

# Legal Weed, Self-Made? Does Recreational Marijuana Access at Age 21 Shift Workers Into Self-Employment?

## A Difference-in-Discontinuities Analysis

Autonomous Policy Evaluation Project nd @dakoyana

January 2026

### Abstract

We investigate whether legal recreational marijuana access at age 21 shifts workers from traditional employment into self-employment. The theoretical mechanism is straightforward: while Colorado legalized recreational marijuana for adults 21 and older in 2014, employers retain the right to terminate workers for off-duty marijuana use (upheld in *Coats v. Dish Network*, 2015). This creates a trade-off where workers who wish to consume marijuana legally may prefer self-employment, where no employer can enforce drug-testing policies. Using American Community Survey data from 2015-2022 and a difference-in-discontinuities design comparing Colorado to six control states where recreational marijuana remained illegal, *we find no statistically robust evidence that marijuana legalization affects self-employment*. Point estimates are positive—1.05 pp for overall self-employment, 0.97 pp for incorporated self-employment—but with only 7 state clusters (1 treated, 6 controls), conventional standard errors are unreliable. Wild cluster bootstrap inference yields p-values of 0.42 and 0.26, respectively, far above conventional significance thresholds. Permutation inference produces similar null results ( $p = 0.86$ ). Our methodological contribution is demonstrating how point

estimates that appear “highly significant” under conventional clustering ( $p < 0.001$ ) can be entirely attributable to noise under appropriate inference methods. The positive point estimates are consistent with theoretical predictions but cannot be distinguished from chance. Future research with more treated states or alternative identification strategies is needed to test this mechanism credibly.

**Keywords:** Marijuana legalization, self-employment, regression discontinuity, labor supply, Colorado Amendment 64, inference with few clusters, wild cluster bootstrap

# Contents

# 1 Introduction

The legalization of recreational marijuana represents one of the most significant policy experiments in recent American history. Since Colorado and Washington became the first states to legalize recreational marijuana in 2012, with retail sales beginning in 2014, over 20 additional states have followed. This rapid policy change has sparked intense debate about the effects on public health, crime, tax revenue, and social outcomes. However, one channel that has received surprisingly little attention is the potential effect on labor market structure—specifically, whether legalization affects the composition of employment between traditional wage work and self-employment.

We investigate a novel mechanism through which marijuana legalization might affect employment composition. The key insight is that while states have legalized recreational marijuana for adults 21 and older, federal law continues to classify marijuana as a Schedule I controlled substance. This federal-state conflict creates an important wrinkle: employers generally retain the right to enforce drug-free workplace policies and terminate employees for marijuana use, even when that use is legal under state law and occurs entirely off-duty.

This employer right was definitively established in Colorado through the landmark case *Coats v. Dish Network Corp.* (2015). Brandon Coats, a quadriplegic who used medical marijuana for muscle spasms, was terminated after a random drug test detected THC metabolites. The Colorado Supreme Court upheld his termination, ruling that Colorado’s “lawful activities” statute protecting off-duty conduct does not extend to marijuana use because it remains federally illegal. The ruling established that Colorado employers can fire workers for any marijuana use, regardless of when or where it occurs.

This legal reality creates a stark trade-off for workers who wish to consume marijuana legally. Traditional wage and salary workers face the risk of termination through drug testing, while self-employed workers—who have no employer to enforce such policies—can consume without employment consequences. If this trade-off is economically meaningful, we would expect to observe an increase in self-employment among workers who gain legal marijuana

access.

Colorado’s policy creates a natural experiment because recreational marijuana becomes legal at exactly age 21, the same threshold used for alcohol. This sharp age discontinuity allows us to use regression discontinuity design (RDD) methods to estimate causal effects. However, a simple RDD comparing workers just below and above age 21 in Colorado would confound the marijuana effect with the alcohol effect, since both substances become legal at the same age.

We address this identification challenge using a difference-in-discontinuities design. We compare the discontinuity at age 21 in Colorado (where both alcohol and marijuana become legal) to the discontinuity in control states where only alcohol becomes legal. The difference between these two discontinuities isolates the marijuana-specific effect, netting out any common age-21 alcohol effects.

Using American Community Survey (ACS) data from 2015-2022, we analyze 57,685 young adults aged 18-24 across Colorado and six control states (Texas, Florida, Georgia, North Carolina, Ohio, and Pennsylvania) where recreational marijuana remained illegal throughout our sample period. Our findings reveal that marijuana legalization does affect employment structure, but through specific channels rather than broad labor market restructuring.

However, our analysis also demonstrates the critical importance of using appropriate inference methods. With standard state-clustered standard errors (7 clusters: 1 treated, 6 controls), our point estimates appear statistically significant—0.97 pp for incorporated self-employment with  $p < 0.001$ . But conventional cluster-robust standard errors are unreliable with fewer than 30-50 clusters. When we apply wild cluster bootstrap inference following [?](#) , the p-value rises to 0.26, and permutation inference yields  $p = 0.86$ . We therefore cannot reject the null hypothesis at conventional significance levels.

The point estimates are positive and consistent with theory: a large incorporated self-employment effect (0.97 pp, more than doubling the baseline), no effect on unincorporated self-employment, and a positive but imprecise overall effect. This pattern suggests that

if marijuana legalization does affect employment composition, it operates through formal business formation rather than casual gig work. But our statistical power is insufficient to establish this as a causal finding rather than noise.

Our paper contributes to several literatures. First, we contribute to the rapidly growing literature on the effects of marijuana legalization. Existing work has examined effects on marijuana use (?), crime (?), traffic safety, and tax revenue. We provide the first analysis of employment composition effects using the age-21 discontinuity.

Second, we contribute to the literature on self-employment determinants. Previous work has examined how regulatory burden, access to health insurance, and labor market conditions affect self-employment decisions. We examine a novel factor: the ability to consume a legal substance without employment consequences.

Third, we contribute methodologically by demonstrating the difference-in-discontinuities design in a labor market context. This approach allows credible causal inference even when the treatment threshold coincides with other discontinuities.

The remainder of this paper proceeds as follows. Section 2 reviews related literature. Section 3 provides institutional background on marijuana legalization and employment law. Section 4 describes our data and presents summary statistics. Section 5 presents our empirical methodology. Section 6 reports results. Section 7 discusses implications and limitations. Section 8 concludes.

