Fast and Spurious: A Causal Evaluation of NYPD's Pursuit Policy, 2022–2025

John Hall
Justin Nix

03-22-2026

Abstract

Objectives: This study examines the consequences of the New York Police Department's late-2022 operational shift that dramatically expanded vehicle pursuits without a formal policy change. We evaluate whether the surge in pursuits affected motor vehicle collisions, robberies, and shootings, and whether any crime-prevention benefits offset the associated collision costs.

Methods: ITS segmented regression estimates changes in pursuit-related collisions following the October 2022 escalation; no external comparison group is required because collisions are mechanically linked to pursuit volume. For crime outcomes, a precinct-level heterogeneous-effects DD design exploits pre-treatment crime-level variation across all 77 NYPD precincts, with event-study diagnostics assessing the parallel trends assumption. A cost–benefit analysis estimates the threshold of prevented crimes needed for pursuit escalation to break even.

Results: Monthly pursuits surged from roughly 8 to over 160 events — a rise exceeding 2,000%. Pursuit-related collisions increased proportionally with pursuit volume and declined sharply following the February 2025 partial re-restriction that reduced pursuits by approximately three-quarters. Crime outcomes provide no evidence of deterrence: the robbery DD estimate is positive and significant, while the shooting DD estimate, though nominally negative, rests on non-parallel pre-trends. The February 2025 reduction in pursuits was not followed by any increase in crime.

Conclusions: Pursuit escalation produced clear and attributable collision harms while providing no credible evidence of crime reduction. Cost–benefit estimates suggest the threshold of prevented shootings required to offset these harms was not achieved. Agencies should require affirmative evidence of crime-prevention benefits before expanding pursuit authority.

Introduction

Vehicle pursuits are among the most dangerous activities in American policing. Each year, an estimated 300 people die in pursuit-related motor vehicle crashes in the United States, roughly one-third of whom are bystanders uninvolved in the pursuit (Rivara and Mack 2004). Pursuits frequently occur at high speeds, on local roads, and at night, creating substantial risk to officers, suspects, and the public (Alpert and Anderson 1986; Hill 2002). Police departments have been making and remaking those choices for decades, under substantial public pressure and with remarkably little empirical guidance (Alpert et al. 1997b; Hicks 2003).

Despite the stakes, empirical evidence on the effects of pursuit policy changes remains thin. A handful of case studies have documented sharp changes in pursuit frequency following policy shifts: Dade County, Florida saw an 82% reduction in pursuits after restricting them to violent felons, while Omaha, Nebraska experienced a 600% increase after relaxing its policy (Alpert et al. 1997a; Hoffmann and Mazerolle 2005). Yet no study has rigorously evaluated the downstream consequences of such changes on crime, collisions, and public safety outcomes using modern causal identification methods. The question that Rivara and Mack (2004) posed two decades ago — whether "the trade-off of fewer pursuit-related crashes and deaths [is] offset by a higher number of fatal crimes" — remains unanswered.

The New York Police Department's experience beginning in late 2022 offers a rare opportunity to address this question. Under Mayor Eric Adams's quality-of-life enforcement mandate and new Chief of Patrol John Chell, the NYPD dramatically escalated its use of vehicle pursuits without any formal revision to the department's written Patrol Guide procedure on pursuits. Monthly pursuit events surged from an average of 8 in the pre-escalation period to over 160 by mid-2023, an increase exceeding 2,000%. This operational shift was followed by two additional regime changes: a compliance memo from Chief of Department Jeffrey Maddrey in August 2023, and a formal policy rewrite by Commissioner Jessica Tisch in February 2025 that restricted pursuits to felonies and violent misdemeanors only. The result is a multi-phase natural experiment in which pursuit intensity varied dramatically across distinct policy regimes within a single jurisdiction.

We exploit this variation to estimate the effects of pursuit escalation on motor vehicle collisions, robberies, shootings, and shots fired incidents. Collision outcomes are mechanically linked to pursuit volume, so we model them with an interrupted time series (ITS) segmented regression using Newey-West HAC standard errors — no external comparison group is required. Crime outcomes are a harder problem. Within-city trends may reflect secular forces unrelated to pursuit policy, so our primary identification strategy is a precinct-level heterogeneous-effects DD design that classifies precincts by their pre-October 2022 crime rate (a predetermined, exogenous characteristic) and estimates the excess post-intervention crime change in high- versus low-crime precincts. Precinct and month fixed effects absorb stable precinct differences and city-wide temporal shocks, respectively; the DD coefficient captures the differential crime trajectory associated with pre-existing crime level. We further examine outcomes across four policy regimes to trace how crime and collisions moved as pursuit intensity escalated, stabilized, and was subsequently re-restricted. The multi-regime design provides a built-in falsification test: if pursuit escalation deters crime through a credible apprehension threat, then the February 2025 re-restriction should be associated with crime increases.

Background

Vehicle Pursuit Policy in American Policing

Police vehicle pursuits have been a subject of scholarly and legal scrutiny since the mid-1980s, when Alpert and Anderson (1986) described the police vehicle as “the most deadly force” in an officer’s arsenal. Despite the risks, pursuit policy remained largely unregulated through the 1970s and early 1980s, with most departments permitting officers broad discretion to initiate and continue chases (Kennedy et al. 1992). The first wave of reform came in response to mounting civil liability under 42 U.S.C. § 1983, which exposed departments to lawsuits when pursuits resulted in injuries to bystanders or suspects (Alpert et al. 1997a; O’Connor and Norse 2005).

By the late 1990s, pursuit policies had coalesced around three general types (Alpert et al. 1997b; Kenney and Alpert 1997). *Judgmental* policies grant officers wide discretion to engage in pursuits based on their assessment of the situation. *Restrictive* policies limit pursuits to specific offense categories (typically violent felonies) and impose conditions related to speed, traffic density, and weather. *Discouragement* policies formally permit pursuits but establish a strong institutional presumption against them, coupled with mandatory supervisory review and potential discipline. A 1997 national survey of 436 agencies found that 48% had revised pursuit policies within the preceding two years (87% toward greater restrictions), yet variation persisted: some agencies reported zero annual pursuits, others over 800 (Kenney and Alpert 1997).

The empirical literature is consistent: the typical pursuit involves a young male driver fleeing a traffic or stolen-vehicle stop, occurs at night on local roads, lasts under five minutes, and ends in suspect termination or a crash (Alpert and Dunham 1988, 1989; Crew and Hart 1999; Hoffmann and Mazerolle 2005). About 20–40% end in a crash, 12–41% result in injury, and roughly 1% in a fatality (Alpert and Dunham 1988; Kenney and Alpert 1997). Between 1994 and 2002, Rivara and Mack (2004) documented 2,654 fatal pursuit-related crashes in the United States (an average of 295 per year) resulting in 3,146 deaths. Of these, 1,088 (35%) were bystanders or occupants of uninvolved vehicles.

The most recent comprehensive practitioner synthesis, produced jointly by the Police Executive Research Forum, the National Highway Traffic Safety Administration, and the U.S. Department of Justice’s Office of Community Oriented Policing Services (Police Executive Research Forum 2023), distills these decades of findings into 65 recommendations organized around a central principle: pursuit policy should balance the risk of immediate apprehension failure against the danger to officers, suspects, and uninvolved bystanders. PERF’s (2023) recommended default standard — restricting pursuits to violent crimes and situations presenting an imminent threat to public safety — mirrors the restrictive policy model that emerged from the liability-driven reforms of the 1990s but grounds it explicitly in a risk-management framework.

The offense triggering a pursuit rarely justifies the risks incurred. Hoffmann and Mazerolle (2005) found that “very few apprehended drivers were charged with crimes more serious than what was known at the time the pursuit was initiated,” yet Alpert and Dunham (1988) showed that while many pursuits begin with traffic violations, a non-trivial share of apprehended suspects are subsequently charged with serious felonies unrelated to the pursuit itself. This suggests that pursuits sometimes uncover criminality that would not otherwise have been detected, albeit at substantial cost.

Theoretical Framework

The case for expanded pursuit authority rests primarily on deterrence theory. General deterrence holds that the credible threat of apprehension discourages potential offenders from fleeing and, more broadly, from committing crimes that might precipitate police contact (Senese and Lucadamo 1996). Specific deterrence operates through incapacitation: apprehended suspects are removed from the crime-committing population through arrest and detention.

Though these mechanisms seem straightforward, there are at least three potential complications that must be considered. First, deterrence through pursuit requires that potential offenders are aware of the prevailing policy and update their behavior accordingly. While formal policy changes may receive limited public attention, a dramatic operational shift, such as a visible, sustained increase in vehicle pursuits, may communicate enforcement intensity through experiential channels even without formal announcements (Homant et al. 1994). Second, the deterrence benefit of pursuits depends critically on the apprehension rate. If pursued suspects frequently escape, as they do in roughly 30% of cases (Alpert and Dunham 1988), the deterrence signal is weakened. Third, deterrence theory assumes rational calculation by offenders, yet the population of fleeing drivers is disproportionately young, male, intoxicated, and unlicensed (Rivara and Mack 2004), characteristics associated with impaired risk assessment. Research on perceptual deterrence provides more specific theoretical mechanisms for this limitation. Pogarsky et al. (2017) demonstrated that potential offenders rely on cognitive heuristics (e.g., anchoring, availability, and question substitution) rather than coherent probability estimates, especially under time pressure. Barnum et al. (2021) further showed that person-specific cognitive biases can decouple perceived apprehension risk from actual enforcement intensity, a decoupling most acute for the young, impulsive, intoxicated population disproportionately represented among fleeing drivers.

The argument to constrain officer discretion to pursue emphasizes the potential for collateral harms. Pursuits impose externalities on third parties who had no involvement in the precipitating offense. Rivara and Mack (2004) found that 30% of pursuit fatalities were occupants of vehicles entirely uninvolved in the pursuit, and an additional 3% were pedestrians or cyclists. Hutson et al. (2007) documented similar patterns in a review of pursuit fatalities from 1982 through 2004. The public has expressed ambivalence about this tradeoff: MacDonald and Alpert (1998) found that while citizens support police pursuits in principle, support dwindles when pursuits involve minor offenses or result in bystander injuries. Indeed, in a recent survey experiment, Mourtgos et al. (2026) demonstrated that public judgments of pursuit appropriateness are sensitive to both the severity of the triggering offense and the risk to bystanders.

Prior Research on Pursuit Policy Effects

Rigorous evaluation of pursuit *policy-level* effects remains scarce. The most frequently cited evidence comes from a pair of natural experiments in the 1990s. When Dade County, Florida restricted pursuits to violent felons in 1992, pursuit frequency dropped by 82% and pursuit-related injuries decreased correspondingly, with no observed increase in crime rates or the number of suspects fleeing from police (Alpert and Madden 1994; Hoffmann and Mazerolle 2005). Conversely, when Omaha, Nebraska relaxed its pursuit policy in 1993, pursuit frequency increased by 600% (Alpert et al. 1997a). These cases are suggestive, but they predate modern causal inference methods and lack rigorous comparison groups.

Hicks (2006) surveyed state-level pursuit policies and documented substantial variation in restrictiveness but did not link policy variation to outcomes. Christie (2020) reviewed pursuit management strategies across multiple jurisdictions and concluded that “the available evidence suggests that restrictive policies reduce pursuit-related harm without measurably increasing crime,” but acknowledged that the evidence base was largely descriptive. Wade (2015) developed a pursuit risk classification system but focused on individual pursuit characteristics rather than policy-level effects. Crew and Hart (1999) attempted a cost-benefit framework for pursuit decisions but lacked the outcome data to conduct a full welfare analysis. Alpert and Lum (2014) recently synthesized the policy-evaluation literature and arrived at a sobering conclusion: despite decades of research documenting pursuit characteristics and outcomes, rigorous causal evidence linking pursuit policy changes to downstream crime rates remains scarce.

