social sciences

MDPI

Article

Aggregate Growth, Divergent Local Trajectories: Recorded Firearm Dealership Counts in 307 U.S. Counties, 2010–2019

Keshav Bhattarai

Department of Physical Sciences, University of Central Missouri, Warrensburg, MO 64093, USA; bhattarai@ucmo.edu

Abstract

This study examines changes in recorded firearm dealership counts and their demographic and commercial correlates in a sample of 307 U.S. counties from 2010 to 2019. The descriptive panel includes 3070 county–year observations. Across the sampled counties, the total number of recorded dealerships increased from 6995 to 7303 (4.40%), while the county median declined from 17 to 16. Dealership counts increased in 142 counties, decreased in 134, and remained unchanged in 31, demonstrating that aggregate growth can coexist with local declines. Poisson pseudo-maximum-likelihood (PPML) models with county- and year-fixed effects include 306 counties and 3060 observations after excluding one county with zero recorded dealerships throughout the study period. Population and establishment counts are positively associated with dealership counts when modeled separately. When included together, their respective coefficients are 0.499 (95% confidence interval: −0.027 to 1.024) and 0.117 (95% confidence interval: −0.266 to 0.500). Both intervals include zero, leaving their distinct conditional associations uncertain. These findings apply to the selected panel and do not estimate national or border-region trends. Counties were included based on data availability in the reconstructed panel and agreement between internal identifiers. Residual spatial patterning also means that the reported cluster-adjusted confidence intervals depend on the assumed structure of dependence among observations. Finally, the study measures recorded dealership counts rather than verified outlet activity, transactions, trafficking, or policy effects.

Keywords: firearm dealerships; county panel; population; establishments; Poisson pseudo-maximum likelihood; United States

1. Introduction

Regional growth in firearm dealership counts can mask contractions in many of the communities that contribute to that total. Across the selected counties examined here (Figure 1), recorded counts increased by 4.40% during 2010–2019, even as 134 counties lost dealerships and the median county count fell. Understanding this divergence requires examining the distribution of local change alongside the aggregate trajectory. These counts indicate recorded institutional presence; they do not, by themselves, measure transactions, travel access, or firearm movement.

This distinction is important for research on the geography of firearm retail institutions because a rising regional total does not establish that retail presence expanded in a typical county or that all communities followed the same trajectory. Existing studies examine related but distinct dimensions of dealer geography. Pear et al. (2023) investigate dealer density in relation to community firearm outcomes, while Stansfield et al. (2023) examine dealer

Check for updates

Academic Editor: Tina Freiburger

Received: 22 October 2026
Revised: 20 November 2026
Accepted: 22 November 2026
Published: 25 November 2026

Copyright: © 2026 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

compliance and local homicide. Transaction-based studies connect purchaser, firearm, and retailer characteristics to subsequent recovery (Robinson et al. 2024; Laqueur and Smirniotis 2025). Together, these studies highlight the importance of examining where firearm dealers are located, although they focus on outcomes other than the changes in county-level dealership counts examined in this study.

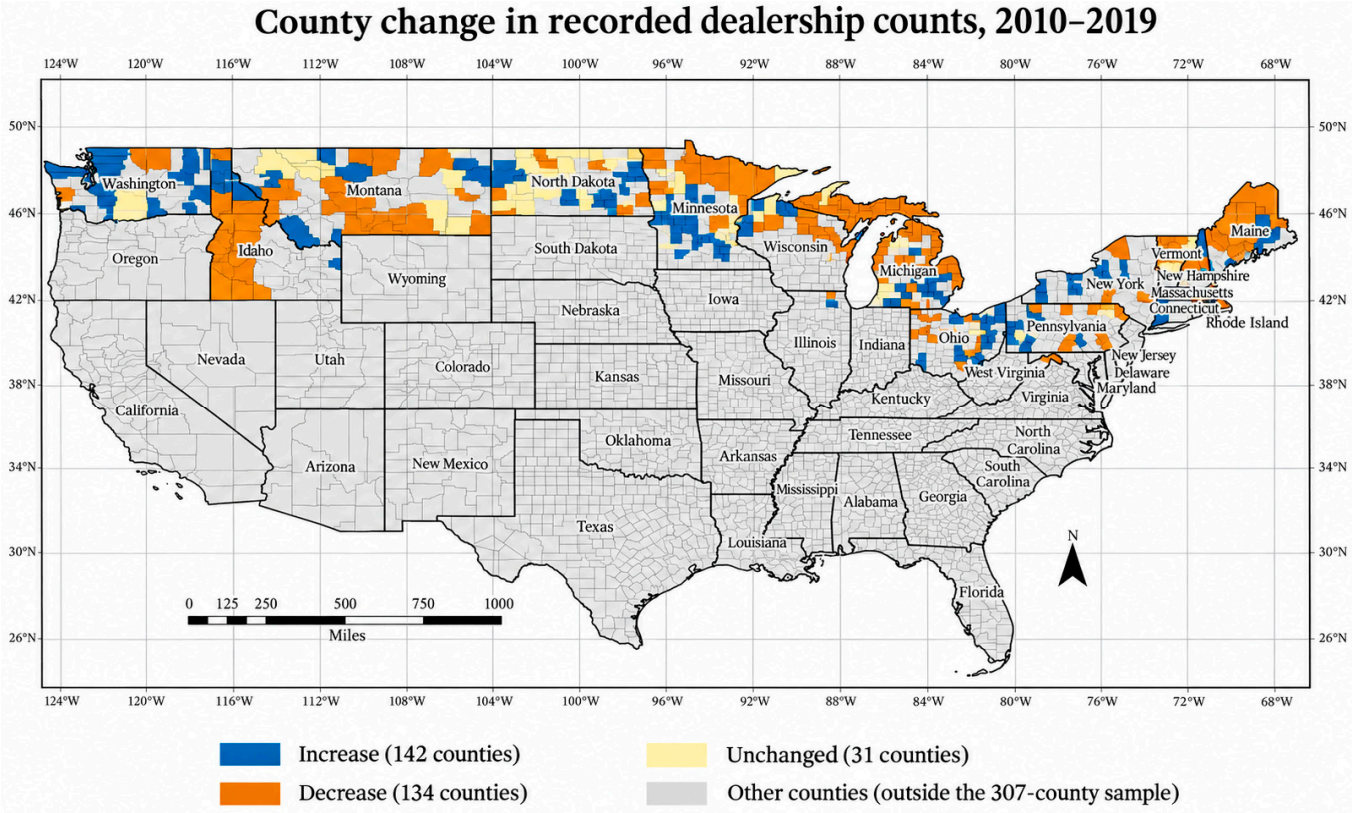

Figure 1. County-level changes in recorded dealership counts from 2010 to 2019. Blue indicates an increase (142 counties), orange a decrease (134), and tan no change (31); gray counties lie outside the 307-county descriptive sample. All 307 sample identifiers matched a single 2010 census county polygon. The join uses five-digit Federal Information Processing Standards (FIPS) county codes; supplied point coordinates are not used for placement. Categories show the sign of endpoint change, not its magnitude, per-capita change, spatial significance, or causal effects. Small counties may be difficult to distinguish at this scale. Source: County–year panel and U.S. Census Bureau (2010), 1:20,000,000 cartographic boundaries. Projection: contiguous U.S. Albers equal-area projection.

Building on this motivation, the present study treats recorded dealership counts as an outcome and provides a longitudinal account of how those counts changed in the selected counties. It examines annual totals, county distributions, and gains, losses, and unchanged counts between 2010 and 2019, then evaluates population and business establishments as demographic and commercial correlates. Its empirical contribution is to bring these comparisons together within a consistent panel. The study does not claim that county-level change has never been studied or that the selected counties represent the United States.

Population and establishments capture related dimensions of local scale. Examining them separately estimates their respective adjusted associations with dealership counts, while including both assesses whether each retains a distinct association conditional on the other. The principal models use county and year effects to examine within-county change while accounting for stable county characteristics and common annual influences. Neither these adjustments nor positive coefficients in the separate models establish a causal mechanism.

The findings show modest aggregate growth alongside substantial variation in local trajectories: 142 counties gained dealerships, 134 lost dealerships, and 31 retained the same number of dealerships. Population and establishments each have positive associations with dealership counts in separate models, but the intervals for both coefficients include zero when the two are modeled jointly. The distribution of local change and the uncertainty surrounding these mutually adjusted associations are therefore central to the study's contribution. Figure 1 presents the county-level geography of endpoint change.

Two research questions organize the main analysis. RQ1 asks how aggregate totals, county distributions, and county-level endpoint changes changed from 2010 to 2019. RQ2 asks whether within-county changes in dealership counts accompany changes in population and establishments, examined separately and jointly with county and year effects. Objectives 1 and 2 correspond to these respective questions. Appendix C presents the exploratory geographic question and all Canadian model results as a separate data-linkage exercise.

The study thus contributes a sample-specific account of domestic dealership trajectories and their demographic and commercial correlates. Section 2 defines the available county sample and conceptual framework, and Section 3 explains panel construction. Sections 4 and 5 present the methods and results, while Sections 6–8 address interpretation, limitations, and conclusions. Appendices A and B provide supporting sample and audit details, and Appendix C documents the inherited Canadian linkage exercise.

2. Study Area and Conceptual Framework

2.1. Study Area Selection

The study examines aggregate growth alongside divergent county trajectories using a nonrandom convenience sample derived from a reconstructed county–year panel. The panel initially contains 412 county identifiers. However, mismatches arose in the candidate list, inclusion decisions, and any thresholds for proximity or data availability. Thus, this study does not establish that the counties were originally selected on theoretical grounds.

Sample construction proceeds through internal identifier screening and verification of annual record completeness for 2010–2019. A county is classified as identifier-consistent (Identifier_consistent = TRUE) when its original five-digit county code matches the first five digits of its recorded 11-digit geographic identifier (GEOID). Of the 412 initial identifiers, 105 fail this check and are excluded. The remaining 307 counties all have complete records for the ten-year period, yielding a final descriptive sample of 3070 county–year observations. The final sample retains 74.5% of the initial identifiers, a measure of internal screening retention rather than national representativeness or border coverage. Appendix A lists the retained counties, and Appendix B summarizes the screening counts.

The final sample spans 18 states: Connecticut, Idaho, Illinois, Iowa, Maine, Maryland, Massachusetts, Michigan, Minnesota, Montana, New Hampshire, New York, North Dakota, Ohio, Pennsylvania, Vermont, Washington, and Wisconsin. Its geographic scope includes inland counties and extends beyond both counties adjoining Canada and states sharing an international boundary with Canada. Examples include Fairfield County, Connecticut (09001), and McHenry County, Illinois (17111). Likewise, the inclusion of counties in Great Lakes states does not establish a systematic selection rule based on water boundaries. Alaska, Hawaii, and Puerto Rico are absent from the reconstructed panel because of their geographic separation.

This geographic distribution does not establish a defined border-based sampling design. Such a design would require explicit criteria distinguishing counties that directly adjoin the international boundary, counties located nearby, and counties connected through transport corridors. Direct adjacency would require a boundary-intersection rule specifying how to treat water boundaries; proximity would require a distance threshold and a reference

geometry; and corridor membership would require a defined transportation or border-crossing connection. These criteria do not apply to Alaska, Hawaii, and Puerto Rico.

Without an independently defined universe of northern border counties, border coverage cannot be verified or calculated. Coverage would equal the number of retained counties belonging to that universe divided by the total number of counties in that universe. The 74.5% identifier-retention rate therefore remains an internal screening statistic. Similarly, the supplementary geographic sample retains 285 of the 307 counties (92.8%) and excludes 22 (7.2%) through linkage screening; these percentages describe linkage retention and exclusion, not border coverage.

