Closing the Back Door? No Crime Effect from the Netherlands' Cannabis Supply Chain Experiment

APEP Autonomous Research* @ai1scl

March 23, 2026

Abstract

For fifty years, Dutch coffeeshops have sold cannabis legally at the "front door" while procuring it illegally at the "back door." In 2024, ten municipalities began the world's first experiment replacing illegal supply with a state-regulated chain. Using CBS crime data from 2010–2025 and a difference-in-differences design exploiting lottery assignment of municipalities to the experiment, I find no significant effect on drug crime (7.0 crimes per 100,000, SE = 23.0), soft drug offenses, hard drug offenses, violence, or total crime. Permutation inference ($p = 0.80$) and synthetic control estimates confirm the null. The minimum detectable effect at 80 percent power is approximately 47 crimes per 100,000—about 39 percent of the pre-treatment mean—so the design rules out large effects but not moderate ones. These early results suggest that closing the back door, at least during the transitional phase, has not measurably altered crime.

JEL Codes: K42, I18, H75

Keywords: cannabis legalization, supply chain regulation, crime, Netherlands, wietexperiment

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 20m).

1. Introduction

The Netherlands has tolerated cannabis retail since 1976, yet every gram sold in a Dutch coffeeshop arrives through an illegal supply chain. This paradox—legal at the front door, criminal at the back door—has defined global cannabis policy debates for half a century. Advocates of full legalization argue that the criminal supply chain sustains organized crime, corrupts growers, and degrades product safety (MacCoun and Reuter, 2001; Caulkins et al., 2012). Opponents warn that formalizing supply will expand availability, lower prices, and increase drug-related offenses (Kleiman et al., 2011). Until now, no country has run a controlled experiment to test these claims.

In June 2024, the Netherlands did exactly that. Ten municipalities entered the *Experiment gesloten coffeeshopketen* (closed coffeeshop supply chain experiment, or *wietexperiment*), which replaces illegal wholesale procurement with state-licensed cultivation, seed-to-sale tracking, and mandatory quality testing (van Laar and van der Pol, 2022). Ten additional municipalities, matched by WODC and RAND Europe on population size and coffeeshop density, serve as designated controls. The treatment was assigned by lottery from a pool of 26 volunteer municipalities, generating quasi-random variation in exposure to legal supply.

This paper estimates the effect of the cannabis supply chain experiment on crime. I use annual municipality-level crime data from Statistics Netherlands (CBS) covering 2010–2025, combined with a two-way fixed effects difference-in-differences (DiD) design that compares the 10 treatment municipalities to the 10 matched controls before and after the June 2024 rollout. I examine five outcomes: total drug crime, soft drug offenses, hard drug offenses, violent crime, and total registered crime.

The headline finding is a precisely centered null. The estimated effect on total drug crime is 7.0 additional offenses per 100,000 residents (SE = 23.0, $p = 0.77$), or 5.8 percent of the pre-treatment treatment-group mean. None of the five crime outcomes shows a statistically significant effect. These results survive a battery of robustness checks: restricting the pre-period to 2016–2025 (where a joint pre-trend test yields $p = 0.96$), adding municipality-specific linear trends, excluding the COVID years of 2020–2021, and replacing the DiD with a synthetic control method drawing from 316 non-experiment municipalities as donors. Permutation inference—randomly reassigning treatment among the 20 experiment municipalities 1,000 times—yields a two-sided p-value of 0.80, firmly in the null.

This paper contributes to three literatures. First, it provides the first causal evidence on the Dutch cannabis supply chain experiment, ahead of the government's own evaluation scheduled for 2028–2029. Prior empirical work on the experiment is limited to descriptive baseline reports (RAND Europe, 2024). Second, it contributes to the economics of cannabis legalization and crime. Gavrilova et al. (2019) show that medical marijuana laws in US border states reduced violent crime by undercutting Mexican drug trafficking organizations. Adda et al. (2014) find that depenalizing cannabis possession in London's Lambeth district had no effect on drug offenses but increased other crimes through a policing reallocation channel. Dragone et al. (2019) document that Italy's 2016 "cannabis light" legalization reduced organized crime offenses by disrupting black market revenues. Anderson and Rees (2023) provide a comprehensive review, concluding that the evidence on cannabis and crime remains mixed. My contribution is to study supply-side legalization in a setting where retail was already legal, isolating the supply chain mechanism that prior work cannot separate from demand-side liberalization.

