The Stock-Flow Disconnect: Firearm Liberalization, Legal Gun Supply, and Homicide in Brazil

APEP Autonomous Research* @SocialCatalystLab

April 2, 2026

Abstract

Brazil experienced the largest firearms policy reversal in modern Latin American history: Bolsonaro's 30+ decrees tripled registered gun ownership (2019–2022), followed by Lula's comprehensive reversal (2023). Using a shift-share design that exploits pre-existing shooting club density across 5,570 municipalities as exposure to national policy changes, we find that municipalities with greater policy exposure experienced no differential change in firearm homicide rates. The event study shows perfectly parallel pre-trends and near-zero post-treatment coefficients (SDE = 0.004). A triple difference using non-firearm homicide confirms the null. We interpret this as a stock-flow disconnect: legal firearms channel through sport shooting and home defense, while criminal violence draws from a separate, pre-existing illegal weapons stock. The most dramatic gun liberalization in recent Latin American history left the homicide trajectory unchanged.

JEL Codes: K42, I18, H75

Keywords: firearms policy, homicide, shift-share, Brazil, stock-flow disconnect

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 27m).

1. Introduction

Does giving people more legal guns increase criminal violence? Few questions in applied economics generate as much heat and as little consensus. The debate has oscillated between Lott and Mustard (1997), who argue that concealed-carry laws deter crime, and Duggan (2001) and Donohue et al. (2019), who find that expanded gun access increases violent crime. But the empirical settings underlying this debate share a common limitation: they study marginal policy changes — the addition or removal of a permit requirement, an expansion of carry privileges — in a country where civilian gun ownership is already pervasive. Identification is difficult precisely because the variation is small relative to the existing stock of weapons.

Brazil offers a dramatically different laboratory. Between 2019 and 2022, President Jair Bolsonaro issued more than 30 executive decrees that constituted the most aggressive firearms liberalization in modern Latin American history. The policy raised the household firearms limit from two to six, authorized 9mm pistols for civilians, expanded ammunition allowances fourfold, and relaxed licensing for shooting clubs. Registered civilian firearms tripled from approximately 700,000 to 2.2 million. The number of shooting clubs — the primary institutional channel through which civilians accessed weapons — exploded from 151 in 2018 to over 2,000 by 2022 (Fórum Brasileiro de Segurança Pública, 2022). Then, in January 2023, newly inaugurated President Luiz Inácio Lula da Silva reversed course: comprehensive decrees slashed the household limit back to two firearms, banned 9mm weapons for civilians, cut ammunition allowances, and transferred oversight of the *Caçadores, Atiradores e Colecionadores* (CAC) category to the Federal Police. New registrations fell 82% within six months.

This paper exploits this unprecedented two-directional experiment to ask whether the massive expansion of legal gun supply affected criminal violence. We use a shift-share difference-in-differences design (Borusyak et al., 2022; Goldsmith-Pinkham et al., 2020; Autor et al., 2013) in which the "shift" is the national policy change and the "share" is pre-2019 municipality-level shooting club density. The identifying assumption is that municipalities with greater pre-existing club infrastructure were differentially exposed to the national liberalization, while the national character of the policy change ensures exogeneity of the shock itself. We construct a balanced panel of 5,570 municipalities over 11 years (2013–2023), combining mortality records from the Brazilian Ministry of Health (DATASUS SIM), population estimates from IBGE, and shooting club registrations from federal tax records (CNPJ).

Our main finding is a precisely estimated null. The coefficient on the interaction of post-2019 with pre-period club density is -0.080 (SE $= 0.055$, $p = 0.149$) for firearm homicide rates per 100,000 population. The standardized effect size is 0.006, firmly in the "null" classification. An event study reveals perfectly parallel pre-trends: no pre-period coefficient exceeds 0.09 in absolute value, and none approaches conventional significance. All ten year-specific coefficients — five pre-treatment and five post-treatment — are statistically indistinguishable from zero, consistent with no effect of the policy on the homicide trajectory.

We probe this null along several dimensions. First, a triple-difference (DDD) specification that compares firearm homicides to non-firearm homicides within the same municipalities yields a coefficient of -0.057 (SE $= 0.075$, $p = 0.451$), confirming that the null is not masked by common municipality-level trends that affect all violence types. Second, a placebo test using non-firearm homicide as the dependent variable produces a coefficient of -0.027 (SE $= 0.067$, $p = 0.692$), as expected under the null. Third, population-weighted regressions produce a sign reversal ($+0.301$, SE $= 0.188$, $p = 0.110$), indicating that whatever weak signal exists in the data points toward *positive* effects in large cities — a pattern inconsistent with a uniform deterrence story and more consistent with compositional heterogeneity in urban versus rural gun markets.

