

Reducing Violent Crime Without New Budget, New Staff, or More Arrests

A Pragmatic Guide for City Leaders

Joe Eichenbaum, Partner

17A

February 2026



Contents

1. Executive Summary	3
2. Where Violence Is, and How It Moves	5
3. What City Leaders Can Do: Environmental Interventions	11
4. Case Study: Coordinated Action in Dallas	13
5. Getting Started: A Practical Roadmap	18
Appendix	23

1. Executive Summary

Violent crime concentrates geographically, and the concentration is extreme. Across nine cities analyzed for this paper, roughly 3–5% of city geography accounts for 20% of all violent crime. The pattern holds everywhere: a small number of places carry a vastly disproportionate share of violence. Code violations, 311 complaints, and illegal dumping cluster in much of the same territory. For residents in these areas, citywide statistics showing declining crime can feel meaningless, and they're right to say so.

Violent crime's geographic concentration is not static. It shifts substantially from year to year. The measured concentration churns: in a typical city, a third to half of the highest-crime micro-areas are different from one year to the next. Some of that turnover reflects genuine shifts in conditions — a corridor deteriorates, a problem property attracts activity. Some of it reflects the inherent volatility of small-number crime data, where a handful of additional incidents can push a neighborhood across a threshold. But from the standpoint of a city trying to direct resources, the distinction doesn't matter much. Either way, last year's priority list is wrong. The geography that shows up in your data moves, and your response has to move with it.

The good news: the broad trend is favorable, and the real target is small. Violent crime is declining nationally and in most cities, including in chronically high-crime neighborhoods. Across nine cities, roughly 90% of geography is stable from year to year, and where crime is genuinely shifting, improvement outpaces deterioration by nearly five to one. The places trending in the wrong direction represent just 1–3% of city geography in any given year — roughly 20 to 30 half-mile areas in a typical city. That's not an overwhelming problem. It's a list that fits on a single page. But each of those areas, once you're on the ground, contains dozens of specific problems at specific addresses, each one owned by a different city department. The targeting tells you where to look. The harder part is coordinating the response: getting the right agencies to the right addresses, tracking follow-through, and sustaining engagement over time.

Cities can respond rapidly using capacity they already have. Environmental interventions — cleaning vacant lots, fixing streetlights, citing blighted properties, removing illegal dumping — can be directed to specific blocks within weeks and redirected as conditions shift. The departments that deliver this work — code enforcement, sanitation, public works, parks — already exist, already have staff and budgets, and already do this work every day. What they typically lack is a system that tells them where it matters most this week. Policing reform and large-scale social programs matter immensely but are slow, expensive, or politically contentious. Environmental interventions occupy a different space: fast, adaptive, and deliverable with existing resources. And because geographic concentration characterizes not just violence but also 311 requests, code violations, and other indicators of neighborhood distress, the same coordination model serves multiple city priorities at once.

These interventions work, and they work best when paired with a back-office engine that keeps them targeted. Randomized trials have shown that greening vacant lots reduces nearby gun violence by

29% and that improved street lighting cuts outdoor nighttime crime by 39%. But the evidence also suggests that isolated, one-off deployments lose their effect. What sustains impact is a coordination system that combines data-driven targeting with operational follow-through: identifying the small number of places trending the wrong way, assigning the right agencies, and retargeting as conditions shift. In our own local government public safety work, we have seen firsthand how this coordinated approach gets results. This is especially true in Dallas. In 2024–2025, the city deployed coordinated environmental interventions across departments, retargeting weekly. At every risk level, areas that received intervention outperformed comparable areas that did not — and the effect was strongest in areas at the tipping point between stability and escalation. Section 4 presents the full case study.

Any city can start in 90 days with existing staff and data. This approach does not require new programs, new hires, or new funding. It requires a coordinator, a short list of priority locations drawn from frequently updated data, and a commitment to regular cross-agency meetings. Cities that sustain this work treat it as an operating rhythm — not a one-time campaign. The goal is not a single program but a way of directing city government toward wherever concentrated need emerges, adapting as that geography shifts.

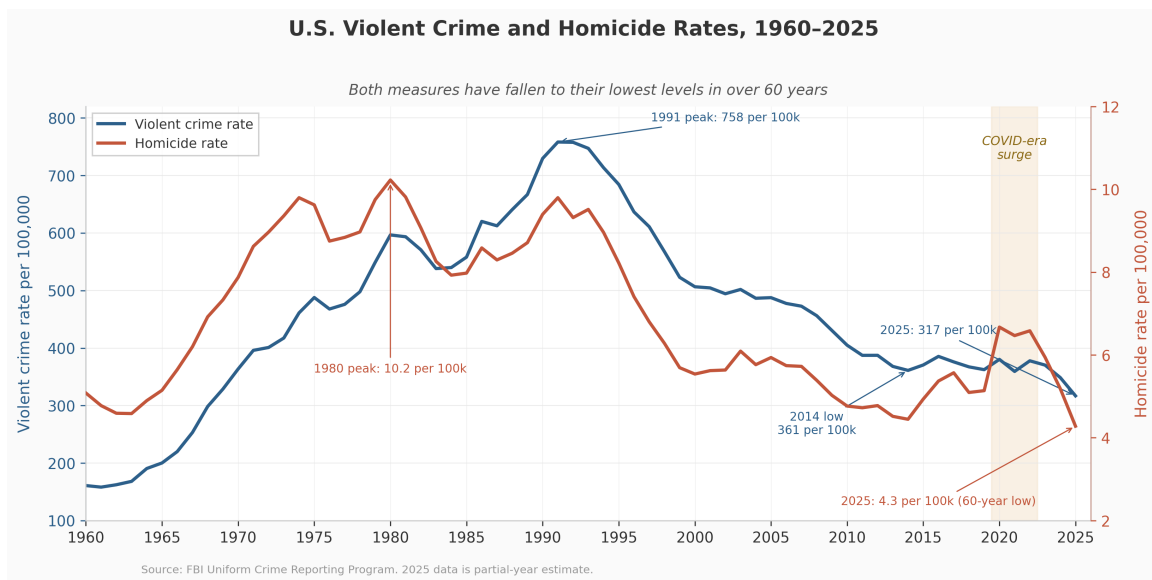
2. Where Violence Is, and How It Moves

Violent crime devastates individuals, families, and communities. Beyond the direct harm to victims, persistent violence erodes trust in government and makes future violence more likely. Where violent crime is common and sustained, the failure to reduce it is widely perceived as a failure of government itself.

Understanding where violence concentrates, and how that concentration changes over time, is the foundation for any place-based strategy.

The broad decline is real

On average, Americans today are roughly half as likely to be the victim of a violent crime as they were in the early 1990s. This decline is one of the most significant and undercelebrated public safety achievements of the last fifty years. After a sharp spike during the COVID-19 pandemic (homicides surged roughly 30% in 2020), violent crime has resumed its downward trajectory. By 2025, national homicide rates had fallen below pre-pandemic levels to their lowest point in over sixty years.



This isn't just a national average. Eight of nine cities analyzed for this paper saw violent crime decline from 2024 to 2025, with citywide drops ranging from 3% to 24%.

Concentration is devastating where it lands

But that broad decline coexists with a well-documented local reality: violent crime concentrates geographically. What criminologists call the “law of crime concentration” — formalized by David Weisburd in 2015 — holds that a small share of places accounts for a disproportionate share of crime. A study of six major U.S. cities found that 3% to 10% of street segments accounted for half of all crime.

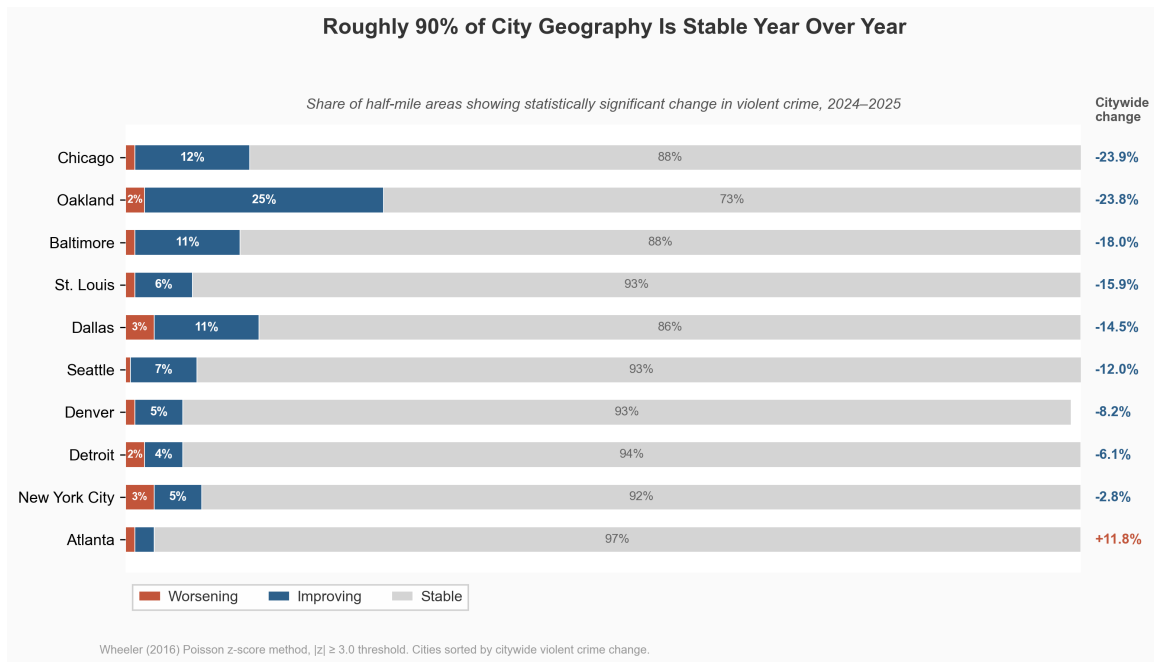
This concentration is extreme. In Dallas, the geography containing the top 20% of violent crime represents just 3.4% of the city. In New York, 2.5%. In Denver, 1.6%. For residents living in these areas, the national decline can feel invisible. The danger is real and the frustration is legitimate.

Former Baltimore Mayor and Maryland Governor, and current 17A Senior Advisor Martin O’Malley saw this dynamic firsthand. “When I was mayor of Baltimore, I’d walk neighborhoods where three or four shootings had happened in a month, and someone would always say, ‘They keep telling us crime is going down. Where?’” he recalls. “They weren’t wrong. Citywide numbers were improving, but violence had moved onto their block that year, and that was all that mattered to them.”