Third, this paper speaks to the broader question of whether legalizing prohibited markets reduces the violence and disorder associated with black market enforcement (Miron and Zwiebel, 1995; Becker et al., 2006). The null result is informative: if the criminal supply chain were a major source of drug-related crime in experiment municipalities, replacing it should have produced a detectable decline. Instead, the evidence suggests that the "back door problem" generates less measurable crime than commonly assumed, or that the transitional phase—where both legal and illegal supply coexist—has not yet disrupted established criminal networks.

The paper proceeds as follows. Section 2 describes the institutional background. Section 3 presents the data. Section 4 details the empirical strategy. Section 5 reports results, and Section 6 discusses implications.

2. Institutional Background

The Dutch tolerance policy. The Netherlands has operated a *gedoogbeleid* (tolerance policy) for cannabis since 1976. Under this regime, coffeeshops—licensed cannabis retail establishments—may sell up to 5 grams per customer per transaction, hold a maximum stock of 500 grams, and are prohibited from advertising, selling hard drugs, or admitting minors (MacCoun, 2011). As of 2023, approximately 570 coffeeshops operated across the country, concentrated in urban centers (Trimbos Institute, 2024).

The core paradox is supply: while retail sale is tolerated, cultivation and wholesale distribution remain criminal offenses under the Opium Act. Coffeeshop operators must procure their entire inventory from illegal producers, creating a "back door" (*achterdeur*) through which organized criminal networks supply the legal market (Grund and Breeksema, 2016). Police and prosecutors have long recognized this contradiction as a driver of organized crime involvement in cannabis cultivation, with associated violence, environmental damage from illegal grow operations, and product contamination risks.

Design of the experiment. In response to these concerns, the Dutch parliament authorized the Controlled Cannabis Supply Chain Experiment (*Experiment gesloten coffeeshopketen*) through the 2019 Act on the Experiment with a Closed Coffeeshop Chain. The experiment's design, developed with WODC (the Ministry of Justice's research center) and RAND Europe, involved selecting 10 municipalities from 26 volunteers by lottery in August 2019 (van Laar and van der Pol, 2022). Ten designated growers were subsequently selected through a lottery and integrity screening process in December 2020.

The experiment operates in three phases. The *preparatory phase* (2020–2023) involved licensing growers, establishing quality control systems, and building the regulatory infrastructure. The *transitional phase* began on December 15, 2023 in Breda and Tilburg (the first two municipalities ready to launch) and expanded to all 10 treatment municipalities on June 17, 2024. During the transitional phase, coffeeshops may sell both regulated products from designated growers and existing (illegally sourced) inventory, allowing time for legal supply to scale. The *experimental phase* began on April 7, 2025, after which coffeeshops in participating municipalities may sell only legally produced cannabis.

Treatment and control municipalities. The 10 treatment municipalities—Almere, Arnhem, Breda, Groningen, Heerlen, Maastricht, Nijmegen, Tilburg, Voorne aan Zee, and Zaanstad—represent substantial geographic and demographic diversity, from large university cities (Groningen, Nijmegen) to smaller industrial towns (Heerlen, Voorne aan Zee). They collectively contain approximately 80 coffeeshops. The 10 control municipalities—Enschede, Haarlem, Helmond, Leeuwarden, Leiden, Lelystad, Roermond, Tiel, Utrecht, and Zutphen—were selected by WODC/RAND to match treatment municipalities on population size, urbanization level, and coffeeshop density (RAND Europe, 2024). Because the experiment municipalities were drawn from *volunteers*, selection on unobservables (e.g., political preferences, local drug policy orientation) is a concern. The matched control municipalities, also volunteers, partially address this, but I cannot rule out that volunteering municipalities differ systematically from non-volunteers.

3. Data

I use two datasets from Statistics Netherlands (CBS), the Dutch national statistical office.