The screening procedure makes the final sample reproducible within the reconstructed panel. Consequently, the findings describe the counties retained and do not support population-wide inference.

2.2. Conceptual Framework

The framework examines how changes in population and business establishments are associated with changes in recorded firearm dealership counts within counties. Population and establishments represent distinct demographic and commercial dimensions of local scale. Separate and joint models therefore address complementary descriptive questions: how dealership counts relate to each dimension individually and how each association changes when the other dimension is held constant. The contribution concerns county trajectories and domestic correlates within the available sample.

The outcome, D_{it}, is the integer recorded in the supplied annual county field "dorp," a firearm dealership count. It is interpreted as a recorded count of dealerships rather than a verified measure of operating retail outlets, sales, capacity, or unique businesses (Adcock and Collier 2001). The available data do not provide license-level identifiers, information on included license types, extraction dates, or establishment classifications. Consequently, the series supports analysis of recorded counts, but its underlying business and licensing composition cannot be established.

This measurement limitation matters because the Bureau of Alcohol, Tobacco, Firearms and Explosives (ATF) distinguishes among Type 01 dealers, which include gunsmiths; Type 02 pawnbrokers; and Type 03 collectors. Collector status does not authorize the holder to engage in the business of dealing in firearms (Bureau of Alcohol, Tobacco, Firearms and Explosives n.d.). Without the underlying records, these categories cannot be assigned retrospectively to the aggregate "dorp" series. Moreover, a license count does not necessarily represent distinct premises or firms, and a valid license neither measures transactions nor establishes continuous retail activity.

The explanatory variables capture two dimensions of local scale. Domestic population, P_{it}, measures the number of residents in the county where dealerships are counted; population growth may support additional retail institutions. Business establishments, E_{it}, recorded as "est," represent the local commercial base, whose expansion may also accompany dealership growth. E_{it} is interpreted as a correlate of commercial scale.

Models M1 and M2 examine population and establishments separately, while M3 includes both variables to estimate their mutually adjusted associations. All three models include county- and year-fixed effects, focusing the domestic analysis on within-county change while accounting for time-invariant county characteristics and shared annual influences. Because population and commercial activity can change together, their conditional associations may be difficult to distinguish precisely. Comparing the separate and joint models helps assess this overlap, while the uncertainty around the estimates indicates their precision. An auxiliary collinearity statistic alone cannot explain the joint-model findings (O'Brien 2007).

M1 and M2 are retained because population and establishments represent distinct theoretical dimensions of local scale. However, each separate-model coefficient may also reflect shared variation with the omitted correlate and therefore cannot be interpreted as an independent effect. M3 is preferable for estimating mutually adjusted associations when conditioning on both variables is substantively appropriate. Model choice thus follows the descriptive question being addressed rather than whether particular p-values fall below 0.05.

3. Data Preparation and Analytical Sample

Source Panel and Variable Definitions

The county–year panel uses dealership data from Johnson and Robinson (2024). Section OA1 of Johnson and Robinson's (2024) online appendix documents the upstream federal firearms license (FFL) records from the U.S. Bureau of Alcohol, Tobacco, Firearms and Explosives (ATF). The authors submitted a Freedom of Information Act request in May 2019 and received annual license lists for 2003–2013, generally starting in January, that included license-holder names, addresses, and, where applicable, business names. They obtained the 2014–2018 lists from the ATF website. Both sets of lists excluded Type 03 collector licenses. The license number encodes geographic and license-type information that can support county aggregation.

The metadata documents the ATF inputs, but does not provide a record-level crosswalk to the annual county field "dorp." Within the main study period, they document how the upstream lists were obtained for 2010–2018. They do not identify the source extracts for 2019 or 2020. The 2020 data appear only in a domestic sensitivity analysis. The exact annual snapshot dates, the included license types, and the rules for counting multiple licenses at a single address remain unverified for "dorp." The Type 03 exclusion described for the source lists therefore cannot on its own establish the composition of the supplied series. Throughout the analysis, "recorded dealership counts" refer to the values in that series.

The original dataset reports 33,892 rows across 3090 counties after excluding missing dealership values. The reconstructed county–year panel contains 4532 observations for 412 county identifiers over 2010–2020. Restricting this panel to 2010–2019 yields 4120 observations for 412 identifiers. Requiring Identifier_consistent = TRUE excludes 105 discordant identifiers and their 1050 observations, leaving 307 counties and 3070 observations. Each retained county has ten annual observations, positive U.S. population and establishment values, and nonnegative integer dealership counts, with no duplicate county–year keys or blank required domestic fields. Canadian linkage eligibility does not determine inclusion in the domestic analysis; Appendix C documents the separate linkage screening.

Recorded zero dealership counts are retained in descriptive analyses. Billings County, North Dakota (38007), has zero recorded dealerships in every study year and is therefore excluded from county fixed-effects Poisson pseudo-maximum likelihood (PPML) estimation, leaving a principal domestic estimation sample of 306 concordant counties and 3060 observations. Billings County remains in descriptive statistics, ordinary least squares (OLS) analyses, and state fixed-effects models where the required data are available. The supplements document the reconstructed and linkage sensitivity samples. Appendix B reports internal consistency checks and model refits documented in the inherited analytical record.

4. Statistical Methods

4.1. Poisson Pseudo-Maximum Likelihood Estimation

The principal multiplicative conditional-mean specification is given in Equation (1) (Santos Silva and Tenreyro 2006; Correia et al. 2020).

$$E[D_{it} \mid X_i, \alpha_i, \tau_t] = exp[\alpha_i + \tau_t + \beta_P \, ln(P_{it}) + \beta_E \, ln(E_{it})] \quad (1)$$

Here, D_{it} is the dealership count in county i in year t; α_i and τ_t are county- and year-fixed effects; T_i is the set of observed years for county i; and X_i contains its covariate histories over those years. The conditional-mean assumption requires outcome disturbances to be uncorrelated with past, present, and future included covariates, conditional on the fixed effects. This is a maintained identifying assumption rather than an established property of the data.

The coefficients β_p and β_E are elasticities of the conditional expectation of the dealership count. A 1% increase in a covariate corresponds approximately to a β% change in the expected count, holding other model terms constant. Population and establishments must be positive because they are logged; dealership counts may equal zero. ln denotes the natural logarithm, and exp denotes the exponential function.

M_1 includes log population alone, M_2 includes log establishments alone, and M_3 includes both. The principal concordant-sample PPML models use 3060 observations from 306 counties after exclusion of the all-zero county. County effects capture stable county characteristics, while year effects capture influences shared across counties within each year.

PPML accommodates zero dependent-variable observations and does not require a log transformation of D_{it} (Santos Silva and Tenreyro 2006; Correia et al. 2020). A model would exclude any row that is missing a required value. It would also exclude rows with nonpositive population or establishment counts because those covariates are logged. No such rows occur in the attached domestic sample for the main study period. Only the ten rows for county 38007 are omitted from the county fixed-effects PPML analysis because that county's outcome is zero every year. The other 17 zero-outcome rows are retained. Neither overdispersion nor an individual zero dealership count is itself an exclusion criterion.

PPML relies on the specified conditional expectation rather than a full Poisson outcome distribution. As Anderson et al. (2018, pp. 2754–55) explain, additive and multiplicative disturbances can represent the same conditional mean. This feature supports modeling expected dealership counts while leaving the exogeneity assumption and the choice of covariates to substantive assessment.

Population enters with an estimated coefficient rather than as an offset. An offset $\ln(P_{it})$ would impose a population elasticity of one; the count model allows this association to differ from proportionality.

4.2. County Totals and Profiled County Fixed Effects

$$D_{i+} = \sum_{t \in T} D_{it} \tag{2}$$

Equation (2) defines D_{i+}, the sum of annual recorded dealership counts over the set T_i of years observed for county i. It is measured in dealership years and does not count unique businesses over the decade. Equation (3) writes the same multiplicative mean using a vector of covariates.

$$\mu_{it} = exp(\alpha_i + x_{it}^T \beta) \tag{3}$$

In Equation (3), x_{it} contains the included covariates and year indicators; the superscript T denotes the transpose of a vector. β is their coefficient vector. Equation (4) follows by solving the county-intercept score condition for a county with a positive outcome total.

For a county with a finite intercept estimate, the intercept score sets the sum over observed years of $D_{it} - \hat{\mu}_{it}$ to zero. Substituting the mean in Equation (3) yields the fitted

county multiplier $\exp(\hat{\alpha}_i) = D_{i+}/\Sigma_{s \in Ti}\, exp(x_{is}^T \hat{\beta})$. Multiplying this ratio by the fitted index term for year t gives Equation (4).

$$\hat{\mu}_{it} = D_{i+} \frac{exp\left(x_{it}^T \hat{\beta}\right)}{\sum_{s \in T} exp\left(x_{is}^T \hat{\beta}\right)} \quad (4)$$

Equation (4) makes the fitted county total equal the observed total. The denominator sums over observed years s in T_i. If $D_{i+} = 0$, the county intercept approaches negative infinity and the fitted means approach zero; after profiling, this group contributes no information for the common slopes. This separation property, rather than a prohibition on zero outcomes in PPML, explains the exclusion of county 38007 (Correia et al. 2020).

The profiling in Equations (2)–(4) is an algebraic consequence of the PPML county-intercept first-order conditions. Anderson et al. (2018) provide related PPML fixed-effects aggregate-matching results; these results provide methodological context rather than evidence about firearm dealerships. However, matching fitted and observed county totals does not independently validate the counts or geography.

4.3. Clustered Inference and Finite-Sample Adjustment

Reported one-way cluster covariance matrices use the finite-sample multiplier in Equation (5), and the corresponding confidence intervals use a t distribution with G − 1 degrees of freedom (Cameron and Miller 2015). N is the number of fitted observations, G is the number of clusters, and K is the rank of the full estimable design, including fixed effects. This convention reproduces the supplied principal estimates and intervals. It does not guarantee adequate inference with 18 state clusters.

$$q = [G/(G-1)] \times [(N-1)/(N-K)] \quad (5)$$

For county-effects models with I counties, T years, and p slopes, Equation (6) gives the full-rank parameter count. County and year indicators include one intercept and omit redundant columns. M1 and M2 have K = 316 and M3 has K = 317. Appendix A records counts for all retained sensitivity specifications.

$$K = I + T - 1 + p \quad (6)$$

The domestic models primarily use county clustering, which accommodates temporal dependence and county-level heteroskedasticity. It assumes sufficiently weak dependence across counties. State-clustered results provide a sensitivity comparison that allows dependence within states; only 18 state clusters are available. Appendix B reports residual spatial diagnostics. Section 7 discusses these inferential limitations. The analysis uses locality-clustered inference only for the linkage exercise in Appendix C (Cameron and Miller 2015).

No bootstrap for a small number of state clusters or a spatial heteroskedasticity- and autocorrelation-consistent covariance estimator has been implemented. A suitable Poisson estimating-equation bootstrap would require a justified treatment of nuisance fixed effects and state-level scores; the methodological discussion in Roodman et al. (2019) does not itself supply an implemented correction here.

4.4. Quadratic Functional-Form Check

The functional-form check adds squared mean-centered log-population and log-establishment terms to M3. Define $z_{p,it} = \ln(P_{it}) - \text{mean}[\ln(P)]$ and $z_{e,it} = \ln(E_{it}) - \text{mean}[\ln(E)]$, with means taken over the estimation sample.