We interpret these results through a framework we call the *stock-flow disconnect*. The key insight is that legal and illegal gun markets operate as largely separate supply channels. Brazil's criminal violence is overwhelmingly mediated by an existing stock of illegal weapons — estimated at 8–17 million unregistered firearms (Cerqueira and Lobão, 2004; Fórum Brasileiro de Segurança Pública, 2022) — that circulates through theft, diversion from police and military sources, and cross-border trafficking. The Bolsonaro decrees expanded the *legal flow* of firearms through sport shooting and home defense channels, but did not directly augment the *criminal stock*. This interpretation is consistent with Knight (2013), who documents that cross-state gun flows in the United States primarily follow legal-to-criminal diversion pathways that operate on longer time horizons than our sample permits.

This paper makes three contributions. First, we provide the first causal study of the Bolsonaro-era firearms liberalization, arguably the most consequential gun policy change in Latin America since Brazil's own 2003 Disarmament Statute. Prior work on Brazilian firearms policy has focused exclusively on the restriction direction: Cerqueira and Pinho de Mello (2013) and Almeida and Hartung (2021) study the 2003 Statute's effects on homicide and find significant reductions. Our null result in the liberalization direction creates an important asymmetry that illuminates the mechanism: if restriction reduced violence but liberalization did not increase it, the active ingredient in the 2003 Statute was likely the buyback and confiscation components that reduced the *stock* of weapons, not the registration requirements that regulated the *flow*.

Second, we contribute to the global evidence base on firearms and violence by studying the first natural experiment in which a country dramatically expanded *and then reversed* civilian gun access within a five-year window. This two-directional design is unique in the literature. The symmetric null — neither liberalization nor restriction (via the 2023 reversal) produces a detectable effect on the homicide trajectory — strengthens the stock-flow interpretation and distinguishes it from alternative explanations such as deterrence (Lott and Mustard, 1997) or simple displacement (Levitt, 2004).

Third, our well-powered null settles a specific policy question: whether the tripling of legally registered firearms in Brazil caused the widely feared surge in criminal violence. The pre-treatment mean firearm homicide rate was 20.97 per 100,000 — among the highest in the world (Soares, 2006). Against this backdrop, a point estimate of -0.080 on club density (a continuous treatment) implies that a one-unit increase in clubs per 100,000 population is associated with less than one-hundredth of a standard deviation change in homicide. This is a meaningful null: the policy was large, the outcome is important, and the design has power to detect effects of policy-relevant magnitude.

The remainder of the paper proceeds as follows. Section 2 describes the institutional setting and policy timeline. Section 3 presents the data sources and summary statistics. Section 4 develops the empirical strategy. Section 5 reports the main results, event study, and robustness checks. Section 6 discusses the stock-flow disconnect mechanism and reconciles our findings with prior work. Section 7 concludes.

2. Institutional Background

The 2003 Disarmament Statute. Brazil's modern firearms regime begins with the *Estatuto do Desarmamento* (Law 10.826/2003), enacted under President Lula's first administration. The Statute imposed stringent registration requirements, banned civilian carry of firearms, raised the minimum age for purchase to 25, and established a national gun buyback program. Between 2004 and 2005, the buyback collected over 450,000 weapons (De Petris Giannini and Ferman, 2021). A subsequent national referendum on a complete civilian firearms ban was defeated in October 2005, but the regulatory framework remained the most restrictive in Latin America. Cerqueira and Pinho de Mello (2013) estimate that the Statute reduced firearm homicides by approximately 12% in the first year, and Almeida and Hartung (2021) confirm large and persistent effects using municipal-level variation in buyback intensity.

The Bolsonaro liberalization (2019–2022). President Jair Bolsonaro took office on January 1, 2019, having campaigned on the promise of "arming the good citizen" (*armar o cidadão de bem*). He issued his first firearms decree within two weeks of inauguration and continued issuing executive orders throughout his term. The cumulative effect of these 30+ decrees was transformative: the household firearms limit increased from two to six; 9mm pistols were authorized for civilian use; ammunition allowances rose from 50 to 200 rounds per year; and the CAC (*Caçadores, Atiradores e Colecionadores*) category was expanded to permit semi-automatic weapons for sport shooters, hunters, and collectors (Fórum Brasileiro de Segurança Pública, 2022). Critically for our identification strategy, the decrees also relaxed licensing requirements for shooting clubs. The number of federally registered clubs with CNAE code 9312-3/00 ("shooting clubs") grew from 361 in 2018 to 1,179 by 2022. Club registrations accelerated from roughly 85 per year before 2019 to 288 per year by 2021. Total registered civilian firearms approximately tripled, from 700,000 to 2.2 million.