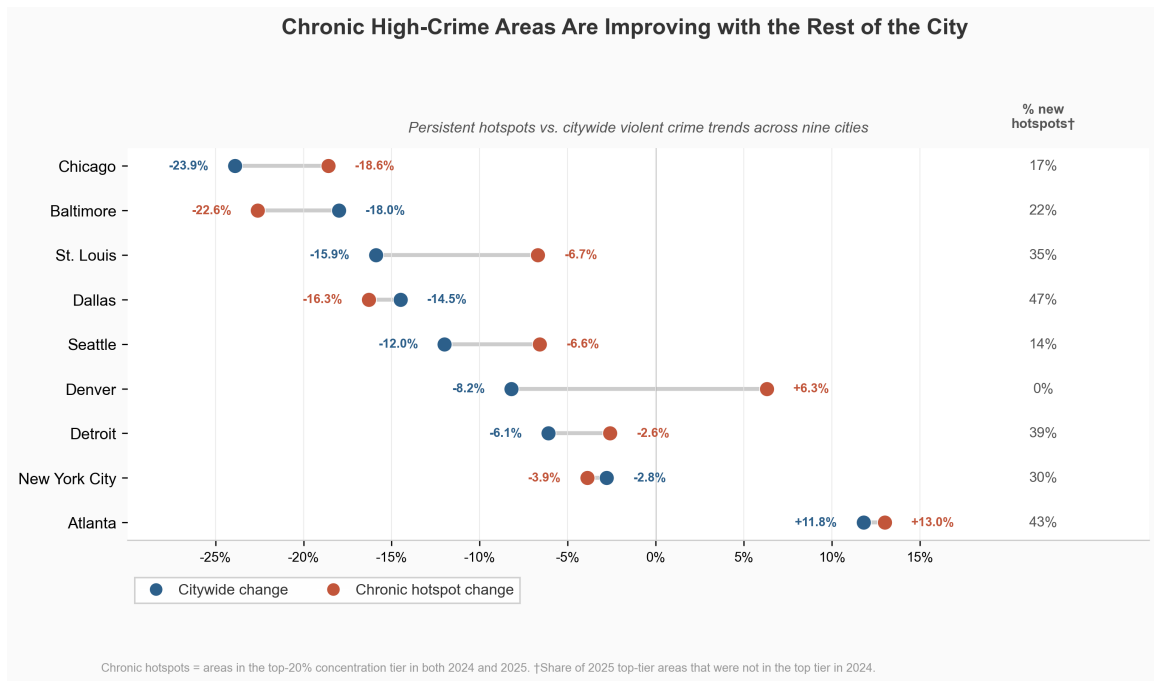
What's actually changing

The common assumption is that the same neighborhoods stay dangerous year after year while the rest of the city improves. The data tells a different story.

When you test whether year-over-year changes in each half-mile area reflect an emerging trend — rather than the normal ups and downs you’d expect from small numbers — a clear picture emerges across ten cities.



This is what “crime is going down” actually looks like at the neighborhood level. It’s not a uniform decline everywhere. There is broad stability across most of most cities, concentrated improvement in a relatively small number of places, and a much smaller number of places moving in the wrong direction. That last group — the 1–3% where violence is escalating — is where targeted intervention can have the greatest impact.












There are exceptions. Denver’s chronic hotspots worsened even as the city overall improved. Atlanta’s worsened alongside a citywide increase. But the dominant pattern — in seven of nine cities — is that persistently high-crime areas are declining. They are not being left behind.

The real policy target is small. Even as overall concentration persists — roughly 3–5% of city geography continues to account for 20% of violent crime, year after year — the specific locations shift. In Dallas, nearly half of the cells in the highest-concentration tier in 2025 were not in that tier in 2024. This is not a Dallas anomaly. Across nine cities, the geography of concentration reshuffles substantially every year — with year-over-year overlap in the highest tier ranging from just 28% (Dallas) to 79% (Denver). Wheeler and Reuter (2021) found the same pattern in their analysis of Dallas hotspots: traditionally defined boundaries fail to capture how concentration moves over time.

Geographic Churn: Where Violence Concentrates Reshuffles Every Year

Each square = one half-mile grid cell in the top-20% concentration tier, 2024 vs. 2025.

■ Newly concentrated (entered top tier in 2025)
 ■ Cooled off (left top tier in 2025)
 ■ Persistent (top tier both years)



Half-mile grid cells. Overlap = Jaccard similarity (persistent + union of both years' top-tier cells). See Appendix C-D.

Understanding this rotation is important. Some of it reflects real shifts — a corridor deteriorates, a problem property attracts activity, a previously stable area starts trending the wrong way. Some of it is simply the nature of small numbers: an area with 3 violent crimes one year and 7 the next can cross a concentration threshold without anything fundamental having changed. The stability analysis helps distinguish between the two. When you focus on the areas where the change is large enough to indicate a genuine trend, the picture sharpens considerably: the places where violence is escalating in a given year represent roughly 1–3% of city geography. In Dallas, that was 26 half-mile areas. That’s not

an overwhelming problem. That’s a list that fits on a single page — and it’s where targeted intervention can do the most good.

Why this matters for city leaders

This reframes the challenge. The question for a new mayor or deputy mayor isn’t “how do we fix every high-crime neighborhood?” Most of the city is stable, chronic hotspots are generally improving, and the broad trend is your ally. The question is: can you identify the small number of places — roughly 20 to 30 at any given time — where violence is escalating, and can you get there fast enough to make a difference?

That’s not a resource problem. It’s a detection and coordination problem. You don’t need new programs, new staff, or new funding to respond to 25 locations. You need a system that can spot the emerging trouble spots and redirect existing city capacity — code enforcement, public works, street-lights, property maintenance — toward those places on a rolling basis. And you need that system to be adaptive, because the list will change.

This pattern — geographic concentration that persists structurally even as the specific locations shift — is not unique to violent crime. Cities see similar dynamics in 311 service requests, code violations, illegal dumping complaints, and other indicators of neighborhood distress. The places generating the most service demand rotate over time, even as the overall level of concentration stays roughly constant. Violence is the highest-stakes version of a more general phenomenon: city resources need to follow shifting geographic need, and most cities aren’t set up to do that. We focus on violence in this paper because the consequences of getting it wrong are most severe and the data is most granular — but the coordination model described in Sections 4 and 6 applies wherever concentrated need moves faster than the city’s response.

The next section explains why environmental interventions — the kind of work those city departments already do — are effective at reducing violence. The section after that shows how one city built the coordination system to deploy them.

3. What City Leaders Can Do: Environmental Interventions

Most public safety strategies fall into two familiar categories. The first is political: policing reform, criminal justice policy, sentencing changes. These debates are important, but they are slow, contentious, and often outside the direct control of city government. The second is expensive: housing, behavioral health, workforce development, large-scale social programs. These investments matter immensely, but they require significant new resources and take years to show results.

There is a third category that is neither political nor expensive: environmental interventions. These strategies focus on the physical conditions of places where violence concentrates, and can often be deployed quickly using existing city capacity.

What the research shows

Environmental interventions address what a block or corridor looks like, how it is maintained, and what behaviors it enables or discourages. A substantial body of research shows that improving physical environments in high-violence locations can meaningfully reduce crime:

- A randomized controlled trial in Philadelphia found that cleaning and greening vacant lots reduced nearby gun violence by 29% (Branas et al., 2018).
- A study of New York City public housing found that improved street lighting led to a 39% reduction in outdoor nighttime crimes, including violent offenses (Chalfin et al., 2021).
- Additional studies show consistent benefits from graffiti removal, trash cleanup, and securing abandoned buildings.

These effects occur not simply because places look better, but because environmental conditions shape opportunity, visibility, and social norms. Well-maintained spaces are harder to use for illicit activity and signal that someone is paying attention.

Governor O'Malley describes a similar dynamic from his time in Baltimore. "Residents always knew which vacant lot was the problem, which streetlight had been out for months. When the city finally showed up and addressed those things, people noticed overnight. You'd hear it at the next community meeting. And the guys who had been using that corner noticed too."

The agencies that can deliver

The departments that can execute environmental interventions already exist in every city government. They already do this work, just not necessarily in the places where it would have the greatest impact on violence.

Code Enforcement is often the anchor agency. Code officers can issue citations, conduct property

assessments, and pursue nuisance abatement for persistently problematic addresses. The approach may vary depending on the property owner: a large institutional landlord with a pattern of neglect warrants a different enforcement posture than an elderly longtime homeowner who may need support navigating repairs.

Sanitation handles trash removal and illegal dumping cleanup, visible signs of disorder that shape how a space feels and whether it attracts further neglect.

Public Works maintains lighting, signage, and infrastructure. Broken streetlights and missing signage create opportunities for crime and signal that no one is watching.

Transportation can install barriers, adjust signal timing, add crosswalks, and improve street signage to improve safety and visibility.

Parks and Recreation can green vacant land, maintain public spaces, and activate underused areas to increase legitimate foot traffic and community presence.

The gaps: targeting and coordination

These agencies already have staff. They already have budgets. They already do cleanups, inspections, and repairs every day. What they typically lack are two things:

Targeting: A clear, data-driven answer to the question: *which blocks should we prioritize?* This requires combining crime data with local knowledge, identifying not just where violence has occurred but where conditions are ripest for it to continue. And because the geography of violence shifts over time, targeting must be updated frequently enough to keep pace with changing conditions. As Section 2 describes, the number of areas trending in the wrong direction at any given time is small, roughly 20 to 30 in a typical city. That's a manageable list. But it's only the first scale of the problem.

Coordination: Each of those 20 to 30 areas, once you're on the ground, contains dozens of specific problems at specific addresses: a vacant lot that needs clearing, a streetlight that's been out for months, a blighted property that needs a code citation, illegal dumping that needs cleanup, a bus stop that's become a gathering point. A single corridor might involve four or five different departments. The targeting tells you where to look. Coordination is what gets the right agencies to the right addresses, tracks whether the work is actually happening, and ensures repeat engagement over time. Without that management structure, even well-targeted efforts fade after the first visit.

When city leadership provides both—clear priorities and a structure to execute on them—routine city functions become a coordinated response to the places where multiple forms of need concentrate. Violence reduction is the highest-stakes application, but the same coordination structure addresses blight, illegal dumping, lighting failures, and the 311 backlogs that erode residents' confidence in city government.