To be sure, the consistent pattern across these studies (i.e., that restrictive policies reduce pursuit frequency and pursuit-related harm) constitutes a meaningful signal; meanwhile, the absence of crime spikes in those same settings is suggestive but falls short of causal. Prior work cannot distinguish a policy effect from a coincident secular crime decline. Recent evaluations have extended this evidence base with more systematic documentation and, in some cases, pre-post comparative designs. The Policing Project at New York University School of Law et al. (2021) conducted a formal cost-benefit analysis of policy changes in two Virginia jurisdictions, finding that Roanoke County’s 2013 adoption of a restrictive policy reduced pursuit frequency and harm without measurable crime increases, while arrest rates for serious offenses were largely unchanged. Hosking (2021) examined two waves of reform in Queensland, Australia, documenting that restrictive policies reduced pursuit-related injuries without generating the crime consequences that critics predicted, though initial officer resistance partially attenuated implementation. Butler (2018) reported that Milwaukee's 2017 policy liberalization produced a 155% surge in pursuits the following year, with rising pursuit-related collisions accompanying the increase. Rivara et al. (2025) recently documented pursuit outcomes from 64 law enforcement agencies in Washington State from 2019 to 2024, finding persistently high rates of pursuit initiation for traffic violations and substantial injury risk under relatively permissive enforcement standards. Aronie and Alpert (2020) synthesized the cumulative evidence to argue that the data "strongly favor a restrictive pursuit policy," noting that approximately 91% of pursuits are initiated for non-violent offenses and roughly 30% result in a crash.

To our knowledge, this is the first study to apply modern causal identification methods — interrupted time series with robust standard errors and a precinct-level heterogeneous-effects DD design — to evaluate the downstream consequences of a major U.S. pursuit policy change. The NYPD's multi-regime structure, spanning escalation, stabilization, and re-restriction within three years, provides an unusually rich setting.

The NYPD Policy Change

The NYPD's pursuit escalation unfolded in three distinct phases. The first phase began in late 2022, coinciding with the appointment of John Chell as Chief of Patrol in December 2022 and Mayor Adams's broader quality-of-life enforcement mandate. Monthly pursuit events increased from 32 in November 2022 to 53 in December, then to 133 in January 2023 — more than all of 2018 and 2019 combined. Crucially, no formal change to the department's written Patrol Guide procedure on vehicle pursuits (PG 212-39) accompanied this operational shift. Rather, the increase reflected an enforcement posture emphasizing aggressive pursuit of "ghost cars," ATVs, and vehicles with illegal plates. By the end of the first quarter of 2023, the NYPD recorded 304 pursuits, a roughly 600% increase over the same period the prior year. Chief Chell publicly defended the surge, stating: "People thinking they can take off on us? Those days are over" (Brosnan 2023).

The second phase began in August 2023, when Chief of Department Jeffrey Maddrey circulated an internal memo reminding officers to follow existing Patrol Guide guidelines and emphasizing that pursuits "must be terminated when risks outweigh the danger of delayed apprehension" (Gonen 2023). The memo enumerated risk factors (e.g., nature of offense, time of day, weather, population density) and stressed supervisory responsibility and possible discipline for noncompliance. However, it did not narrow the offense categories that could trigger a pursuit, functioning more as compliance reinforcement than policy reform. Pursuit volume remained elevated throughout 2024, with quarterly counts approaching 600.

The third phase began in February 2025, when Commissioner Jessica Tisch implemented the first structural rewrite of the pursuit procedure since the pursuit surge began. The new policy restricted pursuits to felonies and violent misdemeanors only, explicitly excluding traffic infractions, violations, and non-violent misdemeanors. Officers would not face discipline for good-faith pursuit termination, and heightened risk-assessment requirements were introduced for residential areas, schools, and playgrounds. The NYPD disclosed that 67% of 2024 pursuits began when drivers fled traffic stops, scenarios that would largely fail to meet the new pursuit criteria. Initial data showed a 66% decline in pursuits in the first two weeks following implementation.

Data and Methods

Data Sources

We draw on five administrative data series from NYC Open Data. *Calls for service* records (October 2018 to present) identify pursuit-related radio broadcasts, from which we aggregate monthly pursuit event counts. *Motor vehicle collision* records (2016 to present) report vehicle-level crash involvement; we collapse to the incident level and identify pursuit-related crashes using the pre-crash condition field, which designates "Police Pursuit" as the precipitating circumstance (an administrative classification applied by officers in the collision reporting system). *NYPD complaint data* (January 2014 to August 2025) provide monthly robbery complaint counts. *NYPD shooting incident data* (through September 2025) record victim-level shooting events; we deduplicate to the incident level using unique incident keys. *Shots-fired complaints* (beginning January 2018) capture incidents of gunfire in which no victim was struck. Table 1 reports annual totals for each outcome across the full study period.

Measures

All variables are measured at the monthly level. *Vehicle pursuit events* reflect counts of pursuit-related calls for service broadcasts, capturing the treatment itself. *Pursuit-related motor vehicle collisions* capture incidents involving vehicles in pursuit situations. *Robberies* are drawn from complaint-level robbery reports. *Shooting incidents* are from victim-level records. Finally, *total gun violence events* reflect the sum of shooting incidents and shots-fired reports. We focus on robbery and shooting incidents as our primary crime outcomes following both empirical convention and practitioner emphasis. NYPD Chief of Patrol John Chell explicitly linked expanded pursuit enforcement to these outcomes, stating that illegal vehicles "cause street violence — robberies and shootings" (Brosnan 2023).

Analytic Strategy

We face two distinct inferential problems that require different designs. For pursuit-related collisions, the causal pathway from pursuit volume to crash risk is direct and theoretically unambiguous — no external comparison group is needed, and ITS provides well-identified estimates of the collision burden. For crime outcomes, ITS describes the within-NYC trajectory before and after the intervention and provides the upper bound on potential deterrence benefits in our cost-benefit framework (under the assumption that all observed crime changes are causally attributable to the policy), but it cannot isolate the policy effect from national secular crime trends. Our primary causal identification strategy for crime uses all 77 NYPD precincts as the unit of analysis. Note that the 105th Precinct was split into two (the 105th and the 116th) in December 2024; for this analysis, we treated it as single unit throughout the study period. Precinct fixed effects absorb stable cross-sectional differences in crime, time fixed effects absorb city-wide shocks, and pre-existing crime levels provide the identifying contrast. This design classifies precincts by their pre-October 2022 mean crime rate — predetermined and exogenous — and estimates the excess post-intervention crime change in high- versus low-crime precincts (SEs clustered at the precinct level, G = 77). Borough-level ITS provides supporting within-city heterogeneity evidence, and the four-regime event study tests temporal consistency across all policy phases. The February 2025 Tisch re-restriction serves as the paper's strongest internal falsification test of the deterrence mechanism.

Interrupted Time Series (ITS)

The segmented regression model is estimated on monthly data from January 2018 through September 2025:

$$Y_t = \beta_0 + \beta_1 T + \beta_2 D_t + \beta_3 P_t + \sum_{m=2}^{12} \gamma_m \,\text{Month}_m + \delta \text{COVID}_t + \varepsilon_t$$

where T is a linear time trend (months from start), D_t is a binary indicator equal to 1 for months on or after October 2022, P_t counts months elapsed since the intervention (0 in the pre-period), Month_m are seasonal fixed effects, and COVID_t flags the pandemic period (March 2020 through June 2021). The coefficient β_2 captures the immediate level shift at the intervention, while β_3 captures any change in the post-intervention slope. We estimate Newey-West heteroskedasticity- and autocorrelation-consistent (HAC) standard errors with a lag of 6 months, computed via `NeweyWest(model, lag = 6, prewhite = FALSE)` in R's `sandwich` package.

Borough-Level Interrupted Time Series

To examine within-NYC heterogeneity in crime responses, we estimate the same segmented regression model separately for each of the five NYC boroughs:

$$Y_{bt} = \beta_0 + \beta_1 T + \beta_2 D_t + \beta_3 P_t + \sum_{m=2}^{12} \gamma_m \,\text{Month}_m + \delta \text{COVID}_t + \varepsilon_{bt}$$

where subscripts b and t denote borough and month, respectively. The specification is identical to the city-level ITS model, and Newey-West HAC standard errors are computed with lag = 6 throughout for comparability. This provides a heterogeneity test of deterrence: if apprehension threat is the mechanism, boroughs with larger pursuit escalations should show larger crime reductions. Because all boroughs share the same national trends, within-city heterogeneity correlated with pursuit intensity is not confounded by city-level forces.

Precinct-Level Panel Analysis

The policy change created city-wide variation in pursuit exposure at the precinct level: when the department shifted enforcement posture in October 2022, which precincts actually used the new pursuit authority — and how intensively — depended on local crime patterns and operational circumstances. We exploit this cross-sectional variation through a heterogeneous-effects DD design using a balanced panel of 77 NYPD precincts over 93 months (January 2018–September 2025).

Heterogeneous-effects DD design. We classify precincts into high- and low-crime groups based on their mean monthly crime rate during the pre-October 2022 period (January 2018–September 2022), split at the median. This classification is constructed entirely from pre-intervention data and is therefore exogenous to the policy change. Critically, using post-treatment pursuit counts as the grouping variable would be endogenous: pursuit intensity itself is an outcome of the policy, and precincts with higher crime would be expected to generate more pursuits under any enforcement posture. By fixing the grouping on pre-intervention crime rates, we ensure that the identifying contrast cannot be contaminated by post-treatment changes in pursuit behavior or crime dynamics. Each outcome uses its own pre-treatment baseline: the shooting analysis groups precincts by pre-treatment shooting rate (median: 0.73 incidents/month) and the robbery analysis by pre-treatment robbery rate (median: 13.1 incidents/month). A supplementary quartile dose-response specification assigns precincts to quartiles Q1–Q4 and tests whether the post-treatment crime change varies linearly across the quartile ranking. The intervention date is October 2022 for all precincts — a common break reflecting the citywide policy change.

The DD scalar estimate is:

$$Y_{pt} = \alpha_p + \lambda_t + \beta \cdot (\text{Post}_t \times \text{HighCrime}_p) + \varepsilon_{pt}$$

where α_p are precinct fixed effects, λ_t are month-year fixed effects, and β is the DD estimate: the excess post-intervention change in crime for high-crime precincts relative to low-crime precincts. Standard errors are clustered at the precinct level (G = 77).

The DD event study specification replaces the scalar interaction with a set of relative-time interactions:

$$Y_{pt} = \alpha_p + \lambda_t + \sum_{k \neq -1} \delta_k \cdot \mathbf{1}[t - t_0 = k] \times \text{HighCrime}_p + \varepsilon_{pt}$$

where $t_0 =$ October 2022 and k indexes quarterly bins relative to the intervention — each bin aggregates three consecutive calendar months — referenced at $k = -1$ (July–September 2022). The Tisch reversal period (January 2025 onward, with the policy officially taking effect in February 2025) is pooled into a single lump bin ($k = 9$). Pre-period δ_k

coefficients near zero would confirm that high- and low-crime precincts were on parallel crime trajectories before the policy change — the key identifying assumption for this design. Post-period δ_k coefficients trace the divergence between groups over time. Standard errors are clustered at the precinct level (G = 77) throughout.

Event Study / Multi-Period Analysis

We define four policy regimes to examine how outcomes evolved across distinct phases. The first is the *pre-escalation* period, which spans January 2018 to September 2022 and serves as our baseline. The second is the *escalation* period from October 2022 through July 2023, when pursuits surged. Third is the *post-Maddrey* phase, which begins with Chief of Department Jeffrey Maddrey's compliance memo in August 2023 and runs through December 2024. Finally, the *post-Tisch* phase begins with Commissioner Tisch's formal policy rewrite in February 2025 and runs through the end of our study period (September 2025).