The augmented linear predictor adds the two centered squared terms in Equation (7). The test uses the augmented model's county-clustered covariance matrix to evaluate the

joint null hypothesis that both added coefficients are zero. This tests only these two terms; it does not validate the conditional mean or independence assumptions.

$$\text{Additional index terms} = \eta_p \, z_{p,it}^2 + \eta_e \, z_{e,it}^2 \quad (7)$$

In Equation (7), η_p and η_e are the coefficients on the squared mean-centered log-population and log-establishment terms, respectively. The means in the definitions of $z_{p,it}$ and $z_{e,it}$ are calculated across the retained county–year estimation observations. The null is $\eta_p = \eta_e = 0$.

5. Results

5.1. Dealership Trends and Descriptive Context

Across the 307 concordant counties, the number of recorded dealerships increased from 6995 in 2010 to 7303 in 2019, a 4.40% increase (Table 1). The mean rose from 22.79 to 23.79, while the median fell from 17 to 16; the annual total peaked at 7513 in 2017. These results address RQ1 and Objective 1 by showing aggregate growth alongside a declining median. They do not establish uniform county growth or greater concentration.

Table 1. Annual recorded dealership counts in the concordant county sample.

Year	Counties	Total	Mean	Median	Zeros
2010	307	6995	22.79	17	4
2011	307	6911	22.51	16	4
2012	307	6914	22.52	16	3
2013	307	7118	23.19	16	2
2014	307	7417	24.16	17	2
2015	307	7480	24.36	17	1
2016	307	7463	24.31	17	2
2017	307	7513	24.47	16	2
2018	307	7439	24.23	17	3
2019	307	7303	23.79	16	4

Endpoint changes reveal the local variation behind the total: counts rose in 142 counties, fell in 134, and were unchanged in 31. The median change was zero, with a 25th percentile of −2 dealerships, a 75th percentile of 3 dealerships, and a range of −14 to 40 dealerships. Quartiles of county dealership counts shifted from 9, 17, and 28 in 2010 to 9, 16, and 30 in 2019. These summaries include all 307 concordant counties, including counties with zero dealerships, and describe heterogeneous change without identifying its mechanisms.

The mapped categories give a spatial description of the same endpoint comparison as the text. Counties with increases gained a total of 779 recorded dealerships, while counties with decreases lost 471, for a net gain of 308. The map distinguishes counties outside the sample from counties with no change.

Table 2 summarizes 3070 county–years from 307 counties before the all-zero county was excluded from the county-effects PPML. The within-county standard deviation (Within SD) is the root-mean-square deviation from each county's mean, with N in the denominator. Canadian input and distance summaries have been moved to Appendix C.

Table 2. Variable summaries by analytical sample.

Sample and Variable	Mean	SD	Min	Median	Max	Within SD
Domestic dealership count	23.63	23.20	0.00	17.00	186.00	2.80
Domestic U.S. population	89,069.78	173,656.98	475.00	36,971.50	1,280,122.00	3499.54
Domestic establishments	2194.81	4595.21	8.00	821.00	39,250.00	97.80

In Table 2, SD denotes standard deviation; Min and Max denote minimum and maximum. In the regression tables, CI stands for confidence interval.

5.2. Domestic Population and Establishments

Table 3 addresses RQ2 and Objective 2 by relating changes in dealership counts within counties to changes in population and establishments, with county- and year-fixed effects. In separate models, both estimated associations are positive: the population coefficient is 0.606 (95% CI [0.230, 0.983]; $p = 0.002$) in M1, and the establishment coefficient is 0.409 (95% CI [0.126, 0.691]; $p = 0.005$) in M2. Mutual adjustment in M3 yields coefficients of 0.499 (95% CI [−0.027, 1.024]; $p = 0.063$) and 0.117 (95% CI [−0.266, 0.500]; $p = 0.548$), respectively. Both joint-model intervals include zero, leaving each conditional association uncertain. These within-county elasticities describe expected recorded counts and do not identify causal effects.

Table 3. Principal county- and year-fixed-effects PPML models with county-clustered inference.

Model	Term	Estimate	95% Confidence Interval	*p*-Value
M1	ln(U.S. population)	0.606	[0.230, 0.983]	0.002
M2	ln(establishments)	0.409	[0.126, 0.691]	0.005
M3	ln(U.S. population)	0.499	[−0.027, 1.024]	0.063
M3	ln(establishments)	0.117	[−0.266, 0.500]	0.548

To illustrate the scale of these associations, Equation (8) converts a coefficient b to the model-implied percentage change in the expected count for a hypothetical 10% increase in the corresponding covariate. For the M3 population coefficient, this gives 4.87%, with a transformed coefficient interval of approximately −0.26% to 10.25%. For establishments, it gives 1.12%, with an interval of approximately −2.50% to 4.88%. These are conditional model illustrations, not observed changes or causal forecasts.

$$\text{Model-implied percentage change} = 100 \times [(1.10)^{b} - 1] \quad (8)$$

In the 306-county M3 sample, the median 2010–2019 population change is −0.59% (interquartile range: −3.95% to 3.37%); the median establishment change is −1.28% (−5.38% to 2.96%). Median annual changes are −0.07% and −0.08%, respectively. Thus, the hypothetical 10% increases exceed the upper quartile of observed changes over the decade for both variables.

All models include county- and year-fixed effects and use 3060 observations from 306 counties. Inference uses 306 county clusters and confidence intervals based on Student's t distribution with 305 degrees of freedom. The full-rank convention implies K = 316 for M1/M2 and K = 317 for M3.

5.3. Collinearity Diagnostics

Table 4 reports auxiliary variance inflation factors (VIFs) after removing county and year means from the two log covariates in the 306-county M3 estimation sample. Their adjusted correlation is 0.671, and each VIF is 1.818. These values indicate overlapping variation but do not indicate severe collinearity or explain the change in precision when both covariates enter the Poisson model (O'Brien 2007).

Table 4. Collinearity diagnostics for the principal model sample.

Sample and Transformation	Covariate	VIF
M3 adjusted for county and year	ln(U.S. population)	1.818
M3 adjusted for county and year	ln(establishments)	1.818

Equation (9) calculates the auxiliary two-variable VIF using the county- and year-adjusted correlation ρ. It is a linear-design diagnostic, not a diagnostic of the full nonlinear PPML information matrix. Precision also depends on outcome information, weighting, specification, and the variation used to identify the coefficients (O'Brien 2007).

The separate models are theoretically useful as complementary descriptions rather than replacements for M3. Adding a correlated covariate changes both the estimand and its uncertainty; the joint-model intervals alone do not identify collinearity as the cause. Principal components would replace population and establishments with a composite; ridge regularization would introduce a penalty and require different choices for estimation and inference. Neither is required by a VIF of 1.818, and neither has been estimated here. Retaining M1–M3 preserves interpretable covariates and makes the conditional uncertainty visible rather than removing a variable to obtain significance.

$$\mathrm{VIF} = 1/(1-\rho^2) \tag{9}$$

5.4. Domestic Sensitivity Analyses

Table 5 retains domestic sensitivity specifications for covariance estimation, study period, source reconstruction, and outcome transformation. The table reports coefficients and 95% cluster-robust confidence intervals. All models include county- and year-fixed effects. Canadian-linked sample restrictions and Canadian terms are reported exclusively in Appendix C.

Table 5. Domestic sensitivity analyses with county- and year-fixed effects.

Specification/Clustering	N/Counties	U.S. Population ln(P)	Establishments ln(E)
1. Baseline PPML/county	3060 306	0.499 [−0.027, 1.024]	0.117 [−0.266, 0.500]
2. Baseline PPML/state	3060 306	0.499 [−0.141, 1.138]	0.117 [−0.594, 0.828]
3. Include 2020/county	3366 306	0.558 [0.049, 1.068]	0.026 [−0.344, 0.396]
4. OLS: ln(1 + D)/county	3070 307	0.568 [0.129, 1.007]	0.044 [−0.316, 0.403]
5. All reconstructed/county	4110 411	0.521 [0.101, 0.942]	0.132 [−0.177, 0.441]

5.4.1. Sample Reconstruction and Transformed Outcome

The sensitivity analysis using ordinary least squares (OLS) with a log-transformed outcome retains 307 counties and 3070 observations, with population and establishment coefficients of 0.5682 (standard error = 0.2230; $p = 0.0113$) and 0.0438 (standard error

= 0.1827; $p = 0.8106$). The PPML sensitivity analysis with county fixed effects and all reconstructed identifiers retains 411 counties and 4110 observations after excluding the all-zero county. Inclusion of identifier-discordant records tests sensitivity to reconstruction; it does not validate those records.

Equation (10) defines the supplementary OLS model. The intercepts a_i and year effects τ_t adjust the conditional mean of $\ln(1 + D_{it})$; θ_p and θ_e are its domestic slopes, and ε_{it} is the additive error. These coefficients are not PPML elasticities of expected dealership counts. The reported within R^2 is 0.039. The main period is 2010–2019; a separate 2010–2020 model uses 306 counties and 3366 observations.

$$\ln(1 + D_{it}) = a_i + \tau_t + \theta_p \ln(P_{it}) + \theta_e \ln(E_{it}) + \varepsilon_{it} \quad (10)$$

5.4.2. Domestic Sensitivity Estimates

Population point estimates remain positive across the reported specifications, but their uncertainty changes with clustering, period, and sample. In M3, the population p-value is 0.063 with county clustering and 0.118 with state clustering; the establishment p-value under state clustering is 0.733. Extending the panel through 2020 yields population and establishment coefficients of 0.558 ($p = 0.032$) and 0.026 ($p = 0.891$). Including all reconstructed identifiers yields coefficients of 0.521 ($p = 0.015$) and 0.132 ($p = 0.401$).

Using the same M3 sample, state clustering leaves both point estimates unchanged and widens both 95% intervals. The population interval widens from $[-0.027, 1.024]$ to $[-0.141, 1.138]$, and the establishment interval widens from $[-0.266, 0.500]$ to $[-0.594, 0.828]$ (Table 5). State clustering allows counties in the same state to share shocks, but geographic dependence may cross state lines. In Appendix B, the M3 residuals show positive spatial association under the graph based on the four nearest neighbors in seven of ten years at an unadjusted significance level of 0.05. The annual diagnostics and the two clustering schemes show that the reported precision depends on the covariance assumptions. They do not provide spatially robust intervals for M1–M3. The separate-model associations are therefore suggestive, and the evidence for either mutually adjusted association remains inconclusive.

5.5. Functional Form and Numerical Diagnostics

Profiling the county intercepts reproduced the M3 coefficients of 0.498643 and 0.116958, with a maximum absolute full-design score component of 2.6×10^{-7}. The audit also reports that M1 and M2 were reproduced to the precision shown. It independently verifies sample flow, endpoint totals, zero entries, and the adjusted covariate correlation; those checks are distinct from regression refitting. Together, these results demonstrate numerical agreement. They do not establish substantive validity or the validity of statistical inference.

The inherited quadratic-term test provides an additional, limited assessment of model specification. It reports a joint Wald statistic of 0.513 with two restrictions and a p-value of 0.774 (Appendix A), indicating that the test did not detect the additional curvature specified in Equation (7). This result, however, does not validate the full conditional mean.

Assessing the reliability of statistical inference therefore also requires examining the assumptions about covariance. County-clustered covariance estimation accommodates heteroskedasticity and temporal dependence within counties without requiring equality of the conditional mean and variance. Its validity nevertheless depends on sufficient independence across counties and adequate information across clusters. The state-clustered sensitivity calculation allows dependence across counties within the same state, but inference remains sensitive to the limited number of state clusters (18).