4. Case Study: Coordinated Action in Dallas

In 2024–2025, Dallas tested what this approach looks like in practice. The city’s leadership, in close partnership with Child Poverty Action Lab (“CPAL”), committed to treating violence reduction as a broader quality of life problem in the city. City leaders reframed violent crime as a cross-agency coordination problem, not a policing problem, and directed departments that shape physical environments toward the places where violence was concentrating.

The effort was part of a long-running partnership between the City of Dallas and CPAL. CPAL’s working team built an analytical infrastructure that made targeting possible, translated data into operational priorities on a weekly cycle, and kept the work accountable to outcomes. City departments (code enforcement, transportation, and public works) executed the interventions. CPAL’s program personnel ensured the initial list of priorities was aligned with neighborhood realities, and built a process for site validation that the City is now carrying forward.

The results offer the clearest available evidence that coordinated, place-based environmental intervention works with existing city resources.

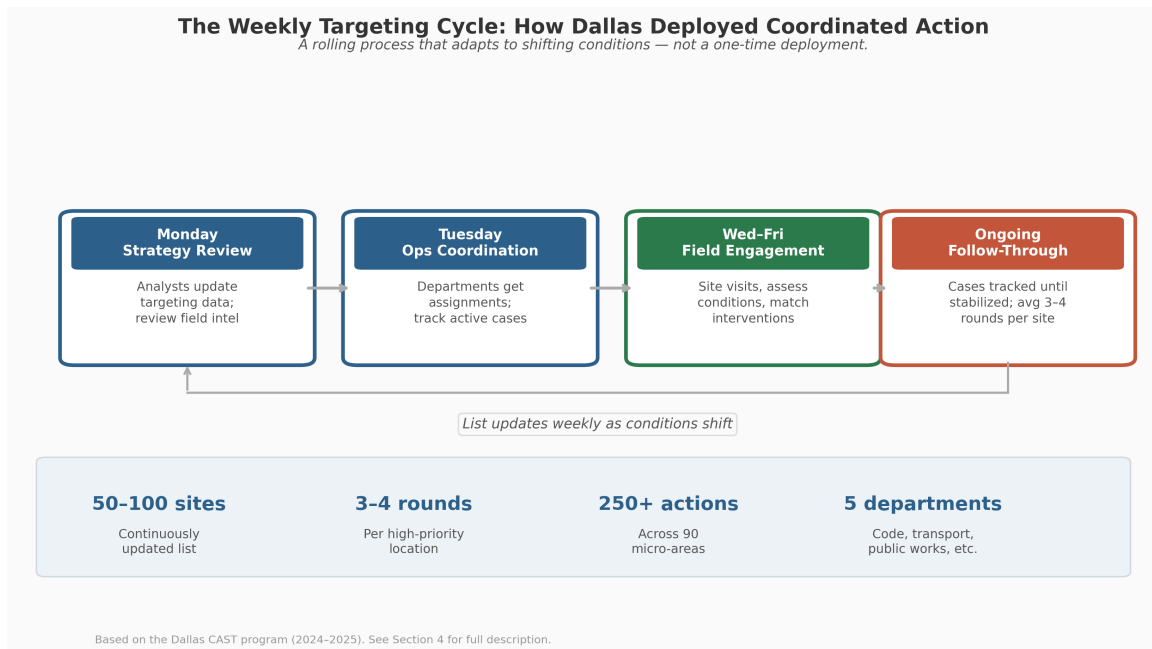
What Dallas did

The city organized the work around four ideas.

First, city leadership made it clear that violence reduction was a cross-agency priority. Department heads heard directly from the City that they would be expected to prioritize specific high-violence locations and report on progress. This was a core expectation from management.

Second, the working team maintained a short, continuously updated list of priority locations. Our analysts combined crime data, code violations, 311 complaints, and environmental indicators to identify roughly 50–100 sites at any given time. Critically, the list wasn’t static. As conditions shifted (some areas stabilizing, others deteriorating), the team shifted the targeting with them, typically on a weekly cycle. This is exactly the adaptive approach Section 2 argues is necessary: targeting that moves as the geography of violence moves.

Third, coordination happened on a regular weekly cadence. The team reviewed targeting data and field intelligence every Monday, updating priorities based on the latest conditions. Separate meetings with operational departments on Tuesdays, facilitated by CPAL, tracked active cases, assigned follow-ups, and closed out completed work. Throughout the week, field staff visited sites, assessed conditions, and gathered the context that CPAL fed back into the next cycle’s targeting. CPAL functioned as the connective tissue between data and action, ensuring that what the numbers showed actually translated into work on the ground.



Fourth, every site had clear ownership and follow-through. Departments tracked their cases in shared systems, updating progress as work moved. Sites stayed on the active list until conditions stabilized, not until the first intervention was complete. On average, high-priority locations received multiple rounds of engagement over several months.

What the work looked like

The team's targeting identified where to focus. But each priority area, once you're on the ground, contained dozens of specific problems at specific addresses, each one the responsibility of a different department. Ninety micro-areas produced over 250 completed interventions spanning vegetation clearing, streetlight repairs, trash and dumping cleanup, property citations, and infrastructure fixes. CPAL's Neighborhood's team was the core translation layer between data-driven prioritization and work assignment.

Some sites had clear solutions. After a murder near an apartment in Old East Dallas, the team assessed a vacant lot sandwiched between a cottage and a modern three-story building. The lot had ankle-high grass, graffiti on the fence, and old streetlights. The diagnosis was straightforward: cut back vegetation, paint over graffiti, convert streetlights to LED bulbs. Within two weeks, the grass was cut. The remaining items were in progress. When people feel safe walking around, there are more eyes and ears to report issues, and more community investment in maintaining the improvement.

Other sites didn't have a two-week fix. A house near the intersection of North Buckner Boulevard

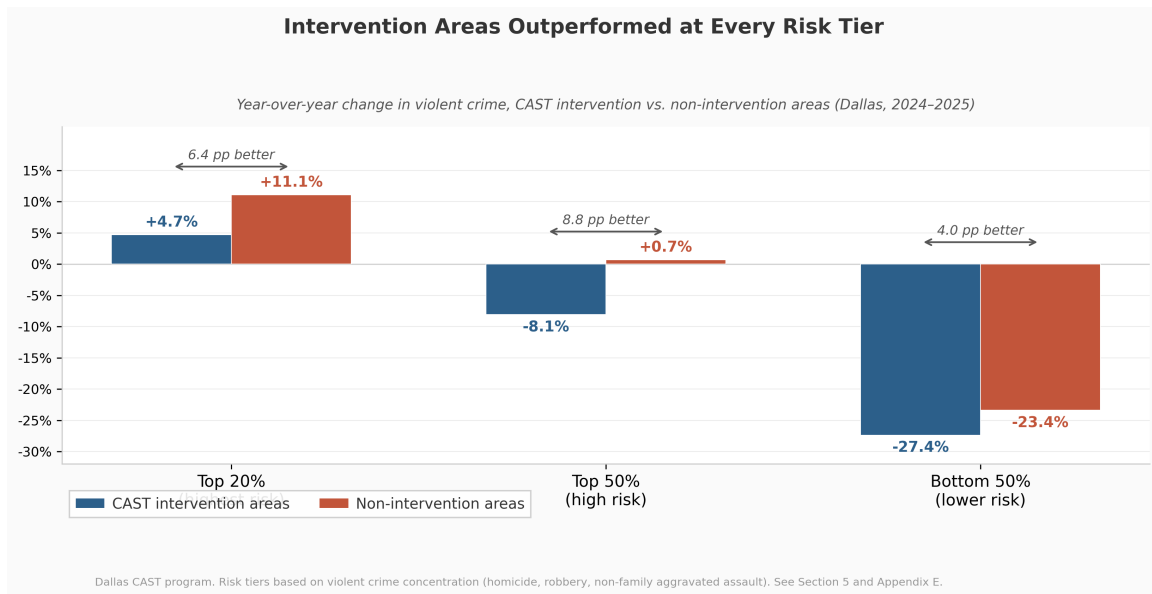
and Peavy Road in Far East Dallas presented a different challenge. The property's out-of-state owner had let it become a camp for people experiencing homelessness. Across the street, a beer store clerk described constant drug activity and 911 calls. A bus stop nearby had become a gathering point for people who never boarded buses. A surveillance camera tower installed months earlier hadn't deterred drug dealing beneath it. The team considered options: work with Dallas Area Rapid Transit to remove the bench, increase patrols, pursue the absentee landlord through community prosecution. They took the clerk's information and promised to return. This area would likely require months of repeat visits across dozens of interventions.

As [Kevin Oden](#), Dallas's director of Emergency Management and Crisis Response, put it: "Not every site needs the same things. It's a sudoku for them to figure out."

Agencies are already excellent at deploying interventions. The City can provide those agencies the discipline of concentrating them. City government often defaults to spreading resources on a first-in, first-out basis, but violence concentrates geographically, and the response has to concentrate in return. Fifty locations with repeated engagement will have more impact than 500 locations with one-off visits. And because environmental conditions deteriorate (vacant lots refill with trash, code violations recur), the work didn't end after a first pass. The goal wasn't a clean block for a photo. It was sustained change in the conditions that enable violence.

What happened

We know that Dallas improved. Violent crime declined 14.5% citywide, part of a national trend. The question is whether areas that received coordinated intervention fared better than comparable areas that didn't. The team's evaluation framework, comparing intervention areas to non-intervention areas within the same risk tier, provides the clearest test. At every risk level, the answer is yes.



The stabilization effect was clearest in the middle tier: areas with enough crime to be at risk of escalating, but not so entrenched that deeper structural forces dominate. Here, intervention areas declined 8% while comparable non-intervention areas were flat. These are precisely the areas Section 2 identifies as the narrow policy target: the 1–3% of geography at risk of tipping into higher concentration. Intervention pulled them back from that edge.

Even in the highest-risk areas, where crime increased regardless, intervention areas saw less than half the increase of non-intervention areas. Dampening the magnitude of surges matters — it’s what keeps a temporary spike from becoming an entrenched hotspot.

The results for homicide specifically were even stronger. At every risk tier, the gap between intervention and non-intervention areas was larger for homicide than for violent crime overall. Appendix E presents the full intervention results, and Appendix F tests their statistical significance using methods appropriate for crime count data.