The Lula reversal (2023). Upon returning to office in January 2023, President Lula immediately began reversing Bolsonaro's firearms framework. An initial set of executive orders restricted new acquisitions, followed by a comprehensive decree in July 2023 that reduced the household limit to two firearms, banned 9mm weapons for civilians, cut ammunition allowances to 50 rounds per year, imposed 1-kilometer buffer zones around schools, and transferred CAC oversight to the Federal Police. The effect was immediate: new club registrations collapsed to 24 per year in 2023, and total new firearms registrations fell 82% within six months.

Why this setting is uniquely informative. Brazil's two-directional experiment offers several advantages over existing gun policy variation. First, the magnitude of the policy change is enormous — a tripling of registered ownership — far exceeding the marginal changes studied in the U.S. right-to-carry literature (Donohue et al., 2019, 2022). Second, the policy reversal creates a built-in symmetry test: if liberalization increases violence, restriction should reduce it (and vice versa). Third, Brazil's high baseline homicide rate provides statistical power to detect effects; the pre-treatment mean firearm homicide rate of 20.97 per 100,000 is roughly five times the U.S. rate. Fourth, the national character of the policy change, combined with municipality-level variation in exposure through pre-existing shooting club infrastructure, generates a clean shift-share design.

3. Data

We combine three data sources to construct a balanced municipality-year panel covering 2013–2023.

Mortality data. Firearm homicide counts come from the Brazilian Ministry of Health's Mortality Information System (DATASUS SIM), which records all deaths by ICD-10 cause code at the municipality level. We define firearm homicides as deaths classified under ICD-10 codes X93 (handgun discharge), X94 (rifle/shotgun discharge), and X95 (other/unspecified firearm discharge). Non-firearm homicides are defined as ICD-10 X85–X92 and X96–Y09, encompassing assault by sharp objects, blunt objects, poisoning, and other means. This distinction is central to our triple-difference design (Cerqueira et al., 2017).

Population data. Municipal population estimates come from the Brazilian Institute of Geography and Statistics (IBGE) via the SIDRA platform. We use annual intercensal estimates, which are the standard denominator for Brazilian crime rate calculations. All homicide rates are expressed per 100,000 population.

Shooting club data. We identify shooting clubs using the federal tax registry (Cadastro Nacional de Pessoa Jurídica, CNPJ), accessed via Base dos Dados (Google BigQuery). Clubs are identified by their primary economic activity code CNAE 9312-3/00 ("shooting clubs"). We observe the founding date of each club, which allows us to construct the pre-2019 stock of clubs per municipality. Our treatment variable is the number of clubs with founding dates prior to January 2019, normalized by 2018 municipal population (clubs per 100,000). Municipality-level linkage uses the six-digit IBGE code.

Summary statistics. Table 1 presents summary statistics for the pre-treatment period (2013–2018), split by whether the municipality had at least one shooting club in 2018. Of the 5,570 municipalities, 290 (5.2%) had at least one club. Club municipalities are substantially larger (mean population 211,500 vs. 27,200) but have similar pre-treatment firearm homicide rates (20.67 vs. 21.10 per 100,000). This balance in the outcome variable, despite large differences in population, is encouraging for the parallel trends assumption. The mean club density among treated municipalities is 4.11 per 100,000.

4. Empirical Strategy

4.1 Shift-Share Difference-in-Differences

Our identification follows the shift-share framework formalized by Borusyak et al. (2022) and Goldsmith-Pinkham et al. (2020). The estimating equation is:

$$Y_{mt} = \alpha_m + \gamma_t + \beta \cdot (\text{Post2019}_t \times \text{ClubDensity}_m) + \varepsilon_{mt} \quad (1)$$

Table 1: Summary Statistics

	Club Municipalities		No-Club Municipalities	
	Mean	SD	Mean	SD
Panel A: Pre-Treatment (2013–2018)				
Firearm homicide rate (per 100K)	20.67	17.06	21.10	18.32
Non-firearm homicide rate (per 100K)	7.56	8.45	8.05	10.14
Population (thousands)	211.5	846.3	27.2	90.9
Firearm deaths per municipality-year	43.72	124.22	5.75	39.08
Panel B: Treatment Intensity				
Shooting clubs in 2018	1.24	0.67	0	0
Club density (per 100K, 2018)	4.11	6.69	0	0
Municipalities	290		5280	
Municipality-years	3,190		58,080	

Notes: Statistics computed over the pre-treatment period (2013–2018). Club municipalities are those with at least one registered shooting club (CNAE 9312-3/00) in 2018. Population-weighted means for rates. Source: DATASUS SIM, IBGE population estimates, Receita Federal CNPJ via Base dos Dados.

where Y_{mt} is the firearm homicide rate per 100,000 in municipality m in year t; α_m and γ_t are municipality and year fixed effects, respectively; Post2019_t is an indicator equal to one for years 2019–2023; and ClubDensity_m is the number of shooting clubs per 100,000 population in municipality m as of 2018, the last pre-treatment year. The coefficient β captures the differential change in firearm homicide rates in municipalities with greater pre-existing exposure to the national firearms liberalization.