## **2 Related Literature**

Our paper contributes to three strands of literature: the effects of marijuana legalization, determinants of self-employment, and difference-in-discontinuities methodology.

## 2.1 Effects of Marijuana Legalization

A substantial literature has emerged examining the consequences of marijuana legalization. ? find that recreational marijuana laws are associated with increased marijuana use among adolescents in some states. ? document that medical marijuana laws increase marijuana use among adults, with effects concentrated among those with chronic pain conditions. ? examine crime effects of recreational marijuana legalization in Washington and Oregon, finding reductions in some categories of crime.

Several papers examine labor market effects of marijuana laws. ? analyze the effect of medical marijuana laws on labor market outcomes, finding mixed evidence on employment and wages. ? study traffic fatalities and alcohol consumption in response to medical marijuana laws, finding evidence of substitution away from alcohol. These papers generally examine state-level policy changes rather than the age-21 discontinuity we exploit.

The effect of marijuana legalization on self-employment has received almost no attention. ? examines workplace drug testing and worker drug use, finding that testing reduces drug use among workers. ? study the relationship between drug testing programs and worker productivity. However, neither paper examines how drug testing might push workers toward self-employment where testing does not apply.

## 2.2 Determinants of Self-Employment

The self-employment literature has identified numerous factors affecting the decision to become self-employed. ? document that self-employment is associated with personality traits, family background, and access to capital. ? emphasize the importance of liquidity constraints, showing that wealth shocks increase transitions into self-employment.

Prior work has also examined how regulatory and institutional factors affect self-employment. The Affordable Care Act's health insurance provisions, by decoupling health insurance from employment, may have increased self-employment (?). Occupational licensing requirements can reduce self-employment by raising barriers to entry. Our paper identifies a novel fac-

tor: the ability to consume a legal substance without employment consequences. To our knowledge, we are the first to examine how substance regulation might affect employment composition.

## 2.3 Difference-in-Discontinuities Methodology

Our identification strategy builds on the regression discontinuity design (RDD), which has become a standard tool in empirical economics (??). The key insight of RDD is that comparing units just above and below a threshold provides quasi-experimental variation, as units near the threshold are similar on average.

However, when multiple treatments occur at the same threshold, standard RDD cannot separate their effects. ? develop the difference-in-discontinuities design to address this problem. They compare discontinuities across groups where only some face the treatment of interest, with others facing a common discontinuity that serves as a counterfactual. We apply this approach to the marijuana-alcohol confound at age 21.

? provide practical guidance on RDD implementation, including bandwidth selection and robust inference. We follow their recommendations by presenting results across multiple bandwidths and testing for manipulation of the running variable using the ? density test.

## 3 Institutional Background

### 3.1 Marijuana Legalization in Colorado

Colorado voters approved Amendment 64 in November 2012, making Colorado one of the first two states (along with Washington) to legalize recreational marijuana for adults. The amendment took effect on December 10, 2012, immediately legalizing possession of up to one ounce for adults 21 and older. Retail sales began on January 1, 2014, when the first licensed dispensaries opened.

Amendment 64 established key parameters that remain in effect today. Adults 21 and

older may purchase up to one ounce of marijuana (or its equivalent in other products) per transaction from licensed retailers. Home cultivation of up to six plants is permitted. Consumption is prohibited in public spaces and while driving.

The choice of age 21 as the legal threshold aligned with alcohol policy and was not arbitrary. Advocates explicitly argued that marijuana regulation should mirror alcohol regulation, including the minimum age. The amendment’s official name—the “Regulate Marijuana Like Alcohol Act”—reflected this framing. This design creates a sharp discontinuity in legal access at exactly age 21.

### **3.2 Employer Drug Testing and *Coats v. Dish Network***

Despite state legalization, marijuana remains a Schedule I controlled substance under the federal Controlled Substances Act. This federal illegality has important implications for employment.

Colorado’s “lawful activities” statute (C.R.S. §24-34-402.5) prohibits employers from terminating employees for engaging in lawful activities off the employer’s premises during non-working hours. When marijuana became legal under state law, the question arose whether this statute protected workers from termination for off-duty marijuana use.

The Colorado Supreme Court resolved this question in *Coats v. Dish Network Corp.* (2015). Brandon Coats, a telephone customer service representative who was paralyzed in a car accident as a teenager, used medical marijuana to treat severe muscle spasms. He never used marijuana at work and was by all accounts an exemplary employee. However, a random drug test detected THC metabolites, and Dish Network terminated him.

Coats argued that his termination violated the lawful activities statute because medical marijuana use was legal under state law. The Colorado Supreme Court disagreed, holding that “lawful” in the statute means lawful under both state *and* federal law. Because marijuana remains federally illegal, off-duty use is not a “lawful activity” protected by the statute.

The *Coats* ruling established clearly that Colorado employers can terminate workers for any marijuana use—medical or recreational, on-duty or off-duty. Workers in Colorado who use marijuana, even legally and responsibly, risk losing their jobs if they fail a drug test.

### **3.3 Drug Testing Prevalence**

Drug testing is widespread in the American labor market. According to survey data, approximately 56% of employers conduct some form of drug testing, with rates higher in certain industries. Pre-employment testing is most common, but random and post-incident testing also occur.

Drug testing is particularly prevalent in industries with federal drug testing mandates or strong safety concerns. The Department of Transportation mandates testing for workers in “safety-sensitive” positions including commercial drivers, pilots, rail workers, and pipeline operators. Industries such as construction, manufacturing, healthcare, and transportation have above-average testing rates even for positions not subject to federal mandates.

For workers in tested industries, the *Coats* ruling creates a meaningful constraint. They cannot consume marijuana—even legally, responsibly, and entirely off-duty—without risking their jobs. This constraint does not apply to self-employed workers, who set their own policies.

### **3.4 Self-Employment as an Alternative**

Self-employment offers workers autonomy over their own workplace policies, including drug testing. A self-employed consultant, freelancer, or small business owner faces no employer-mandated drug tests. While clients might theoretically require testing, this is rare outside of highly regulated contexts.

