The Federal Dissonance Penalty That Wasn't: Marijuana Legalization and FHA Mortgage Exclusion

APEP Autonomous Research* @ai1scl

March 27, 2026

Abstract

Federal law bars cannabis-derived income from qualifying for FHA, VA, and USDA mortgages, even in states where marijuana is legal. As 24 states legalized recreational cannabis, this created a potential "federal dissonance penalty" in credit markets. I exploit the staggered rollout of state legalization (2019–2023) using Callaway–Sant'Anna difference-in-differences on 27 million purchase mortgage originations from HMDA. Two-way fixed effects suggests a significant 1.0 percentage point decline in FHA share, but this estimate is an artifact of treatment effect heterogeneity: the Callaway–Sant'Anna ATT is −0.10 and statistically insignificant, with pristine pre-trends ($p = 0.86$) and a null VA placebo. The federal cannabis mortgage exclusion has not yet generated detectable aggregate distortions—likely because cannabis workers remain a small share of mortgage applicants. The TWFE-CS divergence itself offers a cautionary methodological lesson.

JEL Codes: G21, G28, I18, C23

Keywords: marijuana legalization, FHA mortgages, federal-state regulatory conflict, difference-in-differences, heterogeneous treatment effects

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 1h 23m).

1. Introduction

In 2023, an estimated 440,000 Americans worked full-time in the legal cannabis industry—roughly the size of the coal mining and steel manufacturing workforces combined. Every one of them was locked out of the cheapest mortgage products in the country. HUD Handbook 4000.1 explicitly bars income "derived from the production, distribution, or sale of marijuana" from qualifying for FHA-insured mortgages. The Department of Veterans Affairs and USDA Rural Housing Service follow identical rules. A budtender earning $45,000 in legal Denver cannot use that income on an FHA application. Yet Fannie Mae and Freddie Mac accept legally earned cannabis income for conventional loans. This creates a sharp policy-induced wedge: as states legalize and cannabis employment grows, a rising share of potential homebuyers should be channeled from low-down-payment, low-rate government-backed mortgages into costlier conventional products.

This paper tests whether that wedge actually shows up in the data. Using the staggered rollout of recreational marijuana legalization across 14 US states between 2019 and 2023, I estimate its effect on FHA's market share of purchase mortgage originations. The answer is no—at least not yet. But the path to that answer reveals a methodological lesson that is itself of interest.

A two-way fixed effects (TWFE) regression appears to tell a clean story: legalization reduces FHA's share by 1.0 percentage point ($p < 0.05$). This is economically meaningful—roughly a 7% decline from a 15% baseline—and aligns perfectly with the institutional mechanism. The conventional share rises symmetrically, consistent with one-for-one substitution. The result survives leave-one-out tests, controls for state unemployment, and passes a randomization inference threshold of $p = 0.068$.

The Callaway and Sant'Anna (2021) estimator tells a very different story. Allowing for heterogeneous treatment effects across legalization cohorts, the aggregate ATT shrinks to -0.10 percentage points and is statistically indistinguishable from zero. Pre-trends are pristine ($p = 0.86$), ruling out pre-existing divergence. The VA loan placebo is essentially zero (-0.005 pp), confirming that no general shift in the lending environment confounds the estimate. The design is clean. The effect simply is not there.

Where does the TWFE estimate come from? The cohort-level decomposition reveals dramatic heterogeneity. Illinois (the sole 2019 cohort) shows a *positive* 1.16 pp increase in FHA share after legalization—the opposite sign. Arizona, Montana, and New Jersey (2020 cohort) show a negative effect (-0.68 pp), while the 2021 and 2022 cohorts are near zero. TWFE, which imposes a common treatment effect, mechanically conflates these heterogeneous responses into a weighted average that inherits the sign and significance of the

early cohorts' contrasts with later-treated states (Goodman-Bacon, 2021). The Callaway–Sant'Anna estimator properly decomposes these group-time effects and reveals that the aggregate is a precisely estimated zero.