6. Discussion

6.1. Interpretation of the Research Questions

The central finding is a divergence between aggregate and local change in recorded firearm retail presence. Across the 307 selected counties, the total number of dealerships recorded rose by 4.40%, from 6995 in 2010 to 7303 in 2019. Yet increases in 142 counties were nearly balanced by decreases in 134, while counts remained unchanged in 31 counties (Figure 1). The county median fell from 17 to 16, and the median endpoint change was zero. RQ1 therefore reveals local contraction and stability within a modestly expanding total—variation that the aggregate growth rate alone conceals.

This distinction matters because retail institutions serve and operate within communities. A rising total does not necessarily imply growth in the typical county or uniform expansion across the region. County change is therefore a useful outcome alongside dealer density and downstream firearm outcomes. The panel identifies where recorded presence changed and supports questions about accompanying local characteristics, although endpoint counts cannot distinguish individual openings, closures, relocations, or changes in concentration.

RQ2 places these trajectories within the counties' changing demographic and commercial context. In separate models, the population elasticity is 0.606 (95% CI [0.230, 0.983]), and the establishment elasticity is 0.409 (95% CI [0.126, 0.691]). These positive associations support considering both dimensions of local scale, without identifying the mechanisms through which retail institutions change.

Mutual adjustment provides a more qualified interpretation of the associations. In M3, the population coefficient is 0.499 (95% CI [−0.027, 1.024]), and the establishment coefficient is 0.117 (95% CI [−0.266, 0.500]). Both confidence intervals include zero, indicating uncertainty about the direction and magnitude of each association after adjustment for the other variable. The confidence interval for establishments widens from M2 to M3, consistent with reduced precision in the conditional estimate. These results warrant caution in drawing conclusions about the association of each variable with dealership counts after adjustment for the other.

These results complement studies that use dealer geography to explain firearm outcomes. Pear et al. (2023) examine dealer density in community-level injury research, while Stansfield et al. (2023) connect licensed dealers and compliance to local homicide. Here, recorded county counts are the outcome and are examined in relation to demographic and commercial change. Studies of transactions and recoveries connect retail records to later firearm recoveries (Robinson et al. 2024; Laqueur and Smirniotis 2025). The present contribution thus concerns the changing institutional geography that can precede a more direct study of acquisition and movement.

The sensitivity results qualify the precision of the domestic findings. Population point estimates remain positive across the reported main and sensitivity specifications, but their uncertainty varies with sample period and clustering. State clustering allows a different dependence structure, whereas extending the study period through 2020 changes the observations and historical conditions represented. With only 18 state clusters, conventional inference warrants particular care (Cameron and Miller 2015). Substantive interpretation should therefore follow the estimates, intervals, and design rather than a single p-value threshold (Wasserstein and Lazar 2016).

6.2. Policy and Research Implications

For monitoring and research, the immediate implication is to examine county trajectories alongside aggregate counts. Regional totals alone obscure the increases, decreases, and stability observed here. Considering population and establishments together provides

a fuller context for comparing counties, although uncertainty in the joint model warrants cautions about treating either factor alone as a sufficient explanation for differences among counties. The results describe retail institutions and do not estimate the effects of licensing, enforcement, or zoning.

The most useful next research step is to verify the recorded dealership unit and sample construction, then link validated retail histories to more direct measures of activity and accessibility. Purchaser residence, transaction volume, and routes through authorized crossings would help distinguish local commercial context from actual cross-border interaction. Linked sales and recovery records would address firearm movement on its own evidentiary basis, with attention to tracing coverage and selection (Bureau of Alcohol, Tobacco, Firearms and Explosives 2026a, 2026b; Robinson et al. 2024). Policy, theft, and ownership attitudes become testable mechanisms when corresponding measures are available.

6.3. Academic Contributions

The main academic contribution is a longitudinal account of how aggregate dealership growth can coexist with divergent county trajectories. The total count increased by 4.40%, while nearly equal numbers of counties gained and lost dealerships and the median declined. These findings show the value of examining the distribution of change alongside totals. This contribution is specific to the selected sample rather than an estimate of national change.

A second contribution is the comparison of demographic and commercial correlates within one panel. Separate-model associations are positive, while both mutually adjusted intervals include zero. This comparison highlights overlap between the two measures of local scale and the limits of attributing dealership change to either measure alone. By treating institutional presence as the outcome, the study complements research on dealer density and downstream firearm outcomes. The supplementary linkage material is included for transparency and methodological scrutiny.

7. Limitations

Section 2.1 states the reproducible inclusion criteria and state coverage and identifies what remains unknown about the original selection process. Internal identifier concordance cannot establish coordinate accuracy, original-source validity, national representativeness, or coverage across all border counties. The available data also lack license-level records, extraction dates, information on which license types were included, and a complete variable provenance record. Consequently, recorded dealership counts cannot be equated with verified active retail premises.

Identifier concordance, unique county–year keys, and integer-count checks establish internal consistency only. An annual difference may reflect changes in licensed establishments, license composition, or the compilation of the series. The documented change from FOIA lists to website lists in 2014 warrants particular attention, as does the unverified source for 2019, the endpoint used to calculate the reported 4.40% increase. The direction or magnitude of any measurement error cannot be determined from the aggregate file. Year effects absorb common annual shifts on the model's log-mean scale, while county-specific changes in coverage may still affect the slopes. A source audit would require the dated ATF extracts, explicit license-type and counting rules, and reconciliation to the county–year values, including a documented extract for 2019 and, separately, 2020.

Moving the Canadian exercise to the supplement limits the main claims but does not resolve its candidate-selection, assignment-validity, or population-provenance problems. The proposed multi-locality index in Section S.3 remains unestimated; validating that extension requires original geographic identifiers, harmonized annual populations, and crossing-route inputs.

Thus, the geographic-matching concern is only partially addressed: flagged observations are screened and the remaining exercise is explicitly exploratory, but the flagged matches have yet to be verified or corrected against authoritative sources, and the models have yet to be refitted with validated matches. Likewise, residual diagnostics and alternative clustering expose dependence concerns without resolving them. These limitations remain substantive conditions for stronger claims.

Appendix C retains exploratory exercises and sensitivity analyses based on recorded values. Their measurement issues are documented there rather than repeated throughout the main text. Time-varying omitted variables, including income, policy, recreation, and transport access, may affect the domestic estimates. County and year effects do not establish causality (Bell et al. 2019). Appendix B's exploratory residual diagnostics indicate spatial association in seven of the ten years covered by M3 at an unadjusted significance level of 0.05. County clustering may therefore miss dependence across counties, while conventional state-cluster intervals rely on only 18 clusters. Neither a spatial covariance correction nor a suitable few-cluster bootstrap has been implemented. Reported confidence intervals remain conditional on the stated dependence assumptions.

Equation (11) defines the Pearson residual used in the newly implemented exploratory spatial diagnostic. Appendix B describes the spatial weights, the permutation procedure, annual results, and inferential limits. These diagnostics evaluate the supplied point geography and do not validate the point identities.

$$r_{it}^{P} = \frac{D_{it} - \hat{\mu}_{it}}{\sqrt{\hat{\mu}_{it}}} \tag{11}$$

In Equation (11), r_{it}^{P} denotes the Pearson residual for county i in year t.

D_{it} is the observed dealership count, and $\hat{\mu}_{it}$ is its fitted conditional mean, including the fixed effects.

The denominator in Equation (11) is a working Poisson standard deviation, not an assumption that the observed counts are equidispersed. Appendix B describes exploratory annual permutations and their limits. These diagnostics do not supply spatially robust coefficient intervals.

8. Conclusions

Recorded growth in firearm dealerships was uneven across the 307 selected counties during 2010–2019. The total rose from 6995 to 7303 (4.40%), while the county median declined from 17 to 16. Increases in 142 counties coexisted with decreases in 134 and no change in 31 (Figure 1). Examining county trajectories alongside totals therefore provides a more informative account of this sample than aggregate growth alone. The pattern does not establish a national trend or a change in concentration.

The domestic models further show why demographic and commercial change should be considered together. Both variables have positive associations in separate PPML models with county- and year-fixed effects. In the joint model, population has a coefficient of 0.499 (95% CI [−0.027, 1.024]), and establishments have a coefficient of 0.117 (95% CI [−0.266, 0.500]). Both intervals include zero, leaving each conditional association uncertain. A formal coefficient-difference test would be needed to compare their magnitudes.

The Canadian linkage exercise is excluded from the main research questions and substantive conclusions; Appendix C preserves the calculations and unresolved input issues for transparency.

The contribution is a sample-specific longitudinal description combining aggregate growth, divergent county trajectories, and domestic correlates. It motivates validation and future research rather than national or cross-border extrapolation.

Progress beyond this baseline requires verification of license coverage, sample construction, reference points, and annual population series. Transaction and recovery records, purchaser residence, and crossing-access measures could then connect validated retail histories more directly to acquisition and movement.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Data will be made available upon request.

Acknowledgments: The author gratefully acknowledges David Johnson (Department of Economics, University of Central Missouri) for providing the original data and for his valuable guidance and insightful ideas, which helped shape this research. The author used Grammarly 14.1349.0 for language editing and ChatGPT 6 to assist with coding. The author is responsible for checking source definitions, analyses, citations, and the final submitted text. The author conducted the research.

Conflicts of Interest: The author declares no conflicts of interest.

Appendix A

County identifiers in the descriptive sample

The following 307 county identifiers define the descriptive panel. The principal PPML models with county fixed effects exclude 38007, the all-zero county. Geographic models also exclude the 22 records listed below. Geographic exclusions

09001, 16009, 16013, 16017, 16021, 16035, 16049, 16055, 16057, 16061, 16069, 16079, 16081, 16085, 17111, 19059, 19063, 23001, 23003, 23011, 23017, 23019, 23021, 23023, 23027, 23031, 24001, 25009, 26001, 26003, 26005, 26007, 26009, 26011, 26013, 26015, 26019, 26023, 26031, 26033, 26035, 26045, 26049, 26051, 26053, 26055, 26057, 26059, 26061, 26063, 26065, 26067, 26069, 26071, 26073, 26077, 26079, 26083, 26087, 26089, 26091, 26095, 26097, 26099, 26101, 26103, 26107, 26109, 26111, 26113, 26117, 26119, 26121, 26123, 26125, 26127, 26129, 26131, 26135, 26137, 26141, 26143, 26145, 26151, 26153, 26155, 26157, 26161, 26165, 27001, 27005, 27007, 27025, 27035, 27057, 27059, 27061, 27065, 27069, 27071, 27087, 27095, 27097, 27107, 27111, 27113, 27115, 27123, 27125, 27129, 27135, 27137, 27159, 27167, 30001, 30003, 30007, 30009, 30013, 30015, 30019, 30025, 30027, 30029, 30035, 30037, 30039, 30041, 30045, 30055, 30059, 30063, 30065, 30069, 30073, 30075, 30081, 30085, 30087, 30089, 30093, 30095, 30097, 30101, 30107, 30109, 30111, 33001, 33005, 33007, 33009, 33013, 33015, 33017, 36007, 36009, 36011, 36015, 36017, 36027, 36029, 36043, 36063, 36067, 36069, 36073, 36075, 36083, 36089, 36091, 36093, 36095, 36099, 36107, 36109, 36111, 36121, 36123, 38003, 38005, 38007, 38009, 38011, 38015, 38019, 38021, 38023, 38025, 38029, 38031, 38035, 38039, 38041, 38043, 38053, 38061, 38065, 38067, 38073, 38079, 38081, 38093, 38095, 38097, 38099, 38101, 38105, 39005, 39007, 39009, 39011, 39013, 39025, 39029, 39031, 39035, 39037, 39039, 39053, 39055, 39059, 39071, 39073, 39089, 39093, 39099, 39103, 39105, 39117, 39119, 39129, 39135, 39141, 39143, 39145, 39155, 39157, 39161, 39163, 39173, 39175, 42003, 42005, 42009, 42013, 42023, 42027, 42031, 42033, 42043, 42053, 42061, 42063, 42067, 42069, 42079, 42081, 42087, 42093, 42097, 42105, 42107, 42109, 42115, 42121, 42129, 50001, 50003, 50005, 50007, 50013, 50015, 50021, 50027, 53003, 53007, 53009, 53015, 53019, 53027, 53029, 53035, 53037, 53039, 53047, 53049, 53051, 53057, 53059, 53061, 53063, 53071, 53073, 53075, 53077, 55005, 55007, 55013, 55029, 55061, 55067, 55075, 55083, 55099, 55113, 55125, 55129