Generalized Synthetic Control (Appendix B)

As a supplementary robustness check for the shootings outcome, we apply the generalized synthetic control method (Xu 2017) to a panel of large U.S. cities (Sharkey 2026). Unlike TWFE difference-in-differences, gsynth estimates an interactive fixed effects model that explicitly allows latent factors to load differently across units, accommodating the heterogeneous time trends that invalidated simpler national comparison designs for these data. NYC is the single treated unit; the donor pool excludes 24 cities with documented pursuit policy changes between 2018 and 2024 to avoid violations of the Stable Unit Treatment Value Assumption (SUTVA). The number of latent factors is selected by leave-one-out cross-validation. Inference uses 500 parametric bootstrap iterations. Results are reported in Appendix B.

Results

Pursuit Activity

Monthly pursuit events increased sharply following the department's enforcement posture shift in late 2022. In the pre-escalation period (January 2018–September 2022), the NYPD averaged 7.7 pursuit events per month. Pursuit volume rose to a mean of 87.9 during the initial escalation phase (October 2022–July 2023) and continued climbing to 159.3 per month during the Post-Maddrey period (August 2023–January 2025), as the August 2023 compliance memo reinforced existing procedures without restricting eligible offenses. Annually, pursuit events totaled 1,592 in 2023 and 1,930 in 2024, compared with 214 in 2022 (Table 1). The ITS model estimates an immediate level shift of b_2 = 120.9 events per month (SE = 37.8, p = .00201).

The February 2025 Tisch re-restriction reversed the escalation. Monthly pursuits fell from a pre-Tisch mean of 133.8 to 38.4 (-71.3%), consistent with the department's disclosure that approximately two-thirds of 2024 pursuits — most initiated from fled car stops — would not qualify under the new felony-and-violent-misdemeanor threshold. Figure 1 illustrates pursuit volume across all four policy regimes.

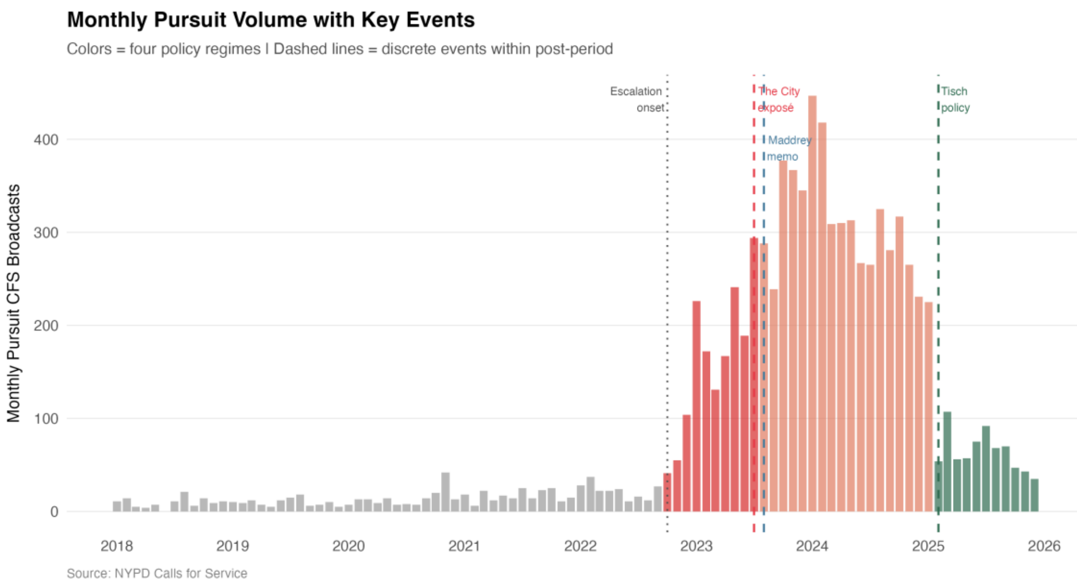

Figure 1: *Monthly NYPD pursuit events across four policy regimes: Pre-escalation (January 2018–September 2022), Escalation (October 2022–July 2023), Post-Maddrey (August 2023–January 2025), and Post-Tisch (February 2025–September 2025). Vertical dashed lines indicate intervention dates.*

Table 1 provides annual totals for each outcome across the full study period, offering a pre/post summary of how the key indicators changed alongside pursuit escalation.

Table 1: Annual Summary of Key Outcomes, 2018–2025

Year	Pursuits	Δ Pursuits	Crashes	Δ Crashes	Robberies	Δ Robberies	Shootings	Δ Shootings	Gun Violence	Δ Gun Violence
2018	61	—	58	NA%	12,965	NA%	754	—	2,017	—
2019	64	+4.9%	73	+25.9%	13,434	+3.6%	777	+3.1%	2,324	+15.2%
2020	99	+54.7%	69	-5.5%	13,187	-1.8%	1,531	+97%	4,478	+92.7%
2021	109	+10.1%	62	-10.1%	13,860	+5.1%	1,562	+2%	5,028	+12.3%
2022	214	+96.3%	84	+35.5%	17,439	+25.8%	1,294	-17.2%	4,162	-17.2%
2023	1,592	+643.9%	318	+278.6%	16,920	-3%	974	-24.7%	3,109	-25.3%
2024	1,930	+21.2%	421	+32.4%	16,566	-2.1%	904	-7.2%	2,917	-6.2%
2025*	424	-78%	89	-78.9%	11,366	-31.4%	553	-38.8%	2,056	-29.5%

*2025 data through September.

Pursuit-Related Collisions

Pursuit-related collisions increased substantially following the escalation, representing the most direct and causally interpretable impact of the policy change (see Table 2). The ITS model estimates an immediate level shift of b_2 = 22.6 additional collisions per month (SE = 8.45, p = .00924), with no significant change in the post-intervention slope (p = .583). Prior to the escalation, the NYPD averaged 5.5 pursuit-related collisions per month; this rose to 34.4 per month during the Post-Maddrey period — a more than fivefold increase in total crash burden. Although the per-pursuit crash rate fell from 71.5% to approximately 21.6% as lower-risk pursuits were added, the absolute collision toll increased substantially given the scale of the volume increase. The pre-escalation per-pursuit crash rate may also partly reflect administrative underreporting in earlier periods when pursuit tracking was less systematic, which would make the observed decline appear steeper than the underlying behavioral shift.

The per-pursuit rate decline reflects compositional shift, not reduced harm. Pre-escalation pursuits were heavily selected — only situations serious enough to overcome institutional norms — and consequently ended in collisions at an elevated rate. As volume expanded after October 2022, the marginal pursuit increasingly involved lower-risk situations: fled traffic stops, ATV operators, vehicles with illegal plates. These added pursuits carry lower per-event crash probability, pulling the average rate down. The policy question turns on aggregate harm: a rate of 21.6% applied to roughly 5,500 Post-Maddrey pursuits generates far more collisions than 71.5% applied to approximately 800 pre-escalation pursuits.

The multi-period event study corroborates this pattern closely. Collision counts trended upward throughout the escalation phase (b = 1.58 per month, p = .00486), jumped at the start of the Post-Maddrey period (b = 17.1, p = <.001), and declined sharply following the Tisch re-restriction (b = -15.2, p = <.001). This close correspondence across all four policy regimes is consistent with a direct causal mechanism.

Crime Outcomes

The central question is whether pursuit escalation reduced crime through deterrence. We address this through three complementary within-NYC designs: city-level ITS (describing the crime trajectory), borough-level ITS (testing whether boroughs with larger escalations show larger reductions), and the precinct-level DD (primary causal identification). Across all designs, the evidence for deterrence is weak.

City-Level Interrupted Time Series

Within-NYC crime trends changed substantially after October 2022: during the pre-intervention window, shooting incidents were trending upward at b_1 = 1.3 incidents per month (p = <.001); the counterfactual projection continues this upward trajectory, so the post-intervention estimates reflect deviations from an already-rising baseline. Shooting incidents declined relative to this counterfactual: the ITS estimates an immediate level shift of b_2 = -33.2 incidents per month (SE = 9.25, p = <.001), with a sustained slope decline of b_3 = -2.32 per month thereafter (p = <.001). Aggregate gun violence (shootings plus shots-fired reports) declined similarly — an immediate level shift of b_2 = -137.4 (p = <.001) and a sustained slope of b_3 = -7.1 per month (p = <.001).

Table 2: Interrupted Time Series Estimates (Newey-West HAC Standard Errors, Lag = 6)

Outcome	Term	Estimate	95% CI	p-value
Pursuits	Level shift	120.9**	[46.8, 195]	.00201
Pursuits	Slope change	-1.3	[-4.8, 2.2]	.47191
Collisions	Level shift	22.6**	[6, 39.1]	.00924
Collisions	Slope change	-0.2	[-1.1, 0.6]	.58345
Robberies	Level shift	-26.3	[-146.2, 93.6]	.66887
Robberies	Slope change	-12.2***	[-16.3, -8]	< .001
Shootings	Level shift	-33.2***	[-51.4, -15.1]	< .001
Shootings	Slope change	-2.3***	[-2.9, -1.7]	< .001
Gun Violence	Level shift	-137.4***	[-205, -69.8]	< .001
Gun Violence	Slope change	-7.1***	[-9.1, -5.2]	< .001

Robbery showed no immediate change (b_2 = -26.3, p = .669), but the post-intervention slope turned sharply negative (b_3 = -12.2 per month, p = <.001), indicating a gradual trajectory shift. Intervention date sensitivity analyses confirm that the key estimates are robust to shifting the treatment date by ±2 months. Pre-period placebo tests (within-NYC temporal placebos) produce no comparable signals, ruling out spurious detection within the analysis window; these tests do not, however, rule out national crime trend confounding. Three concurrent confounders are most plausible: the April 2023 NYS bail reform amendments that restored pretrial detention authority for certain offenses, post-COVID crime normalization following NYC's 2021 shooting peak, and NYPD strategy shifts across the Sewell, Maddrey, and Tisch administrations. The ITS crime estimates should therefore be read as an upper bound on potential deterrence benefits rather than causal effect estimates; the generalized synthetic control analysis (Appendix B) provides the most direct test of whether the ITS crime decline outpaced national trends, and it does not. Figure 2 displays the segmented regression fits for all four outcomes.

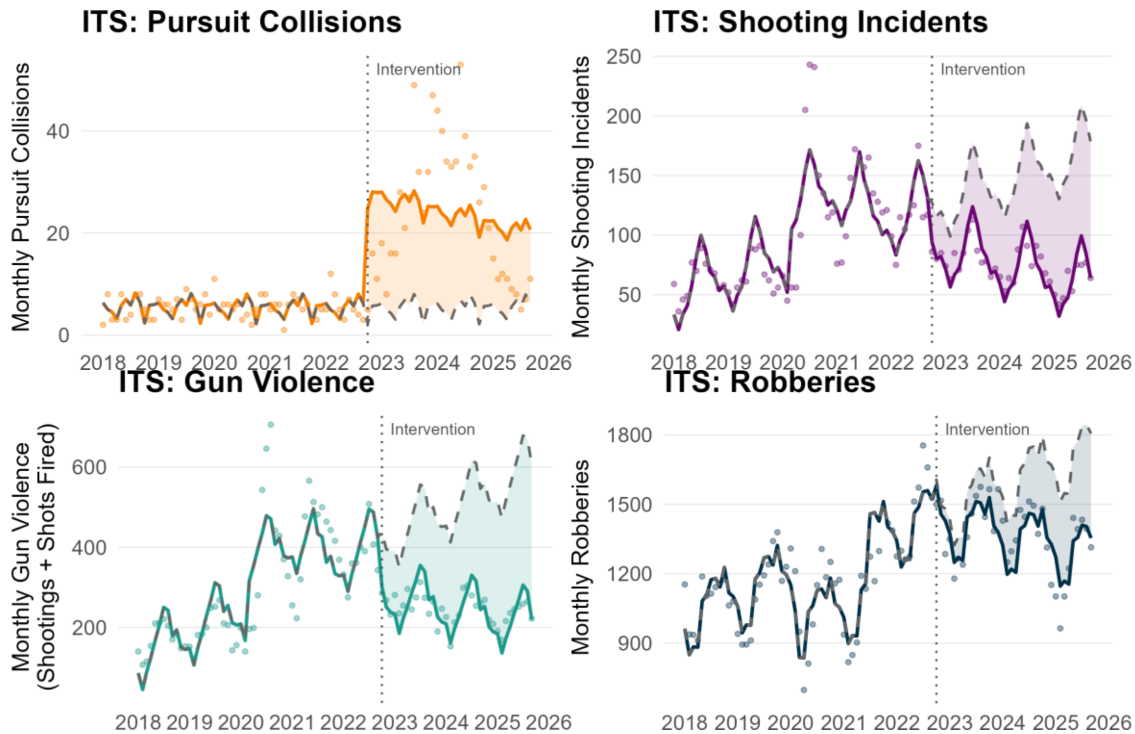

Figure 2: *Interrupted time series segmented regression fits for four outcomes, January 2018–September 2025. Top row: pursuit-related collisions (left) and shooting incidents (right). Bottom row: aggregate gun violence — shootings plus shots-fired (left) and robberies (right). Shaded bands show the estimated post-intervention effect (fitted minus counterfactual). Vertical dashed lines mark the October 2022 intervention. Newey-West HAC standard errors (lag = 6). Points show monthly observed counts.*