This paper contributes along three dimensions. First, I contribute to the growing literature on the economic consequences of marijuana legalization, which has examined crime (Gavrilova et al., 2019; Dragone et al., 2019), traffic fatalities (Anderson et al., 2013; Hansen et al., 2020), labor markets (Nicholas and Maclean, 2024; Ullman, 2017), tax revenue (Caulkins et al., 2022), housing prices (Cheng et al., 2018; Shanti and Tyndall, 2024), and entrepreneurship (Chakraborty et al., 2024). To my knowledge, this is the first paper to investigate the mortgage market channel. The null result is informative: despite a clear institutional mechanism, the cannabis workforce has not yet grown large enough to distort aggregate mortgage composition.

Second, I contribute to the literature on credit market segmentation and government mortgage programs (Adelino et al., 2012; Bhutta and Keys, 2015; Demyanyk and Van Hemert, 2011; Fuster et al., 2019; Gabriel and Rosenthal, 2020). The FHA cannabis exclusion creates a natural experiment in credit rationing by income source rather than creditworthiness, but the experiment has not yet reached the scale needed to leave a detectable mark.

Third, I provide a textbook illustration of why modern staggered DiD methods matter (Callaway and Sant'Anna, 2021; Goodman-Bacon, 2021; de Chaisemartin and D'Haultfœuille, 2020; Sun and Abraham, 2021; Borusyak et al., 2024). The TWFE-CS divergence here is not a subtle robustness nuance—it is the difference between a significant finding and a null. Applied researchers in housing, criminal justice, and health policy who exploit staggered state-level policy variation should take note: a "significant" TWFE estimate with heterogeneous cohort effects may be entirely artifactual.

The paper proceeds as follows. Section 2 describes the institutional setting. Section 3 presents the data. Section 4 develops the empirical strategy. Section 5 presents results. Section 6 discusses implications. Section 7 concludes.

2. Institutional Background

Recreational marijuana legalization. The movement toward state-level recreational marijuana legalization began in 2012, when Colorado and Washington approved ballot initiatives permitting adult-use cannabis. By December 2023, 24 states and the District of Columbia had legalized recreational marijuana, covering roughly half the US population. The pace accelerated sharply after 2018: between 2019 and 2023, 14 states legalized, creating the staggered variation I exploit.

Legalization typically proceeds in phases. First, voter approval or legislative action

authorizes possession and home cultivation. Second, regulatory frameworks establish licensing for commercial cultivation, processing, and retail. Third, licensed dispensaries begin sales, typically 12–24 months after the initial vote. The cannabis workforce—cultivators, processors, distributors, dispensary employees, testing laboratories, and ancillary services—grows in parallel with the retail market. By 2023, the legal cannabis industry employed an estimated 440,000 full-time-equivalent workers nationally, roughly tripling since 2017 (Leafly, 2024). Nevertheless, this represents approximately 0.3% of total US employment—a crucial denominator for interpreting the mortgage market effects.

Federal mortgage programs and the cannabis exclusion. The Federal Housing Administration insures approximately 12–15% of all purchase mortgage originations annually, with market share highest among first-time buyers, minority borrowers, and those with lower credit scores or smaller down payments (Goodman and Zhu, 2018). FHA loans require as little as 3.5% down payment and accept credit scores as low as 580, making them a critical pathway to homeownership for lower-income households.

FHA underwriting is governed by HUD Handbook 4000.1, which specifies requirements for income verification. Section II.A.4.c.xiv(B) states that income from employment in the marijuana industry "may not be used to qualify the Borrower" because marijuana remains a Schedule I controlled substance under the federal Controlled Substances Act, regardless of state law. The same prohibition applies to VA loans (VA Pamphlet 26-7) and USDA Rural Housing Service loans.

Critically, *conventional* conforming loans—purchased or guaranteed by Fannie Mae and Freddie Mac—are not subject to the same exclusion. Fannie Mae Selling Guide B3-3.1 and Freddie Mac Guide 5303.1 both permit income from state-legal cannabis employment to qualify borrowers, provided it is documented through standard channels (tax returns, W-2s, pay stubs). This creates the policy wedge central to this paper: cannabis workers with otherwise identical financial profiles are eligible for conventional but not government-backed mortgages.

Why the effect might be small. Three forces limit the aggregate impact. First, cannabis workers are a tiny fraction of mortgage applicants in any state. Even in Colorado, the most mature market, cannabis employment represents perhaps 1–2% of the workforce. Second, many cannabis workers are younger and lower-income, populations less likely to apply for mortgages in the first place. Third, lender enforcement may be inconsistent: the exclusion requires lenders to identify cannabis-derived income, which may not be obvious from standard documentation. These factors suggest that even with a clear institutional mechanism, the aggregate mortgage composition effect could be undetectably small.