Suspected location mismatch involving Billings: 30001 Beaverhead County → Billings (TP), Ontario; 30003 Big Horn County → Billings (TP), Ontario; 30007 Broadwater County → Billings (TP), Ontario; 30009 Carbon County → Billings (TP), Ontario; 30025 Fallon County → Billings (TP), Ontario; 30027 Fergus County → Billings (TP), Ontario; 30037

Golden Valley County → Billings (TP), Ontario; 30045 Judith Basin County → Billings (TP), Ontario; 30059 Meagher County → Billings (TP), Ontario; 30065 Musselshell County → Billings (TP), Ontario; 30069 Petroleum County → Billings (TP), Ontario; 30075 Powder River County → Billings (TP), Ontario; 30087 Rosebud County → Billings (TP), Ontario; 30093 Silver Bow County → Billings (TP), Ontario; 30095 Stillwater County → Billings (TP), Ontario; 30097 Sweet Grass County → Billings (TP), Ontario; 30107 Wheatland County → Billings (TP), Ontario; 30111 Yellowstone County → Billings (TP), Ontario.

Assignments requiring review of historical Canadian geography: 23003 Aroostook County → Clair (VL), New Brunswick; 27007 Beltrami County → Long Sault 12 (IRI), Ontario; 27057 Hubbard County → Long Sault 12 (IRI), Ontario; 27097 Morrison County → Long Sault 12 (IRI), Ontario.

RQ1 and complete variable descriptions

Section 5.1 reports county changes and quartiles, and Table 2 describes domestic distributions and within-county variation. The domestic PPML models with county fixed effects use 306 counties after exclusion of the all-zero county, compared with 307 in the descriptive sample.

Missingness and exclusion audit

The selected panel of 4532 rows for 2010–2020 has no blank dealership, U.S. population, establishment, Canadian population, or distance fields and no duplicate county–year keys. This internal audit does not establish the validity of zeros or recover the unavailable national source panel. Appendix B records the independently recomputed sample flow and corrected descriptive minimum.

Model output and inference

Locality clustering and all Canadian coefficients are discussed only in Appendix C. Main-text Table 5 reports domestic sensitivity specifications; Table A7 reports the Canadian-linked county fixed-effects specifications. Section 4.3 explains domestic inference conventions.

Collinearity and joint functional-form test: Model-specific parameter counts

The adjusted covariate correlation is 0.671, and the auxiliary VIF is 1.818. These values do not establish severe collinearity or explain the uncertainty in the joint model estimates. The quadratic Wald statistic is $\chi^2(2) = 0.513$, $p = 0.774$; its asymptotic reference distribution does not resolve few-cluster or spatial-inference concerns.

For I counties, T years, and p slopes, the full design rank is $I + T - 1 + p$ when the retained columns are independent. Table A1 gives the implied counts. K includes the intercept and all estimable fixed effects; G denotes the number of covariance clusters rather than the fixed-effect dimension. Conventional t intervals use $G - 1$ degrees of freedom, and $q = [G/(G-1)] \times [(N-1)/(N-K)]$.

Table A1. Model-specific parameter counts.

Model	N	Fixed Effects	Slopes	K
M1 and M2	3060	306 counties and 10 years	1	316
M3 and Table 5 rows 1–2	3060	306 counties and 10 years	2	317
Include 2020	3366	306 counties and 11 years	2	318
OLS log outcome	3070	307 counties and 10 years	2	318

Table A1. *Cont.*

Model	N	Fixed Effects	Slopes	K
All reconstructed	4110	411 counties and 10 years	2	422
Geographic domestic	2840	284 counties and 10 years	2	295
Add Canadian population	2840	284 counties and 10 years	3	296
G1	2850	18 states and 10 years	3	30
G2	2850	18 states and 10 years	4	31
M3 quadratic check	3060	306 counties and 10 years	4	319

Appendix B

Appendix B documents the refitting of models M1–M3, G1–G2, and the Canadian county-effects model, confirming agreement with the reported estimates to the precision shown. It also independently verifies sample flow, endpoint totals, zero entries, and the adjusted covariate correlation using the supplied data.

Table A2. Sample flow and exclusions.

Stage	Years	Counties	Rows
Full reconstructed panel	2010–2020	412	4532
Reconstructed main period	2010–2019	412	4120
Identifier-concordant descriptive panel	2010–2019	307	3070
Principal county-effects PPML	2010–2019	306	3060
Geographic descriptive and G1/G2 panel	2010–2019	285	2850
Geographic county-effects PPML	2010–2019	284	2840
G2 sensitivity excluding counties with zero Canadian-population entries	2010–2019	259	2590
County-effects sensitivity excluding counties with zero Canadian-population entries	2010–2019	258	2580

The first screening step for the main study period excludes 105 discordant county identifiers, or 1050 rows. Geographic screening then drops 22 concordant counties, or 220 rows. County 38007, which has zero dealerships in every year, contributes ten rows and is excluded only from the PPML models with county fixed effects. Of the 27 zero-dealership outcomes in the concordant main-period panel, 17 remain in the principal PPML models. All required domestic covariates are positive; there are no additional missing-value exclusions. All reconstructed and identifier-concordant domestic counts are positive integers or zero.

All 412 great-circle distances were recomputed from the supplied coordinates; the maximum absolute difference from the cached corrected distances is below 3×10^{-13} miles. The recomputed distances differ from the original distance field by more than ten miles in 31 records. This arithmetic agreement does not validate coordinates, the nearest-locality choice, or travel access. The corrected distances range from 0.5563 to 209.3854 miles. The

retained 71 locality labels have internally consistent locality–year population values and fixed county assignments.

The "Gun dealership by year" sheet labels an establishment-total column "Canada Population." For example, its 2010 value of 672,003 equals the U.S. establishment total, not a Canadian population total. This header should not be interpreted as referring to Canadian population. Table 2 and the domestic models use the county–year panel; the supplementary Canadian models use its Canadian-population field. The same locality population is recorded for each county assigned to that locality and cannot be summed to form a unique Canadian population. Static model-output and VIF sheets contain multiple specifications; their sample labels must be used rather than treating every entry as the principal M3 result.

Table A3. Zero-population exclusions by assigned locality.

Assigned Locality	Zero Rows	Counties
Agency 1 (IRI), Ontario	10	1
Blood 148A (IRI), Alberta	30	3
Cockburn Island (TP), Ontario	63	7
Dixville (MÉ), Quebec	10	1
East Hereford (MÉ), Quebec	50	5
Franklin (MÉ), Quebec	10	1
Improvement District No. 4 Waterton (ID), Alberta	10	1
Saint-Armand (MÉ), Quebec	50	5
Saint-Georges-de-Clarenceville (MÉ), Quebec	10	1
Saint-Théophile (MÉ), Quebec	10	1

All 26 counties with any zero Canadian-population entry are removed only for this sensitivity analysis, including years with positive population entries. This removes 260 rows and ten locality labels. The G2 sensitivity model with state fixed effects includes 17 states, ten years, four slopes, and K = 30; the corresponding county-effects model has K = 270. Both use 61 locality clusters. The G2 coefficient for Canadian population is 0.0228 (95% CI −0.0048 to 0.0505; $p = 0.104$), and the coefficient for distance is 0.0230 (−0.0420 to 0.0881; $p = 0.481$). In the county-effects sensitivity, the Canadian-population coefficient is 0.1533 (−0.0312 to 0.3378; $p = 0.102$). These conditional intervals include zero and do not resolve source accuracy or spatial dependence.

For each fitted model, a fixed neighborhood graph is constructed using the supplied U.S. point coordinates for the counties included in the estimation. Great-circle distance determines the four nearest neighbors of each county. All links between distinct counties are made reciprocal, and the resulting rows are standardized to sum to 1. Each model uses its own retained county set. Equation (11) defines Pearson residuals, centered separately within each year. Moran's statistics are given by Equation (A1) (Moran 1950).

$$I_t = (n/S_0) \times [\Sigma_i \Sigma_j w_{ij} u_{it} u_{jt} / \Sigma_i u_{it}^2] \quad (A1)$$

Here, n is the number of counties, u_{it} is the annual mean-centered Pearson residual, w_{ij} is a spatial weight, and S_0 is the sum of all weights, equal to n under row standardization. The graph is fixed across years within a model. For each model, the random number generator is initialized with the seed 4576597 and draws 999 independent permutations per

year in chronological order. Equation (A2) calculates a two-sided pseudo-p-value using the departure from the expectation under exchangeability, $I_0 = -1/(n-1)$.

$$p_t = \{1 + \Sigma_\beta \, 1[|I_t^{(b)} - I_0| \geq |I_t - I_0|]\}/1000 \quad \text{(A2)}$$

The sum in Equation (A2) is over the 999 permutations b; $1[\cdot]$ is an indicator, and $I_t^{(b)}$ is the statistic after permuting residuals among county points. The permutations are exploratory residual randomizations. They do not refit the regression, reproduce the clustered null, adjust for multiple annual tests, or constitute a formal test of the coefficient covariance estimator. The choice of neighbors and the results depend on the unvalidated identities of the input points.

At an unadjusted significance level of 0.05, seven of ten annual diagnostics for M3 and six of ten supplementary Canadian county-effect diagnostics indicate spatial association; none of the annual diagnostics does so for G1 or G2 using this graph. The latter result is insufficient to establish independence across localities or states. The spatial patterns in the county fixed-effects models underscore that the reported one-way clustered intervals depend on assumptions about spatial dependence. Neither a spatial heteroskedasticity- and autocorrelation-consistent (HAC) coefficient covariance estimator nor a PPML bootstrap suitable for a small number of state clusters has been implemented. Re-estimation with verified points and a defensible dependence model remains necessary before stronger inferential claims.

Table A4. Annual exploratory residual spatial diagnostics.