Crime data. CBS StatLine table 83648NED provides counts of registered crimes by type and municipality for 2010–2025. I extract five crime categories: total drug offenses (Opium

Table 1: Pre-Treatment Summary Statistics (2010–2023)

	Treatment (10 municipalities)	Control (10 municipalities)
Mean population	165,733	120,434
Crime rates per 100,000:		
Drug crime (total)	138.38	103.24
Soft drugs	67.76	54.97
Hard drugs	63.11	46.27
Violence	742.38	672.68
Total crime	7,306.09	6,910.69
Municipality-years	127	140

Notes: Crime rates are registered crimes per 100,000 residents from CBS StatLine (table 83648NED). Treatment municipalities: Almere, Arnhem, Breda, Groningen, Heerlen, Maastricht, Nijmegen, Tilburg, Voorne aan Zee, Zaanstad. Control municipalities selected by WODC/RAND: Enschede, Haarlem, Helmond, Leeuwarden, Leiden, Lelystad, Roermond, Tiel, Utrecht, Zutphen.

Act violations), soft drug offenses (cannabis-related), hard drug offenses, violent crimes, and total registered crimes. I convert counts to rates per 100,000 residents using CBS population data (table 70072NED).

Sample. The main analysis sample consists of the 20 experiment municipalities (10 treatment, 10 control) observed annually over 16 years, yielding 307 municipality-year observations. For the synthetic control robustness check, I expand the sample to include all 480 Dutch municipalities with complete data over the period.

Table 1 presents pre-treatment summary statistics. Treatment municipalities are larger on average and have higher drug crime rates, consistent with their urban character and larger coffeeshop presence. The pre-treatment mean drug crime rate in treatment municipalities is 138 per 100,000, compared to 103 in controls—a level difference absorbed by municipality fixed effects in the DiD. The treatment group has fewer municipality-years (127 vs. 140) because Voorne aan Zee (GM1992) was created by a 2022 merger; pre-merger years have missing data in the CBS panel.

4. Empirical Strategy

4.1 Identification

The identifying assumption is parallel trends: absent the experiment, treatment and control municipalities would have experienced the same changes in crime rates over time. I estimate:

$$Y_{mt} = \alpha_m + \gamma_t + \beta \cdot (\text{Treat}_m \times \text{Post}_t) + \varepsilon_{mt} \quad (1)$$

where Y_{mt} is the crime rate per 100,000 in municipality m in year t, α_m and γ_t are municipality and year fixed effects, Treat_m indicates the 10 experiment municipalities, and $\text{Post}_t = 1$ for 2024–2025. The coefficient β captures the average treatment effect on the treated.

Standard errors are clustered at the municipality level. With only 20 clusters (10 treated, 10 control), asymptotic cluster-robust inference may be unreliable (Cameron et al., 2008). I therefore complement clustered standard errors with two non-parametric approaches: (i) permutation inference, which randomly reassigns treatment status among the 20 municipalities 1,000 times and computes the share of permutation estimates exceeding the observed estimate in absolute value; and (ii) a synthetic control method (Abadie et al., 2010) using the full panel of Dutch municipalities as potential donors, with Abadie-type placebo inference.

4.2 Event Study

To assess pre-trends, I estimate an event study specification:

$$Y_{mt} = \alpha_m + \gamma_t + \sum_{k \neq -1} \delta_k \cdot (\text{Treat}_m \times \mathbf{1}[t - 2024 = k]) + \varepsilon_{mt} \quad (2)$$

with 2023 ($k = -1$) as the reference year. A joint F-test of the pre-treatment coefficients $\{\delta_k\}_{k<-1}$ evaluates the parallel trends assumption.

4.3 Threats to Validity

Three concerns warrant discussion. First, treatment municipalities *volunteered* for the experiment, raising selection concerns. The WODC/RAND matching partially addresses this by constructing controls from the same volunteer pool. Second, the transitional phase allows both legal and illegal supply, potentially diluting treatment effects if the criminal supply chain remains active alongside legal growers. This is a feature, not a bug: the transitional phase is the policy-relevant treatment. Third, with 10 treated units and only 2 post-treatment years, statistical power is limited. I report 95 percent confidence intervals and discuss what effect