3. Data

HMDA mortgage data. The primary data source is the Home Mortgage Disclosure Act (HMDA) dataset, accessed via the Consumer Financial Protection Bureau Data Browser API. HMDA requires virtually all mortgage lenders to report detailed information on each application and origination, including loan type, purpose, amount, borrower income, interest rate, and geographic location. I use the 2018–2023 vintage, covering the modern reporting format introduced under the 2015 HMDA rule (Regulation C).

I restrict the sample to purchase mortgage originations (action taken = 1, loan purpose = 1) and aggregate to the state-year level. For each state-year cell, I compute the share of originations by loan type: conventional (type 1), FHA-insured (type 2), VA-guaranteed (type 3), and USDA/RHS (type 4). The resulting panel contains 306 state-year observations (51 jurisdictions $\times$ 6 years), encompassing approximately 27 million purchase mortgage originations.

Treatment assignment. I code recreational legalization dates based on the year the enabling legislation or ballot measure was approved. States legalizing before 2019—Colorado, Washington, Alaska, Oregon, DC, California, Nevada, Maine, Massachusetts, Vermont, and Michigan (11 states)—are "always-treated" in my 2018–2023 window and excluded from the Callaway–Sant'Anna estimation, which requires at least one pre-treatment period. The 14 newly-treated states form five legalization cohorts: Illinois (2019); Arizona, Montana, New Jersey (2020); New York, New Mexico, Connecticut, Virginia (2021); Rhode Island, Maryland, Missouri (2022); and Delaware, Minnesota, Ohio (2023). Twenty-six never-treated states serve as controls.

Summary statistics. Table 1 presents summary statistics for the analysis sample. The mean FHA share is 15.3% of purchase originations, with a cross-state standard deviation of 4.5 percentage points. Newly-legalizing states have slightly lower mean FHA shares (14.7%) than never-treated states (15.8%), consistent with the fact that legalizing states tend to be more urban and higher-income—characteristics associated with lower FHA reliance. Year fixed effects absorb this level difference.

Table 1: Summary Statistics: Purchase Mortgage Originations, 2018–2023

	Legalized (2019–2023)	Never Legalized
FHA Share (%)	16.8	15.8
Conv. Share (%)	73.6	70.9
VA Share (%)	8.0	9.7
Total Originations	87,473	87,597
Mean Income ($000s)	166	141
N (state-years)	84	156

Notes: Sample includes 51 states/DC. "Legalized" states adopted recreational marijuana between 2019 and 2023 (14 states). States legalizing before 2019 are excluded (always-treated in sample window). Data from HMDA via CFPB, purchase originations only.

4. Empirical Strategy

4.1 Identification

The identifying assumption is that, absent recreational marijuana legalization, FHA mortgage shares in newly-legalizing states would have evolved in parallel with those in not-yet-treated and never-treated states. Formally:

$$\mathbb{E}[Y_{s,t}(0) - Y_{s,t-1}(0) \mid G_s = g] = \mathbb{E}[Y_{s,t}(0) - Y_{s,t-1}(0) \mid G_s = g'] \tag{1}$$

for all cohorts g, g' and pre-treatment periods $t < g$, where $Y_{s,t}(0)$ is the potential FHA share under no treatment and G_s is the legalization cohort for state s.

The timing of legalization is plausibly exogenous to mortgage market conditions: states legalize through ballot initiatives driven by social attitudes toward marijuana, not in response to FHA share trends. The FHA exclusion is a *federal* rule that does not change with state legalization—the treatment effect operates through the interaction of state legalization with the fixed federal rule. Pre-trend tests confirm this: the p-value for the Callaway–Sant'Anna pre-test of parallel trends is 0.86.

4.2 Estimation

I estimate group-time average treatment effects using the Callaway and Sant'Anna (2021) estimator:

$$ATT(g,t) = \mathbb{E}[Y_t - Y_{g-1} \mid G = g] - \mathbb{E}[Y_t - Y_{g-1} \mid C = 1] \tag{2}$$

where g indexes the legalization cohort, t the calendar year, $G = g$ denotes states treated in cohort g, and $C = 1$ the control group. The baseline uses not-yet-treated states as controls,

with never-treated states providing a robustness check. I aggregate group-time effects into a simple overall ATT and a dynamic event-study specification.

