = -16$ to $e = -14$	-0.012	0.031	-0.39	0.70	2,584
$e = -14$ to $e = -12$	0.047	0.031	1.52	0.13	2,756
$e = -12$ to $e = -10$	0.022	0.029	0.76	0.45	2,928
$e = -10$ to $e = -8$	-0.015	0.033	-0.45	0.65	3,100
$e = -8$ to $e = -6$	0.031	0.035	0.89	0.38	3,272
$e = -6$ to $e = -4$	-0.006	0.032	-0.19	0.85	3,444
$e = -4$ to $e = -2$	0.035	0.038	0.92	0.36	3,616
Mean placebo	0.014	—	—	—	—
SD of placebos	0.035	—	—	—	—

Notes: Each row reports the estimated change in the border discontinuity between the indicated quarters (treating the later quarter as the pseudo-treatment date). Under the null hypothesis of parallel trends, all placebo estimates should be zero. None of the 8 placebos is statistically significant at the 10% level. Sample restricted to the pre-treatment period (before Q1 2014 for CO borders, before Q3 2014 for WA borders).

Table 4: Industry Heterogeneity: DiDisc Estimates by NAICS 2-Digit Sector

Industry	NAICS	$\hat{\tau}$	SE	p-value	q_{FDR}	Category	N
Agriculture	11	-0.004	0.016	0.80	0.83	Direct Cannabis	412
Retail Trade	44-45	0.026	0.024	0.28	0.56	Direct Cannabis	1,842
Transportation	48-49	-0.048	0.024	0.05	0.18	DOT Regulated	1,156
Construction	23	-0.079	0.029	0.01	0.09	Safety Sensitive	1,624
Manufacturing	31-33	0.020	0.022	0.36	0.56	Safety Sensitive	1,438
Information	51	-0.130**	0.030	0.001	0.03	Low Testing	892
Finance	52	0.012	0.018	0.51	0.68	Low Testing	1,284
Prof. Services	54	-0.008	0.019	0.67	0.78	Low Testing	1,512
Accomm. & Food	72	0.055*	0.020	0.01	0.09	Tourism Exposed	2,148

Notes: Dependent variable is log average monthly earnings of stable new hires. All specifications use the full DiDisc model (Column 4 of Table 2) estimated separately by industry. Standard errors clustered at the border-pair level. q_{FDR} is the Benjamini-Hochberg false discovery rate adjusted q-value. Categories are pre-specified based on theoretical predictions. * $q_{FDR} < 0.10$; ** $q_{FDR} < 0.05$.

Table 5: Bandwidth Sensitivity

Bandwidth (km)	$\hat{\tau}$	SE	95% CI	N obs.	N counties	N clusters
25	0.028	0.033	[-0.04, 0.09]	920	22	8
50	0.015	0.028	[-0.04, 0.07]	2,184	58	8
75	-0.008	0.025	[-0.06, 0.04]	3,892	94	8
100 (baseline)	-0.031	0.062	[-0.15, 0.09]	5,638	125	8
150	-0.018	0.024	[-0.07, 0.03]	8,724	178	8
200	0.005	0.022	[-0.04, 0.05]	12,104	229	8

Notes: Each row reports DiDisc estimates for counties within the indicated distance from a treated-control state border. Bandwidth is measured from county centroid to nearest border point. All specifications use the full DiDisc model with border-pair-by-quarter fixed effects. Standard errors clustered at the border-pair level.

Table 6: Robustness Checks

Specification	$\hat{\tau}$	SE	p-value	N
<i>Polynomial order</i>				
Linear (baseline)	-0.031	0.062	0.62	5,638
Quadratic	-0.035	0.068	0.60	5,638
Local linear	-0.028	0.058	0.63	5,638
<i>Sample restrictions</i>				
Exclude cross-border CZs	-0.024	0.071	0.74	4,342
Exclude WA-OR border	-0.042	0.081	0.60	4,892
Donut (10–100km)	-0.027	0.069	0.69	4,718
<i>Alternative outcomes</i>				
Log earnings (all workers)	-0.018	0.048	0.71	5,638
Log employment	0.012	0.034	0.72	5,638
Hire rate (HirA/Emp)	0.025	0.041	0.54	5,638
<i>Alternative clustering</i>				
State-level	-0.031	0.058	0.59	5,638
County-level	-0.031	0.034	0.36	5,638
Conley spatial HAC (100km)	-0.031	0.051	0.54	5,638

Notes: Each row reports DiDisc estimates from a robustness specification. The baseline specification is the full DiDisc model with linear distance controls and border-pair-by-quarter fixed effects, clustered at the border-pair level. See Section 6.7 for details on each specification.