Borough-Level Interrupted Time Series

Estimating the ITS model separately for each of the five NYC boroughs reveals substantial heterogeneity in crime responses. For shooting incidents, all five boroughs show negative level shifts (b_2) at October 2022, but the magnitude varies considerably. The Bronx — which experienced the highest post-escalation pursuit volume relative to baseline (approximately a 21-fold increase) — shows the largest immediate reduction (b_2 = -13.1, SE = 3.4, p = <.001), while Staten Island — with the lowest pursuit escalation — shows a near-zero and non-significant reduction (b_2 = -0.5, SE = 0.5, p = .348). However, the relationship between pursuit intensity and crime reduction is not monotonically ordered: Queens recorded the highest raw pursuit volume in the post-escalation period yet achieved a shooting reduction (b_2 = -6.4, p = .00198) comparable to Manhattan (b_2 = -5.8, p = .00161), despite Manhattan having roughly one-third the pursuit activity. The slope-change coefficient (b_3) is negative for all five boroughs for shooting incidents, reflecting a sustained post-intervention decline.

For robbery, the pattern is more heterogeneous and does not track pursuit intensity: Manhattan shows the largest negative level shift (b_2 = -43, SE = 8.6, p = <.001) despite moderate pursuit activity, while Staten Island — the lowest-pursuit borough — registers a statistically significant *positive* level shift (b_2 = 8.4, SE = 2.1, p = <.001). Neither the level-shift pattern nor the slope-change pattern for robbery is consistent with a deterrence mechanism operating through borough-level pursuit intensity.

Figure 3 displays the borough-level ITS fits for all four outcomes across all five boroughs. The 5 × 4 grid (rows = boroughs, columns = outcomes) reports the immediate level-shift coefficient (b_2) and Newey-West p-value in each panel.

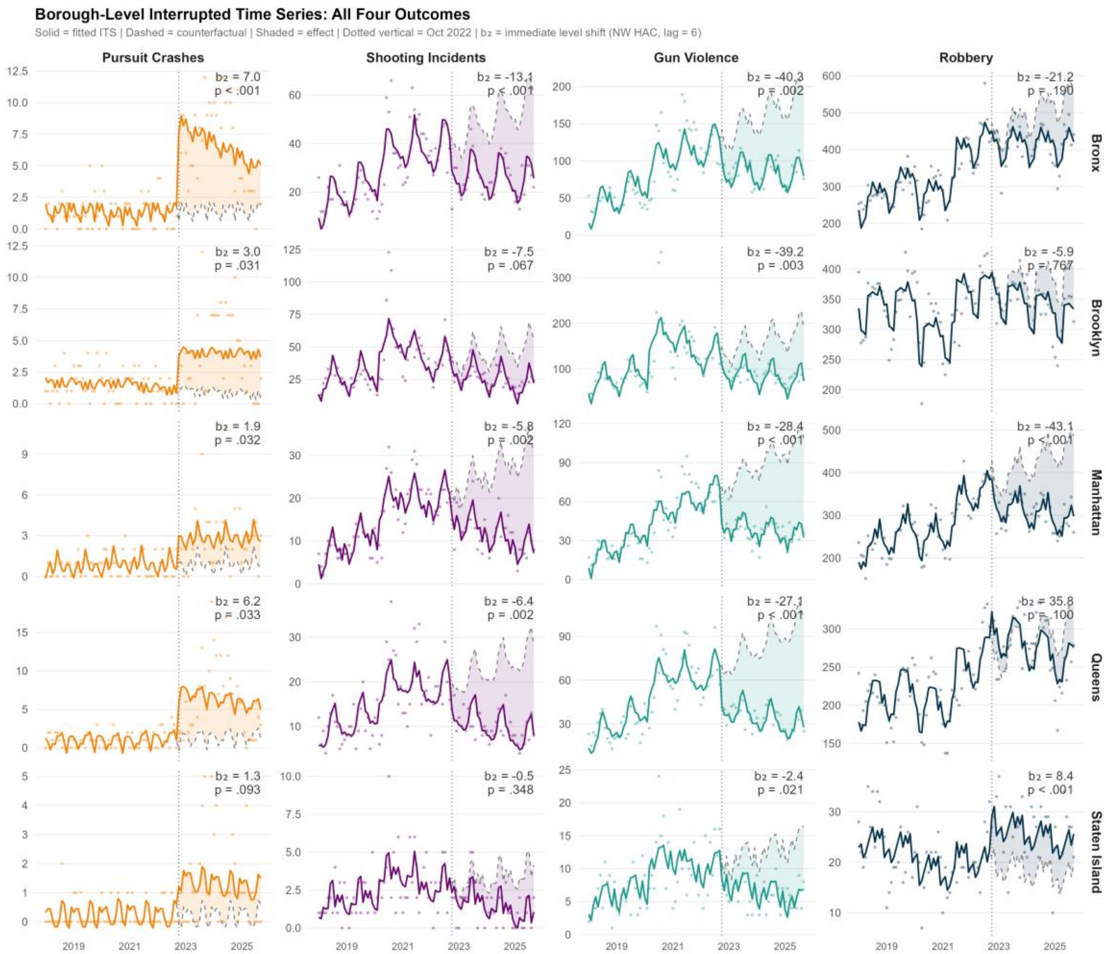

Figure 3: *Borough-level interrupted time series fits for all four outcomes, January 2018–September 2025. Rows correspond to boroughs (Bronx, Brooklyn, Manhattan, Queens, Staten Island); columns correspond to outcomes (pursuit-related collisions, shooting incidents, gun violence, robbery). Each panel shows the observed monthly counts (points), segmented regression fit (solid line), counterfactual projection from the pre-intervention trend (dashed line), and shaded post-intervention effect band. Each panel's upper right corner reports the immediate level-shift estimate b_2 and its Newey-West p-value (lag = 6). Vertical dotted line marks October 2022.*

This cross-borough heterogeneity underscores the contrast: collision outcomes track pursuit volume closely, while crime reductions do not scale with borough-level pursuit intensity. If apprehension threat were the mechanism, the Bronx (~21-fold pursuit increase) should show substantially larger reductions than Manhattan (~7-fold), but the gap in crime reductions does not match the gap in escalation. This motivates the precinct-level DD, which tests the same question at finer resolution.

Precinct-Level Analysis

The 77-precinct panel is the study's sharpest identification lens. Each outcome uses its own pre-treatment median to split precincts into high and low groups (Figure 4); the split is fixed in the pre-intervention window (January 2018–September 2022) and cannot be contaminated by post-treatment dynamics.

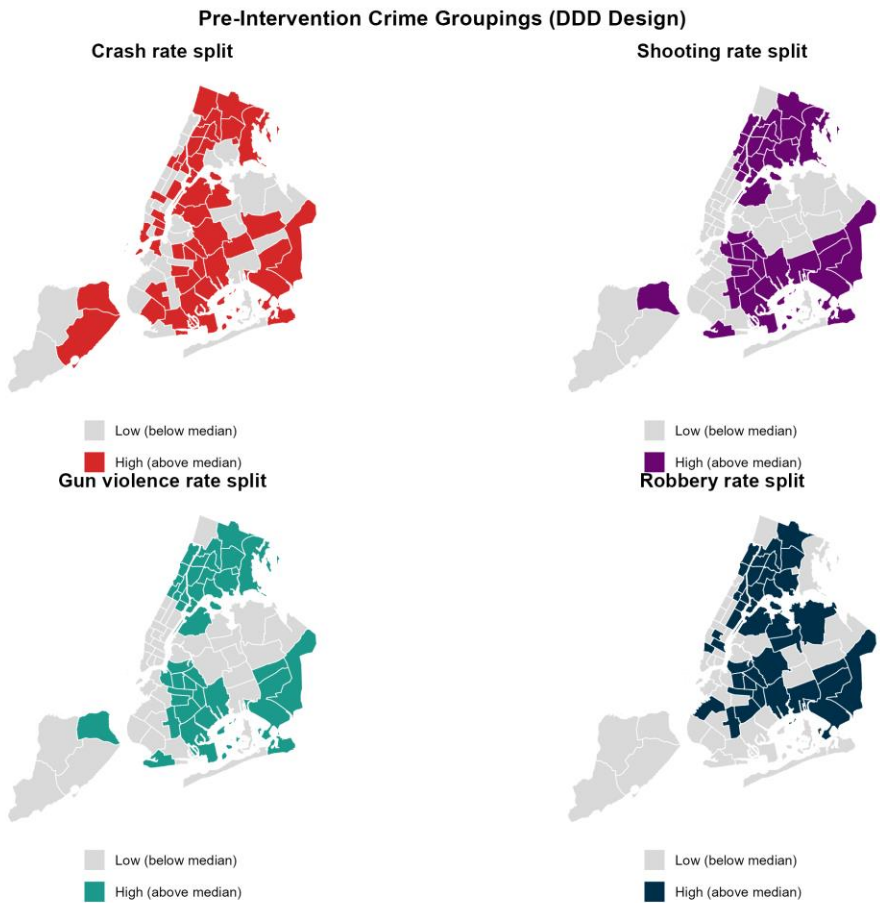

Figure 4: *Pre-intervention groupings for the heterogeneous-effects DD design across all four outcomes. Top row: pursuit-related collisions (left) and shooting incidents (right); bottom row: gun violence (left) and robberies (right). Dark color = high pre-treatment rate (above median); light gray = low pre-treatment rate. Each outcome uses its own outcome-specific grouping variable. The panel covers all 77 NYPD precincts (G = 77 clusters in the DD model). The gun violence and shooting groupings assign the majority of precincts to the same category, reflecting the strong within-precinct correlation between shooting incidents and shots-fired reports.*

Heterogeneous-Effects DD Results. The DD scalar estimate compares the post-October 2022 change in crime between high- and low-crime precincts, net of all precinct-level baselines and city-wide monthly shocks. For shooting incidents, the DD estimate is b = -0.559 (SE = 0.076, p = <.001), indicating that high-crime precincts experienced nominally larger shooting reductions than low-crime precincts in the post-intervention period. For robbery, the DD estimate is b = 2.39 (SE = 0.94, p = .0134), a positive and statistically significant result directly inconsistent with deterrence. Notably, even if high-crime precincts were experiencing faster secular crime declines than low-crime precincts — the most plausible confound, given post-pandemic mean-reversion — this would bias the DD coefficient *downward* (toward more negative values) and the robbery coefficient *downward* (toward zero or negative values). The fact that the robbery DD estimate is positive and significant despite this expected bias direction makes it especially difficult to explain as a measurement artifact; it is the opposite direction from what deterrence theory predicts. The quartile dose-response specifications confirm these patterns: each one-quartile step up in pre-treatment shooting rate is associated with b = -0.279 additional shooting reduction per month (p = <.001), while each robbery-quartile step is associated with b = 1.45 additional robberies per month (p = .0013). For collisions, the highest-crash-quartile precincts show a positive and significant post-intervention DD (b = 0.246, p = .00259), consistent with the expected mechanism: precincts that already had higher collision rates saw the largest absolute increases as pursuit volume expanded department-wide. This collision quartile gradient is directionally informative, not anomalous — it reflects higher pursuit activity in precincts that had already demonstrated higher crash risk under prior pursuit procedures. The positive robbery DD estimate should be read alongside the city-wide ITS slope: the ITS captures a genuine NYC-wide robbery decline, but this design shows this decline was not concentrated in higher-crime precincts as a deterrence mechanism would predict — if anything, the cross-precinct gradient ran the wrong way.