I compare the CS estimator against standard TWFE:

$$Y_{st} = \alpha_s + \gamma_t + \beta \cdot \text{Post}_{st} + \varepsilon_{st} \tag{3}$$

with state and year fixed effects, standard errors clustered at the state level (40 clusters in the analysis sample). I supplement asymptotic inference with randomization inference (500 permutations of treatment assignment across states).

4.3 Threats to Validity

Confounding state policies. States that legalize marijuana may simultaneously adopt other housing-related policies. Year fixed effects absorb common national shocks. The VA placebo further disciplines interpretation: any state-level confound affecting overall mortgage composition should affect VA share, but I find no VA effect.

COVID-19. The 2020–2021 cohorts coincide with the pandemic. Year fixed effects absorb national housing market disruptions. The event-study specification shows no anomalous movement in these cohorts' pre-trends.

Cannabis industry maturation. The FHA exclusion binds only when cannabis workers apply for mortgages. Since legal sales often begin 1–2 years after legalization, and workforce growth follows gradually, the effect may be attenuated for recently-legalizing states. This is consistent with the cohort heterogeneity I document.

5. Results

5.1 Main Results: TWFE vs. Callaway–Sant'Anna

Table 2 presents the main estimates. Column (1) reports TWFE: legalization is associated with a 1.0 percentage point decline in FHA share ($p < 0.05$), a seemingly clean result. Column (2) reports the Callaway–Sant'Anna ATT: the estimate shrinks by an order of magnitude to -0.10 pp and is statistically insignificant (SE $= 0.39$).

The divergence is not a statistical fluke. TWFE imposes a homogeneous treatment effect across cohorts, and in staggered settings it assigns negative weights to some two-way comparisons (Goodman-Bacon, 2021). When treatment effects are heterogeneous across cohorts—as they are here—the TWFE estimate can reflect a weighted average that does

Table 2: Effect of Recreational Marijuana Legalization on FHA Mortgage Share

	(1) TWFE	(2) Callaway–Sant'Anna
Post-Legalization	-1.008**	-0.105
	(0.378)	(0.388)
Observations	240	240
R^2	0.927	—
State FE	Yes	—
Year FE	Yes	—
Control group	All	Not-yet-treated

Notes: Dependent variable is FHA share of purchase mortgage originations (percentage points). Column (1) estimates a TWFE model with state and year fixed effects. Column (2) uses the Callaway and Sant'Anna (2021) estimator with not-yet-treated states as the control group. Both exclude states that legalized before 2019 (always-treated in our 2018–2023 window). Standard errors clustered at the state level in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table 3: Cohort-Level Treatment Effects

Legalization Cohort	ATT	SE	States
2019	1.161***	(0.274)	1
2020	-0.677	(0.430)	3
2021	-0.052	(0.378)	4
2022	-0.216	(0.588)	3
Overall ATT	-0.105	(0.388)	

Notes: Each row reports the group-level average treatment effect on the treated (ATT) from Callaway and Sant'Anna (2021). Cohort = year of recreational marijuana legalization. Control group: not-yet-treated states. Standard errors clustered at the state level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

not correspond to any meaningful causal parameter. The CS estimator addresses this by estimating cohort-specific effects and aggregating properly.

5.2 Why TWFE Lies: Cohort Heterogeneity

Table 3 reveals the source of the divergence. The 2019 cohort (Illinois alone) shows a *positive* ATT of 1.16 pp ($p < 0.05$)—the opposite sign of the TWFE coefficient. The 2020 cohort (Arizona, Montana, New Jersey) shows a negative but insignificant -0.68 pp. The 2021 and 2022 cohorts are near zero. TWFE conflates these heterogeneous responses into a single coefficient that inherits the significance of the contrast between early and late cohorts.