What this doesn't prove

This is not a randomized controlled trial. Intervention areas weren’t assigned randomly; they were chosen because they had high crime and conditions that seemed amenable to environmental intervention. It’s possible that something about these areas, beyond the intervention itself, explains part of the difference.

When tested using Poisson-based methods appropriate for crime count data (see Appendix F), the intervention effect is statistically significant in the tier where the theory predicts it should be strongest:

areas between the 20th and 50th percentile of crime concentration. In the highest-risk and lowest-risk tiers, the differences go in the same direction but are not statistically significant, largely because the sample sizes are too small to detect the observed effects with confidence. The consistent direction across all tiers is itself notable.

But the consistency of the pattern, and the size of the differences, suggests that targeted, repeated, cross-departmental action made a real contribution. Dallas is one city and one year of data. But it's evidence that this approach can work — and that cities don't have to wait for perfect proof before starting.

5. Getting Started: A Practical Roadmap

The previous sections describe what coordinated, place-based action looks like and the evidence that it works. This section is about how to begin.

The data makes the case for starting: across ten cities, roughly 90% of geography is stable, the broad trend is favorable, and the places where violence is trending the wrong way represent just 1–3% of city geography in any given year. That’s roughly 20 to 30 half-mile areas. The challenge is not the size of the problem. It’s building the system to detect those areas and get city resources there before conditions entrench.

Any city can launch this approach in 90 days and see measurable outcomes within 180, using existing staff and existing data. This is not a new program. It’s a way of directing what city government already does toward the places where it will have the greatest impact, starting with violent crime but extending naturally to the code violations, 311 backlogs, and quality-of-life conditions that concentrate in many of the same places.

What you need

Mayoral commitment. This work only moves if city leadership makes clear that violence reduction is a cross-agency priority, not just a police department problem. Department heads need to hear, directly, that they will be expected to prioritize specific high-violence locations and report on progress. Without that signal from the top, coordination meetings become optional and follow-through fades.

A coordinator. One person (existing staff, not a new hire) who owns the rhythm: scheduling meetings, tracking assignments, flagging stalled cases, and connecting analytical targeting with operational execution. In Dallas, this role sat in the Emergency Management and Crisis Response office. In other cities, it could sit in the mayor’s office, the city manager’s office, or a public safety coordinator role. The title doesn’t matter. What matters is that someone wakes up every Monday thinking about which sites need attention this week.

Two to three invested, capable agency leads at the table. You don’t need every relevant department on day one. To start, you need code enforcement and one or two other agencies—sanitation, public works, or transportation—with leads who are willing to show up and do the work. Over time, you can expand to parks, additional public works divisions, and other departments. Within each agency, the right configuration will emerge: in Dallas, code enforcement brought representatives from multiple divisions (field inspections, nuisance abatement, multi-tenant) because the work touched all of them. Start with a small, committed group. The table grows as the work proves its value.

Data to identify priority locations. At minimum, incident-level crime data for homicide, robbery, and aggravated assault with addresses or coordinates—anything geocodable. Most cities already collect this through NIBRS. But crime data alone understates the opportunity. Code violations, 311 com-

plaints, vacancy rates, and lighting outage reports often concentrate in overlapping geography, and they give operational departments a direct reason to engage. Cities that layer these data sources into the targeting process find that the priority list serves multiple departments' missions, not just public safety. A basic crime map is enough to start; richer data sources make the coordination more effective and more durable over time.

A short list of priority locations. Don't try to cover the whole city, and the data says you don't have to. Ninety percent of the city is stable. The Dallas model worked because it concentrated resources on roughly 50–100 locations at a time, few enough to allow repeat engagement, enough to cover the places where violence was most acute. The discipline of a short list is what separates targeted action from business as usual.

What you don't need

New funding. The departments involved already have budgets and staff. Code enforcement already conducts inspections. Sanitation already cleans up dumping. Public works already fixes streetlights. The intervention is redirecting a share of existing capacity toward specific locations, not building something from scratch. The primary upfront investment is analytical: someone needs to pull the crime data, identify priority areas, and match them with appropriate interventions. In most cities, this can be done with existing analysts or a short-term partnership.

Perfect analytics. You don't need a predictive model, a risk-terrain analysis, or a proprietary platform to get started. A map of where violent crime concentrates, built from the data your police department already collects, is sufficient. Analytical sophistication can grow over time, but it should never be the reason you haven't started.

A new department or program. This is a coordination mechanism, not an initiative with its own staff and letterhead. It lives in the space between existing departments, connecting what they already do to where it's most needed. Cities that treat this as a new program risk building something that's easy to defund. Cities that embed it as an operating rhythm make it harder to undo.

The timeline

The framing that guided Dallas, and that we'd recommend for any city, is: 90 days to launch, 180 days to measurable outcomes.

Weeks 1--4: Build the foundation.

Secure explicit commitment from the mayor or city manager. Designate a coordinator. Identify the first two to three agency leads who will participate. Pull crime data and generate an initial list of candidate priority locations. Hold an initial convening to explain the approach, set expectations, and

establish the meeting cadence. The goal by the end of month one is a functioning coordination structure with a first round of priority sites.

Weeks 5--8: Assess and deploy.

Conduct field assessments of priority locations: walk the sites, photograph conditions, talk to nearby residents and businesses. Match specific interventions to specific sites: which locations need code enforcement action? Where is the lighting problem? Which lots need clearing? Begin first-round interventions and establish the live tracking system (a shared spreadsheet or project tracker is fine to start). Hold the first full cycle of weekly coordination meetings.

Months 3--6: Build the rhythm.

This is where the model either takes hold or fades. The weekly cadence should be routine by now: strategy meetings to identify new sites and review field work, department meetings to track cases and assign follow-ups. The key indicator isn't how many cleanups you've done; it's whether high-priority locations are receiving repeat engagement. Sites should stay on the list until conditions have stabilized, not until the first intervention is complete. By month six, the coordinator should be able to report on which sites have been closed out, which are still active, and how intervention areas are performing relative to comparable non-intervention areas.

Month 6 and beyond: Assess, refine, institutionalize.

Evaluate results. Are intervention areas outperforming comparable non-intervention areas? Which types of interventions seem most effective in which contexts? Refine the priority list based on updated crime data. Identify what's working and expand it; identify what isn't and adjust. The goal is to shift from a pilot to an institutional rhythm, something that happens every week because it's how the city operates, not because someone is championing a project.

Common pitfalls

Cities that have attempted place-based coordination tend to encounter the same failure modes. Knowing them in advance makes them easier to avoid.

Spreading too thin. The instinct in city government is to serve the whole city equally. It feels politically safer than concentrating resources in specific areas. But even distribution is the default that these neighborhoods have been living under, and it isn't working. The data shows that roughly 20–30 areas in a typical city are trending toward escalation at any given time. That's a manageable list. Fifty locations with sustained, repeat engagement will have more impact than 500 locations with a single visit. The discipline of saying “these blocks first” is what makes this approach different from business as usual.

Confusing activity with impact. Counting cleanups, inspections, and lighting repairs feels productive. But the measure of success isn't how many work orders you've completed; it's whether violence is declining in the places you're targeting. Track interventions, but evaluate outcomes. If a location has received multiple rounds of engagement and conditions haven't improved, that's a signal to reassess the approach, not to keep doing the same thing.

Skipping the data. This model depends on knowing *where*. Without crime data to identify priority locations, you're guessing, and in most cities the guess will be wrong. The areas that show up in 311 complaints or council constituent calls aren't always the areas with the highest violence. Data doesn't need to be perfect, but it needs to exist.

Letting the cadence slip. Weekly meetings sound easy until competing priorities push them to bi-weekly, then monthly, then "as needed." The coordination rhythm is the engine. When it stops, so does follow-through. Protect the meeting cadence the way you'd protect any critical management function. If the coordinator is out, someone else runs the meeting. It doesn't get canceled.

Treating this as a one-time campaign. Environmental conditions deteriorate. Vacant lots refill with trash. Code violations recur. Absentee landlords don't change behavior after a single citation. A cleanup that makes a block look better for a week isn't a violence reduction strategy. Sustained change requires sustained attention. The cities that get results are the ones that commit to being in these places repeatedly, over months, not the ones that launch a visible push and move on.

About the Author

Joe Eichenbaum

Partner, 17A

Joe Eichenbaum is a partner at 17A, a management consulting firm that works exclusively with state and local government. His practice focuses on public safety analytics, place-based strategy, and operational coordination for cities. He has led analytical and advisory engagements for multiple city governments, with particular depth in crime concentration analysis, cross-agency coordination, and the use of administrative data to direct city resources toward concentrated need.

The analysis in this paper draws on work conducted in partnership with the City of Dallas and the Child Poverty Action Lab (CPAL), as well as multi-city research using open data from nine U.S. cities.

Contributor

Martin O'Malley

Senior Strategic Advisor, 17A

Martin O'Malley served as Mayor of Baltimore from 1999 to 2007 and as Governor of Maryland from 2007 to 2015. As mayor, he led Baltimore through a sustained reduction in violent crime using data-driven, place-based strategies, including the pioneering CitiStat performance management system. Governor O'Malley contributed to this paper as a 17A Senior Strategic Advisor, drawing on his experience leading public safety operations in one of America's most challenging urban environments.

Appendix

A. Definitions

Violent crime in this paper refers to homicide, robbery, and non-family aggravated assault. We exclude family-related aggravated assaults because these offenses reflect dynamics (domestic violence, intimate partner violence) that are less responsive to environmental intervention and are typically addressed through different policy channels.

This definition is narrower than the FBI's Uniform Crime Reporting (UCR) standard, which includes all aggravated assaults and forcible rape in its violent crime category. It is also narrower than many cities' internal definitions, which may include additional offense types or use local classification systems. We use the narrower definition because it more precisely captures the place-based, stranger-involved violence that environmental interventions are most likely to affect.

Homicide data in the national trends section (Section 2) uses the FBI UCR definition: murder and non-negligent manslaughter. City-level homicide counts are drawn from local police department data via the AH Datalytics Real-Time Crime Index (RTCI), which tracks reported homicides across major U.S. cities using official department sources.

Geographic concentration is measured by dividing a city into uniform spatial units (half-mile grid cells in the Dallas analysis) and ranking those units by the volume of violent crime incidents. "Top 20%" refers to the cells that collectively account for 20% of the city's total violent crime; "bottom 50%" refers to the cells that account for the lowest 50% of violent crime. Because crime is highly concentrated, the top 20% of violent crime typically occurs in a very small share of city geography (3.4% in Dallas), while the bottom 50% of crime is spread across a large geographic majority (86.2% in Dallas).