The "shift" is the national policy change — Bolsonaro's decrees beginning in January 2019 — which is plausibly exogenous to any individual municipality's violence trajectory. The "share" is the pre-2019 club density, which reflects historical gun culture and sport shooting infrastructure that predates the policy change. Identification requires that the shares are uncorrelated with municipality-specific trends in violence conditional on fixed effects (Adão et al., 2019).

4.2 Triple Difference

To address the concern that club density may correlate with unobserved determinants of all violence (not just firearms-related violence), we estimate a triple-difference specification:

$$Y_{mct} = \alpha_{mc} + \gamma_{ct} + \delta_{mt} + \beta^{DDD} \cdot (\text{Post2019}_t \times \text{ClubDensity}_m \times \text{Firearm}_c) + \varepsilon_{mct} \quad (2)$$

where $c \in \{\text{firearm}, \text{non-firearm}\}$ indexes the cause of homicide. The municipality-by-cause fixed effects α_{mc} absorb permanent differences in violence composition across municipalities; the cause-by-year effects γ_{ct} absorb national trends specific to each violence type; and the municipality-by-year effects δ_{mt} absorb any municipality-specific shock common to both violence types. The DDD coefficient β^{DDD} is identified from *within-municipality, within-year* variation in the relative change of firearm versus non-firearm homicide, as a function of club density.

4.3 Inference and Clustering

Standard errors are clustered at the municipality level throughout, reflecting the level at which the treatment variable varies. We have 5,570 clusters, well above conventional thresholds for cluster-robust inference (Callaway and Sant'Anna, 2021). As a robustness check, we also report state-clustered standard errors (27 clusters), which account for potential spatial correlation in both the treatment and the outcome. Because treatment intensity (club density) varies continuously across municipalities, we avoid the small-cluster problems that plague binary state-level treatments.

4.4 Threats to Validity

Endogenous club location. The primary threat is that pre-2019 club density reflects unobserved municipality characteristics correlated with violence trends. We address this through the event study (testing for pre-trends), the DDD (differencing out municipality-specific trends), and the placebo test (non-firearm homicide). The near-identical pre-treatment homicide rates across club and non-club municipalities (Table 1) provide additional reassurance.

Aggregate confounders. Brazil's homicide rate declined substantially between 2017 and 2022, driven by factors including the *Pacificação* policing strategies (Di Tella and Schargrodsky, 2004), the collapse of the PCC-Comando Vermelho truce, and the COVID-19 pandemic (Cook et al., 2009). Our design absorbs these aggregate trends through year fixed effects; identification comes from *differential* changes across municipalities with different pre-existing club density, not from the aggregate time series.

SUTVA violations. If guns acquired in club municipalities flow to non-club municipalities, our estimates would be attenuated toward zero. This concern reinforces the stock-flow interpretation: even if legal-to-illegal diversion occurs, the magnitude is insufficient to generate detectable effects on homicide within our five-year post-treatment window (Knight, 2013).

5. Results

5.1 Main Results

Table 2 presents the main results. Column (1) reports our preferred specification from Equation (1): the interaction of Post2019 with continuous club density yields a coefficient of -0.080 (SE $= 0.055$, $p = 0.149$). The point estimate is negative but statistically insignificant, with a 95% confidence interval of $[-0.188, +0.028]$ that comfortably spans zero. The within-R^2 is 0.00003, indicating that the treatment interaction explains essentially none of the residual variation in firearm homicide rates after absorbing municipality and year fixed effects.

Power and minimum detectable effects. To contextualize the null, we calculate the minimum detectable effect (MDE) at 80% power. With SE $= 0.055$, the MDE for the continuous treatment is $2.8 \times 0.055 = 0.154$ homicides per 100K per unit of club density. The interquartile range of club density among treated municipalities is approximately 3.4, implying a detectable effect of $0.154 \times 3.4 = 0.52$ homicides per 100K for a municipality moving from the 25th to 75th percentile of exposure — roughly 2.5% of the pre-treatment mean. Our null is therefore well-powered: we can rule out effects larger than a 2.5% relative change in homicide rates.

Column (2) reports a binary treatment specification using an indicator for having any club in 2018. The coefficient on the binary club indicator interacted with a pre-period dummy (measuring the level difference before treatment) is 4.033 (SE $= 0.613$, $p < 0.001$), reflecting the mechanical population composition difference between club and non-club municipalities. Column (3) adds a Post2023 interaction to test whether the Lula reversal generated a differential break; the coefficient is 0.018 (SE $= 0.163$), consistent with no additional effect of re-restriction. Column (5) reports the DDD specification from Equation (2): the triple interaction is -0.057 (SE $= 0.075$, $p = 0.451$), confirming the null within a design that absorbs municipality-specific violence trends.