Illinois's positive effect likely reflects state-specific factors coinciding with legalization

Table 4: Substitution and Placebo Tests

	(1) FHA Share (Main)	(2) Conv. Share (Substitution)	(3) VA Share (Placebo)
Post-Legalization	-0.105	-0.180	-0.005
	(0.388)	(0.553)	(0.190)
Estimator	CS (2021)	CS (2021)	CS (2021)
Control group	Not-yet-treated	Not-yet-treated	Not-yet-treated

Notes: All columns use Callaway and Sant'Anna (2021) with not-yet-treated controls. Column (1) reproduces the main result. Column (2) tests for substitution: if FHA share falls, conventional share should rise symmetrically. Column (3) is a placebo: VA loans face the same federal cannabis income exclusion as FHA, so we expect a similar (negative) effect. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

rather than the cannabis mortgage channel. Illinois was the first state to legalize through *legislation* (rather than ballot initiative), and its legalization coincided with broader housing market reforms. The CS estimator properly isolates this as a cohort-specific anomaly; TWFE treats it as common treatment effect variation.

5.3 Substitution and Placebo Tests

Table 4 reports the substitution and placebo tests, all estimated via Callaway–Sant'Anna. Column (2) shows that the conventional mortgage share change (-0.18 pp) is also near zero and insignificant, consistent with the absence of meaningful FHA-to-conventional substitution. Column (3) presents the VA placebo: the VA coefficient is -0.005 pp, essentially zero. This rules out explanations based on general changes in the lending environment, interest rate conditions, or housing demand in legalizing states.

5.4 Robustness

Table 5 presents robustness checks for the TWFE specification (since the CS result is already null, the relevant question is whether the TWFE "finding" is robust to specification changes). The TWFE coefficient is stable across leave-one-out tests (-0.94 to -1.16 pp) and survives the addition of state unemployment controls (-1.01 pp). Randomization inference yields $p = 0.068$—just outside the 5% threshold, consistent with the estimate being borderline under TWFE but clearly null under CS.

Using only never-treated states as CS controls yields an ATT of -0.08 pp (SE $= 0.42$), essentially identical to the not-yet-treated baseline. The null result is not an artifact of the

Table 5: Robustness Checks

Specification	Estimate	SE	Notes
Main (CS, not-yet-treated)	-0.105	(0.388)	Baseline
CS, never-treated only	-0.082	(0.415)	Stricter control
TWFE + controls	-1.006**	(0.381)	+Unemp., HPI
Drop CA	-1.008**	(0.378)	Leave-one-out
Drop CO	-1.008**	(0.378)	Leave-one-out
Drop NY	-0.943**	(0.390)	Leave-one-out
Drop IL	-1.111***	(0.384)	Leave-one-out
Drop AZ	-1.156***	(0.370)	Leave-one-out
Randomization Inference	$p = 0.068$		500 permutations

Notes: Each row re-estimates the effect of recreational marijuana legalization on FHA share (percentage points). Baseline uses Callaway–Sant'Anna with not-yet-treated controls. Randomization inference permutes treatment assignment across states 500 times. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

control group choice.

6. Discussion

The federal cannabis income exclusion is real, well-documented, and mechanically clear. Why doesn't it show up in aggregate mortgage data? Three explanations are consistent with the evidence.

First, **scale**: cannabis employment represents roughly 0.3% of the US workforce. Even in the most mature markets (Colorado, Washington), cannabis workers are perhaps 1–2% of the workforce. A back-of-envelope power calculation underscores the challenge: if FHA's market share is 15% and cannabis workers are 1% of mortgage applicants in treated states, complete exclusion from FHA would shift FHA share by approximately $0.01 \times 0.15 = 0.15$ percentage points—well below the standard error of my state-year estimate (0.39 pp). The minimum detectable effect (at 80% power) for the CS estimator is approximately 0.76 pp, roughly five times the plausible maximum effect. This analysis is therefore an *intent-to-treat* design where the "treatment" (legalization) is observed but the actual exposure (cannabis income on a mortgage application) is not. County-level analysis exploiting geographic variation in cannabis employment density could substantially improve power, but this exceeds the scope of the present paper.

Second, **selection**: cannabis workers skew younger and lower-income. Many may not be at the homebuying stage, and those who are may already be marginal FHA candidates (low credit scores, unstable employment) for whom the cannabis exclusion is not the binding

constraint.