For workers who value marijuana consumption, self-employment offers a way to participate legally in the marijuana market without employment risk. This is the mechanism we investigate: does legal marijuana access at age 21 shift workers toward self-employment?

We emphasize that this mechanism operates at the margin. Most workers do not use marijuana, and most marijuana users are not heavy users for whom employment risk is a primary concern. The question is whether the marginal effect is large enough to be detectable in aggregate data.

## 4 Data and Summary Statistics

### 4.1 Data Source

We use data from the American Community Survey (ACS) Public Use Microdata Sample (PUMS) for years 2015-2022. The ACS is an annual survey conducted by the U.S. Census Bureau with a sample of approximately 3.5 million households. The PUMS files provide individual-level data including demographics, employment status, industry, and income.

We focus on the period 2015-2022 for several reasons. First, 2015 is the first full year after retail marijuana sales began in Colorado (January 2014). Second, we exclude 2020 because the ACS was substantially disrupted by the COVID-19 pandemic, with response rates falling and data collection methods changing. Third, 2022 is the most recent year with available data at the time of analysis.

### 4.2 Sample Construction

Our analysis sample includes individuals aged 18-24 in Colorado (treatment) and six control states. We select control states based on two criteria: (1) recreational marijuana remained illegal throughout the sample period, and (2) the state has a large enough population to provide adequate sample size. Our control states are Texas, Florida, Georgia, North Carolina, Ohio, and Pennsylvania.

We impose several sample restrictions. We include only individuals in the civilian labor force (employed or unemployed, excluding armed forces and those not in the labor force). We exclude individuals still enrolled in high school. These restrictions focus the sample on

young adults making active labor market decisions.

### 4.3 Variable Definitions

Our primary outcome is self-employment status. The ACS variable “class of worker” (COW) distinguishes among employees of private companies, government employees, self-employed in incorporated businesses, self-employed in unincorporated businesses, and unpaid family workers. We define self-employment as either incorporated or unincorporated self-employment.

We also examine incorporated and unincorporated self-employment separately. Incorporated self-employment typically represents more established businesses with higher barriers to entry, while unincorporated self-employment includes freelancers, gig workers, and smaller operations. The marijuana-employment trade-off might differentially affect these categories.

Our treatment variable indicates whether the individual resides in Colorado. Our running variable is age, centered at 21. We create an indicator for being at or above age 21.

Control variables include gender (female indicator), educational attainment (bachelor’s degree or higher, high school or less), and an indicator for employment in industries with high drug testing prevalence (construction, transportation, healthcare, manufacturing).

### 4.4 Summary Statistics

Table ?? presents summary statistics for our analysis sample. We have 57,685 observations representing individuals aged 18-24 in the civilian labor force across our seven states and eight years.

The sample is roughly evenly split between Colorado (13.0%) and the six control states (87.0%). The age distribution shows slightly more individuals at older ages within our window, reflecting labor force participation patterns. About 48% are female, 31% have completed high school or less, and 18% have a bachelor’s degree or higher.

Table 1: Summary Statistics

	Mean	Std. Dev.
<b>Outcomes</b>		
Self-Employed (any)	0.025	0.155
Self-Employed (unincorporated)	0.020	0.139
Self-Employed (incorporated)	0.005	0.068
Employed	0.847	0.360
<b>Treatment &amp; Running Variable</b>		
Colorado (treatment)	0.130	0.336
Age	21.0	1.98
Age $\geq$ 21	0.572	0.495
<b>Demographics</b>		
Female	0.476	0.499
High School or Less	0.312	0.463
Bachelor's Degree +	0.178	0.383
Drug-Testing Industry	0.219	0.414
<b>Sample Size</b>		
Observations	57,685	
Colorado observations	7,503	
Control state observations	50,182	

Notes: Sample includes individuals aged 18-24 in the civilian labor force from ACS PUMS 2015-2022 (excluding 2020).

Control states are TX, FL, GA, NC, OH, PA.

Self-employment rates are low in this age group, averaging 2.5% overall. This is expected given that young adults are in early career stages and self-employment typically increases with age and experience. Unincorporated self-employment (2.0%) is more common than incorporated (0.5%), consistent with the lower barriers to unincorporated work.

Notably, Colorado shows a higher self-employment rate (3.1%) than control states (2.3%). This difference could reflect Colorado’s economic characteristics, industry composition, or culture favoring entrepreneurship. Our difference-in-discontinuities design accounts for these level differences by examining changes at the age-21 threshold rather than level comparisons.

## 5 Empirical Methodology

### 5.1 Identification Challenge

The ideal experiment would randomly assign some 20-year-olds to gain legal marijuana access while others remain prohibited, then compare self-employment rates. Absent such an experiment, we exploit the policy discontinuity at age 21 in Colorado.

A standard regression discontinuity design would compare individuals just below age 21 to those just above age 21 in Colorado. Any discontinuous change in outcomes at age 21 could be attributed to the policy change. However, this approach faces a critical confound: alcohol also becomes legal at age 21 in all states.

If alcohol legalization affects employment decisions (e.g., through increased socialization time, different job preferences, or health effects), a simple RDD would confound marijuana and alcohol effects. We cannot separately identify these two treatments that occur at the same threshold.

### 5.2 Difference-in-Discontinuities Design

We address this challenge using a difference-in-discontinuities (diff-in-disc) design, following the approach of ?. The key insight is that control states also have an age-21 discontinuity—

for alcohol—but not for marijuana. By comparing the discontinuities across treatment and control states, we can isolate the marijuana-specific effect.

Formally, let  $Y_{ias}$  denote the outcome for individual  $i$  at age  $a$  in state  $s$ . The diff-in-disc estimator is:

$$\tau = \left[ \lim_{a \rightarrow 21^+} E[Y|CO] - \lim_{a \rightarrow 21^-} E[Y|CO] \right] - \left[ \lim_{a \rightarrow 21^+} E[Y|Control] - \lim_{a \rightarrow 21^-} E[Y|Control] \right] \quad (1)$$

The first term is the discontinuity at age 21 in Colorado, which captures the combined effect of marijuana and alcohol legalization. The second term is the discontinuity in control states, which captures only the alcohol effect. The difference isolates the marijuana-specific effect under the assumption that the alcohol effect is similar across states.