Year	M3 I (p)	G1 I (p)	G2 I (p)	Canadian FE I (p)
2010	0.1352 (0.001)	0.0383 (0.280)	0.0335 (0.326)	0.1255 (0.001)
2011	0.1360 (0.001)	0.0148 (0.638)	0.0130 (0.677)	0.1183 (0.003)
2012	0.1114 (0.002)	0.0061 (0.800)	0.0060 (0.815)	0.0671 (0.054)
2013	0.0241 (0.437)	0.0043 (0.876)	0.0052 (0.854)	0.0224 (0.497)
2014	0.1032 (0.002)	0.0023 (0.868)	0.0053 (0.820)	0.1207 (0.002)
2015	0.0342 (0.347)	−0.0026 (0.974)	0.0016 (0.884)	0.0256 (0.426)
2016	0.0025 (0.845)	−0.0099 (0.872)	−0.0071 (0.945)	−0.0095 (0.871)
2017	0.0980 (0.005)	0.0003 (0.917)	0.0032 (0.845)	0.0849 (0.018)
2018	0.1112 (0.001)	0.0060 (0.802)	0.0080 (0.757)	0.0959 (0.009)
2019	0.1361 (0.001)	0.0398 (0.236)	0.0438 (0.190)	0.1027 (0.002)

These diagnostics use supplied reference points, whereas Figure 1 places county counts on Census county polygons using FIPS codes. The polygon join supports cartographic placement of the recorded county identifiers; it does not validate the distinct points used for residual neighborhoods. Diagnostics based on verified county reference geometry and a prespecified alternative neighborhood would be needed to assess sensitivity to that choice.

Appendix C. Exploratory Canadian Data Linkage

Scope and interpretation. This appendix contains all Canadian model results. They are retained solely as an exploratory data-linkage exercise. An inherited single-locality assignment is not a verified nearest-neighbor search, geographic catchment, border-access measure, or measure of cross-border demand.

Unresolved zero entries. C denotes the data recorded in the Canadian population input, not independently validated Canadian population. The 253 zero entries cannot be interpreted as a measured absence of residents. Full-panel Canadian coefficients are

retained only as diagnostics of this unvalidated input. The existing sensitivity excludes all 26 counties with any zero entry, including years with positive entries, leaving 259 counties and 2590 observations for G2 and 258 counties and 2580 observations for county-effects PPML. Appendix B reports the coefficients and the use of 61 locality clusters for this sensitivity analysis. Positive recorded values are not independently validated either. No undocumented zero is imputed, treated as a known true population, or replaced with an invented value.

The domestic analysis uses county population and establishment counts without reference to Canadian assignments. The supplementary coefficients are retained to show how results from the inherited linkage change with the fixed effects, clustering, and exclusion of counties with undocumented zero inputs. Their interpretation is limited to the recorded inputs and the fitted specifications. A substantive interpretation would require verification of locality identifiers and annual populations, the candidate universe for matching, and routes through border crossings.

A supplementary geographic question concerns the population of Canadian localities assigned to counties in the inherited dataset. Research on interstate sourcing and official tracing describes connections beyond local boundaries (Knight 2013; Coates and Pearson-Merkowitzz 2017; Bureau of Alcohol, Tobacco, Firearms and Explosives 2026a; Royal Canadian Mounted Police 2026). Such evidence motivates the study of geographic context but does not validate this dataset's Canadian assignments or imply cross-border demand. The assigned population and point-to-point distance are exploratory contextual proxies whose construction and limitations are examined explicitly.

The embedded reference-point display and its state-name key are retained as Figure A1. Figure 1 is the separate county-polygon map of domestic endpoint change. Neither display verifies the inherited Canadian assignments.

Canadian locality population $C_{it} = P_{j(i),t}$ is the population of locality $j(i)$ assigned to county i. The assignment is fixed over time. Distance d_i is the great-circle separation between the supplied U.S. and Canadian reference points. These measures describe external geographic context. They do not measure Canada's national population or access to arms through an authorized crossing. Shared assignments also create dependence because multiple counties receive the same locality population series.

Prior research provides context for this geographic extension. Studies of students and handgun owners suggest that firearm attitudes and ownership motives can differ across national samples (Baxter-Moore et al. 2016; Stroebe et al. 2021). Research on interstate sourcing and regulatory spillovers examines connections shaped by acquisition conditions and geography (Coates and Pearson-Merkowitzz 2017; Johnson and Suwanprasert 2026, working paper), while theft offers another pathway from lawful possession to illegal supply (Donohue et al. 2025). These studies underscore the distinction between population, proximity, and observed interaction.

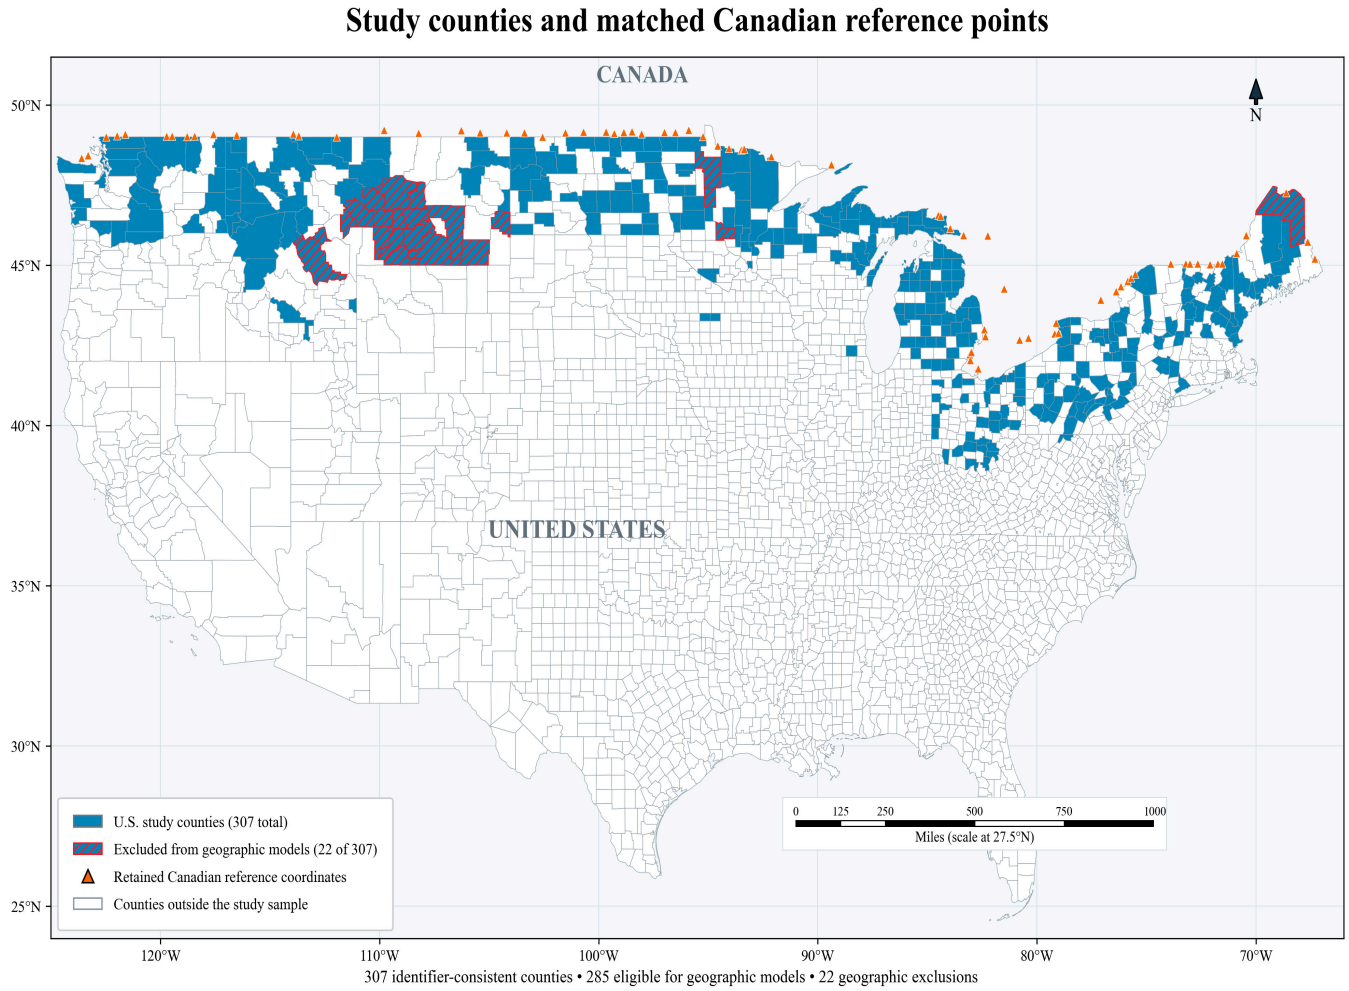

Figure A1. Supplied reference points for the 307 identifier-concordant U.S. counties and distinct Canadian coordinate pairs associated with the 285-county exploratory sample. Red rings identify the 22 counties excluded from geographic models. State abbreviations are keyed below the figure. This coordinate display replaces the earlier broad regional map; it does not depict county boundaries, all Canadian settlements, or a validated nearest-locality search. Source: Location coordinates in the Distances sheet of the Canadian dataset. The display reproduces supplied coordinates without certifying their identities.

Appendix C.1. County Identifiers and Geographic Reference Points

The Distances sheet contains 412 rows, each with its own GEOID and reconstructed county prefix. It contains 412 distinct original county codes, of which 307 match the reconstructed prefix in the corresponding rows. The 575 location records refer to an unavailable upstream file and cannot be reproduced from the attachment. For an 11-digit Census tract GEOID, the first five digits encode state and county, and the remaining six identify the tract (U.S. Census Bureau n.d.). A tract-based reference point is not automatically a county centroid. Prefix agreement is an internal identifier check, not independent verification of a coordinate's location, geographic vintage, or county membership.

The supplied U.S. and Canadian coordinate pairs were used exactly as recorded for the distance audit. Their identities and coordinate reference metadata have not been externally verified; the U.S. points are not established as population-weighted county centers. Google Earth Pro 7.3.7.1327 (64-bit) and ArcGIS Pro 3.7 were used for data processing. Published reference points and mean centers have distinct definitions (U.S. Census Bureau n.d., n.d.; Esri n.d.). The audit therefore distinguishes the recomputation of the distance between two points from the validation of either point or its assignment.

Appendix C.2. Distance and Canadian Locality Assignment

For the supplied U.S. point i and Canadian point j, Equation (A3) calculates the great-circle distance on a sphere using the haversine formula (Veness and Hijmans n.d.; see their documentation of Sinnott 1984).

$$d_{ij} = 2R\,arcsin\left(\sqrt{sin^2\left(\frac{\varphi_j - \varphi_i}{2}\right) + \cos(\varphi_i)\cos\left(\varphi_j\right)sin^2\left(\frac{\lambda_j - \lambda_i}{2}\right)}\right) \quad \text{(A3)}$$

For each county i, d_i in Equations (A6), (A7) and (A9) is d_{ij} for its assigned locality $j = \mathrm{j(i)}$.

Where φ denotes latitude, λ denotes longitude, all angles are in radians, and R = 3958.7613 miles. The resulting distance measures surface separation between the two reference points. It does not measure a road route, travel through a border crossing, or observed movement. The inherited point pairing does not establish that the assigned locality is nearest or demonstrate interaction between a U.S. county and residents of that locality.

Each county retains the locality label recorded in the data workbook throughout the study period. C_{it} denotes the annual population value attached to that locality. The audit found no changes in county assignments and no conflicting population values for any locality–year combination. These internal checks do not establish that the location is nearest among all eligible Canadian localities. Esri (n.d., n.d.) provides guidance on candidate coverage, coordinate systems, geodesic distance, and tie handling. A future matching exercise designed for independent replication should specify these elements.

The supplied materials do not identify a reproducible population or geographic product, release, variable identifier, download date, annual interpolation rule, or historical boundary crosswalk. Statistics Canada's Table 17-10-0155-01 documents estimates on 2021 census-subdivision boundaries (Statistics Canada 2026). The Census county intercensal series and County Business Patterns documentation describe potential official sources. Original extracts and a variable-level provenance record are needed to validate the spatial linkages.