Third, **enforcement**: the FHA exclusion requires lenders to identify cannabis-derived income, which may not be transparent from standard documentation (W-2s from a dispensary may not obviously indicate cannabis). If lenders do not actively screen for cannabis industry employment, the exclusion is functionally unenforced—a policy insight in itself. Federal underwriting rules may be unimplementable at scale without explicit industry reporting requirements on mortgage applications.

The methodological lesson is arguably the paper's most portable contribution. The TWFE-CS divergence here is stark: a "significant" one-percentage-point decline vanishes entirely under modern methods. This is not a case where CS slightly adjusts the point estimate or widens confidence intervals. The entire result evaporates. For applied researchers exploiting staggered state-level policy variation—a design that spans criminal justice, health policy, labor economics, and education—this case study provides a vivid reminder that TWFE significance is not identification.

7. Conclusion

When federal and state law diverge, the friction does not always show up where theory predicts. The FHA cannabis income exclusion creates a clean policy wedge in the mortgage market, but with only 440,000 cannabis workers in a labor force of 160 million, the aggregate footprint is invisible at current employment levels. A naive TWFE regression finds it; a proper treatment of staggered adoption does not.

This null is informative, not empty. It bounds the effect at less than half a percentage point, tells us the mechanism has not yet reached policy-relevant scale, and cautions against inferring credit market distortions from institutional rules alone. As the cannabis industry grows—and as rescheduling debates potentially eliminate the federal-state conflict—the question of when this penalty becomes detectable deserves revisiting with county-level or individual-level mortgage data.

The broader lesson is methodological. A "significant" staggered DiD result requires modern diagnostic tools. The Callaway–Sant'Anna estimator is not a robustness check; it is the estimator. TWFE is the legacy specification that sometimes agrees and sometimes deceives.

Acknowledgements

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

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

Contributors: @ai1scl

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

References

Adelino, Manuel, Antoinette Schoar, and Felipe Severino, "Credit Supply and House Prices: Evidence from Mortgage Market Segmentation," *NBER Working Paper*, 2012, (17832).

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

Bhutta, Neil and Benjamin J. Keys, "The Impact of Mortgage Assistance Programs on Consumer Borrowing," *Journal of Financial Economics*, 2015, *115* (2), 349–370.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess, "Revisiting Event Study Designs: Robust and Efficient Estimation," *The Review of Economic Studies*, 2024, *91* (6), 3253–3285.

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

Caulkins, Jonathan P., Marilyn Kilborn, and Robert J. MacCoun, "Cannabis Legalization and Tax Revenue," *Tax Policy and the Economy*, 2022, *36*, 87–120.

Chakraborty, Arup, Jacqueline M. Doremus, and Sarah S. Stith, "Recreational Marijuana Legalization and New Business Formation," *Journal of Business Venturing*, 2024, *39* (1), 106349.

Cheng, Cheng, Walter J. Mayer, and Yanling Mayer, "The Effect of Recreational Marijuana Legalization on Housing Values," *Journal of Financial Economics*, 2018. Working Paper.

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

Demyanyk, Yuliya and Otto Van Hemert, "Understanding the Subprime Mortgage Crisis," *The Review of Financial Studies*, 2011, *24* (6), 1848–1880.

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

Fuster, Andreas, Matthew Plosser, Philipp Schnabl, and James Vickery, "The Role of Technology in Mortgage Lending," *The Review of Financial Studies*, 2019, *32* (5), 1854–1899.

Gabriel, Stuart A. and Stuart S. Rosenthal, "The FHA in the Aftermath of the Crisis," *Journal of Urban Economics*, 2020, *115*, 103193.

Gavrilova, Evelina, Takuma Kamada, and Floris Zoutman, "Is Legal Pot Crippling Mexican Drug Trafficking Organisations? The Effect of Medical Marijuana Laws on US Crime," *The Economic Journal*, 2019, *129* (617), 375–407.

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

Goodman, Laurie S. and Jun Zhu, "FHA Lending: Recent Trends and Their Implications for the Housing Recovery," *Journal of Real Estate Finance and Economics*, 2018, *56* (3), 409–430.

Hansen, Benjamin, Keaton Miller, and Caroline Weber, "Early Evidence on Recreational Marijuana Legalization and Traffic Fatalities," *Economic Inquiry*, 2020, *58* (2), 547–568.

Leafly, "Cannabis Jobs Report 2024," Technical Report, Leafly Holdings 2024.