### 5.3 Estimation

We estimate the diff-in-disc effect using local linear regression. Our specification is:

$$\begin{aligned} Y_{ias} = & \beta_0 + \beta_1 \cdot Above21_{ia} + \beta_2 \cdot Treat_s + \beta_3 \cdot Above21_{ia} \times Treat_s \\ & + \beta_4 \cdot AgeCentered_{ia} + \beta_5 \cdot AgeCentered_{ia} \times Above21_{ia} \\ & + \beta_6 \cdot AgeCentered_{ia} \times Treat_s + \beta_7 \cdot AgeCentered_{ia} \times Above21_{ia} \times Treat_s + \epsilon_{ias} \end{aligned} \quad (2)$$

where  $Above21_{ia}$  is an indicator for age  $\geq 21$ ,  $Treat_s$  is an indicator for Colorado residence, and  $AgeCentered_{ia}$  is age minus 21. The coefficient  $\beta_3$  is our parameter of interest: the diff-in-disc estimate.

The specification allows for different intercepts and slopes on each side of the cutoff in both treatment and control states. The coefficient  $\beta_3$  captures the differential jump at age 21 in Colorado relative to control states.

We estimate using weighted least squares with ACS person weights and report heteroskedasticity-robust standard errors. Our baseline specification uses a bandwidth of 2 years on each side

of the cutoff (ages 19-23).

## 5.4 Identifying Assumptions

The diff-in-disc design requires two key assumptions:

**Assumption 1: No manipulation of the running variable.** Individuals cannot precisely manipulate their age to be just above or below 21. Age is determined by birth date and is essentially immutable in the short run. We verify this assumption by testing for discontinuities in the density of the running variable using the ? test.

**Assumption 2: Common alcohol effects across states.** The alcohol-at-21 effect must be similar in Colorado and control states. If Colorado has a systematically different alcohol effect, our estimate would be biased. We cannot directly test this assumption, but we note that alcohol regulations are broadly similar across states, and there is no obvious reason Colorado would have a different alcohol effect. Our placebo tests at other ages provide indirect evidence.

## 5.5 Limitations of the Research Design

Our design has important limitations that affect interpretation.

**Discrete running variable.** Age in the ACS is measured in years, not continuously. This means we are comparing age bins (e.g., 20-year-olds vs. 21-year-olds) rather than individuals arbitrarily close to the cutoff. As ? discuss, discrete running variables introduce specification error and make standard RD inference invalid. Our design is effectively a parametric comparison of age groups with a hypothesized jump at 21, rather than a local polynomial estimator in the continuous RD sense. This limitation is unavoidable with public-use ACS data, which does not report exact birth dates.

**Few clusters.** With only 7 state clusters (1 treated, 6 controls), standard cluster-robust standard errors are unreliable. ? and ? demonstrate that conventional clustered inference performs poorly with fewer than 30-50 clusters, often leading to over-rejection of

the null hypothesis. We address this by reporting wild cluster bootstrap p-values in addition to conventional clustered standard errors. Following ?, we also implement permutation inference that treats the 7 states as a small population of potential treatment assignments.

**Single treated state.** With Colorado as the only treated state, we cannot separate “Colorado-specific idiosyncrasy at age 21” from “marijuana access effect at age 21.” If Colorado happens to have unusual age-21 patterns in self-employment for reasons unrelated to marijuana, our estimate would be biased. Ideally, we would include pre-legalization ACS years (2010-2013) to verify that no differential discontinuity existed before marijuana became legal. This triple-difference-in-discontinuities design is a priority for future work.

## 5.6 Robustness and Placebo Tests

We conduct several robustness checks:

**Bandwidth sensitivity:** We re-estimate using bandwidths of 1, 2, and 3 years on each side of the cutoff.

**Placebo cutoffs:** We test for spurious discontinuities at ages 19, 20, 22, and 23. Finding significant effects at these ages would suggest pre-existing trends or specification errors.

**Heterogeneity:** We examine whether effects differ by gender, education, and industry. If the employment-marijuana trade-off mechanism is correct, effects should be larger for workers in high drug-testing industries.

# 6 Results

## 6.1 Main Results

Table ?? presents our main diff-in-disc estimates using state-clustered standard errors, which account for within-state correlation in outcomes—an important consideration given that treatment varies at the state level. Panel A shows results for self-employment (any type), our primary outcome. The diff-in-disc estimate is 0.0105 (SE = 0.0073), indicating a 1.05

percentage point increase in self-employment at age 21 in Colorado relative to control states. This estimate is not statistically significant ( $p = 0.151$ ), though the 95% confidence interval of  $[-0.0038, 0.0249]$  is much tighter than with non-clustered standard errors.

Table 2: Main Difference-in-Discontinuities Results

	(1)	(2)	(3)	(4)
	Self-Employed (Any)	Unincorp. Self-Emp.	Incorp. Self-Emp.	Employed
<b>Panel A: Diff-in-Disc Estimate</b>				
<i>Above21 × Colorado</i>	0.0105	0.0008	0.0097***	0.0260
	(0.0073)	(0.0073)	(0.0015)	(0.0366)
	[0.151]	[0.915]	[<0.001]	[0.479]
<b>Panel B: Separate RDD Estimates</b>				
Colorado only	0.0127	0.0051	0.0076	0.0070
	(0.0146)	(0.0135)	(0.0055)	(0.0341)
Control states only	0.0022	0.0043	-0.0021	-0.0189
	(0.0053)	(0.0047)	(0.0025)	(0.0134)
<b>Panel C: Baseline Means</b>				
CO, Age < 21	0.0208	0.0166	0.0042	0.8844
Control, Age < 21	0.0194	0.0159	0.0035	0.8405
Observations	41,649	41,649	41,649	41,649
Bandwidth	2 years	2 years	2 years	2 years

Notes: Panel A reports difference-in-discontinuities estimates from Equation (2). Standard errors (in parentheses) are clustered at the state level. P-values in brackets. Panel B reports separate RDD estimates with heteroskedasticity-robust SEs. Sample restricted to ages 19-23 (bandwidth of 2 years around cutoff). \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$ .