The zero-input audit and interpretation are stated at the beginning of this appendix. Appendix B lists the ten affected locality labels and reports the existing exclusion sensitivity. No zero entry is interpreted as a verified population count.

Appendix C.3. Geographic Linkage Validity

The implemented assignment rule is to retain each county's recorded "Addresses in Canada" (Address_Can) label and attach that locality's annual Canadian_population series. The label remains unchanged across years. The conceptual target is external geographic context, not a demonstrated retail catchment. The materials do not identify the original Canadian candidate universe, its population threshold, geographic extent, locality types, search metric, tie rule, or crossing-access restrictions. Consequently, the recorded assignments cannot be certified as nearest among all eligible places or as behaviorally relevant destinations. Recomputing the distance between assigned points audits arithmetic only.

A single locality provides a compact, reproducible label within the workbook, but it can omit other accessible settlements and assign the same settlement to multiple counties. Its population is therefore a contextual proxy, not a measure of Canadian demand, customers, firearm purchases, or the total population accessible to a county. Great-circle distances do not account for road networks, water barriers, or the locations and operation of border crossings. The U.S. point is not a verified population-weighted county center. These limitations preclude a substantive claim that a single assigned locality is the appropriate catchment.

Exclusions follow the supplied flags after screening for identifier concordance: 18 counties assigned to Billings (TP), Ontario, one assigned to Clair (VL), New Brunswick, and three assigned to Long Sault 12 (IRI), Ontario. Appendix A lists all 22. Billings is a

suspected name/location mismatch; the latter cases are historical-geography review flags. These flags do not establish that all affected assignments are errors, that an independently verified reassignment has been made, or that unflagged pairs are valid. Canadian geographic identifiers and boundary histories are needed to distinguish similarly named places and to maintain comparable population counts over time. Statistics Canada (2019) documents municipal boundary and name changes relevant to longitudinal comparisons; however, this guidance does not validate these particular observations.

Among the 285 retained counties, 71 locality labels are each assigned to between 1 and 32 counties; 31 labels are assigned to one county each, and the median is two counties per label. Locality clustering is primary for Canadian-population specifications, allowing temporal dependence and dependence among counties assigned to the same locality when assignments are fixed. It leaves cross-locality spatial and statewide dependence unresolved. Changing clustering changes covariance estimates, not the fitted coefficients. Replacing county fixed effects with state fixed effects changes the variation used for identification and may change the coefficient. Neither comparison isolates assignment error.

Assignment and index construction could contribute to instability through the omission of accessible populations, inappropriate reference points, repeated exposure, undocumented zeros, or inconsistent geography across years. The available comparisons cannot identify the contribution of each source. G2 estimates the Canadian-population coefficient at −0.0018, with a locality-clustered interval [−0.0210, 0.0173], whereas the supplementary estimate with county fixed effects is 0.133 with an interval of [−0.024, 0.290]. These different conditional estimands and their uncertainty do not establish either a robust geographic relationship or its absence. Appendix B's exclusion of counties with any zero Canadian population entry is a measurement sensitivity check, not an alternative locality assignment.

A multi-locality alternative would require an independently defined candidate set J_i, verified locality–year populations on comparable boundaries, and travel times through authorized crossings. Equation (A4) defines one possible normalized contextual index.

$$A_{it} = \Sigma_j \in J_i \, w_{ij} \, P_{jt}, \; w_{ij} = \exp(-t_{ij}/h) / \Sigma_k \in J_i \exp(-t_{ik}/h) \tag{A4}$$

Here, J_i is the set of eligible Canadian localities for county i; j and k index localities in that set; P_{jt} is locality j's population in year t; t_{ij} is route travel time; h is a prespecified positive decay scale in the same time units; and w_{ij} is a normalized weight. This normalized index measures average locality size, not total accessible population or demand. An unnormalized sum would represent a different construct and require a separate justification. Candidate eligibility, crossing availability, scale selection, the treatment of zero and missing values, and overlapping exposures must be specified before examining results. Equation (A4) is a proposed construction, not an estimated result.

This alternative has not been estimated because the supplied data workbook contains one inherited assignment per county and no validated candidate set, route matrix, or complete candidate-population panel. Treating its assigned localities as the complete candidate universe would reproduce unknown selection and omit unassigned alternatives. A future comparison should hold the county sample, years, domestic controls, and fixed effects constant and report coefficient estimates, intervals, and assignment changes without using statistical significance to select a weighting rule. Overlapping multi-locality exposures would require a method for accounting for dependence beyond a single assigned-locality cluster.

$$\mu_{it} = exp[\alpha_i + \tau_t + \beta_P \, ln(P_{it}) + \beta_E \, ln(E_{it}) + \gamma \, ln(1 + C_{it})] \tag{A5}$$

The coefficient γ multiplies the transformed recorded input $ln(1 + C_{it})$; Section S.6 gives its interpretation with respect to C_{it}. County fixed effects absorb time-invariant distance. This supplementary specification uses 284 counties and 2840 observations (see Table A7 and specifications S2–S3). Equation (A5) gives its conditional mean.

Appendix C.4. Geographic Model G1

Equation (A6) specifies G1. C_{it} is the recorded input assigned to the Canadian locality, as defined at the beginning of this appendix; d_i is the time-invariant separation of supplied points in miles.

$$\mu_{it}^{G1} = exp\left[\delta_{s(i)} + \tau_t + \gamma_P \, ln(P_{it}) + \theta \, ln(1 + C_{it}) + \lambda \, ln(1 + d_i)\right] \quad (A6)$$

Here, δ_{si} is the fixed effect for the state containing county i, τ_t is a year-fixed effect, γ_p is the coefficient on log U.S. population, θ is the coefficient on log(1 + Canadian population), and λ is the coefficient on log(1 + distance). G_1 and G_2 use the reported 285-county, 2850-observation geographic sample. These state fixed-effects models are distinct from the supplementary models with county fixed effects and 284 counties.

In Equations (A6) and (A7), $s(i)$ indexes the state containing county i, and the superscripts G1 and G2 label the two model-specific conditional means. In Equation (A5), α_i, τ_t, β_p, and βE have the same roles as in the domestic model, and γ is the additional Canadian-input coefficient. In Equations (A8) and (A9), εC and εd denote local elasticities rather than regression disturbances. These local elasticities follow from the chain rule; the finite-change expression in Section S.6 follows by taking the ratio of two exponential conditional means.

Appendix C.5. Geographic Model G2

$$\mu_{it}^{G2} = exp\left[\delta_{s(i)} + \tau_t + \gamma_P \, ln(P_{it}) + \gamma_E \, ln(E_{it}) + \theta \, ln(1 + C_{it}) + \lambda \, ln(1 + d_i)\right] \quad (A7)$$

Equation (A7) adds log establishments to Equation (A6), with the coefficient γ_e. State- and year-fixed effects are retained.

State effects account for stable differences between states but leave unobserved county differences uncontrolled. Distance is identified from variation among counties within states. Canadian population varies across assigned localities and over time, so its coefficient in G1 and G2 can reflect both cross-county and temporal information. G1 adjusts for domestic population; G2 also adjusts for establishments. Neither specification removes all potential omitted-variable bias.

G1 and G2 estimate an exponential conditional mean using PPML; the outcome is not log-transformed. The recorded Canadian input and distance enter as ln(1 + C) and ln(1 + d). These transformations are mathematical definitions, not validation of the input values. Their interpretation is stated once below.

Appendix C.6. Interpretation of Transformed Predictors

For positive recorded C, the local elasticity is θC/(1 + C); for positive distance d, it is λd/(1 + d), as recorded in Equations (A8) and (A9). In Equation (A5), γ replaces θ. At C = 0, a conventional percentage-change elasticity is undefined because a percentage change from zero is undefined. A finite change from C_1 to C_2 instead implies a proportional change in the fitted mean of $\exp\{\theta[\ln(1 + C_2) - \ln(1 + C_1)]\} - 1$, holding other terms constant. This algebra does not validate the recorded values.

$$\varepsilon_C = \frac{\theta C_{it}}{1 + C_{it}} \quad (A8)$$

$$\varepsilon_d = \frac{\lambda d_i}{1 + d_i} \quad (A9)$$

The reference scales are one person for the recorded input and one mile for distance. A unit conversion must convert both the variable and its reference scale to preserve the transformation.

Because zero-input provenance is unknown, full-panel coefficients are diagnostic only. The sensitivity analysis in Appendix B excludes every county with any zero entry and provides the relevant comparison for values recorded as positive. It does not remedy gaps in the underlying source documentation.

Appendix C.7. Exploratory Geographic Associations

The primary confidence intervals use clustering by the 71 assigned localities to account for shared exposure. In G1, the Canadian population and distance coefficients are −0.0045 (95% CI [−0.0230, 0.0139]; $p = 0.626$) and −0.0069 (95% CI [−0.0670, 0.0531]; $p = 0.819$). The corresponding G2 estimates are −0.0018 (95% CI [−0.0210, 0.0173]; $p = 0.850$) and 0.0004 (95% CI [−0.0608, 0.0615]; $p = 0.990$). All intervals include zero, leaving the direction of both associations uncertain. Section S.8 presents the supplementary within-county comparison.

G1 and G2 estimates are reported in Table A5 with alternative county- and locality-clustered intervals.

G1 and G2 use 2850 county–year observations from 285 counties during 2010–2019, with state- and year-fixed effects. G1 controls for U.S. population; G2 also controls for establishments. The all-zero county can remain in these models because they do not include county fixed effects, which explains why the sample is larger than the one used for Table A7. Each county receives its assigned Canadian locality's population. Distance between the supplied reference points is measured in miles, and these points have not been verified as population centers. Canadian population and distance enter as $\ln(1 + C_{it})$ and $\ln(1 + d_i)$. Alternative covariance estimates use 285 county clusters or 71 locality clusters; the corresponding specifications retain the same observations and coefficients. Full-rank designs imply K = 30 for G1 and K = 31 for G2. These comparisons assess sensitivity to the assumed dependence structure rather than validating geographic relationships.

Additional checks vary the distance measure, sample, and state coverage. Using the supplied distance field in the 285-county G2 sample yields a distance coefficient of −0.0008 ($p = 0.976$); using all 307 concordant pairs yields −0.0161 ($p = 0.458$). Because sample coverage also changes, the difference cannot be attributed solely to distance measurement. Omitting one state at a time produces Canadian population coefficients from −0.0057 to 0.0054 and distance coefficients from −0.0092 to 0.0204. These point-estimate ranges lack corresponding uncertainty estimates and therefore do not establish statistically meaningful differences. Overall, the checks leave the direction of the geographic associations unresolved rather than establishing their absence.

Table A5. G1 and G2 regression estimates of recorded U.S. dealership counts with county- and locality-clustered inference.