Table 2: Effect of Cannabis Supply Legalization on Crime

	Drug Crime (1)	Soft Drug (2)	Hard Drug (3)	Violence (4)	Total Crime (5)
Treatment × Post	6.98	-2.13	13.31	8.14	-119.12
	(23.05)	(12.59)	(10.99)	(27.93)	(179.88)
Pre-treatment mean	119.96	61.05	54.28	705.83	7,099
Effect (% of mean)	5.8%	-3.5%	24.5%	1.2%	-1.7%
Municipality FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	307	307	307	307	307
Municipalities	20	20	20	20	20

Notes: Each column reports estimates from a separate regression of crime rates per 100,000 residents on the interaction of treatment status with a post-2024 indicator, with municipality and year fixed effects. Standard errors clustered at the municipality level in parentheses. Treatment municipalities entered the controlled cannabis supply chain in June 2024. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

sizes the design can rule out.

5. Results

5.1 Main Results

Table 2 reports the main DiD estimates. Column (1) shows that the experiment is associated with 7.0 additional drug crimes per 100,000 residents ($p = 0.77$), an economically small effect equal to 5.8 percent of the pre-treatment treatment-group mean of 120 per 100,000. Decomposing by drug type, soft drug offenses show a slight decline of 2.1 per 100,000 (Column 2), while hard drug offenses increase by 13.3 (Column 3, $p = 0.24$)—suggestive but far from conventional significance. Violence (Column 4) and total crime (Column 5) show no meaningful effects.

The point estimates are small relative to the variation in the data. The 95 percent confidence interval for total drug crime ranges from approximately -38 to $+52$ per 100,000. To quantify what the design can rule out, I compute the minimum detectable effect (MDE) at 80 percent power using the standard formula MDE $= 2.8 \times$ SE, which yields approximately 47 per 100,000, or 39 percent of the treatment group's pre-treatment mean.[1] The design

[1]This uses MDE $\approx (t_{0.025} + t_{0.20}) \times$ SE $\approx 2.8 \times 23.0 \approx 64$, but with only 10 treated clusters, a more conservative calculation using the t-distribution with 18 degrees of freedom gives MDE $\approx 2.9 \times 23.0 \approx 67$ per 100,000. I report the asymptotic calculation.

Table 3: Event Study: Drug Crime Rate

Year (relative time)	Coefficient	SE
2016 ($t-8$)	-7.38	(22.86)
2018 ($t-6$)	-11.88	(21.47)
2020 ($t-4$)	-20.60	(23.52)
2021 ($t-3$)	-12.55	(17.32)
2022 ($t-2$)	-13.33	(20.19)
2024 ($t+0$)	11.71	(13.24)
2025 ($t+1$)	5.52	(14.08)
2023 ($t-1$)	[Reference]	
Pre-trend F-test	$F = 0.28$, $p = 0.96$	

Notes: Event study coefficients from a regression of drug crime rates per 100,000 on interactions of treatment status with year indicators, using 2023 ($t-1$) as the reference year. Sample restricted to 2016–2025 for clean pre-trends. Municipality and year fixed effects included. Standard errors clustered at the municipality level. The joint F-test fails to reject the null of zero pre-treatment effects ($p = 0.96$). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

therefore rules out large crime increases or decreases but cannot detect moderate effects in the range of 10–30 percent.

Transitional vs. experimental phase. Splitting the post-period into 2024 (transitional: legal and illegal supply coexist) and 2025 (near-experimental: predominantly legal supply), the estimated effects on drug crime are 11.7 per 100,000 in 2024 and 5.5 in 2025, both statistically insignificant. The similar magnitudes suggest that the null is not driven solely by transitional-phase dilution, though only one year of near-experimental data is available.

5.2 Event Study

Table 3 reports event study coefficients from the 2016–2025 specification. The pre-treatment coefficients are uniformly small and statistically insignificant, with a joint F-test of $F = 0.28$ ($p = 0.96$). Parallel trends hold convincingly over the 2016–2023 pre-period. The treatment-period coefficients ($k = 0$ and $k = 1$) are positive but small and insignificant, consistent with the main DiD results.