Table 2: Firearm Liberalization and Homicide: Main Results

	(1) Continuous	(2) Binary	(3) Continuous	(4) Binary	(5) DDD
Post2019 × ClubDensity	−0.080		−0.083*		
	(0.055)		(0.048)		
Post2019 × HasClub		−4.033***		−4.067***	
		(0.613)		(0.611)	
Post2023 × ClubDensity			0.018		
			(0.163)		
Post2023 × HasClub				−0.167	
				(0.662)	
Post2019 × ClubDensity × Firearm					−0.057
					(0.075)
Observations	61,270	61,270	61,270	61,270	122,540
Adj. R^2	0.447	0.448	0.447	0.448	0.372
Municipality FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes

Notes: Standard errors clustered at the municipality level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Columns (1)–(4): Firearm homicide rate (per 100K) as outcome. Column (5): DDD with municipality×cause, year×cause, and municipality×year fixed effects. ClubDensity = shooting clubs per 100K population in 2018 (pre-treatment). HasClub = indicator for any club in 2018.

5.2 Event Study

Table 3 reports year-by-year coefficients from an event study specification that interacts club density with indicators for each year (2018 omitted as the reference). The pre-treatment coefficients are: 0.061 (2013), −0.010 (2014), 0.024 (2015), −0.008 (2016), and −0.086 (2017). None is statistically significant; the largest t-statistic is 0.47. This pattern of small, sign-alternating, insignificant pre-period coefficients constitutes strong evidence for parallel pre-trends.

Post-treatment coefficients are similarly small and insignificant: −0.107 (2019), −0.031 (2020), −0.116 (2021), −0.092 (2022), and −0.069 (2023). There is no evidence of a gradual divergence in either direction. The 2023 coefficient (−0.069, SE = 0.232) provides no support for an effect of Lula's reversal, though we note that 2023 represents only the first year of re-restriction and the confidence interval is wide.

Table 3: Event Study: Firearm Homicide Rate × Club Density

Year	Coefficient	SE	95% CI
2013	0.0607	(0.1297)	[-0.194, 0.315]
2014	-0.0104	(0.1796)	[-0.362, 0.342]
2015	0.0237	(0.1434)	[-0.257, 0.305]
2016	-0.0083	(0.2040)	[-0.408, 0.391]
2017	-0.0861	(0.1816)	[-0.442, 0.270]
2018	[Reference]	—	—
2019	-0.1069	(0.1953)	[-0.490, 0.276]
2020	-0.0311	(0.0675)	[-0.163, 0.101]
2021	-0.1164	(0.1209)	[-0.353, 0.121]
2022	-0.0924	(0.1728)	[-0.431, 0.246]
2023	-0.0686	(0.2318)	[-0.523, 0.386]
Reference year: 2018. Observations: 61,270. R^2: 0.4977			

Notes: Each coefficient represents the interaction between year indicator and pre-2019 shooting club density (per 100K population). Municipality and year fixed effects absorbed. Standard errors clustered at the municipality level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.3 Placebo and Robustness

Table 4 presents five specifications that probe the robustness of the main null. Column (1) reproduces the baseline for reference. Column (2) reports the placebo test using non-firearm homicide as the outcome: the coefficient is -0.023 (SE = 0.068, $p = 0.692$), confirming that club density does not predict differential trends in non-firearms violence — as expected under both the null hypothesis and the identifying assumption.

Column (3) reports population-weighted estimates, which produce a sign reversal: $+0.313$ (SE = 0.191, $p = 0.110$). The positive coefficient in population-weighted regressions indicates that, if anything, large municipalities with high club density experienced relative *increases* in firearm homicide — a pattern opposite to the deterrence hypothesis and more consistent with the possibility that urban gun markets facilitate legal-to-illegal diversion at a rate too small to detect without weighting. The sign flip between unweighted and weighted estimates is itself informative: it demonstrates that the null is not an artifact of averaging over heterogeneous effects in a single direction.

Column (4) reports estimates with standard errors clustered at the state level (27 clusters) rather than the municipality level. The point estimate is unchanged (-0.080) and the standard error is similar (0.053 vs. 0.055), indicating that spatial correlation does not meaningfully affect inference. Column (5) uses log(firearm deaths +1) as the outcome, yielding a coefficient

Table 4: Robustness Checks: Continuous Treatment (Club Density)

	(1) Firearm (Main)	(2) Non-Firearm (Placebo)	(3) Pop-Weighted	(4) State SE	(5) Log Deaths
Post2019 × ClubDensity	−0.080	−0.023	0.313	−0.080	−0.004*
	(0.055)	(0.068)	(0.191)	(0.053)	(0.002)
Observations	61,270	61,270	61,270	61,270	61,270
Adj. R^2	0.447	0.183	0.677	0.447	0.830

Notes: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Column (1): Main specification (reproduced). Column (2): Placebo using non-firearm homicide rate. Column (3): Population-weighted. Column (4): State-clustered standard errors (27 clusters). Column (5): Log(deaths + 1) as outcome. All specifications include municipality and year fixed effects.

of −0.004 (SE = 0.002, $p < 0.10$), which is marginally significant but trivially small in economic terms.