As described in Appendix C, measuring change in concentrated areas requires care about which year's data is used to define the geography. This paper uses a persistent-cell approach that avoids the methodological pitfalls of single-baseline analysis.

B. National Crime Trends: Data and Sources

National violent crime and homicide rates shown in Section 2 are drawn from the FBI Uniform Crime Reporting (UCR) program, which has collected standardized crime data from law enforcement agencies across the United States since 1960.

The data covers 1960–2025 and includes rates per 100,000 population for violent crime, murder, and property crime. Key reference points:

Year	Violent Crime Rate	Homicide Rate	Notes
1960	160.9	5.1	Series baseline
1980	596.6	10.2	First peak
1991	758.2	9.8	Historic high (violent crime)
2000	506.5	5.5	Post-1990s decline
2014	361.1	4.4	Pre-COVID low
2019	362.4	5.1	Pre-pandemic baseline
2020	380.3	6.7	COVID-era surge (+30% homicide)
2021	359.3	6.5	Elevated
2022	377.8	6.6	Elevated
2023	370.3	5.9	Beginning decline
2024	348.6	5.2	Approaching pre-pandemic levels
2025	316.8	4.3	Below pre-pandemic levels (partial year)

Rates are per 100,000 population. Source: FBI Uniform Crime Reporting Program.

City-level homicide rates referenced in Section 2 are drawn from the AH Datalytics Real-Time Crime Index (RTCI), which compiles reported homicide counts from official city police department sources. The RTCI covers major U.S. cities from 2018 to the present and is updated regularly. Source links for individual cities are available in the data files accompanying this paper.

C. Methodology: Crime Concentration Analysis

The concentration analysis used in this paper divides a city’s geography into uniform half-mile (approximately 800-meter) grid cells. Each cell is assigned the count of violent crime incidents (homicide, robbery, non-family aggravated assault) that occurred within its boundaries during the analysis period.

Ranking and tiering. Cells are ranked by total violent crime volume. Tiers are defined by cumulative share of total violent crime:

- *Top 20% tier:* The smallest set of cells that collectively account for 20% of the city’s total violent crime. In Dallas, this corresponds to 3.4% of city geography (35 cells).
- *Top 50% tier:* The smallest set of cells that collectively account for 50% of the city’s total violent crime. In Dallas, this corresponds to 13.8% of city geography (144 cells).
- *Bottom 50% tier:* All remaining cells — those that collectively account for the lowest 50% of violent crime. In Dallas, this corresponds to 86.2% of city geography (897 cells).

The baseline selection problem

When measuring year-over-year change in concentrated areas, results depend heavily on which year defines the geography. There are two natural approaches:

- *Current-year baseline (2025)*: Rank cells by 2025 volume, identify the top tier, then compare those specific cells' 2024 counts to their 2025 counts.
- *Prior-year baseline (2024)*: Rank cells by 2024 volume, identify the top tier, then compare those specific cells' 2024 counts to their 2025 counts.

These approaches select different cell sets, and the results can diverge dramatically. Using Dallas as an example: the top-20% tier defined by 2025 data shows a +6.7% increase in violent crime. The top-20% tier defined by 2024 data shows a **-32.9%** decline. The same city, same time period, same underlying data — a 40 percentage point swing driven entirely by which year defines the geography.

This happens because the 2025 baseline mechanically selects cells that spiked or stayed high, while the 2024 baseline selects cells that were at their peak and are likely to regress. Neither approach is wrong, but neither is bias-free.

The persistent-cell approach

To cut through this problem, the analysis in this paper classifies cells into four groups based on their status in both years:

- *Persistent*: In the top tier in both 2024 and 2025. These are chronic high-crime locations.
- *Rotated in*: In the top tier in 2025 but not 2024. These areas newly entered concentration.
- *Rotated out*: In the top tier in 2024 but not 2025. These areas improved relative to the rest of the city.
- *All other*: Not in the top tier in either year.

This classification avoids the baseline selection bias entirely. Persistent cells are in the top tier regardless of which year you use — they represent the genuinely chronic locations. Their year-over-year change is the cleanest read on whether concentrated areas are improving.

Geographic churn is measured by the Jaccard index: the number of cells in the top tier in both years, divided by the number in the top tier in either year. A Jaccard overlap of 100% would mean the same cells are concentrated in both years; lower values indicate more geographic turnover.

Dallas grid parameters:

- Grid cell size: 0.5 miles x 0.5 miles
 - Total grid cells covering Dallas city limits: 1,041
 - Violent crime definition: murder, robbery, aggravated assault
 - Baseline year: 2024
 - Evaluation year: 2025
-

- Crime data source: Dallas Police Department NIBRS quarterly data, geocoded to grid cells (see Appendix I)

D. Multi-City Concentration Analysis

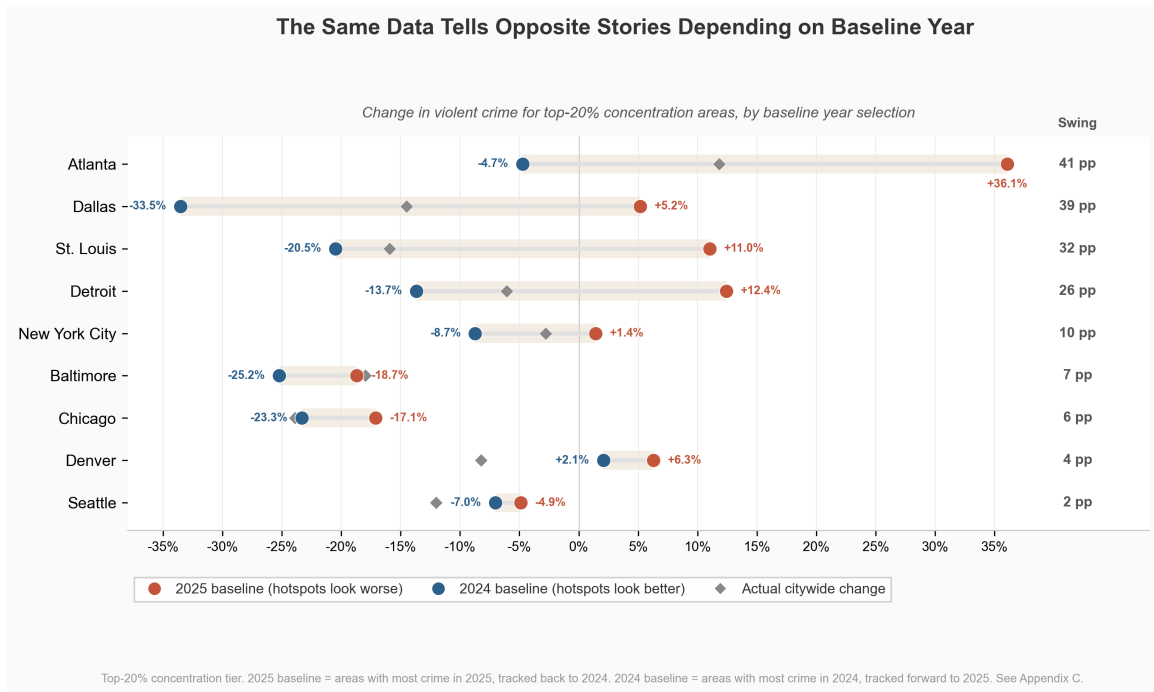
The tables below present the full results of the concentration analysis across nine cities, including the baseline sensitivity comparison, geographic churn measures, and persistent-cell findings discussed in Section 2.

D.1 Baseline Sensitivity: Top-20% Tier

The same data can tell different stories depending on which year defines the concentrated geography. This table shows how the measured change in the top-20% tier varies by baseline year.

City	Citywide Change	Top-20% (2025 base)	Top-20% (2024 base)	Swing	Jaccard Overlap
Dallas	-14.4%	+6.7%	-32.9%	39.6pp	35%
St. Louis	-15.9%	+5.2%	-20.5%	25.7pp	44%
Atlanta	+11.8%	+32.5%	-4.7%	37.3pp	40%
Detroit	-6.1%	+11.1%	-13.7%	24.8pp	43%
New York City	-2.8%	+1.4%	-8.7%	10.1pp	55%
Denver	-8.2%	+6.3%	+2.1%	4.2pp	79%
Chicago	-23.9%	-17.1%	-23.3%	6.2pp	67%
Seattle	-12.0%	-4.9%	-7.0%	2.1pp	67%
Baltimore	-18.0%	-18.7%	-25.2%	6.5pp	67%

pp = percentage points. Jaccard overlap = share of cells in the top-20% tier in both years out of all cells in the tier in either year.



Cities with low Jaccard overlap (Dallas, St. Louis, Atlanta, Detroit) show the largest swings between baselines. Cities with high overlap (Denver, Seattle, Chicago, Baltimore) show consistent results regardless of baseline. This confirms that the divergent readings are an artifact of geographic churn, not a real difference in how concentrated areas are trending.

D.2 Baseline Sensitivity: Top-50% Tier

City	Citywide Change	Top-50% (2025 base)	Top-50% (2024 base)	Swing	Jaccard Overlap
Dallas	-14.4%	-1.8%	-26.0%	24.2pp	50%
St. Louis	-15.9%	+1.4%	-22.8%	24.2pp	45%
Atlanta	+11.8%	+20.1%	-7.1%	27.2pp	53%
Detroit	-6.1%	+5.6%	-12.0%	17.5pp	52%
New York City	-2.8%	0.0%	-5.5%	5.5pp	74%
Denver	-8.2%	+7.3%	-12.7%	20.0pp	51%
Chicago	-23.9%	-19.1%	-27.4%	8.3pp	64%
Seattle	-12.0%	-5.7%	-9.5%	3.9pp	82%
Baltimore	-18.0%	-13.8%	-21.6%	7.9pp	66%

The 50% tier is broader and somewhat more stable, but the baseline sensitivity pattern persists. In volatile cities, the swing remains large (Dallas: 24pp, St. Louis: 24pp, Atlanta: 27pp).