When the full 2010–2025 sample is used, the joint pre-trend test rejects ($p < 0.001$), driven by divergence in the 2010–2013 period. This early-period divergence reflects higher drug crime growth in treatment municipalities during the 2010s that had converged by 2016. I report the shorter pre-period as the preferred specification and the full sample as a sensitivity

Table 4: Robustness: Drug Crime Rate

Specification	Estimate	SE
(1) Baseline (2010–2025)	6.98	(23.05)
(2) Shorter pre-period (2016–2025)	18.25	(19.68)
(3) Municipality-specific trends	39.34	(28.16)
(4) Excluding COVID (2020–2021)	4.00	(24.02)
(5) Permutation p-value	$p = 0.800$	
(6) Synthetic control (aggregate)	-10.75	—

Notes: All specifications use total drug crime per 100,000 as the dependent variable. Rows (1)–(4) report DiD estimates with municipality and year fixed effects, clustered SEs. Row (5) reports the two-sided p-value from 1,000 random reassignments of treatment among the 20 experiment municipalities. Row (6) reports the average post-treatment gap from a synthetic control using 316 non-experiment municipalities as donors.

check.

5.3 Robustness

Table 4 summarizes robustness checks for the drug crime outcome. Row (2) restricts the pre-period to 2016–2025, yielding a point estimate of 18.3 (SE = 19.7)—slightly larger but still insignificant. Row (3) adds municipality-specific linear trends, producing an estimate of 39.3 (SE = 28.2). The larger magnitude under trends specification suggests some upward drift in treatment relative to control, but the estimate remains well within one standard error of zero. Row (4) excludes the COVID-affected years 2020–2021, with no material change ($\hat{\beta} = 4.0$, SE = 24.0). Row (5) reports the permutation p-value of 0.80, confirming that the observed effect is well within the noise generated by random treatment reassignment. Row (6) shows the synthetic control estimate: the aggregate treatment unit had 10.8 *fewer* drug crimes per 100,000 than its synthetic counterpart in the post-period—a slight negative effect, opposite in sign to the DiD and consistent with no true effect.

Placebo crime categories. Violence (the primary secondary outcome) shows no effect in any specification. Total crime shows a non-significant decline of 119 per 100,000 (SE = 180), or 1.7 percent of the pre-treatment mean. These placebo results provide no evidence of crime displacement or spillovers across offense categories. The absence of an effect on violence, in particular, suggests that the experiment has not triggered territorial conflicts between legal growers and existing criminal networks—a concern raised in policy debates. The null on total crime also argues against a policing substitution story, in which police resources freed from cannabis enforcement are redeployed to other crime types.

6. Discussion

The null effect has three possible interpretations. First, the back door problem may generate less crime than commonly assumed. If cannabis supply networks in the Netherlands are non-violent and well-established, replacing them with legal supply would produce little change in observable crime. Decorte (2010) documents that Dutch cannabis cultivation is often small-scale and locally embedded rather than controlled by organized crime networks, consistent with this interpretation.

Second, the transitional phase—in which legal and illegal supply coexist—may not yet have disrupted criminal operations. Legal growers faced significant capacity constraints during the first year, with five of ten designated cultivators failing to meet production requirements by August 2024. If coffeeshops continue to rely partially on illegal supply during the transition, the treatment contrast is diluted. The experimental phase beginning April 2025, which requires exclusively legal supply, will provide a sharper test.

Third, the null may reflect limited statistical power. With 10 treated municipalities and 2 post-treatment years, the design cannot detect effects smaller than approximately 30 percent of the pre-treatment mean (using the standard 80 percent power threshold). The effect of supply legalization on crime may be real but modest—a 10 percent decline in drug offenses, for instance, would be invisible in this design.

These interpretations are not mutually exclusive. The most policy-relevant conclusion is that, eighteen months into the world's most ambitious cannabis supply chain reform, there is no evidence of a crime increase. Fears that legalizing the back door would unleash new criminal activity—by lowering prices, expanding availability, or disrupting enforcement—find no support in the data. This is consistent with Gavrilova et al. (2019)'s finding that legalizing cannabis supply in US border states reduced violence, though the Dutch setting differs in that retail was already tolerated.