Model and Clustering	Term	Estimate	95% CI	*P*
G1 county G = 285	ln(U.S. population)	0.6169	[0.5740, 0.6598]	<0.001
G1 county G = 285	ln(1 + Canadian population)	−0.0045	[−0.0233, 0.0142]	0.634
G1 county G = 285	ln(1 + distance)	−0.0069	[−0.0602, 0.0464]	0.798
G1 locality G = 71	ln(U.S. population)	0.6169	[0.5724, 0.6614]	<0.001
G1 locality G = 71	ln(1 + Canadian population)	−0.0045	[−0.0230, 0.0139]	0.626

Table A5. *Cont.*

Model and Clustering	Term	Estimate	95% CI	*P*
G1 locality G = 71	ln(1 + distance)	−0.0069	[−0.0670, 0.0531]	0.819
G2 county G = 285	ln(U.S. population)	0.4415	[0.2193, 0.6636]	<0.001
G2 county G = 285	ln(establish-ments)	0.1689	[−0.0455, 0.3832]	0.122
G2 county G = 285	ln(1 + Canadian population)	−0.0018	[−0.0212, 0.0176]	0.854
G2 county G = 285	ln(1 + distance)	0.0004	[−0.0553, 0.0560]	0.989
G2 locality G = 71	ln(U.S. population)	0.4415	[0.2484, 0.6345]	<0.001
G2 locality G = 71	ln(establish-ments)	0.1689	[−0.0119, 0.3497]	0.067
G2 locality G = 71	ln(1 + Canadian population)	−0.0018	[−0.0210, 0.0173]	0.850
G2 locality G = 71	ln(1 + distance)	0.0004	[−0.0608, 0.0615]	0.990

Appendix C.8. Supplementary Within-County Canadian Population Association

The supplementary model with county- and year-fixed effects adds $\ln(1 + C_{it})$ to the domestic controls, using 284 counties and 2840 observations (Table A7, specifications S2–S3). The Canadian population coefficient is 0.1329. It is identified by the variation remaining after adjustment for county and year effects and domestic covariates. It is a conditional within-county association, not a coefficient from an explicitly first-differenced population-growth model.

With locality clustering, the Canadian population coefficient is 0.133, with a 95% interval of [−0.024, 0.290] (71 localities; reported SE = 0.0788; $p \approx 0.096$). The positive estimate remains compatible with a small negative association, zero, or a positive association. County clustering gives a narrower interval of [0.027, 0.239] (284 counties; reported SE = 0.0537; $p \approx 0.014$). This comparison shows sensitivity to dependence assumptions; clustering should follow the exposure and error structure rather than whether an interval includes zero (Cameron and Miller 2015).

Both covariance calculations leave the point estimate unchanged. Locality clustering is primary because multiple counties share the same Canadian population series, whereas the narrower county-clustered interval does not establish greater reliability. Neither estimator addresses all possible dependence across localities, so spatial and influence assessments remain necessary. Time-invariant distance is absorbed by county effects.

The supplementary Canadian-population coefficient describes an association identified from the variation remaining after adjustment for county and year effects and the included controls. It differs from the G1/G2 coefficients, which are estimated with state and year effects and can reflect both cross-county and temporal variation. Neither specification establishes cross-border purchasing or a causal relationship between Canadian population and U.S. dealership counts.

Table A6. Recorded Canadian input and distance summaries.

Sample and Variable	Mean	SD	Min	Median	Max	Within SD
Geographic Canadian population	12,280.29	35,645.08	0.00	1036.00	233,040.00	833.42
Geographic distance (miles)	97.21	55.96	0.56	99.55	209.39	0.00

Table A7. Canadian-linked county fixed-effects sensitivity results.

Specification/Clustering	N/Counties	U.S. Population ln(P)	Establishments ln(E)	Canadian Population ln(1 + C)
S1. Geographic sample/county	2840 284	0.548 [0.004, 1.091]	0.095 [−0.314, 0.504]	-
S2. Canadian population/county	2840 284	0.626 [0.084, 1.168]	0.042 [−0.368, 0.452]	0.133 [0.027, 0.239]
S3. Same as S2/Canadian locality	2840 284	0.626 [0.036, 1.216]	0.042 [−0.488, 0.572]	0.133 [−0.024, 0.290]

References

Adcock, Robert, and David Collier. 2001. Measurement validity: A shared standard for qualitative and quantitative research. *American Political Science Review* 95: 529–46. [CrossRef]

Anderson, James E., Mario Larch, and Yoto V. Yotov. 2018. GEPPML: General equilibrium analysis with PPML. *The World Economy* 41: 2750–82. [CrossRef]

Baxter-Moore, Nick, Munroe Eagles, Dupinder Aheer, Racquel Maxwell, Lisa-Anne Pilkey, and Kimmy Samra. 2016. Explaining Canada–US differences in attitudes toward crime and justice: An empirical test of S.M. Lipset's account. *American Review of Canadian Studies* 46: 430–51. [CrossRef]

Bell, Andrew, Malcolm Fairbrother, and Kelvyn Jones. 2019. Fixed and random effects models: Making an informed choice. *Quality & Quantity* 53: 1051–74. [CrossRef]

Bureau of Alcohol, Tobacco, Firearms and Explosives. 2026a. Firearms Trace Data: Canada—2020–2025. Available online: https://www.atf.gov/firearms/report/firearms-trace-data/firearms-trace-data-canada-2020-2025 (accessed on 21 October 2026).

Bureau of Alcohol, Tobacco, Firearms and Explosives. 2026b. National Firearms Commerce and Trafficking Assessment (NFCTA): Firearms Trafficking Investigations—Volume Three. Available online: https://www.atf.gov/firearms/national-firearms-commerce-and-trafficking-assessment-nfcta-firearms-trafficking-investigations-volume-three (accessed on 21 October 2026).

Bureau of Alcohol, Tobacco, Firearms and Explosives. n.d. Federal Firearms Licenses. Available online: https://www.atf.gov/firearms/federal-firearms-licenses (accessed on 21 November 2026).

Cameron, A. Colin, and Douglas L. Miller. 2015. A practitioner's guide to cluster-robust inference. *Journal of Human Resources* 50: 317–72. [CrossRef]

Coates, Michael, and Shanna Pearson-Merkowitzz. 2017. Policy spillover and gun migration: The interstate dynamics of state gun control policies. *Social Science Quarterly* 98: 500–12. [CrossRef]

Correia, Sergio, Paulo Guimarães, and Tom Zylkin. 2020. Fast Poisson estimation with high-dimensional fixed effects. *The Stata Journal* 20: 95–115. [CrossRef]

Donohue, John J., Samuel V. Cai, Matthew V. Bondy, and Philip J. Cook. 2025. Why do right to carry laws increase violence? Effects on gun theft and clearance rates. *Journal of Urban Economics* 147: 103761. [CrossRef]

Esri. n.d.a. Generate Near Table (Analysis). ArcGIS Pro Documentation. Available online: https://doc.esri.com/en/arcgis-pro/latest/tool-reference/analysis/generate-near-table.html (accessed on 10 April 2026).

Esri. n.d.b. Geodesic Versus Planar Distance. ArcGIS Pro Documentation. Available online: https://doc.esri.com/en/arcgis-pro/latest/tool-reference/spatial-analyst/geodesic-versus-planar-distance.html (accessed on 10 April 2026).

Esri. n.d.c. Mean Center (Spatial Statistics). ArcGIS Pro Documentation. Available online: https://doc.esri.com/en/arcgis-pro/latest/tool-reference/spatial-statistics/mean-center.html (accessed on 10 April 2026).

Johnson, David B., and Joshua J. Robinson. 2024. Online appendix to Gun dealer density and its effect on homicide (Supplemental material). *Journal of Law and Economics* 67: 1–30. [CrossRef]

Johnson, David Blake, and Wisarut Suwanprasert. 2026. Regulatory spillovers and the interstate flow of crime guns [Working paper]. *SSRN*. [CrossRef]

Knight, Brian. 2013. State gun policy and cross-state externalities: Evidence from crime gun tracing. *American Economic Journal: Economic Policy* 5: 200–29. [CrossRef]

Laqueur, Hannah S., and Colette Smirniotis. 2025. Identifying high-risk firearms dealers: A machine learning study of rapidly diverted firearm sales in California. *Criminology & Public Policy* 24: 363–403. [CrossRef]

Moran, P. A. P. 1950. Notes on continuous stochastic phenomena. *Biometrika* 37: 17–23. [CrossRef]

O'Brien, Robert M. 2007. A caution regarding rules of thumb for variance inflation factors. *Quality & Quantity* 41: 673–90. [CrossRef]

Pear, Veronica A., Garen J. Wintemute, Nicholas P. Jewell, Magdalena Cerdá, and Jennifer Ahern. 2023. Community-level risk factors for firearm assault and homicide: The role of local firearm dealers and alcohol outlets. *Epidemiology* 34: 798–806. [CrossRef]

Robinson, Sonia L., Christopher D. McCort, Colette Smirniotis, Garen J. Wintemute, and Hannah S. Laqueur. 2024. Purchaser, firearm, and retailer characteristics associated with crime gun recovery: A longitudinal analysis of firearms sold in California from 1996 to 2021. *Injury Epidemiology* 11: 8. [CrossRef]

Roodman, David, Morten Ørregaard Nielsen, James G. MacKinnon, and Matthew D. Webb. 2019. Fast and wild: Bootstrap inference in Stata using boottest. *The Stata Journal* 19: 4–60. [CrossRef]

Royal Canadian Mounted Police. 2026. 2025 Commissioner of Firearms Report. Available online: https://rcmp.ca/en/corporate-information/publications-and-manuals/2025-commissioner-firearms-report (accessed on 25 October 2026).

Santos Silva, João Manuel Caravana, and Silvana Tenreyro. 2006. The log of gravity. *The Review of Economics and Statistics* 88: 641–58. [CrossRef]

Sinnott, Roger W. 1984. Virtues of the haversine. *Sky & Telescope* 68: 159.

Stansfield, Richard, Daniel Semenza, Jie Xu, and Elizabeth Griffiths. 2023. Licensed firearm dealers, legal compliance, and local homicide: A case study. *Criminology & Public Policy* 22: 323–45. [CrossRef]

Statistics Canada. 2019. Interim List of Changes to Municipal Boundaries, Status, and Names Up to January 1, 2019. Available online: https://www150.statcan.gc.ca/n1/pub/92f0009x/92f0009x2019001-eng.htm (accessed on 26 October 2026).

Statistics Canada. 2026. Population Estimates, July 1, by Census Subdivision, 2021 Boundaries (Table 17-10-0155-01) (Data Set). Available online: https://www150.statcan.gc.ca/t1/tbl1/en/tv.action?pid=1710015501 (accessed on 17 September 2026).

Stroebe, Wolfgang, Jannis Kreienkamp, N. Pontus Leander, and Maximilian Agostini. 2021. Do Canadian and U.S. American handgun owners differ? *Canadian Journal of Behavioural Science/Revue Canadienne des Sciences du Comportement* 53: 221–31. [CrossRef]

U.S. Census Bureau. 2010. 2010 Cartographic Boundary Files (Data Set). Available online: https://www.census.gov/geographies/mapping-files/2010/geo/carto-boundary-file.html (accessed on 25 February 2026).

U.S. Census Bureau. n.d.a. Gazetteer File Record Layouts. Available online: https://www.census.gov/programs-surveys/geography/technical-documentation/records-layout/gaz-record-layouts.html (accessed on 25 February 2026).

U.S. Census Bureau. n.d.b. Gazetteer Files. Available online: https://www.census.gov/geographies/reference-files/time-series/geo/gazetteer-files.html (accessed on 17 September 2026).

U.S. Census Bureau. n.d.c. Understanding Geographic Identifiers (GEOIDs). Available online: https://www.census.gov/programs-surveys/geography/guidance/geo-identifiers.html (accessed on 17 September 2026).

Veness, Chris, and Robert Hijmans. n.d. distHaversine: Haversine Great Circle Distance (R Documentation). CRAN. Available online: https://search.r-project.org/CRAN/refmans/geosphere/html/distHaversine.html (accessed on 5 November 2026).

Wasserstein, Ronald L., and Nicole A. Lazar. 2016. The ASA statement on p-values: Context, process, and purpose. *The American Statistician* 70: 129–33. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.