The event study plots (Figure 5) trace the quarterly DD coefficients — the evolving difference between high- and low-crime precincts relative to the quarter before intervention (bin −1 = July–September 2022) — for all four outcomes. Pre-period coefficients reveal that the parallel trends assumption is not satisfied for shooting incidents or robbery: 57% of pre-period shooting coefficients and 71% of pre-period robbery coefficients are statistically distinguishable from zero. Pre-period divergence for gun violence follows a pattern similar to shootings, reflecting the strong correlation between the two outcomes at the precinct level. For crashes, the pre-period coefficients are generally flat, consistent with parallel trends between high- and low-crash precincts before October 2022. This pre-period divergence — evident before any pursuit policy change for crime outcomes — indicates that high- and low-crime precincts were already on non-parallel trajectories, most likely reflecting differential rates of secular crime decline. The DD crime estimates therefore cannot be given a causal interpretation under the standard parallel trends assumption.

What accounts for the nominally negative shooting DD estimate, then? The more plausible explanation is that high-crime precincts experienced faster mean-reversion in the post-pandemic crime cycle than low-crime precincts — not a deterrence effect of pursuit escalation. For both crime outcomes, the pre-period divergence runs in the same direction as a deterrence prediction would, making the nominally negative shooting DD estimate uninformative as evidence for deterrence. The robbery DD estimate is positive and significant despite this, which runs in the opposite direction from deterrence regardless of the parallel trends issue.

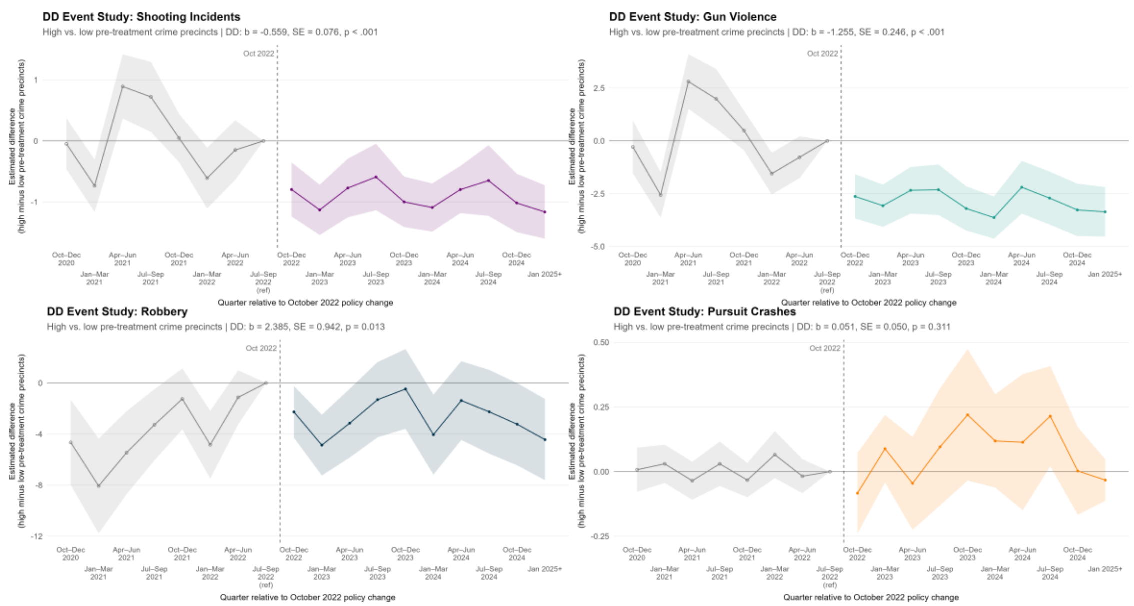

Figure 5: *Precinct-level heterogeneous-effects DD event study, all four outcomes. Top row: shooting incidents (left) and gun violence (right); bottom row: robbery (left) and pursuit-related collisions (right). X-axis shows three-month (quarterly) bins relative to the October 2022 policy change; bin −1 (July–September 2022) is the reference quarter. Points show the estimated DD coefficient for each quarterly bin. The Maddrey compliance memo falls in bin 3 (July–September 2023); bin 9 covers January 2025 onward; the Tisch policy rewrite officially took effect in February 2025 but is included in this lump bin as she announced it in mid-January. Pre-period coefficients reveal non-parallel trends for crime outcomes before the intervention, undermining the parallel trends assumption required for causal interpretation; collision pre-period coefficients are consistent with parallel trends. Standard errors clustered by precinct (G = 77). Note for the pursuit-related collisions panel: high-crash precincts are expected to exhibit larger collision increases after the policy change (reflecting higher pursuit activity), not crime reductions; the directional prediction for collisions differs from the crime outcome panels.*

Tisch Falsification and the Multi-Period Event Study

The February 2025 Tisch re-restriction provides the strongest internal falsification test in this paper: it is a clean reduction in pursuit intensity within the same jurisdiction and institutional context, occurring well after the escalation period, allowing a direct comparison of crime trajectories before and after the reversal without the identification challenges that attend between-jurisdiction comparisons. We estimate the post-Tisch effects using a multi-period segmented regression model that extends the city-level ITS specification to accommodate four distinct slope regimes, with the Tisch indicator entering as both a level shift and a slope change from February 2025 onward. The Tisch reversal was publicly announced (New York City Police Department 2025), strengthening this falsification: a public reversal should accelerate decay of deterrence beliefs faster than a change communicated only through observed behavior — making any resulting crime upturn more detectable. The four-regime event study reveals heterogeneous trajectories. For collisions, the pattern tracks pursuit volume across all phases (as described above). For crime outcomes, the most informative comparison is the Tisch reversal: if pursuit escalation deters crime through credible apprehension threat, then the re-restriction that reduced monthly pursuits by -71.3% should produce a corresponding crime increase.

Table 3: Mean Monthly Outcomes by Policy Regime

Regime	Months	Pursuits	Crashes	Crash Rate[a]	Robberies	Shootings
Pre-escalation (Jan 2018–Sep 2022)	57	7.7	5.5	71.5%	1,168.2	99.4
Escalation (Oct 2022–Jul 2023)	10	87.9	16.6	18.9%	1,368.6	84.6
Post-Maddrey (Aug 2023–Jan 2025)	18	159.3	34.4	21.6%	1,400.1	74.1
Post-Tisch (Feb 2025–Sep 2025)	8	38.4	9.2	24.1%	1,283.0	62.9

[a]Pursuit-related crashes as a percentage of pursuit events.

The Tisch re-restriction did not increase crime (see Table 3). Following the February 2025 re-restriction, robberies declined -7.6% from their pre-Tisch mean — the opposite direction from what a deterrence model predicts. Shooting incidents fell -19.2%, and aggregate gun violence declined -6.8%. The multi-period model confirms that the post-Tisch shooting trajectory continued downward (b = -7.2 per month, p = .00603), while the post-Tisch level shift for shootings was not statistically distinguishable from zero (b = 8, p = .657).

Cost-Benefit Analysis

We conduct a partial cost-benefit analysis from a societal perspective — valuing all costs and benefits regardless of who bears them (see Table 4). This analysis is best understood as a sensitivity exercise that quantifies the minimum causal attribution share required for benefits to exceed costs, not as a claim that observed crime reductions equal the pursuit-deterrence benefit. The benefit numerator under full attribution ($\theta = 1$) uses the ITS-estimated crime trajectory as its upper bound: the within-NYC shooting decline provides the maximum plausible crime reduction that could be attributed to pursuit policy if the entire change were causally driven by the escalation. Costs comprise excess pursuit-related collision externalities: crashes above the pre-escalation monthly baseline, summed over the study period and valued using Federal Highway Administration unit cost schedules. Fatalities (estimated at 1.0% of pursuit-related crashes for the central scenario, with 0.5% and 1.5% as lower and upper bounds; Rivara and Mack (2004); Kenney and

Alpert (1997)) are valued at the value of a statistical life (VSL = \$12.5 million, consistent with current USDOT guidance); injuries (estimated at 35% of crashes resulting in injury, approximately 1.5 injuries per crash; Alpert and Dunham (1988); Kenney and Alpert (1997)) at a blended serious/minor rate; and property damage at actuarial cost. Three scenarios (conservative, central, high) vary the VSL by ±30%. Each prevented shooting incident is valued at approximately \$1.46 million (a victim-cost-based share of the VSL). We introduce an attribution share parameter $\theta \in [0,1]$ and report the break-even threshold θ^* as the decision-relevant summary.

Pursuit-related collision costs are substantial and directly attributable to the policy change. Across the post-escalation period, the model estimates approximately 667 excess pursuit-related collisions above the pre-escalation baseline, generating an estimated 6.7 additional fatalities and 350 additional injuries. Valued at central VSL assumptions, total collision costs amount to \$166,810,705, comprising approximately \$84 million in fatality costs, \$70 million in injury costs, and \$13 million in property damage. These estimates may be conservative: investigative reporting by Gonen (2025) documented approximately 20 fatalities directly linked to NYPD pursuits over the study period (a figure higher than the actuarial estimate used here), suggesting that the central-scenario fatality cost of approximately \$84 million is plausibly a lower bound.

Table 4: Cost-Benefit Analysis: Collision Costs and Break-Even Attribution by Scenario

Scenario	Crash Costs	Break-Even Attribution (%)[a]
Low	\$125,108,029	3.4%
Central	\$166,810,705	4.5%
High	\$208,513,381	5.7%

[a] Minimum causal attribution share at which crime benefits equal collision costs. Values vary across scenarios because crash costs scale with VSL assumptions while full-attribution crime benefits are fixed. Break-even threshold (central): ~119.7 prevented shootings (3.3% of pre-escalation monthly baseline).

The break-even threshold is low in absolute terms — approximately 120 prevented shootings over the study period, equivalent to a causal attribution share of approximately 4.5% of the ITS-estimated crime reduction under central VSL assumptions — yet the available evidence does not establish that even this minimal threshold is met through the causal mechanism of pursuit deterrence. Under full causal attribution — the theoretical upper bound — crime benefits would substantially exceed collision costs (central BCR: 22.1:1); but this assumption is directly contradicted by the convergent evidence above. The generalized synthetic control analysis (Appendix B) is the most direct evidence on this question: NYC's post-2022 shooting trajectory is statistically indistinguishable from the synthetic counterfactual, which means the ITS crime decline tracks what the national trend predicted anyway, placing the likely causal attribution share near zero. The null result does not prove that deterrence is entirely absent; rather, it establishes that the crime reduction required to break even (approximately three to four prevented shootings per month) *should be detectable* given the data resolution available. The absence of a detectable signal at this threshold is informative.

Discussion

Three findings emerge. First, pursuit-related collisions increased substantially and in direct proportion to pursuit volume. The ITS estimates an immediate increase of approximately 23 collisions per month at the intervention onset, and the multi-period event study confirms that collisions tracked pursuit activity closely across all four policy regimes — rising steadily during escalation, peaking in the Post-Maddrey period, and declining sharply following the Tisch re-restriction. The mechanism is direct and theoretically unambiguous: more pursuits create more collision risk for officers, suspects, and uninvolved road users. As documented in the Results, the per-pursuit crash rate fell during the escalation as lower-risk situations were added, but the sheer volume increase produced a more than fivefold increase in total monthly collision burden — a compositional shift that reshapes the per-event rate without reducing aggregate harm.