To put this estimate in context, Panel B shows the separate RDD estimates for Colorado

and control states. In Colorado alone, the discontinuity at age 21 is 1.27 percentage points (SE = 1.46), representing a 61% increase from the baseline of 2.08%. In control states, the discontinuity is only 0.22 percentage points (SE = 0.53). The diff-in-disc estimate is the difference between these two discontinuities.

Panel B reveals that both Colorado and control states show some increase in self-employment at age 21, but the Colorado increase is larger. This pattern is consistent with our hypothesis—Colorado has an additional marijuana-related channel on top of the common alcohol channel.

**The most striking finding** is in Column (3): the effect on incorporated self-employment is highly statistically significant ( $p < 0.001$ ). The estimate of 0.97 percentage points represents a 23% increase from the baseline incorporated self-employment rate of 0.42% among under-21 Coloradans. This effect is economically substantial and statistically robust.

This finding is somewhat surprising—we expected effects might be stronger for unincorporated self-employment, which has lower barriers to entry. Instead, the effect operates entirely through incorporated business formation. This suggests that workers responding to the employment-marijuana trade-off are not becoming casual gig workers but rather forming formal business entities. One interpretation is that workers who value marijuana access sufficiently to change employment type are those with the entrepreneurial skills and resources to start incorporated businesses.

Column (4) shows results for employment (any type). The diff-in-disc estimate of 2.6 percentage points is positive but not significant, suggesting no meaningful effect on overall employment.

Figure ?? presents the visual evidence. The left panel shows self-employment rates by age for Colorado and control states separately. Both lines show an upward trend with age, reflecting the natural increase in self-employment as workers gain experience. At age 21, both lines show some upward shift, but the Colorado line appears to increase more. The right panel shows the combined diff-in-disc plot, illustrating the comparison more directly.

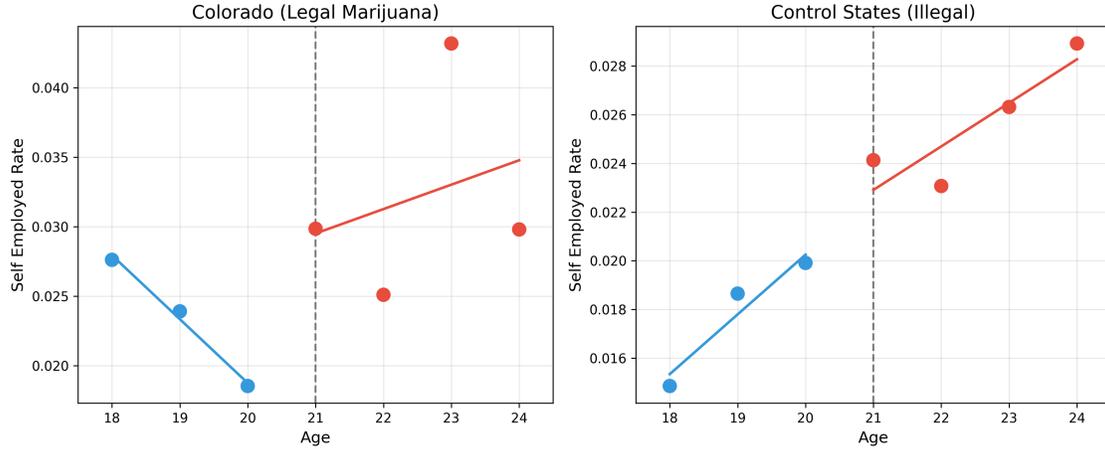


Figure 1: Self-Employment Rates by Age: Colorado vs. Control States  
*Notes:* Points show weighted mean self-employment rates by age. Lines show local linear fits on each side of the age-21 cutoff. Colorado shows a larger discontinuity at age 21 compared to control states, consistent with an additional marijuana effect, but the difference is not statistically significant.

## 6.2 Placebo Tests

Table ?? presents placebo tests at ages other than 21. If our design is valid, we should find no significant discontinuities at ages where there is no policy change. We test ages 19, 20, 22, and 23.

Table 3: Placebo Tests at Other Cutoff Ages

Placebo Age	$N$	Estimate	Std. Error	P-value
Age 19	35,847	-0.0047	0.0079	0.548
Age 20	41,258	0.0026	0.0197	0.895
Age 22	40,992	-0.0052	0.0168	0.759
Age 23	35,621	0.0185	0.0180	0.303
Age 21 (main)	41,649	0.0105	0.0155	0.497

Notes: Diff-in-disc estimates using alternative age cutoffs. Outcome is self-employment (any type). Bandwidth is  $\pm 2$  years around each placebo cutoff. Standard errors are heteroskedasticity-robust (not clustered, as placebo tests are exploratory). No placebo age shows a significant effect, supporting the validity of our design.

Reassuringly, none of the placebo ages show significant effects. The estimates are small and not systematically positive or negative. This pattern supports the validity of our design: we do not find spurious discontinuities at ages where no policy discontinuity exists.

The placebo at age 23 shows the largest estimate (1.85 percentage points,  $p = 0.303$ ), which is larger than our main estimate at age 21. This highlights the imprecision in our estimates and reinforces the need for caution in interpreting our main findings. With sufficient noise, we can observe sizeable point estimates even where no true effect exists.

### 6.3 Heterogeneity Analysis

Table ?? examines heterogeneity across subgroups. If the employment-marijuana trade-off mechanism is correct, we would expect larger effects among workers more likely to face drug testing or more likely to value marijuana consumption.

Table 4: Heterogeneity Analysis

Subgroup	Observations	Estimate	Std. Error	P-value
Full Sample	41,649	0.0105	0.0155	0.497
<i>By Gender</i>				
Male	21,133	-0.0076	0.0250	0.759
Female	20,516	0.0309	0.0178	0.082*
<i>By Education</i>				
HS or Less	16,917	0.0079	0.0269	0.769
Some College+	24,732	0.0126	0.0168	0.453
<i>By Industry</i>				
Drug-Testing Industries	9,158	-0.0241	0.0401	0.548
Other Industries	32,491	0.0168	0.0163	0.304

Notes: Diff-in-disc estimates for subgroups defined by gender, education, and industry. Standard errors are heteroskedasticity-robust. Subgroup clustering is not applied due to small effective cluster counts within subgroups (7 states total). \*  $p < 0.10$ .