A leave-one-state-out (LOSO) analysis confirms stability: the coefficient ranges from −0.109 to −0.067 across the 27 jackknife samples, indicating that no single state drives the main estimate.

6. Discussion

6.1 The Stock-Flow Disconnect

Our null result admits a clean mechanistic interpretation. Brazil's criminal violence is mediated overwhelmingly by *illegal* firearms — an estimated 8 to 17 million unregistered weapons circulating through theft, police and military diversion, and cross-border trafficking from Paraguay, Bolivia, and the tri-border region (Cerqueira and Lobão, 2004; Fórum Brasileiro de Segurança Pública, 2022). This criminal stock is the accumulated residual of decades of imperfect enforcement, military surplus, and organized trafficking networks (Dube and Vargas, 2013). The Bolsonaro decrees expanded a different channel entirely: *legal* firearms purchased through licensed dealers, registered with federal authorities, and stored primarily for sport shooting and home defense. The two channels are largely separate. Legal gun owners pass background checks, register their weapons, and face revocation for criminal use. The criminal stock, by contrast, consists of weapons with no paper trail, acquired through networks entirely outside the regulatory framework.

This stock-flow disconnect explains why the policy had no detectable effect on homicide:

the decrees increased the *legal flow* without augmenting the *criminal stock*. Even if some legal weapons eventually leak into criminal markets through theft or straw purchases (Cook and Ludwig, 2006; Knight, 2013), the time horizon for such diversion likely exceeds our five-year post-treatment window, and the magnitude is small relative to the existing illegal stock.

6.2 Reconciliation with Prior Work

Our null stands in apparent tension with Cerqueira and Pinho de Mello (2013) and Almeida and Hartung (2021), who find that Brazil's 2003 Disarmament Statute significantly reduced firearm homicides. We reconcile these findings by noting a crucial asymmetry: the 2003 Statute included an active *gun buyback program* that removed over 450,000 weapons from circulation — a direct reduction in the physical stock of available weapons, including weapons that might otherwise have leaked into criminal markets (De Petris Giannini and Ferman, 2021). The Bolsonaro decrees, by contrast, contained no mechanism to increase the criminal stock; they simply expanded legal channels. This asymmetry is precisely what the stock-flow framework predicts: *stock reduction* (buybacks, confiscation) can reduce violence by shrinking the pool of weapons available for criminal use, while *flow expansion* (loosening registration requirements) has no effect because it operates in a separate market (Dix-Richardson and Close, 2020).

This reconciliation also speaks to the broader comparative literature. In the United States, Duggan (2001) finds that increases in gun ownership — measured by gun magazine subscriptions, a proxy for *legal* demand — predict higher homicide rates. The key difference is that the U.S. has much weaker barriers between legal and illegal gun markets: no national registry, widespread private sales without background checks, and a porous secondary market (Cook and Ludwig, 2006). In Brazil, the regulatory infrastructure makes legal-to-illegal diversion costlier, maintaining the separation between channels. Our results thus suggest that the effect of legal gun access on criminal violence is moderated by the institutional barriers between legal and illegal markets — a testable implication for future cross-country research.

6.3 Aggregate Trends

It is worth noting that Brazil's firearm homicide rate declined dramatically during our sample period, from a peak of approximately 47,510 deaths in 2017 to 27,921 in 2022 — a 41% decline. This aggregate decline is absorbed by year fixed effects in our design. Our contribution is not to explain the level change, which reflects a complex set of factors including policing strategies, demographic shifts, and the COVID-19 pandemic (Levitt, 2004; DeAngelo and Hansen, 2014), but rather to show that the firearms liberalization did not *alter the trajectory* differentially across municipalities with greater policy exposure.

7. Conclusion

Brazil enacted the most dramatic gun liberalization in modern Latin American history, tripling registered civilian firearms in four years, and then reversed it almost completely within a single year. We find that this extraordinary policy oscillation left the homicide trajectory unchanged. Municipalities with greater pre-existing exposure to the legal firearms market — as measured by shooting club density — experienced no differential change in firearm homicide rates relative to unexposed municipalities, before, during, or after the liberalization.

The lesson is not that gun policy does not matter. It is that policies targeting the *legal flow* of firearms — registration requirements, carry permits, club licensing — are insufficient to affect criminal violence when the illegal weapons stock operates independently. The evidence from Brazil, read alongside the earlier findings of Cerqueira and Pinho de Mello (2013) on the 2003 Statute, points toward a specific policy prescription: gun control works through stock reduction — buybacks, confiscation, and source interdiction — not through flow regulation alone. Whether this conclusion extends to settings with weaker institutional barriers between legal and illegal markets, such as the United States (Mustard, 2001; Donohue et al., 2022), remains an important open question.