Limitations. The most important limitation is timing. The transitional phase is an inherently partial treatment, and the experimental phase has barely begun. A definitive evaluation requires data through at least 2027. Additionally, crime statistics reflect *registered* offenses, not actual criminal activity; changes in reporting behavior or enforcement priorities could mask true effects. Finally, the volunteer-based design limits external validity: municipalities that chose to participate may differ systematically from non-volunteers in ways that affect treatment response.

7. Conclusion

The Netherlands has embarked on the world's first controlled experiment with legal cannabis supply. In the first eighteen months—primarily under the transitional phase where legal and illegal supply coexist—this reform has not detectably changed drug crime, violence, or total crime in participating municipalities. The design rules out large effects (greater than 39 percent of the pre-treatment mean) but not moderate ones, making these results a preliminary assessment rather than a definitive verdict. The experimental phase, which began in April 2025 and requires exclusively legal supply, will provide a sharper test. For now, the evidence offers cautious reassurance: closing the back door has not opened new ones.

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

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller, "Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program," *Journal of the American Statistical Association*, 2010, *105* (490), 493–505.

Adda, Jérôme, Brendon McConnell, and Imran Rasul, "Crime and the depenalization of cannabis possession: Evidence from a policing experiment in London," *Journal of Political Economy*, 2014, *122* (5), 1130–1202.

Anderson, D. Mark and Daniel I. Rees, "The public health effects of legalizing marijuana," *Journal of Economic Literature*, 2023, *61* (1), 86–143.

Becker, Gary S., Kevin M. Murphy, and Michael Grossman, "The market for illegal goods: The case of drugs," *Journal of Political Economy*, 2006, *114* (1), 38–60.

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.

Caulkins, Jonathan P., Angela Hawken, Beau Kilmer, and Mark A.R. Kleiman, *Marijuana Legalization: What Everyone Needs to Know*, Oxford University Press, 2012.

Decorte, Tom, "The case for small-scale domestic cannabis cultivation," *International Journal of Drug Policy*, 2010, *21* (4), 271–275.

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.

Gavrilova, Evelina, Takuma Kamada, and Floris Zoutman, "Is legal pot crippling Mexican drug trafficking organisations? The effect of medical marijuana laws on US crime," *Economic Journal*, 2019, *129* (617), 375–407.

Grund, Jean-Paul C. and Joost Breeksema, "Beyond the back door: Closing the gap between policy and practice in Dutch cannabis regulation," in John Collins, ed., *Collapse of the War on Drugs*, London School of Economics, 2016, pp. 113–133.

Kleiman, Mark A.R., Jonathan P. Caulkins, and Angela Hawken, *Drugs and Drug Policy: What Everyone Needs to Know*, Oxford University Press, 2011.

MacCoun, Robert J., "What can we learn from the Dutch cannabis coffeeshop system?," *Addiction*, 2011, *106* (11), 1899–1910.

_ **and Peter Reuter**, *Drug War Heresies: Learning from Other Vices, Times, and Places*, Cambridge University Press, 2001.

Miron, Jeffrey A. and Jeffrey Zwiebel, "The economic case against drug prohibition," *Journal of Economic Perspectives*, 1995, *9* (4), 175–192.

RAND Europe, "Baseline Report: Closed Coffeeshop Supply Chain Experiment," Technical Report, RAND Corporation and WODC 2024. Prepared for the Dutch Ministry of Justice and Security.

Trimbos Institute, "Nationale Drug Monitor: Kerncijfers en Trends 2024," Technical Report, Netherlands Institute of Mental Health and Addiction, Utrecht 2024.

van Laar, Margriet W. and Peggy van der Pol, "Cannabis policy in the Netherlands: Rationale and design of an experiment with a controlled legal ('closed') cannabis supply chain," *International Journal of Drug Policy*, 2022, *112*, 103950.

A. Data Appendix

Crime data construction. I obtain municipality-level crime statistics from CBS StatLine table 83648NED, which reports registered crimes (*geregistreerde misdrijven*) by crime type and region for all Dutch municipalities. The data cover 2010–2025 at annual frequency. I extract five crime categories using CBS crime type codes: total drug offenses (Opium Act violations, code prefix CRI6000), soft drug offenses (CRI6200), hard drug offenses (CRI6100), violent crimes (CRI3000), and total registered crimes (T001161). Crime counts are converted to rates per 100,000 residents using population data from CBS table 70072NED. Municipalities with zero or missing population data are excluded.