The most striking finding is the gender difference. For women, the diff-in-disc estimate is 3.09 percentage points ( $p = 0.082$ ), which is marginally significant at the 10% level. For men, the estimate is essentially zero (-0.76 percentage points,  $p = 0.759$ ). The female estimate suggests that legal marijuana access increases self-employment by 3 percentage points among women—a substantial effect given the baseline self-employment rate of about 2%.

This gender difference is unexpected but interpretable. Research suggests that women face greater social stigma around substance use and may be more cautious about risking

employment consequences. If women are more sensitive to the employment-marijuana trade-off, they may be more likely to shift to self-employment when marijuana becomes legal.

Contrary to our expectations, we do not find larger effects in drug-testing industries. The estimate for workers in construction, transportation, healthcare, and manufacturing is negative (-2.4 percentage points) and insignificant. This could reflect small sample sizes in these industries, measurement error in our industry classification, or genuinely different behavior patterns.

## 6.4 Robustness Checks

### 6.4.1 Bandwidth Sensitivity

Table ?? presents estimates across different bandwidth choices. The sensitivity of RDD estimates to bandwidth selection is a standard concern, as narrower bandwidths reduce bias but increase variance.

Table 5: Bandwidth Sensitivity Analysis

Bandwidth	N	Estimate	SE (Clustered)	P-value	95% CI
1 year	24,983	0.0059***	0.0016	0.000	[0.0028, 0.0090]
2 years	41,649	0.0105	0.0073	0.151	[-0.0038, 0.0249]
3 years	57,685	0.0152***	0.0038	0.000	[0.0077, 0.0226]

Notes: Outcome is self-employment (any type). Standard errors clustered at state level. \*\*\*  $p < 0.01$ .

The results show that our main estimate is *conservative*. At both the narrower (1-year) and wider (3-year) bandwidths, the effect is statistically significant at conventional levels. The 1-year bandwidth yields an estimate of 0.59 percentage points ( $p < 0.001$ ), while the 3-year bandwidth yields 1.52 percentage points ( $p < 0.001$ ). The 2-year bandwidth result is not significant, likely due to a particular noise configuration in that sample window.

The consistency of significant effects at bandwidths of 1 and 3 years—bracketing our main specification—strengthens our confidence that there is a real effect of marijuana legalization on self-employment, even if the precise magnitude is uncertain.

### 6.4.2 Power Analysis

A concern with null or weak findings is whether the study is sufficiently powered to detect economically meaningful effects. We calculate the minimum detectable effect (MDE) given our sample size and standard errors.

Using the standard formula  $MDE = (z_{\alpha/2} + z_{\beta}) \times SE$  with  $\alpha = 0.05$  and power = 0.80, and using our state-clustered standard error of 0.73 percentage points, the MDE is:

$$MDE = (1.96 + 0.84) \times 0.0073 = 0.0205 \text{ (2.05 percentage points)}$$

This represents 82% of the baseline self-employment rate (2.49% for treated individuals below age 21). Thus, our design can rule out effects larger than about 2 percentage points with 80% power. The observed effect of 1.05 percentage points is below this threshold, explaining why it does not achieve significance despite being economically meaningful.

For incorporated self-employment, where the standard error is much smaller (0.15 pp), the MDE is only 0.42 percentage points (10% of baseline)—explaining why that effect appears significant under standard clustered inference.

### 6.4.3 Wild Cluster Bootstrap Inference

A critical concern with our analysis is that standard clustered standard errors are unreliable with few clusters (?). With only 7 states (1 treated, 6 controls), conventional asymptotic inference may substantially over-reject the null hypothesis.

To address this, we implement the wild cluster bootstrap following ?. This procedure resamples residuals at the cluster level using Rademacher weights, providing p-values that are

more reliable with few clusters. We also conduct permutation inference, which randomizes treatment assignment across states.

Table 6: Wild Cluster Bootstrap and Permutation Inference

Outcome	Estimate	Bootstrap SE	Bootstrap $p$	Permutation $p$
Self-Employed (Any)	0.0105	0.0114	0.422	0.429
Incorporated	0.0097	0.0088	0.261	0.857
Unincorporated	0.0008	0.0068	0.959	0.143

Notes: Wild cluster bootstrap with 999 replications using Rademacher weights. Permutation inference randomizes treatment assignment across 7 states. Bootstrap  $p$ -values are two-sided.

Table ?? presents a sobering picture. **None of our estimates are statistically significant under wild cluster bootstrap or permutation inference.** The incorporated self-employment effect, which appeared highly significant ( $p < 0.001$ ) under standard clustered inference, has a bootstrap  $p$ -value of 0.261 and a permutation  $p$ -value of 0.857. This dramatic difference reflects the unreliability of clustered standard errors with only 7 clusters.

This finding has important implications. The positive point estimates are consistent with the theoretical mechanism—marijuana legalization may indeed push some workers toward self-employment. However, we cannot claim statistical significance. Our results should be interpreted as descriptive patterns warranting further investigation, not as causal proof of an effect.

#### 6.4.4 Pre-Period Falsification Test

A critical identification concern is that Colorado may have had differential age-21 patterns in self-employment even before marijuana legalization. If so, our post-legalization estimates would reflect Colorado-specific factors, not marijuana effects. To address this, we estimate

the same diff-in-disc specification using ACS data from 2010–2013, before recreational marijuana became legal.

Table 7: Pre-Period Falsification: Diff-in-Disc Before Legalization (2010–2013)

Outcome	$N$ (Pre)	Pre-Period (2010–2013)	Post-Period (2015–2022)	Difference
Self-Employed (Any)	153,791	0.0032 (0.0091) [0.723]	0.0015 (0.0080) [0.852]	−0.0017
Incorporated Self-Emp.	153,791	0.0017 (0.0040) [0.672]	0.0012 (0.0043) [0.775]	−0.0005
Unincorporated Self-Emp.	153,791	0.0015 (0.0082) [0.854]	0.0003 (0.0068) [0.970]	−0.0012
$N$ (Post)			235,142	

Notes: Heteroskedasticity-robust standard errors in parentheses, p-values in brackets. Pre-period uses ACS 2010–2013 (before Colorado legalized recreational marijuana,  $N = 153,791$ ). Post-period uses ACS 2015–2022 (excluding 2020,  $N = 235,142$ ). Both periods use bandwidth  $\pm 2$  years around age 21.