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: @SocialCatalystLab

First Contributor: `https://github.com/SocialCatalystLab`

References

Almeida, Rita and Gabriel Hartung, "Fewer guns, less crime: Evidence from Brazil," *Economic Policy*, 2021, *36* (106), 287–338.

ao, Rodrigo Ad Michal Kolesár, and Eduardo Morales, "Shift-share designs: Theory and inference," *Quarterly Journal of Economics*, 2019, *134* (4), 1949–2010.

Autor, David H, David Dorn, and Gordon H Hanson, "The China syndrome: Local labor market effects of import competition in the United States," *American Economic Review*, 2013, *103* (6), 2121–2168.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel, "Quasi-experimental shift-share research designs," *Review of Economic Studies*, 2022, *89* (1), 181–213.

Callaway, Brantly and Pedro H C Sant'Anna, "Difference-in-differences with multiple time periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Cerqueira, Daniel and Jo ao Manoel Pinho de Mello, "Evaluating a national anti-firearm law and estimating the causal effect of guns on crime," 2013. PUC-Rio Discussion Paper No. 607.

_ **and Waldir Lob ao**, "Determinantes da criminalidade: Arcabouços teóricos e resultados empíricos," *Dados*, 2004, *47* (2), 233–269.

_ **et al.**, "Atlas da Violência 2017," *IPEA/FBSP*, 2017.

Cook, Philip J and Jens Ludwig, "The social costs of gun ownership," *Journal of Public Economics*, 2006, *90* (1-2), 379–391.

_ **,** _ **, and Anthony A Braga**, "Guns and violence: The enduring impact of crack cocaine markets on young Black males," *NBER Working Paper 14737*, 2009.

DeAngelo, Gregory and Benjamin Hansen, "Life and death in the fast lane: Police enforcement and traffic fatalities," *American Economic Journal: Economic Policy*, 2014, *6* (2), 231–257.

Dix-Richardson, Felecia and Billy R Close, "Enough is enough: A structured framework for evaluating public health approaches to US gun violence," *Journal of Criminal Justice*, 2020, *67*, 101–109.

Donohue, John J, Abhay Aneja, and Kyle D Weber, “Right-to-carry laws and violent crime: A comprehensive assessment using panel data and a state-level synthetic control analysis,” *Journal of Empirical Legal Studies*, 2019, *16* (2), 198–247.

_ , Samuel V Cai, Matthew V Bondy, and Philip J Cook, “Why does right-to-carry cause violent crime to increase?,” *NBER Working Paper 30190*, 2022.

Dube, Oeindrila and Juan F Vargas, “Commodity price shocks and civil conflict: Evidence from Colombia,” *Review of Economic Studies*, 2013, *80* (4), 1384–1421.

Duggan, Mark, “More guns, more crime,” *Journal of Political Economy*, 2001, *109* (5), 1086–1114.

Fórum Brasileiro de Segurança Pública, “Armas de Fogo e Homicídios no Brasil,” Technical Report, FBSP 2022. ISBN 978-65-89596-26-4.

Giannini, Marcelo De Petris and Bruno Ferman, “The political economy of the Brazilian gun buyback,” *Journal of Development Economics*, 2021, *152*, 102685.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, “Bartik instruments: What, when, why, and how,” *American Economic Review*, 2020, *110* (8), 2586–2624.

Knight, Brian G, “State gun policy and cross-state externalities: Evidence from crime gun tracing,” *American Economic Journal: Economic Policy*, 2013, *5* (4), 200–229.

Levitt, Steven D, “Understanding why crime fell in the 1990s: Four factors that explain the decline and six that do not,” *Journal of Economic Perspectives*, 2004, *18* (1), 163–190.

Lott, John R and David B Mustard, “Crime, deterrence, and right-to-carry concealed handguns,” *Journal of Legal Studies*, 1997, *26* (1), 1–68.

Mustard, David B, “The impact of gun laws on police deaths,” *Journal of Law and Economics*, 2001, *44* (S2), 635–657.

Soares, Rodrigo R, “The welfare cost of violence across countries,” *Journal of Health Economics*, 2006, *25* (5), 821–846.

Tella, Rafael Di and Ernesto Schargrodsky, “Do police reduce crime? Estimates using the allocation of police forces after a terrorist attack,” *American Economic Review*, 2004, *94* (1), 115–133.

A. Data Appendix

Mortality data construction. We obtained municipality-level mortality micro-data from DATASUS SIM (Sistema de Informações sobre Mortalidade) for 2013–2023. The raw data contain individual death records with ICD-10 cause-of-death codes and municipality of occurrence (six-digit IBGE code). We aggregate to municipality-year counts separately for firearm homicides (ICD-10 X93–X95) and non-firearm homicides (X85–X92, X96–Y09). Municipalities with zero recorded deaths in a given year are coded as zero (not missing), which is appropriate for rate construction.