Second, three convergent lines of evidence cut against a deterrence interpretation. The precinct-level DD (our primary causal design for crime outcomes) groups precincts by their pre-intervention crime rate and compares shooting and robbery trajectories across groups after October 2022. Pre-period event study coefficients reveal non-parallel trends for both crime outcomes (57% of pre-period shooting coefficients and 71% of robbery coefficients distinguishable from zero), indicating that the design's identifying assumption is not met and that the nominally negative shooting DD estimate more likely reflects differential secular crime trajectories than a deterrence effect. More tellingly, the robbery DD estimate is positive and significant — a result that is difficult to attribute to confounding, since the most plausible confounder (faster crime decline in high-crime precincts due to post-pandemic mean-reversion) would bias the robbery coefficient *downward*, toward zero or negative. The fact that robbery moved in the wrong direction despite a bias that should work in deterrence's favor makes the positive coefficient especially probative. A national secular crime decline — documented across 40 large U.S. cities through 2025, including a 21% drop in homicide rates (Lopez and Boxerman 2026) — dominated NYC's aggregate trajectory, independent of pursuit policy.

Borough-level heterogeneity provides little additional support for a deterrence mechanism. If deterrence operates through local apprehension threat, boroughs with the largest pursuit escalation should show the largest crime reductions. The Bronx experienced the highest post-escalation pursuit-to-baseline ratio (approximately 21-fold) and did show the largest shooting reduction; however, the relationship is not monotonically ordered across boroughs. Queens recorded the highest raw pursuit volume but achieved crime reductions comparable to Manhattan, which had roughly one-third the pursuit activity. For robbery, the borough-level pattern does not track pursuit intensity at all. This absence of monotonic ordering cannot rule out borough-specific confounders, but it is consistent across both crime outcomes.

The Tisch falsification provides an independent within-sample test. The February 2025 re-restriction reduced monthly pursuits by -71.3% and was publicly announced in advance — conditions that, if deterrence were operating, should accelerate the decay of apprehension-risk perceptions among would-be offenders. No such crime increase is visible: the post-Tisch data show continued crime declines across all three outcomes (Table 3). The multi-period model finds no evidence of a crime upturn in the post-Tisch period. To be clear, this test is not immune to the same national-trend confound that limits the city-level ITS: continued secular crime decline through 2025 would also mask a deterrence-decay effect of similar magnitude. The Tisch falsification is most informative as evidence against a large, immediate crime upturn; it cannot rule out a modest effect obscured by the national trend.

Finally, the cost-benefit analysis reveals an asymmetry between what can and cannot be demonstrated. The collision costs are real, substantial (approximately $166,810,705 under central assumptions), and directly attributable to the policy change. The crime reduction benefits, while potentially large if causally real, have not been demonstrated. The break-even threshold (approximately 120 prevented shootings over the study period) is low but unmet: the available evidence is consistent with a causal share near zero.

Several limitations qualify these conclusions. The heterogeneous-effects DD design requires that high- and low-crime precincts would have followed parallel crime trajectories absent the policy change. We use pre-treatment crime rates as the grouping variable to avoid the endogeneity of post-treatment pursuit intensity as a classifier; this ensures the grouping variable is exogenous to the policy change and to post-treatment crime dynamics. However, the pre-period event study reveals non-parallel trends for both outcomes, indicating that the parallel trends assumption is not met and the shooting DD estimate most likely reflects pre-existing divergence rather than a treatment effect.

An additional, distinct concern is time-varying endogeneity: even if the pre-treatment grouping is fixed and exogenous, the number of pursuits per precinct may change within the post-treatment period as a function of contemporaneous crime (more crime generates more pursuit-eligible incidents), creating a feedback loop within the post-intervention window. Precinct fixed effects absorb any time-invariant component of group differences, but time-varying endogeneity within the post-treatment period remains a limitation of the design. The borough-level ITS assumes that heterogeneity in crime responses across boroughs reflects variation in pursuit intensity rather than other borough-specific factors; unmeasured borough characteristics (e.g., changes in community policing, socioeconomic shocks) could confound the within-NYC comparison. Borough-level pursuit counts also contain measurement error from cross-boundary pursuits — events that begin in one borough and conclude in another — which attenuates the precision of borough-specific exposure estimates. ITS sensitivity analyses confirm robustness to the intervention date within a ±2-month window, and temporal placebos yield no comparable signals. Finally, individual-level pursuit records linking specific pursuit outcomes to officer, suspect, and bystander characteristics are not publicly available, precluding finer-grained analysis of who bears the costs and who might reap the benefits of pursuit policy.

The usual concerns about administrative data also apply. Our pursuit activity measure captures only incidents that officers broadcast over the radio, as required by Patrol Guide procedure. If officers engaged in vehicle pursuits without broadcasting during the study period, those events would be absent from our counts, making our estimates a conservative lower bound on true pursuit volume. If ghosted pursuits were more common in the pre-escalation period, when command attention to pursuits was lower, our pre-period counts would be underestimates, making the escalation appear larger than it was (Pfeiffer and Alpert 2020). The Tisch falsification evidence, however, does not depend on counting pre-period pursuits and is therefore unaffected by this concern. We have no reason to believe that unreported pursuits would cluster systematically around the intervention dates in a way that would account for the dramatic swings we observe in the data.

The policy implications are asymmetric in the same direction as the evidence. Collision costs are demonstrably real; crime prevention benefits remain unclear. The break-even threshold is low enough in absolute terms — the equivalent of preventing roughly three to four shootings per month over the three-year study period — that a genuine deterrence effect, if it existed, should in principle be detectable. The DD evidence and the Tisch falsification both point away from this. The causal attribution framework developed here offers a portable decision template: compute the break-even causal share under local cost and crime parameters, assess whether local crime trends outperform a credible national or regional benchmark, and treat within-jurisdiction declines as presumptively attributable to national trends absent that comparative evidence. The NYPD's experience — with its four distinct policy regimes and built-in falsification test — is unusually rich, but the framework applies wherever administrative data on pursuit activity and crime outcomes can be assembled. Jurisdictions considering expanding officers' discretion to engage in pursuits should require affirmative evidence of crime prevention benefits — at minimum, crime trajectories that outperform contemporaneous regional trends — rather than assuming them from within-jurisdiction time trends alone.

For agencies currently operating under permissive pursuit policies, these findings point toward investment in two areas: data infrastructure and pursuit alternatives. Police Executive Research Forum (2023) recommends that agencies implement real-time pursuit reporting linked to officer, suspect, and outcome characteristics; such records would enable the individual-level analysis that aggregate administrative data cannot support. Pursuit alternatives (e.g., tire deflation devices, GPS tagging technology, aerial support) offer pathways to apprehension without generating the collision externalities documented here (Police Executive Research Forum 2023). The Tisch partial re-restriction offers a promising natural experiment in precisely this direction: with a longer post-restriction observation window, expanded comparison data, and individual-level pursuit records, future research may be able to estimate the effects of the new policy regime with greater precision.

Conclusion

Three hundred people die each year in pursuit-related crashes in the United States, roughly one-third of them bystanders who had no part in the chase (Rivara and Mack 2004). Whether that toll is offset by crime deterrence is what the NYPD's 2022–2025 experience allows us to address, at least in this institutional and temporal context. Pursuit-related collisions increased substantially and in direct proportion to pursuit volume, a causal effect for which ITS provides well-identified estimates given the direct mechanical link between pursuit volume and collision risk, independently corroborated by the Tisch falsification. For crime outcomes, the evidence does not support deterrence. The precinct-level DD (which holds both time trends and precinct baselines fixed and compares precincts by pre-treatment crime rate) yields a nominally negative shooting estimate, but the pre-period event study reveals non-parallel trends between high- and low-pre-treatment-crime precincts, suggesting that the estimate more likely reflects differential secular crime trajectories than a pursuit-driven treatment effect. Critically, the February 2025 re-restriction that reduced monthly pursuits by three-quarters was not followed by any crime increase — a strong within-jurisdiction test of the deterrence hypothesis that the data do not support.

The analysis speaks directly to the question that Rivara and Mack (2004) posed over two decades ago: whether "the trade-off of fewer pursuit-related crashes and deaths [is] offset by a higher number of fatal crimes." At least in the context of the NYPD escalation, the evidence does not resolve the trade-off in favor of the pursuit policy. The crash costs are demonstrably real, and the crime prevention benefits are not demonstrated. The causal attribution framework offers a principled approach: compute the minimum causal share at which benefits equal costs. In this case, the threshold is low, yet the evidence is consistent with a causal share near zero.

The critical open question is whether the pattern extends: does the Tisch partial re-restriction, which has achieved a three-quarters reduction in pursuits while crime continues declining, hold over a longer post-restriction window, and do similar results obtain in other jurisdictions that have recently recalibrated pursuit authority? Answering that question will require longer observation windows, individual-level pursuit records linking specific chases to officer, suspect, and bystander outcomes, and broader national comparison data. For now, the weight of the available evidence counsels against the assumption that expanded pursuit authority deters crime at a scale sufficient to offset its collision costs, and the break-even framework developed here gives practitioners a principled threshold against which future evidence can be assessed.

References

Alpert GP, Anderson PR (1986) The most deadly force: Police pursuits. Justice Quarterly 3:1–14

Alpert GP, Clarke AC, Smith WC (1997a) The constitutional implications of high-speed police pursuits under a substantive due process analysis: Homeward through the haze. University of Memphis Law Review 27:599–662

Alpert GP, Dunham RG (1988) Research on police pursuits: Applications for law enforcement. American Journal of Police 7:123–132

Alpert GP, Dunham RG (1989) Policing hot pursuits: The discovery of aleatory elements. Journal of Criminal Law and Criminology 80:521–539

Alpert GP, Kenney DJ, Dunham R (1997b) Police pursuits and the use of force: Recognizing and managing "the pucker factor" — a research note. Justice Quarterly 14:371–385. https://doi.org/10.1080/07418829700093371

Alpert GP, Lum C (2014) Police pursuit driving: Policy and research. [Springer Science; Business Media], New York, NY

Alpert GP, Madden T (1994) Police pursuit driving: An empirical analysis of critical decisions. American Journal of Police 13:23–46

Aronie J, Alpert GP (2020) 16 to a dealer's 10: Could blackjack odds help inform police pursuit policies?

Barnum TC, Nagin DS, Pogarsky G (2021) Sanction risk perceptions, coherence, and deterrence. Criminology 59:195–223. https://doi.org/10.1111/1745-9125.12266

Brosnan E (2023) NYPD chief attributes more car chases to rise in "ghost vehicles"

Butler LKW (2018) 2018 City of Milwaukee fire and police commission vehicle pursuit report. City of Milwaukee Fire; Police Commission

Christie N (2020) Managing the safety of police pursuits: A mixed method case study of the Metropolitan Police Service, London. Safety Science 129:Article 104848. https://doi.org/10.1016/j.ssci.2020.104848

Crew Jr Robert E., Hart Jr Robert A. (1999) Assessing the value of police pursuit. Policing: An International Journal of Police Strategies & Management 22:58–73

Gonen Y (2023) NYPD chief circulated internal memo on car chases and safety

Gonen Y (2025) 20 people died in NYPD car chases. Now the new police commissioner is cracking down

Hicks WL (2003) Police vehicular pursuits: An overview of research and legal conceptualizations for police administrators. Criminal Justice Policy Review 14:75–95. https://doi.org/10.1177/0887403402250925

Hicks WL (2006) Police vehicular pursuits: A descriptive analysis of state agencies' written policy. Policing: An International Journal of Police Strategies & Management 29:106–124. https://doi.org/10.1108/13639510610648511

Hill J (2002) High-speed police pursuits: Dangers, dynamics, and risk reduction. FBI Law Enforcement Bulletin 71:14–19

Hoffmann G, Mazerolle P (2005) Police pursuits in queensland: Research, review and reform. Policing: An International Journal of Police Strategies & Management 28:530–545. https://doi.org/10.1108/13639510510614591

Homant RJ, Kennedy DB, Howton JD (1994) Risk taking and police pursuit. The Journal of Social Psychology 134:213–221