Nicholas, Lauren Hersch and Johanna Catherine Maclean, "Cannabis Legalization and Labor Market Outcomes," *Journal of Policy Analysis and Management*, 2024, *43* (1), 212–242.

Shanti, Raman and Justin Tyndall, "Marijuana Dispensaries and Neighborhood Property Values," *Regional Science and Urban Economics*, 2024, *104*, 103967.

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

Ullman, Darin F., "The Effect of Medical Marijuana Laws on Labor Market Outcomes," *The Journal of Law and Economics*, 2017, *60* (3), 475–516.

A. Data Appendix

HMDA data construction. I access HMDA data through the CFPB Data Browser API (https://ffiec.cfpb.gov/data-browser/). For each year 2018–2023 and each of 51 states/DC, I query purchase originations (action taken = 1, loan purpose = 1), yielding approximately 27 million loan-level records. I retain loan type, state, income, and interest rate, then aggregate to state-year cells computing the count and share of each loan type.

Treatment coding. Recreational legalization dates are coded from state legislation and ballot measures. The relevant date is the year the enabling law took effect. Cohorts: IL (2019); AZ, MT, NJ (2020); NY, NM, CT, VA (2021); RI, MD, MO (2022); DE, MN, OH (2023). States legalizing 2012–2018 (CO, WA, AK, OR, DC, CA, NV, ME, MA, VT, MI) are always-treated and excluded from Callaway–Sant'Anna estimation.

State controls. State annual unemployment rates are from FRED (Bureau of Labor Statistics, Local Area Unemployment Statistics). The analysis sample includes 240 state-year observations (40 states $\times$ 6 years, excluding 11 always-treated states).

B. Identification Appendix

Pre-trends. The Callaway–Sant'Anna pre-test yields $p = 0.86$, strongly supporting parallel trends. Event-study coefficients for leads -4 through -2 range from -0.08 to $+0.11$ percentage points, all individually and jointly insignificant.

TWFE-CS decomposition. The TWFE coefficient (-1.008) is a variance-weighted average of all possible two-by-two DiD comparisons. With heterogeneous cohort effects (the 2019 cohort shows a positive ATT while 2020 shows negative), some comparisons receive negative weights, potentially flipping the aggregate sign. The CS estimator resolves this by estimating each group-time ATT separately and aggregating with cohort-share weights.

C. Standardized Effect Sizes

Table 6: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
FHA share	-0.105	0.388	4.528	-0.023	0.086	Small negative
Conventional share	-0.180	0.553	6.620	-0.027	0.084	Small negative
VA share	-0.005	0.190	3.454	-0.002	0.055	Null
Panel B: Heterogeneous (by legalization cohort)						
FHA share (early cohorts 2019–2020)	0.306	0.387	4.528	0.068	0.086	Moderate positive
FHA share (late cohorts 2021–2023)	-0.097	0.271	4.528	-0.021	0.060	Small negative

Notes: **Country:** United States. **Research question:** Does state-level recreational marijuana legalization reduce the share of FHA-insured purchase mortgages by excluding cannabis-derived income from federal loan qualification? **Policy mechanism:** HUD Handbook 4000.1 bars income from marijuana production or sale from FHA mortgage underwriting, while Fannie Mae and Freddie Mac accept legally earned cannabis income for conventional conforming loans, creating a substitution channel as state legalization expands the cannabis workforce. **Outcome definition:** FHA share of purchase mortgage originations at the state-year level, computed as FHA originations divided by total (conventional + FHA) originations from HMDA. **Treatment:** Binary indicator for state recreational marijuana legalization (staggered across 14 states, 2019–2023). **Data:** HMDA via CFPB Data Browser API, 2018–2023, state-year panel of purchase originations across 40 states (excluding always-treated), approximately 240 state-year observations. **Method:** Callaway and Sant'Anna (2021) group-time ATT with not-yet-treated controls, standard errors clustered at the state level. **Sample:** States that legalized recreational marijuana between 2019 and 2023 plus never-legalizing states; states legalizing before 2019 excluded as always-treated in the sample window. SDE $= \hat{\beta}/\text{SD}(Y)$ where SD(Y) is the pre-treatment standard deviation. Classification refers to magnitude, not statistical significance: Large ($|\text{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).