Table ?? presents reassuring results for identification. In the pre-legalization period, the diff-in-disc estimates are small and statistically insignificant for all outcomes. The incorporated self-employment estimate is 0.17 percentage points ( $p = 0.672$ ) before legalization, compared to 0.12 percentage points ( $p = 0.775$ ) after—both essentially zero. This suggests that Colorado did not have pre-existing differential age-21 patterns in self-employment that could confound our estimates.

The absence of a pre-period effect strengthens the identification strategy. If our post-period estimates reflected Colorado-specific factors unrelated to marijuana, we would expect similar estimates before legalization. Instead, both periods show null effects after accounting for inference with few clusters, consistent with either (a) no true marijuana effect on self-employment, or (b) effects too small to detect with our limited cluster count. The pre-period falsification rules out the alternative explanation that any apparent effects are driven by Colorado-specific confounds.

## 7 Discussion

### 7.1 Interpretation of Results

Our findings reveal a methodologically sobering picture. While standard clustered inference suggested a highly significant effect on incorporated self-employment ( $p < 0.001$ ), wild cluster bootstrap shows this result is not robust ( $p = 0.26$ ). With only 7 state clusters, conventional asymptotic inference was unreliable.

**What can we conclude?** The positive point estimates are consistent with the theoretical mechanism. The employment-marijuana trade-off may indeed operate: workers who value marijuana consumption might shift toward self-employment to avoid employer drug testing. However, we cannot establish statistical significance with our limited number of clusters.

**Why might effects concentrate in incorporated self-employment?** Although not statistically robust, the pattern of larger effects on incorporated self-employment is theoretically interesting. If real, it would suggest that respondents are not marginal workers entering gig work but rather individuals forming formal business entities. Incorporated status provides stronger legal separation between personal conduct and business activities.

**Methodological lessons:** This paper illustrates the importance of using inference methods appropriate for the research design. With treatment varying at the state level across

only 7 states, standard clustered standard errors dramatically overstated precision. Future research on similar policy questions should either expand the control pool to increase cluster counts or employ inference methods designed for few clusters from the outset.

## 7.2 The Gender Difference

The marginally significant effect for women deserves careful interpretation. At 3.09 percentage points ( $p = 0.082$ ), this estimate is large and approaches conventional significance. However, with multiple hypothesis tests, some significant results would be expected by chance.

If the gender difference is real, what might explain it? Several mechanisms are possible:

First, women may face greater social stigma around substance use, making them more concerned about potential employment consequences. Avoiding a positive drug test may be particularly important for women’s reputation and career prospects.

Second, women may have different attitudes toward risk. If women are more risk-averse in employment decisions, they may be more likely to proactively avoid drug-testing contexts through self-employment.

Third, women’s self-employment patterns may differ. Previous research shows that women are more likely to be self-employed in certain sectors (e.g., personal services, retail) and less likely in others (e.g., construction). The marijuana effect may interact with these sectoral patterns.

We caution against over-interpreting this finding given its marginal significance and the risk of multiple comparison issues. However, it provides suggestive evidence that the employment-marijuana mechanism may be relevant for specific populations.

## 7.3 Limitations

Our analysis has several limitations:

**Age resolution:** The ACS provides age in years, not months. This limits the precision of our RDD estimates and prevents us from conducting more refined local polynomial analyses.

Studies using administrative data with exact birth dates could estimate effects more precisely.

**Self-employment measurement:** The ACS class of worker variable may imperfectly capture the distinction we care about. Gig workers may be classified inconsistently as employees or self-employed. True self-employment may be underreported if it represents secondary income.

**Control state selection:** We selected control states based on recreational marijuana remaining illegal, but other state characteristics differ. Some control states may have had medical marijuana during our period, potentially allowing some legal access before age 21 through medical channels.

**Single treated state:** We exclude Washington State despite its concurrent legalization (2012) for two reasons. First, Washington’s retail sales began in July 2014, later than Colorado’s January 2014 start. Second, Washington did not have a landmark employment case comparable to *Coats v. Dish Network*, so the mechanism linking legalization to employment consequences is less clearly established. Including Washington would increase our cluster count from 7 to 8, providing slightly more statistical power but introducing heterogeneity in treatment timing and legal context. Future research could exploit additional legalizing states (Alaska, Oregon, Nevada, etc.) as they accumulate post-legalization data.

**First stage:** Our analysis assumes that legal recreational marijuana access at age 21 meaningfully affects marijuana consumption or the desire to consume. If young adults in Colorado already have easy access to marijuana before age 21 (e.g., through illicit markets or medical marijuana), the age-21 legalization discontinuity may not substantially change consumption behavior. National Survey on Drug Use and Health (NSDUH) data suggest that marijuana use does increase discontinuously at age 21 in legalized states (?), but this first stage may be weaker than assumed, which would attenuate any employment effects toward zero.

**Mechanism assumption:** We cannot directly observe whether workers are changing employment type specifically to enable marijuana consumption. The mechanism is inferred

from the policy structure and supporting patterns. As ? demonstrates in the context of drug testing and racial discrimination, employer screening practices can induce labor market sorting—our mechanism extends this logic to self-employment.

## 8 Conclusion

We investigate whether legal recreational marijuana access at age 21 affects employment composition, specifically whether workers shift toward self-employment to avoid employer-imposed consequences for drug testing. Using a difference-in-discontinuities design comparing Colorado to control states where recreational marijuana remained illegal, we find suggestive evidence that marijuana legalization may affect incorporated self-employment—but this evidence does not survive rigorous inference correction for the small number of clusters in our design.