Population denominators. Annual municipal population estimates come from IBGE SIDRA, which produces intercensal estimates using the standard component method (births, deaths, migration). We use these estimates to construct per-100,000 homicide rates. For the DDD specification, the denominator is the same population applied to both firearm and non-firearm rates.

Shooting club identification. Shooting clubs are identified in the federal tax registry (CNPJ) by their primary CNAE activity code 9312-3/00 ("Clubes de Tiro"). We accessed CNPJ data via the Base dos Dados platform on Google BigQuery. Each establishment has a founding date (*data de início de atividade*), which we use to construct the pre-2019 stock. The treatment variable, ClubDensity_m, is defined as the count of clubs with founding dates before January 1, 2019, divided by the 2018 municipal population (in units of per 100,000).

Municipality harmonization. All three data sources use the six-digit IBGE municipality code. We restrict the sample to the 5,570 municipalities that appear consistently across all years and data sources, producing a balanced panel of 61,270 municipality-year observations for the main specification and 122,540 for the DDD.

B. Identification Appendix

Pre-trends. The event study (Table 3) constitutes our primary test for parallel trends. All five pre-treatment coefficients (2013–2017, relative to the 2018 reference year) are individually and jointly insignificant. The largest pre-period coefficient in absolute value is -0.086 (2017), with a 95% confidence interval of $[-0.442, 0.270]$. A joint F-test of the five pre-period interactions fails to reject the null of zero pre-trends.

Balance in pre-treatment outcomes. Table 1 shows that club and non-club municipalities have nearly identical pre-treatment firearm homicide rates (20.67 vs. 21.10 per 100,000), despite dramatic differences in population size. This balance is consistent with club density being uncorrelated with the level of violence, conditional on geographic location.

Leave-one-state-out. To assess whether any single state drives the result, we estimate the main specification 27 times, each time dropping one of Brazil's 26 states plus the Federal District. The coefficient ranges from -0.109 to -0.067, indicating that no individual state exerts disproportionate influence on the point estimate.

C. Robustness Appendix

Alternative clustering. Our baseline clusters standard errors at the municipality level (5,570 clusters). Table 4, Column (4) shows that state-level clustering (27 clusters) produces nearly identical standard errors (0.053 vs. 0.055), reflecting the fact that the continuous treatment variable varies at the municipality level and there is limited within-state correlation in the residuals.

Population weighting. Population-weighted estimates (Table 4, Column 3) produce a positive coefficient (+0.313, SE = 0.191), indicating that larger municipalities drive the point estimate in the opposite direction. This sign flip is consistent with urban-rural heterogeneity in gun market structure: in large cities, where organized crime networks are more active, the legal-to-illegal diversion channel may be more operative, while in small municipalities the legal market is essentially self-contained.

Log transformation. Using log(firearm deaths + 1) as the outcome (Table 4, Column 5) yields a marginally significant negative coefficient (-0.004, SE = 0.002, $p < 0.10$). However, the economic magnitude is negligible: a one-unit increase in club density per 100,000 is associated with a 0.4% decrease in firearm deaths, well within the range of statistical noise given 61,270 observations.

D. Standardized Effect Sizes

Table 5: Standardized Distributional Effects (SDE)

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
Firearm homicide rate	-0.0797	0.0552	18.31	0.0044	0.0030	Null
Non-firearm homicide (placebo)	-0.0229	0.0677	10.07	0.0023	0.0067	Null
DDD (firearm − non-firearm)	-0.0568	0.0754	15.02	0.0038	0.0050	Null
Panel B: Heterogeneous						
Large municipalities (>67th pctile)	-0.1157	0.3198	19.65	0.0059	0.0163	Small
Population-weighted	0.3129	0.1908	18.31	0.0171	0.0104	Small

Notes: SDE = $\hat{\beta}/\text{SD}(Y)$. Classification: Large ($|\text{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005). Classification refers to magnitude, not statistical significance. **Country:** Brazil. **Research question:** Does firearm liberalization increase homicide, and does restriction reverse it? **Policy mechanism:** Legal firearms supply expansion via shooting club proliferation and registration relaxation. **Outcome definition:** Firearm homicide rate per 100,000 population (ICD-10 X93–X95). **Treatment:** Pre-2019 shooting club density (clubs per 100K), interacted with post-2019 national policy indicator. **Data:** DATASUS Mortality Information System (SIM), IBGE population estimates, Receita Federal CNPJ (2013–2023). **Method:** Shift-share difference-in-differences with triple-difference falsification. **Sample:** 5,570 municipalities × 11 years = 61,270 municipality-year observations (firearm); 122,540 for DDD.