Municipality identification. Treatment and control municipalities are identified by their CBS municipality codes (GM codes). I verify all 20 codes against the CBS municipality register. Voorne aan Zee (GM1992) was formed by a 2022 merger of three smaller municipalities; for years prior to the merger, I sum crime counts and population across the predecessor municipalities to maintain a balanced panel.

Sample construction. The raw crime data contain 592,800 records across all crime types, municipalities, and years. After filtering to annual data (excluding quarterly breakdowns), the five crime categories of interest, and years 2010–2025, I obtain 7,904 municipality-year observations for all Dutch municipalities. The experiment sample restricts to the 20 treatment and control municipalities, yielding 307 municipality-year observations (some municipalities have incomplete years due to boundary changes).

B. Identification Appendix

Pre-trend analysis. The event study in Table 3 uses 2023 ($k = -1$) as the reference year. With the full 2010–2025 sample, pre-treatment coefficients for 2010–2013 are large and positive, reflecting higher drug crime growth in treatment municipalities during the early 2010s. This pattern had converged by 2016. The joint pre-trend F-test with the full sample rejects ($p < 0.001$), but restricting to 2016–2025 yields $F = 0.28$ ($p = 0.96$), providing strong evidence for parallel trends in the recent pre-period.

Permutation inference details. I generate 1,000 random permutations by reassigning exactly 10 of the 20 municipalities as "treated." For each permutation, I re-estimate the DiD and record the coefficient. The observed coefficient of 7.0 falls at the 40th percentile of the permutation distribution (mean = 1.0, SD = 21.9), yielding a two-sided p-value of 0.80. The

Table 5: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Drug crime (total)	6.98	23.05	74.72	0.093	0.308	Moderate positive
Soft drug crime	-2.13	12.59	41.92	-0.051	0.300	Moderate negative
Hard drug crime	13.31	10.99	31.44	0.423	0.350	Large positive
Violence	8.14	27.93	177.47	0.046	0.157	Small positive
Total crime	-119.12	179.88	1,824	-0.065	0.099	Moderate negative

Notes: **Country:** Netherlands. **Research question:** Does legalizing the cannabis supply chain reduce drug-related crime in Dutch municipalities participating in the wietexperiment? **Policy mechanism:** The experiment replaces the illegal wholesale supply to coffeeshops (the "back door") with a state-regulated closed supply chain of 10 licensed growers, requiring full seed-to-sale tracking and quality testing while maintaining existing retail coffeeshop infrastructure. **Outcome definition:** Registered crimes per 100,000 residents from CBS StatLine table 83648NED, decomposed by type (total drug, soft drug, hard drug, violence, total). **Treatment:** Binary; municipality participates in the controlled cannabis supply chain experiment (transitional phase from June 2024). **Data:** CBS StatLine registered crime statistics and population counts, 2010–2025, municipality-year panel with 20 municipalities (10 treatment, 10 control) and 307 observations. **Method:** Two-way fixed effects DiD with municipality and year fixed effects; standard errors clustered at municipality level; robustness via permutation inference and synthetic control. **Sample:** 10 experiment municipalities selected by lottery from volunteers, matched with 10 control municipalities by WODC/RAND on pre-treatment characteristics; restricted to annual frequency for consistent CBS reporting. SDE = $\hat{\beta}$/SD(Y) where SD(Y) is the pre-treatment standard deviation. Classification refers to magnitude, not statistical significance: Large (|SDE| > 0.15), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).

5th and 95th percentiles of the permutation distribution are -33.9 and $+35.9$, respectively.

Synthetic control details. I aggregate the 10 treatment municipalities into a single treatment unit (population-weighted average crime rate) and construct a synthetic counterpart from 316 non-experiment municipalities with complete data over 2010–2025. The synthetic control matches the pre-treatment drug crime trajectory with an MSPE of 595.7. The post-treatment gap averages -10.8 (treatment below synthetic), indicating a slight crime *reduction* relative to the synthetic counterfactual—opposite in sign to the DiD estimate and consistent with a true null.

C. Standardized Effect Sizes