Hosking PK (2021) Policy reform and resistance: A case study of police pursuit policy change in Queensland, Australia. PhD thesis, Griffith University

Hutson HR, Rice PL Jr., Chana JK, et al (2007) A review of police pursuit fatalities in the United States from 1982–2004. Prehospital Emergency Care 11:278–283. https://doi.org/10.1080/10903120701385414

Kennedy DB, Homant RJ, Kennedy JF (1992) A comparative analysis of police vehicle pursuit policies. Justice Quarterly 9:227–246

Kenney DJ, Alpert GP (1997) A national survey of pursuits and the use of police force: Data from law enforcement agencies. Journal of Criminal Justice 25:315–323. https://doi.org/10.1016/S0047-2352(97)00016-0

Lopez E, Boxerman B (2026) Crime trends in U.S. Cities: Year-end 2025 update. Council on Criminal Justice, Washington, DC

MacDonald JM, Alpert GP (1998) Public attitudes toward police pursuit driving. Journal of Criminal Justice 26:185–194. https://doi.org/10.1016/S0047-2352(97)00080-9

Mourtgos SM, Adams IT, McLean K, Alpert GP (2026) Risk and public judgments on police pursuits: A nationally representative conjoint experiment. Police Quarterly 0:1–25. https://doi.org/10.1177/10986111251412794

New York City Police Department (2025) Commissioner tisch announces new vehicle pursuit policy

O'Connor PT, Norse WL Jr. (2005) Police pursuits: A comprehensive look at the broad spectrum of police pursuit liability and law. Mercer Law Review 57:511–552

Pfeiffer MA, Alpert GP (2020) Developing methodology for finding ghosted police vehicle pursuits. Translational Criminology

Pogarsky G, Roche SP, Pickett JT (2017) Heuristics and biases, rational choice, and sanction perceptions. Criminology 55:85–111. https://doi.org/10.1111/1745-9125.12129

Police Executive Research Forum (2023) Vehicular pursuits: A guide for law enforcement executives on managing the associated risks. Office of Community Oriented Policing Services, Washington, DC

Policing Project at New York University School of Law, Gillooly J, Owens E, Mueller-Smith M (2021) Measuring the costs and benefits associated with vehicle pursuit policies in Roanoke City and Roanoke County, VA. New York University School of Law

Rivara FP et al (2025) Vehicular pursuits in Washington State. Harborview Injury Prevention; Research Center, University of Washington

Rivara FP, Mack CD (2004) Motor vehicle crash deaths related to police pursuits in the United States. Injury Prevention 10:93–95. https://doi.org/10.1136/ip.2003.004853

Senese JD, Lucadamo T (1996) To pursue or not to pursue? That is the question: Modeling police vehicular pursuits. American Journal of Police 15:55–78.

Sharkey P (2026) *AmericanViolence.org: city, fatal and nonfatal shootings*. Princeton, NJ: www.americanviolence.org.

Wade LM (2015) High-risk pursuit classification: A categorical analysis of variables from Georgia police pursuits. Criminal Justice Policy Review 26:278–292. https://doi.org/10.1177/0887403413516000

Xu Y (2017) Generalized synthetic control method: Causal inference with interactive fixed effects models. Political Analysis 25:57–76. https://doi.org/10.1017/pan.2016.2

Appendix A: Supplemental Regression Tables

Table A1: Borough-Level ITS Coefficients — All Four Outcomes

Newey-West HAC standard errors (lag = 6). b_2 = immediate level shift at October 2022; b_3 = change in post-intervention slope. Significance: *** $p < .001$, ** $p < .01$, * $p < .05$.

Table A1. Borough-Level ITS Coefficients — All Four Outcomes

Outcome	Borough	Level Shift b2 (SE)	Slope Change b3 (SE)
Gun Violence	Bronx	-40.32** [p .002] (SE = 12.85)	-2.04*** [p < .001] (SE = 0.38)
Gun Violence	Brooklyn	-39.17** [p .003] (SE = 12.73)	-2.17*** [p < .001] (SE = 0.39)
Gun Violence	Manhattan	-28.43*** [p < .001] (SE = 7.32)	-1.36*** [p < .001] (SE = 0.21)
Gun Violence	Queens	-27.08*** [p < .001] (SE = 4.76)	-1.36*** [p < .001] (SE = 0.16)
Gun Violence	Staten Island	-2.38* [p .021] (SE = 1.01)	-0.21*** [p < .001] (SE = 0.04)
Pursuit Crashes	Bronx	6.97*** [p < .001] (SE = 1.74)	-0.11 [p .247] (SE = 0.09)
Pursuit Crashes	Brooklyn	2.97* [p .031] (SE = 1.35)	0.01 [p .914] (SE = 0.08)
Pursuit Crashes	Manhattan	1.93* [p .032] (SE = 0.88)	-0.00 [p .996] (SE = 0.04)
Pursuit Crashes	Queens	6.23* [p .033] (SE = 2.87)	-0.07 [p .586] (SE = 0.13)
Pursuit Crashes	Staten Island	1.28 [p .093] (SE = 0.75)	-0.01 [p .712] (SE = 0.03)
Robbery	Bronx	-21.23 [p .190] (SE = 16.07)	-3.49*** [p < .001] (SE = 0.73)
Robbery	Brooklyn	-5.93 [p .767] (SE = 19.91)	-1.87** [p .008] (SE = 0.69)
Robbery	Manhattan	-43.14*** [p < .001] (SE = 8.58)	-4.15*** [p < .001] (SE = 0.37)
Robbery	Queens	35.81 [p .100] (SE = 21.50)	-2.59*** [p < .001] (SE = 0.72)
Robbery	Staten Island	8.36*** [p < .001] (SE = 2.06)	-0.06 [p .543] (SE = 0.10)
Shooting Incidents	Bronx	-13.09*** [p < .001] (SE = 3.39)	-0.61*** [p < .001] (SE = 0.14)
Shooting Incidents	Brooklyn	-7.47 [p .067] (SE = 4.03)	-0.73*** [p < .001] (SE = 0.12)
Shooting Incidents	Manhattan	-5.76** [p .002] (SE = 1.76)	-0.52*** [p < .001] (SE = 0.07)
Shooting Incidents	Queens	-6.45** [p .002] (SE = 2.01)	-0.39*** [p < .001] (SE = 0.06)
Shooting Incidents	Staten Island	-0.48 [p .348] (SE = 0.50)	-0.07*** [p < .001] (SE = 0.02)

Table A2: Precinct-Level DD Event Study Coefficients — Shooting Incidents

Estimates from `feols(shooting_incidents ~ i(rel_time, high_crime_shoot, ref = -1) | pct + month_date, cluster = ~pct)`. $t = 0$ corresponds to October 2022. Standard errors clustered by precinct (G = 77). * = 95% CI excludes zero (note: this CI-based indicator differs from the *p*-value stars used in Table A1).

Table A2. DD Event Study Coefficients — Shooting Incidents

Bin	Quarter	Period	Estimate	Std. Error	95% CI
-8	Oct–Dec 2020	Pre	-0.048	0.2086239	[-0.463, 0.368]
-7	Jan–Mar 2021	Pre	-0.736*	0.2147780	[-1.164, -0.308]
-6	Apr–Jun 2021	Pre	0.892*	0.2641933	[0.366, 1.418]
-5	Jul–Sep 2021	Pre	0.722*	0.2872094	[0.15, 1.294]
-4	Oct–Dec 2021	Pre	0.048	0.1972831	[-0.345, 0.441]
-3	Jan–Mar 2022	Pre	-0.61*	0.2497241	[-1.108, -0.113]
-2	Apr–Jun 2022	Pre	-0.149	0.2441653	[-0.635, 0.337]
0	Oct–Dec 2022	Post	-0.796*	0.2222759	[-1.239, -0.354]
1	Jan–Mar 2023	Post	-1.131*	0.2066896	[-1.543, -0.719]
2	Apr–Jun 2023	Post	-0.771*	0.2412181	[-1.252, -0.291]
3	Jul–Sep 2023	Post	-0.592*	0.2730577	[-1.135, -0.048]
4	Oct–Dec 2023	Post	-1*	0.2071978	[-1.413, -0.588]
5	Jan–Mar 2024	Post	-1.092*	0.1981458	[-1.486, -0.697]
6	Apr–Jun 2024	Post	-0.797*	0.1935409	[-1.182, -0.411]
7	Jul–Sep 2024	Post	-0.65*	0.2906372	[-1.229, -0.071]
8	Oct–Dec 2024	Post	-1.017*	0.2400211	[-1.495, -0.539]
9	Jan 2025+	Post	-1.164*	0.2192625	[-1.601, -0.727]

Table A3: Precinct-Level DD Event Study Coefficients — Gun Violence

Estimates from `feols(gun_violence_count ~ i(rel_time, high_crime_gunviol, ref = -1) | pct + month_date, cluster = ~pct)`. Gun violence is defined as the sum of shooting incidents and shots-fired reports. $t = 0$ corresponds to October 2022. Standard errors clustered by precinct (G = 77). * = 95% CI excludes zero.

Table A3. DD Event Study Coefficients — Gun Violence

Bin	Quarter	Period	Estimate	Std. Error	95% CI
-8	Oct–Dec 2020	Pre	-0.288	0.6262261	[-1.536, 0.959]
-7	Jan–Mar 2021	Pre	-2.563*	0.5374695	[-3.634, -1.493]
-6	Apr–Jun 2021	Pre	2.807*	0.6455061	[1.521, 4.092]
-5	Jul–Sep 2021	Pre	1.988*	0.7090973	[0.576, 3.4]
-4	Oct–Dec 2021	Pre	0.488	0.4690219	[-0.446, 1.422]
-3	Jan–Mar 2022	Pre	-1.552*	0.4932361	[-2.534, -0.57]
-2	Apr–Jun 2022	Pre	-0.776	0.4951810	[-1.762, 0.21]
0	Oct–Dec 2022	Post	-2.631*	0.5241768	[-3.675, -1.587]
1	Jan–Mar 2023	Post	-3.071*	0.5015741	[-4.07, -2.072]
2	Apr–Jun 2023	Post	-2.34*	0.5529216	[-3.442, -1.239]
3	Jul–Sep 2023	Post	-2.315*	0.5988363	[-3.508, -1.122]
4	Oct–Dec 2023	Post	-3.202*	0.5247367	[-4.247, -2.157]
5	Jan–Mar 2024	Post	-3.631*	0.4995108	[-4.626, -2.636]
6	Apr–Jun 2024	Post	-2.195*	0.6251015	[-3.44, -0.95]
7	Jul–Sep 2024	Post	-2.71*	0.6281913	[-3.961, -1.459]
8	Oct–Dec 2024	Post	-3.271*	0.6221502	[-4.51, -2.032]
9	Jan 2025+	Post	-3.359*	0.5880461	[-4.53, -2.188]

Table A4: Precinct-Level DD Event Study Coefficients — Robberies

Estimates from `feols(robbery_count ~ i(rel_time, high_crime_rob, ref = -1) | pct + month_date, cluster = ~pct)`. $t = 0$ corresponds to October 2022. Standard errors clustered by precinct (G = 77). * = 95% CI excludes zero.