D.3 Geographic Churn Summary

City	20% Tier: Jaccard	20% Tier: % Newly Concentrated	50% Tier: Jaccard	50% Tier: % Newly Concentrated
Dallas	35%	49%	50%	35%
St. Louis	44%	35%	45%	35%
Atlanta	40%	43%	53%	35%
Detroit	43%	39%	52%	30%
New York City	55%	30%	74%	15%
Denver	79%	0%	51%	28%
Chicago	67%	17%	64%	21%
Seattle	67%	14%	82%	6%
Baltimore	67%	22%	66%	20%

“% Newly Concentrated” = share of 2025’s top-tier cells that were not in the top tier in 2024.

In the most volatile cities, half or more of the highest-concentration geography turns over in a single year. Even in relatively stable cities, 14–28% of the worst cells are new. This is the churn that demands adaptive targeting.

D.4 Persistent vs. Rotating Cells: Top-20% Tier

This table decomposes each city’s concentration geography into four groups: persistent (high both years), rotated in (newly high), rotated out (no longer high), and all other cells.

City	Persistent Cells	Persistent VC Change	Rotated In: Cells	Rotated In: VC Change	Rotated Out: Cells	Rotated Out: VC Change
Dallas	18	−16.0%	17	+58.6%	17	−55.0%
St. Louis	11	−6.7%	6	+46.7%	8	−42.5%
Atlanta	12	+13.0%	9	+84.5%	9	−32.3%
Detroit	25	−2.6%	16	+49.6%	17	−32.3%
New York City	50	−3.9%	21	+21.5%	20	−23.8%
Denver	11	+6.3%	0	n/a	3	−20.3%
Chicago	39	−18.6%	8	−4.8%	11	−43.6%
Seattle	6	−6.6%	1	+13.0%	2	−9.0%
Baltimore	14	−22.6%	4	+8.0%	3	−41.9%

The persistent cells — the genuinely chronic locations — are generally declining alongside or near the citywide rate. In Dallas, persistent cells declined 16.0% versus 14.4% citywide. In Baltimore, 22.6% versus 18.0%. In Chicago, 18.6% versus 23.9%.

The rotation mechanic is large and symmetric: cells that rotated in typically show +30% to +85% increases (this is essentially definitional — they entered the tier because they spiked). Cells that rotated out show −30% to −55% declines (they left because they cooled). These extreme swings are what create divergent readings between baselines, but they describe the mechanics of geographic

turnover, not a trend in chronic areas.

D.5 Persistent vs. Rotating Cells: Top-50% Tier

City	Persistent Cells	Persistent VC Change	Rotated In: Cells	Rotated In: VC Change	Rotated Out: Cells	Rotated Out: VC Change
Dallas	94	−13.6%	50	+63.6%	44	−60.7%
St. Louis	43	−9.5%	23	+50.0%	29	−50.4%
Atlanta	60	+5.2%	33	+104.5%	21	−54.7%
Detroit	96	−3.4%	41	+44.3%	48	−34.4%
New York City	230	−2.4%	41	+28.6%	39	−32.5%
Denver	49	−1.7%	19	+105.6%	28	−52.0%
Chicago	154	−23.1%	41	+14.4%	44	−50.5%
Seattle	31	−7.2%	2	+52.1%	5	−33.1%
Baltimore	56	−18.5%	14	+29.9%	15	−40.6%

The same patterns hold at the 50% tier. Persistent cells track near citywide declines (Dallas: −13.6% vs. −14.4% citywide). The rotation mechanics are even larger at this tier, with cells entering concentration often doubling their prior-year crime counts.

D.6 City-Specific Parameters

City	Grid Cell Size	Total Cells	Violent Crime Definition	Data Source
Dallas	0.5 mi	1,041	Murder, robbery, aggravated assault	Dallas Police Department (NIBRS)
St. Louis	0.4 mi	416	Homicide, robbery, aggravated assault	SLMPD
New York City	0.33 mi	2,854	Murder & non-negligent manslaughter, robbery, felony assault	NYC Open Data
Atlanta	0.4 mi	614	Homicide, robbery, aggravated assault	Atlanta Police Open Data
Detroit	0.5 mi	595	Homicide, robbery, aggravated assault	Detroit Open Data
Denver	0.4 mi	671	Murder, robbery, aggravated assault	Denver Open Data
Chicago	0.4 mi	1,300	Homicide, robbery, aggravated battery/assault	Chicago Data Portal
Seattle	0.4 mi	480	Homicide, robbery, aggravated assault	Seattle Open Data
Baltimore	0.4 mi	470	Homicide, robbery, aggravated assault	Open Baltimore

All data: full-year 2024 vs. 2025 incident counts. Eight cities use data from municipal open data portals; Dallas uses NIBRS data provided directly by the Dallas Police Department (see Appendix I for context on this data source decision). Grid cell sizes vary by city to account for differences in city area and density. Violent crime definitions vary slightly across cities based on local classification systems; all include homicide, robbery, and aggravated assault/battery.

E. Dallas Intervention Analysis: Detailed Results

Section 4 presents a simplified comparison of intervention versus non-intervention areas. This appendix provides the full data underlying those comparisons.

E.1 Deployment Summary

The Dallas coordinated action model operated from 2024 into 2025. Key deployment metrics:

- Total micro-areas receiving intervention: approximately 90
- Total discrete interventions completed: over 250
- Average rounds of engagement per high-priority site: 3–4
- Intervention types: vegetation/visibility, lighting, trash/dumping, property conditions, infrastructure, code enforcement, service connections

E.2 Intervention vs. Non-Intervention Outcomes by Risk Tier

The analysis compares violent crime change (2024–2025) in cells that received coordinated intervention versus cells at the same risk level that did not.

Table E.2a: Violent Crime Change by Concentration Tier and Intervention Status

Risk Tier	% of City Geography	% of City Violent Crime	Intervention Areas	Non-Intervention Areas	Difference
Top 20% (highest risk)	~3%	~20%	+4.7%	+11.1%	6.4 pp better
Top 50% (high risk)	~14%	~50%	–8.1%	+0.7%	8.8 pp better
Bottom 50% (lower risk)	~86%	~50%	–27.4%	–23.4%	4.0 pp better

pp = percentage points. “Better” means the intervention group had a lower (or less positive) rate of change.

Key observations:

Top 20% tier. Both intervention and non-intervention areas in this tier saw increases in violent crime — this tier did not benefit from citywide tailwinds. However, intervention areas saw less than half the increase of non-intervention areas (+4.7% vs. +11.1%). Dampening the magnitude of surges matters: it is the difference between a temporary spike and an area that becomes entrenched in the highest tier.

Top 50% tier. This is where the stabilization effect was most pronounced. Intervention areas declined 8.1% while comparable non-intervention areas were essentially flat (+0.7%). In a system where concentration persists through geographic churn, the areas that fail to decline are precisely those at risk of rotating into higher-risk tiers. Intervention pulled areas back from that edge.

Bottom 50% tier. Even in areas already benefiting from broader citywide improvement, intervention areas declined faster than non-intervention areas (-27.4% vs. -23.4%).

E.3 Homicide-Specific Results

As noted in Section 4, the intervention effect is even larger when measured by homicide alone rather than violent crime overall. At every risk tier, the gap between intervention and non-intervention areas is wider for homicide than for the broader violent crime category.

[PLACEHOLDER: Detailed homicide table to be added when final data is confirmed]

E.4 Caveats

This analysis is observational, not experimental. Intervention areas were selected based on crime data and environmental conditions, not randomly assigned. Several factors should be considered when interpreting results:

- *Selection effects:* Intervention areas were chosen because they had high crime and conditions amenable to environmental intervention. It is possible that some characteristic of these areas — beyond the intervention itself — contributed to better outcomes.
- *Temporal scope:* Results reflect one year of data (2024–2025). Longer-term sustainability is not yet demonstrated.
- *Attribution:* Multiple factors affect crime trends simultaneously. The analysis cannot isolate the effect of environmental intervention from other concurrent changes (policing activity, seasonal patterns, economic conditions).
- *Dosage variation:* Not all intervention areas received the same intensity or type of intervention. The analysis treats all intervention areas equally rather than differentiating by dosage.

Despite these limitations, the consistency of the pattern — intervention areas outperforming non-intervention areas at every risk level, for both violent crime overall and homicide specifically — suggests that the coordinated approach contributed meaningfully to outcomes.

F. Statistical Significance of Year-Over-Year Changes

The comparisons in this paper — including the citywide trends, concentration tier analysis, and intervention evaluation — rely on year-over-year percent changes. This appendix examines whether those changes are statistically distinguishable from the random variation inherent in crime data, using methods drawn from the Poisson distribution.

F.1 Why This Matters: The Limits of Percent Change

Crime counts at the micro-geographic level are small numbers. A half-mile grid cell in Dallas might see 8 violent crimes one year and 14 the next — a 75% increase that sounds alarming but represents only 6 additional incidents. Whether that increase reflects a genuine shift in conditions or ordinary random fluctuation is a question percent change alone cannot answer.

The Poisson distribution provides a framework for answering it. Under the assumption that crime at a given location follows a Poisson process (a reasonable first approximation for rare events distributed across space and time), the expected random variation in crime counts can be estimated directly from the counts themselves. A statistical test — the Poisson z-score — can then distinguish changes that exceed expected variation from those that fall within it.