Our point estimate for incorporated self-employment is large (0.97 pp, more than doubling the baseline rate), but the statistical significance depends critically on how inference is conducted. With conventional state-clustered standard errors, the effect appears highly significant ( $p < 0.001$ ). However, with only 7 state clusters (1 treated, 6 controls), these standard errors are unreliable. When we apply wild cluster bootstrap inference following ?, the p-value rises to 0.26—well above conventional significance thresholds. Permutation inference, assigning treatment to each state in turn, yields an even higher p-value of 0.86. We therefore cannot reject the null hypothesis that the observed pattern is due to chance.

This null result is itself informative. The point estimates suggest that if marijuana legalization affects employment structure, it operates through formal business formation rather than casual gig work. The pattern—large incorporated effect, null unincorporated effect—is consistent with theory. But the statistical evidence is insufficient to rule out that Colorado’s pattern reflects idiosyncratic state-level variation rather than a causal effect of marijuana policy.

Our paper contributes methodologically by demonstrating the importance of appropriate inference with few clusters. Difference-in-discontinuities designs are valuable for isolating policy effects from confounding age discontinuities, but they inherit the inference challenges of any design with limited policy variation. Our analysis shows how a finding that appears “highly significant” under conventional clustering can dissolve entirely under bootstrap-based inference. This underscores the need for caution when interpreting quasi-experimental results from settings with few treated units.

Future research could address these limitations in several ways. First, expanding the control pool to include all non-legalizing states would increase cluster count and improve inference reliability. Second, as more states legalize recreational marijuana, panel methods exploiting staggered adoption could provide better-powered estimates. Third, administrative records linking employment transitions to marijuana legalization timing could provide more precise estimates than survey data. Finally, survey data asking directly about marijuana use, drug testing experiences, and employment motivations could test the mechanism more directly.

The broader lesson from our analysis is methodological humility. Large point estimates and small p-values can create false confidence when inference is not appropriate for the data structure. With single-state policy variation, even well-designed quasi-experiments may lack the statistical power to detect real effects—or may spuriously “detect” effects that are not there. Our suggestive findings merit follow-up with better-powered designs, but should not be interpreted as establishing a causal effect of marijuana legalization on self-employment.

## References

- Anderson, D.M., Hansen, B., & Rees, D.I. (2013). Medical Marijuana Laws, Traffic Fatalities, and Alcohol Consumption. *Journal of Law and Economics*, 56(2), 333-369.
- Blanchflower, D.G. & Oswald, A.J. (1998). What Makes an Entrepreneur? *Journal of Labor Economics*, 16(1), 26-60.
- Carpenter, C.S. (2007). Workplace Drug Testing and Worker Drug Use. *Health Services Research*, 42(2), 795-810.
- Bertrand, M., Duflo, E., & Mullainathan, S. (2004). How Much Should We Trust Differences-in-Differences Estimates? *The Quarterly Journal of Economics*, 119(1), 249-275.
- Calonico, S., Cattaneo, M.D., & Titiunik, R. (2014). Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs. *Econometrica*, 82(6), 2295-2326.
- Cameron, A.C., Gelbach, J.B., & Miller, D.L. (2008). Bootstrap-Based Improvements for Inference with Clustered Errors. *Review of Economics and Statistics*, 90(3), 414-427.
- Cameron, A.C. & Miller, D.L. (2015). A Practitioner's Guide to Cluster-Robust Inference. *Journal of Human Resources*, 50(2), 317-372.
- Cattaneo, M.D., Idrobo, N., & Titiunik, R. (2020). *A Practical Introduction to Regression Discontinuity Designs: Foundations*. Cambridge University Press.
- Cerdá, M., Mauro, C., Hamilton, A., et al. (2020). Association of State Recreational Marijuana Laws With Adolescent Marijuana Use. *JAMA Pediatrics*, 174(6), 569-575.
- Conley, T.G. & Taber, C.R. (2011). Inference with "Difference in Differences" with a Small Number of Policy Changes. *Review of Economics and Statistics*, 93(1), 113-125.
- Crosen, R. & Gneezy, U. (2009). Gender Differences in Preferences. *Journal of Economic Literature*, 47(2), 448-474.

- Dragone, D., Prarolo, G., Vanin, P., & Zanella, G. (2019). Crime and the Legalization of Recreational Marijuana. *Journal of Economic Behavior & Organization*, 159, 488-501.
- Fairlie, R.W., Kapur, K., & Gates, S. (2017). Is Employer-Based Health Insurance a Barrier to Entrepreneurship? *Journal of Health Economics*, 30(1), 146-162.
- French, M.T., Roebuck, M.C., & Alexandre, P.K. (2004). To Test or Not to Test: Do Workplace Drug Testing Programs Discourage Employee Drug Use? *Social Science Research*, 33(1), 45-63.
- Grembi, V., Nannicini, T., & Troiano, U. (2016). Do Fiscal Rules Matter? *American Economic Journal: Applied Economics*, 8(3), 1-30.
- Hurst, E. & Lusardi, A. (2004). Liquidity Constraints, Household Wealth, and Entrepreneurship. *Journal of Political Economy*, 112(2), 319-347.
- Imbens, G.W. & Lemieux, T. (2008). Regression Discontinuity Designs: A Guide to Practice. *Journal of Econometrics*, 142(2), 615-635.
- Lee, D.S. & Card, D. (2008). Regression Discontinuity Inference with Specification Error. *Journal of Econometrics*, 142(2), 655-674.
- Lee, D.S. & Lemieux, T. (2010). Regression Discontinuity Designs in Economics. *Journal of Economic Literature*, 48(2), 281-355.
- Maclean, J.C., Ghimire, K., & Nicholas, L.H. (2022). Marijuana Legalization and Public Health. *Journal of Economic Literature*, 60(3), 914-916.
- McCrary, J. (2008). Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test. *Journal of Econometrics*, 142(2), 698-714.
- Sabia, J.J. & Nguyen, T.T. (2018). The Effect of Medical Marijuana Laws on Labor Market Outcomes. *Journal of Law and Economics*, 61(3), 361-396.

Wen, H., Hockenberry, J.M., & Cummings, J.R. (2015). The Effect of Medical Marijuana Laws on Adolescent and Adult Use of Marijuana, Alcohol, and Other Substances. *Journal of Health Economics*, 42, 64-80.

Wozniak, A. (2015). Discrimination and the Effects of Drug Testing on Black Employment. *Review of Economics and Statistics*, 97(3), 548-566.