Table A4. DD Event Study Coefficients — Robberies

Bin	Quarter	Period	Estimate	Std. Error	95% CI
-8	Oct–Dec 2020	Pre	-4.653*	1.667686	[-7.975, -1.332]
-7	Jan–Mar 2021	Pre	-8.067*	1.861192	[-11.774, -4.361]
-6	Apr–Jun 2021	Pre	-5.462*	1.626128	[-8.701, -2.223]
-5	Jul–Sep 2021	Pre	-3.277*	1.394739	[-6.055, -0.499]
-4	Oct–Dec 2021	Pre	-1.251	1.194778	[-3.631, 1.128]
-3	Jan–Mar 2022	Pre	-4.838*	1.324231	[-7.475, -2.2]
-2	Apr–Jun 2022	Pre	-1.13	1.060108	[-3.241, 0.982]
0	Oct–Dec 2022	Post	-2.266*	1.015219	[-4.288, -0.244]
1	Jan–Mar 2023	Post	-4.867*	1.193730	[-7.245, -2.49]
2	Apr–Jun 2023	Post	-3.148*	1.350905	[-5.838, -0.457]
3	Jul–Sep 2023	Post	-1.307	1.478814	[-4.253, 1.638]
4	Oct–Dec 2023	Post	-0.471	1.565890	[-3.59, 2.647]
5	Jan–Mar 2024	Post	-4.048*	1.570615	[-7.176, -0.92]
6	Apr–Jun 2024	Post	-1.379	1.546947	[-4.46, 1.702]
7	Jul–Sep 2024	Post	-2.257	1.645647	[-5.535, 1.02]
8	Oct–Dec 2024	Post	-3.231*	1.614545	[-6.446, -0.015]
9	Jan 2025+	Post	-4.437*	1.596722	[-7.617, -1.257]

Table A5: Precinct-Level DD Event Study Coefficients — Pursuit Crashes

Estimates from `feols(collision_count ~ i(rel_time, high_crash, ref = -1) | pct + month_date, cluster = ~pct)`. Note: high-crash precincts are expected to show larger collision increases post-policy (more pursuit activity), not larger crime reductions — the directional prediction for crashes differs from crime outcomes.

Table A5. DD Event Study Coefficients — Pursuit Crashes

Bin	Quarter	Period	Estimate	Std. Error	95% CI
-8	Oct–Dec 2020	Pre	0.008	0.04269679	[-0.077, 0.093]
-7	Jan–Mar 2021	Pre	0.03	0.03702846	[-0.043, 0.104]
-6	Apr–Jun 2021	Pre	-0.035	0.03671805	[-0.108, 0.038]
-5	Jul–Sep 2021	Pre	0.03	0.04294260	[-0.055, 0.116]
-4	Oct–Dec 2021	Pre	-0.033	0.03308874	[-0.099, 0.033]
-3	Jan–Mar 2022	Pre	0.066	0.04549284	[-0.025, 0.156]
-2	Apr–Jun 2022	Pre	-0.018	0.03285328	[-0.083, 0.048]
0	Oct–Dec 2022	Post	-0.083	0.07799513	[-0.239, 0.072]
1	Jan–Mar 2023	Post	0.088	0.06539569	[-0.042, 0.219]
2	Apr–Jun 2023	Post	-0.045	0.09000353	[-0.225, 0.134]
3	Jul–Sep 2023	Post	0.096	0.11363032	[-0.13, 0.322]
4	Oct–Dec 2023	Post	0.22	0.12759992	[-0.034, 0.474]
5	Jan–Mar 2024	Post	0.119	0.09041534	[-0.061, 0.299]
6	Apr–Jun 2024	Post	0.114	0.13201180	[-0.149, 0.377]
7	Jul–Sep 2024	Post	0.215*	0.09720136	[0.021, 0.408]
8	Oct–Dec 2024	Post	0.003	0.08487598	[-0.167, 0.172]
9	Jan 2025+	Post	-0.033	0.04010196	[-0.113, 0.047]

Appendix B: Generalized Synthetic Control — Shooting Rates Across U.S. Cities

Design

We estimate the effect of the October 2022 pursuit escalation on shooting rates using the generalized synthetic control method [gsynth; Xu (2017)], applied to a monthly panel of large U.S. cities from January 2014 through December 2025 (Sharkey 2026). The outcome is the annualized shooting rate per 100,000 population, computed from Gun Violence Archive data combining fatal and nonfatal shooting victims. NYC is the single treated unit; all other cities serve as the donor pool. Leave-one-out cross-validation over $r = 0$–5 latent factors selects the optimal specification; inference uses 500 parametric bootstrap iterations. The covariance matrix was singular in the F-test, which is expected with a single treated unit and does not affect point estimates or bootstrap confidence intervals (Xu 2017).

SUTVA exclusions. To avoid SUTVA contamination — which would occur if donor cities' shooting trends were themselves affected by their own pursuit policy changes during the study period — we exclude 24 cities with documented city-level or binding state-level pursuit policy changes between 2018 and 2024 (Table B1). The resulting donor pool comprises 75 cities with stable pursuit policies, observed across 144 months. This is the primary specification. As a sensitivity check, we replicate the analysis with the top-50-by-population panel (49 donors).

Results

Cross-validation selects $r = 3$ latent factor(s) for the SUTVA-clean panel. The pre-treatment root mean square prediction error (RMSPE) is 3.71, indicating adequate counterfactual fit. The average treatment effect on the treated (ATT) is 1.96 per 100,000 (annualized). None of the 39 post-treatment months produces a statistically significant individual ATT at the 95% confidence level. Figure B1 shows NYC's actual shooting rate alongside the gsynth counterfactual; Figure B2 shows the period-by-period ATT gap with 95% bootstrap confidence intervals, all of which include zero throughout the post-treatment window.

The full-panel robustness check ($r = 1$, 49 donors) yields ATT = -1.83 per 100,000 — also not statistically significant. Both specifications are consistent with the null finding: NYC's post-October 2022 shooting trajectory is statistically indistinguishable from what the latent-factor model predicts based on national trends.

Table B1: Cities Excluded from SUTVA-Clean Donor Pool

Cities excluded because of documented pursuit policy changes (city-level or binding state statute) between 2018 and 2024.

City	Reason for Exclusion
Atlanta	Zero-chase policy, January 2020
Chesapeake	Revised pursuit policy, November 2023
Chicago	Tightened 2020, revised 2022
Cincinnati	Violent felonies only, February 2023
Fort Wayne	IN statewide minimum standards, January 2023
Houston	Banned Class C vehicle pursuits, September 2023
Indianapolis	IN statewide minimum standards, January 2023
Jersey City	NJ AG statewide directive, 2021–2022
Lexington-Fayette	KY statewide written-policy mandate, 2020
Long Beach	Special Order 2023-5, June 2023
Louisville/Jefferson County	Loosened July 2019 + KY state mandate 2020
New Orleans	Policy revised August 2019 and February 2024
Newark	NJ AG statewide directive, 2021–2022
Oakland	50 mph speed cap, December 2022
Oklahoma City	Tightened, June 2022
Portland	Overhaul, January 2024
San Francisco	Prop E loosened, March 2024
San Jose	Dept. Order 2022-037, 2022
Seattle	WA state HB 1054/SB 5352/I-2113, 2021–2024
Spokane	WA state law changes, 2021–2024
Stockton	Voters loosened, late 2024
Tampa	Revised after fatal crash, 2022
Toledo	New policy, mid-2024
Washington	Restricted 2022; loosened 2024

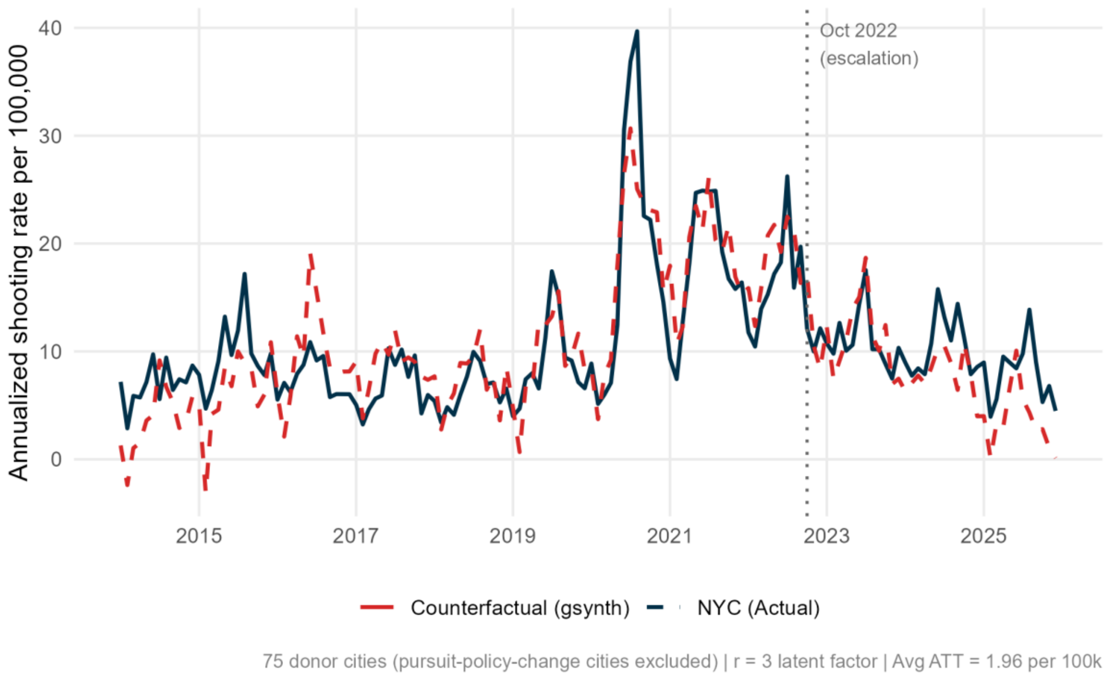

Figure B1: NYC actual shooting rate versus gsynth counterfactual, January 2014–December 2025. Solid line: NYC observed annualized rate per 100,000. Dashed line: gsynth counterfactual constructed from 75 SUTVA-clean donor cities. Vertical dotted line marks the October 2022 pursuit escalation.

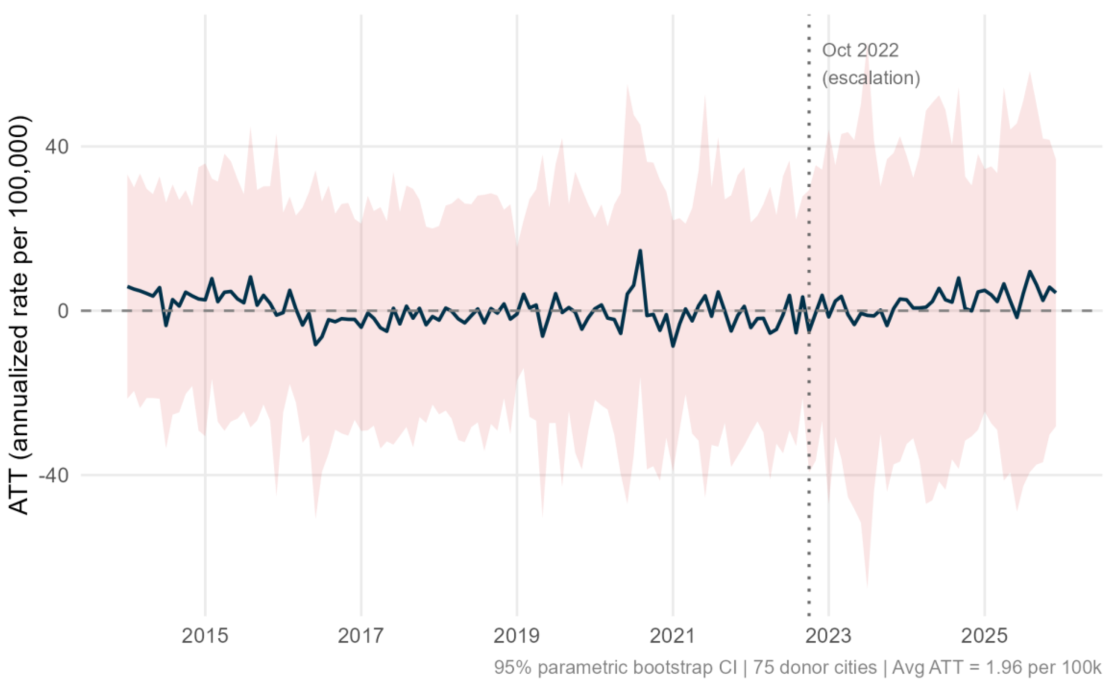

Figure B2: Period-by-period average treatment effect on the treated (ATT) with 95% parametric bootstrap confidence intervals, October 2022–December 2025. The zero line falls within the confidence interval for all 39 post-treatment months, indicating no significant treatment effect in any month.