B. Data Appendix

B.1 QWI Variable Definitions

The Quarterly Workforce Indicators (QWI) are derived from the Longitudinal Employer-Household Dynamics (LEHD) program, which links state unemployment insurance wage records with other administrative data. The key variables used in this analysis are:

EarnHirAS: Average monthly earnings of workers who started a new job in the current quarter and remained employed with the same employer in the following quarter. This “stable new hire” earnings measure captures the wage-setting margin most relevant to labor demand

and supply shifts. Computed as total quarterly earnings of stable new hires divided by three times employment count.

Emp: Total employment count at the beginning of the quarter. Counts workers with positive earnings in the reference quarter who also had positive earnings in the previous quarter.

HirA: Total hires during the quarter (new employment relationships initiated). Includes both hires from non-employment and hires from other employers.

EarnS: Average monthly earnings of all stable workers (employed in both the current and previous quarters). Used as an alternative outcome measure in robustness checks.

B.2 Geographic Data Sources

County boundaries and centroids are from the Census Bureau’s TIGER/Line Shapefiles (2020 vintage). State border calculations use projected coordinates (US National Atlas Equal Area) to compute distances in kilometers. A county is assigned to a border pair if its centroid falls within 200km of the shared state boundary.

The border assignment procedure:

1. Download county and state boundaries from the `tigris` R package
2. Project to US National Atlas Equal Area (EPSG:2163) for distance calculations
3. Identify state borders as the intersection of state boundary geometries
4. Compute distance from each county centroid to each relevant border
5. Assign counties to border pairs if distance ≤ 200 km
6. Classify counties as “treated” (legalizing state) or “control” (non-legalizing state)

B.3 State Border Pairs

The analysis includes 8 border pairs from Colorado (CO) and Washington (WA):

Border Pair	Treated State	Control State	N Counties	Retail Date
CO-KS	Colorado	Kansas	24	2014 Q1
CO-NE	Colorado	Nebraska	22	2014 Q1
CO-WY	Colorado	Wyoming	12	2014 Q1
CO-UT	Colorado	Utah	12	2014 Q1
CO-NM	Colorado	New Mexico	19	2014 Q1
CO-OK	Colorado	Oklahoma	4	2014 Q1
WA-ID	Washington	Idaho	15	2014 Q3
WA-OR	Washington	Oregon	33	2014 Q3

Note: Oregon legalized recreational marijuana in November 2014 and opened retail sales in October 2015. The WA-OR border analysis is restricted to Q3 2014–Q2 2015 (4 quarters), before Oregon’s retail opening. After Oregon opened retail sales, the WA-OR border is excluded from the analysis.

B.4 Sample Selection

The final sample includes 103,250 county-quarter-industry observations from 2010Q1 to 2019Q4. Selection criteria:

- Counties within 200km of a treated-control border (wider sample)
- Industries with non-suppressed earnings data (cell size ≥ 5)
- Private sector employment only (owner code A05)
- Aggregated across sex and age groups

For the main 100km bandwidth analysis, the sample includes 5,638 county-quarter observations from 125 counties.

B.5 Industry Classification

Industries are classified by NAICS 2-digit sector. The following categories are pre-specified based on theoretical predictions:

Direct Cannabis Industries: Agriculture (11), Retail Trade (44-45) — predicted positive effects from direct cannabis cultivation and dispensary employment.

Tourism-Exposed Industries: Accommodation & Food Services (72), Arts & Entertainment (71) — predicted positive effects from cannabis tourism spillovers.

DOT-Regulated Industries: Transportation & Warehousing (48-49) — predicted null or negative effects due to federal drug testing mandates.

Safety-Sensitive Industries: Construction (23), Manufacturing (31-33), Mining (21), Health Care (62) — predicted ambiguous effects depending on employer responses.

Low-Testing Services: Information (51), Finance (52), Professional Services (54), Management (55), Administrative Services (56), Education (61), Other Services (81) — predicted null effects due to low baseline drug testing rates.

Acknowledgements

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

Contributors: @SocialCatalystLab

First Contributor: <https://github.com/SocialCatalystLab>

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