The Poisson z-score is calculated as:

$$z = 2 \times \left[\sqrt{\text{Current}} - \sqrt{\text{Historical}} \right]$$

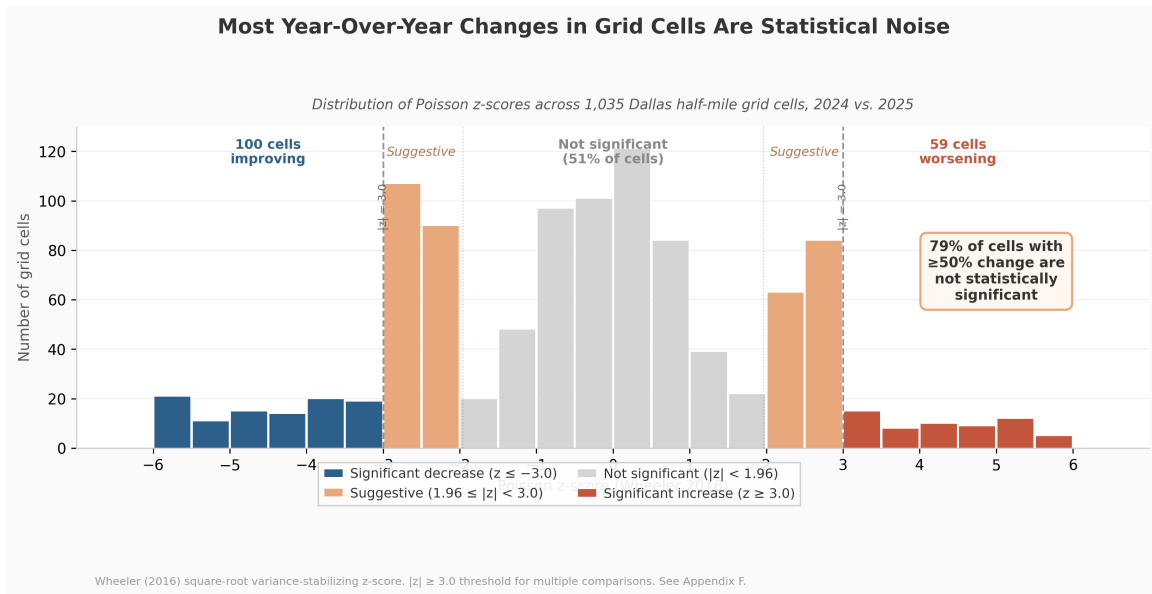
This approximates a standard normal distribution under the null hypothesis that the underlying crime rate has not changed. Positive z-values indicate an increase beyond what randomness would predict; negative z-values indicate a decrease. A threshold of $|z| \geq 3.0$ corresponds roughly to a 1-in-1,000 probability of occurring by chance — appropriate when monitoring many geographic units simultaneously, as the risk of false positives compounds with each additional comparison. This threshold and formula follow Wheeler (2016), who developed the approach specifically for crime analysis applications.

F.2 Cell-Level Analysis: How Much Year-Over-Year Change Is Noise?

Applying the Poisson z-score to all 1,035 Dallas grid cells with any violent crime in either year produces a striking result: the majority of year-over-year changes are not statistically distinguishable from random variation.

Significance Level	Cells	Share
Significant at $p < 0.001$ ($ z \geq 3.0$)	159	15.4%
Suggestive ($1.96 \leq z < 3.0$)	352	34.0%
Not significant ($ z < 1.96$)	524	50.6%

The $|z| \geq 3.0$ threshold is used throughout this appendix, following Wheeler's recommendation for contexts involving multiple simultaneous comparisons. The suggestive category ($1.96 \leq |z| < 3.0$) would be considered significant in a single-comparison context but is treated conservatively here.



Of the 159 cells that cleared the strictest threshold, 100 showed statistically significant *decreases* and 59 showed significant *increases* — a ratio consistent with the overall citywide decline.

The most revealing finding concerns cells with large percent changes. Of the 549 cells showing $\geq 50\%$ year-over-year change in either direction, 431 — 79% — were not statistically significant at the $|z| \geq 3.0$ threshold. A cell that went from 2 crimes to 6 shows a 200% increase, but the Poisson z-score is only 1.7: well within the range of random fluctuation. A cell that went from 10 to 4 shows a 60% decrease, but $z = -2.0$ — suggestive but not definitive.

This has direct implications for how year-over-year comparisons should be interpreted. At the individual cell level, most changes — even dramatic-looking ones — cannot be confidently attributed to anything other than randomness. The patterns become meaningful primarily when aggregated across many cells, where the random noise cancels out and genuine trends emerge.

F.3 Intervention Analysis: Testing the Coordinated Intervention Effect

The intervention evaluation in Section 4 and Appendix E compares percent changes in violent crime between intervention areas and non-intervention areas at each concentration tier. The question this section addresses is whether those differences are statistically significant — that is, larger than what Poisson variation alone would produce.

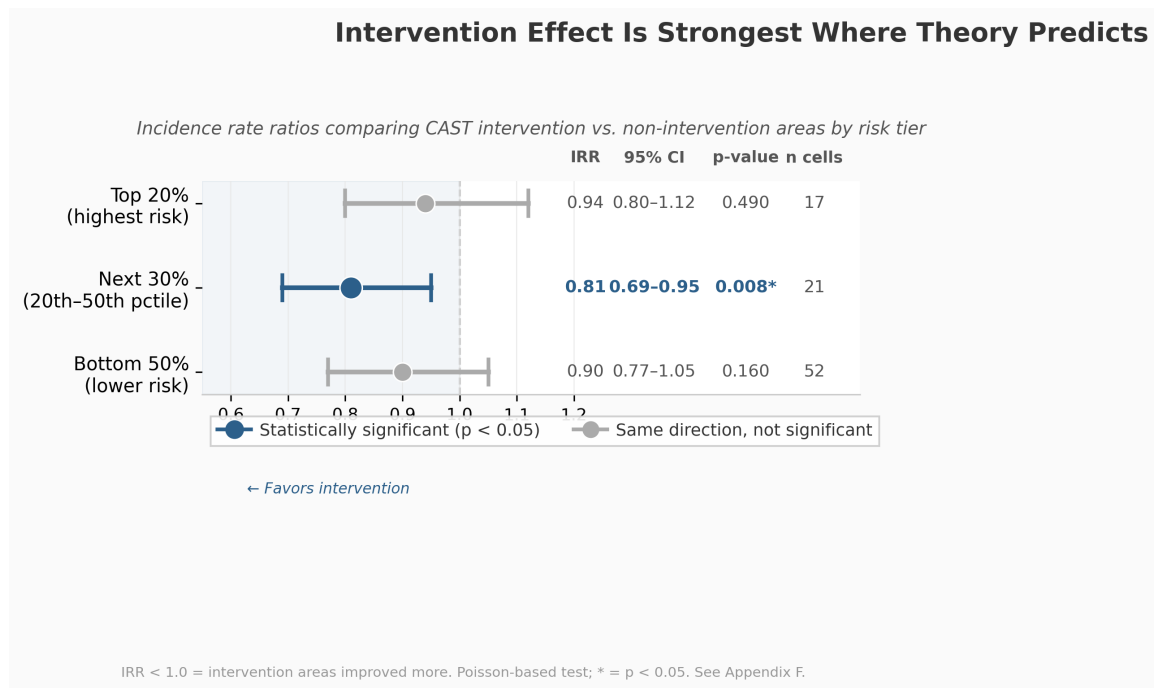
Because the intervention and non-intervention groups differ substantially in size (90 intervention cells versus 931 non-intervention cells across all tiers), the appropriate test is the incidence rate ratio (IRR). The IRR compares the *rate* of change — the ratio of post-intervention to pre-intervention crime —

between the two groups, rather than comparing raw counts. An IRR below 1.0 indicates that the intervention group improved more than the comparison group.

Table F.3: Incidence Rate Ratio Tests by Concentration Tier

Risk Tier	Intervention Cells	Intervention (2024 → 2025)	Non-Intervention (2024 → 2025)	IRR	95% CI	p-value
Top 20%	17	615 → 644 (+4.7%)	441 → 490 (+11.1%)	0.94	0.80–1.12	0.49
Next 30%	21	450 → 339 (–24.7%)	1,461 → 1,363 (–6.7%)	0.81	0.69–0.95	0.008
Bottom 50%	52	432 → 303 (–29.9%)	3,249 → 2,545 (–21.7%)	0.90	0.77–1.05	0.16

IRR < 1.0 indicates intervention areas improved more than non-intervention areas. 95% CI = 95% confidence interval for the IRR. p-values are two-tailed.



The Next 30% tier — areas between the 20th and 50th percentile of crime concentration — shows a statistically significant intervention effect. The IRR of 0.81 (95% CI: 0.69–0.95, $p = 0.008$) indicates that intervention areas experienced a rate of change roughly 19% more favorable than comparable non-intervention areas. In concrete terms: intervention areas in this tier declined 24.7% while non-intervention areas in the same tier declined only 6.7%. This difference is larger than expected from Poisson variation alone.

The Top 20% and Bottom 50% tier differences go in the same direction but are not statistically significant. In the Top 20% tier, the IRR of 0.94 reflects a pattern consistent with the Section 4 narrative —

intervention areas saw less than half the increase of non-intervention areas — but the confidence interval (0.80–1.12) includes 1.0, meaning the difference cannot be confidently distinguished from chance. This is largely a function of sample size: with only 17 intervention cells and 16 non-intervention cells in this tier, a much larger effect would be required to achieve statistical significance. In the Bottom 50% tier, the IRR of 0.90 is similarly favorable but not significant ($p = 0.16$), again reflecting the difficulty of detecting modest effects with limited statistical power.

F.4 What This Means

Three conclusions follow from this analysis.

First, the intervention effect is strongest — and statistically confirmed — in the tier where the paper’s theory predicts it should be. The Next 30% tier represents areas with enough crime to be at risk of escalating into the worst concentration tier, but not so entrenched that deeper structural forces dominate. These are the areas where environmental intervention has the most theoretical leverage, and they are the areas where the data shows the clearest effect.

Second, the absence of statistical significance in the Top 20% and Bottom 50% tiers does not mean the intervention had no effect there. It means the data is insufficient to distinguish the observed differences from random variation. With crime counts this small in individual cells, detecting modest effects requires either larger treatment groups, longer time periods, or both. The consistent direction of the effect across all three tiers — intervention areas outperforming non-intervention areas at every level — is suggestive even where individual tier-level tests are underpowered.

Third, the cell-level Poisson analysis reinforces the central finding of Section 2: much of the apparent geographic churn in crime concentration is noise. When 79% of cells with $\geq 50\%$ year-over-year change are not statistically significant, the implication is that a meaningful share of what appears to be geographic movement is random variation rather than genuine shifts in underlying conditions. This makes the case for adaptive targeting stronger, not weaker: if cities cannot tell in real time which changes are signal and which are noise, frequent reassessment and rapid response become even more important.

F.5 Methodological Notes

The Poisson z-score and incidence rate ratio tests used in this appendix assume that crime counts follow a Poisson distribution — that is, that the variance equals the mean. Crime data sometimes exhibits overdispersion (variance exceeding the mean), which would make the standard errors in these tests too small and the p-values too optimistic. The $|z| \geq 3.0$ threshold used for the cell-level analysis (rather than the conventional 1.96) provides a partial guard against this. For the tier-level IRR tests, overdispersion is less of a concern because aggregation across many cells tends to smooth out the excess variation present at the individual cell level.

The analysis uses 2024 concentration tiers to classify cells, matching the methodology in Appendix E. Crime counts are annual totals for full-year 2024 and 2025.

References: Wheeler, A. P. (2016). Tables and graphs for monitoring temporal crime trends: Translating theory into practical crime analysis advice. International Journal of Police Science & Management, 18(3), 159–172.

G. Multi-City Stability Analysis: Methodology

Section 2 classifies every half-mile area in each city as *stable*, *improving*, or *worsening* based on whether its year-over-year change in violent crime exceeds what random variation would predict. This appendix describes how that classification works.

G.1 The Problem: Distinguishing Signal from Noise in Small-Area Crime Data

Crime counts in half-mile grid cells are small numbers. A cell with 5 violent crimes one year and 9 the next shows an 80% increase — but with counts this low, swings of that size happen routinely by chance. The question is whether a given change reflects something real (a shift in underlying conditions) or falls within the range of normal year-to-year variation.

The Poisson distribution provides a framework for answering this. If crime at a location follows a roughly Poisson process — a standard assumption in criminology for rare events distributed across space and time — then the expected random variation can be estimated directly from the prior year’s count.

G.2 The Wheeler Z-Score

For each grid cell with at least one violent crime in the prior year, a z-score is computed using the variance-stabilizing square-root transform developed by Wheeler (2016):

$$z = 2 \times \left[\sqrt{\text{observed}_{2025}} - \sqrt{\text{expected}_{2024}} \right]$$

where the cell’s own 2024 count serves as the expected value and its 2025 count is the observed value. This formula approximates a standard normal distribution under the null hypothesis that the underlying crime rate has not changed.

Wheeler (2016) specifically recommends this square-root transform over the more common standard Poisson z-score formula $[(\text{observed} - \text{expected}) / \sqrt{\text{expected}}]$ because it produces more reliable significance estimates at the low crime counts typical of half-mile grid cells. The standard formula tends to overstate significance when expected counts are small (e.g., 1–3 crimes), which is common in the lower half of cells. The square-root transform corrects for this.

- Positive $z \rightarrow$ crime increased beyond expected variation (worsening)
- Negative $z \rightarrow$ crime decreased beyond expected variation (improving)
- Cells with $2024 = 0$ receive no z -score (a Poisson expectation cannot be computed from a zero baseline)

G.3 Significance Threshold and Multiple Comparisons

A typical city analysis tests hundreds of grid cells simultaneously. When running many tests at once, the risk of false positives compounds: at the conventional $p < 0.05$ level ($|z| \geq 1.96$), roughly 5% of cells would appear significant by chance alone even if nothing changed anywhere. In a city with 1,000 testable cells, that would produce ~50 false positives.

Wheeler (2016) recommends a $|z| \geq 3.0$ threshold for crime monitoring contexts involving multiple comparisons. This corresponds roughly to $p < 0.001$ per test and substantially reduces the false positive rate. At this threshold, the expected number of false positives in a 1,000-cell analysis drops from ~50 to ~3.

The classifications in Section 2 use this $|z| \geq 3.0$ threshold:

Classification	Criterion	Interpretation
Improving	$z \leq -3.0$	Decrease large enough to indicate an emerging trend, not just fluctuation
Worsening	$z \geq 3.0$	Increase large enough to indicate an emerging trend
Stable	$-3.0 < z < 3.0$, or z undefined	Change within the range of normal year-to-year variation, or insufficient baseline data

G.4 Cross-City Application

The same formula and threshold were applied identically across all ten cities. For each city:

1. Grid cell data was loaded from the multi-city crime explorer dataset, which contains half-mile cells with `crimes_2024` and `crimes_2025` columns (violent crime counts by year).
2. Cells with `crimes_2024 > 0` were included in the z -score computation. Cells with zero 2024 crimes were classified as stable by default (z -score undefined).
3. The Wheeler z -score was computed for each eligible cell.
4. Cells were classified as improving ($z \leq -3.0$), worsening ($z \geq 3.0$), or stable ($|z| < 3.0$).
5. City-level percentages were computed as shares of *all* cells (including those with zero baselines), not just testable cells.

The number of testable cells (those with at least one violent crime in 2024) varies by city, from roughly 200 in smaller cities to over 2,000 in New York. Across all ten cities combined, 7,332 cells were testable.

G.5 Results Summary

City	Total Cells	Testable	Stable	Improving	Worsening	Ratio
Chicago	1,300	892	88%	12%	1%	17:1
Oakland	350	271	73%	25%	2%	11:1
Baltimore	470	357	88%	11%	1%	9:1
St. Louis	416	286	93%	6%	1%	5:1
Dallas	1,041	894	86%	11%	3%	4:1
Seattle	480	301	93%	7%	<1%	17:1
Denver	671	403	93%	5%	1%	4:1
Detroit	595	477	94%	4%	2%	2:1
New York City	2,854	2,164	92%	5%	3%	2:1
Atlanta	614	287	97%	2%	1%	2:1
All cities	8,791	7,332	90%	8%	2%	4.3:1

Ratio = improving cells ÷ worsening cells. Percentages are of total cells (including non-testable). “Stable” includes cells with zero 2024 baseline (z-score undefined).

The cross-city pattern is remarkably consistent: regardless of whether a city’s overall crime declined sharply (Chicago, **-24%**) or increased (Atlanta, **+12%**), roughly 90% of its geography showed no statistically significant change. The variation is primarily in the ratio of improving to worsening cells, which tracks closely with the citywide trend direction and magnitude.

G.6 Where Citywide Improvement Comes From

The stability classification also reveals how citywide declines are distributed across the three groups. In some cities, the majority of the net decline comes from concentrated improvement in the ~8% of cells classified as improving. In others, it comes from the incremental gains across the much larger stable category.

On average, stable cells — while not showing statistically significant individual changes — are gently improving, with roughly 0.5 to 2.5 fewer violent crimes per cell per year. Because stable cells account for ~90% of geography, these small per-cell changes sum to a substantial share of the citywide decline.

The mix varies: in Dallas, 88% of the net violent crime decline came from the concentrated improvement cells; in St. Louis, 73% came from distributed improvement across stable cells. Both pathways contribute to citywide improvement; the stability analysis identifies the small slice of geography where that improvement is not occurring and where conditions may be deteriorating.

G.7 Assumptions and Limitations

- Poisson assumption. The model assumes crime counts are independent events with a stable underlying rate — a simplification, but standard in the criminology of place. Crime data sometimes exhibits overdispersion (variance exceeding the mean), which could produce more significant results than warranted. The conservative $|z| \geq 3.0$ threshold partially mitigates this.
- Single year-over-year comparison. The analysis compares 2024 to 2025. A multi-year analysis (e.g., 2022–2025) would provide a more robust picture of trends but requires consistent multi-year data across all cities, which was not available for all ten.
- Zero-baseline cells. Cells with zero violent crimes in 2024 are excluded from the z-score computation and classified as stable by default. In some cities, this is a large share of cells (e.g., ~45% in Dallas). These cells are disproportionately low-crime areas where the classification is unlikely to be wrong, but a small number may have experienced genuine emergence of violence from a zero baseline.
- Grid cell size. All cities use approximately half-mile grid cells, but exact dimensions vary slightly (0.33 to 0.5 miles) to account for differences in city area and density. The z-score is computed within each city's own grid and is not directly comparable across cities in absolute terms, though the percentage breakdowns are comparable.

Reference: Wheeler, A. P. (2016). Tables and graphs for monitoring temporal crime trends: Translating theory into practical crime analysis advice. International Journal of Police Science & Management, 18(3), 159–172.

H. Data Sources and References

H.1 Data Sources

Source	What It Provides	Coverage	Used In
FBI Uniform Crime Reporting (UCR) Program	National violent crime and homicide rates	1960–present, annual	Section 2 (national trends)
AH Datalytics Real-Time Crime Index (RTCI)	City-level homicide counts and rates	2018–present, multiple cities	Section 2 (city comparisons)
Dallas Police Department (NIBRS)	Incident-level crime data (geocoded), provided directly by DPD	2024–2025	Section 4, Appendix C–E
City open data portals	Incident-level crime data for multi-city analysis	Varies by city	Appendix D, G
FBI NIBRS	Incident-level national crime data	Ongoing transition from UCR	Referenced in Section 5
311 / Service request systems	Code violations, illegal dumping, lighting complaints	Varies by city	Referenced in Sections 3, 4, 6

H.2 References

Branas, C. C., South, E., Kondo, M. C., Hohl, B. C., Bourgois, P., Wiebe, D. J., & MacDonald, J. M. (2018). Citywide cluster randomized trial to restore blighted vacant land and its effects on violence, crime, and safety. *Proceedings of the National Academy of Sciences*, 115(12), 2946–2951.

Chalfin, A., Hansen, B., Lerner, J., & Parker, L. (2021). Reducing crime through environmental design: Evidence from a randomized experiment of street lighting in New York City. *Criminology*, 60(1), 3–44.

Weisburd, D. (2015). The law of crime concentration and the criminology of place. *Criminology*, 53(2), 133–157.

Wheeler, A. P. (2016). Tables and graphs for monitoring temporal crime trends: Translating theory into

practical crime analysis advice. *International Journal of Police Science & Management*, 18(3), 159–172.

Wheeler, A. P., & Reuter, S. (2021). Redrawing hot spots of crime in Dallas, Texas. *Police Quarterly*, 24(2), 159–184.

[Additional references to be added as drafting continues]

I. Data Source Stability and Reproducibility

Open data portals — particularly those built on Socrata — are live systems. Datasets are updated, re-structured, and occasionally re-published in ways that can silently affect downstream analyses. This appendix documents data stability issues encountered during this project and the steps taken to address them.

I.1 The Problem: Mutable Data Sources

The multi-city concentration analysis (Appendix D) draws incident-level crime data from nine municipal open data portals. These portals serve as the public interface for police department records, but they are not static archives. Records may be added, reclassified, deduplicated, or removed after initial publication, and dataset schemas can change without notice.

This creates a reproducibility challenge: the same API query run against the same endpoint on two different dates can return materially different results. Unlike a published dataset with a fixed version, an open data portal is a moving target.

I.2 Dallas: Observed Discrepancy

During validation, we identified a significant discrepancy between two data pulls from the Dallas Police Incidents dataset (Socrata resource `qv6i-rr17`):

Year	Earlier Data Pull	Current API Query (Feb 2025)	Difference
2024	6,648 violent crimes	4,989 violent crimes	–25%
2025	5,684 violent crimes	4,455 violent crimes	–22%

The earlier pull was used to generate the Crime Explorer grid and had been validated against Dallas Police Department reporting. The current API query uses the same offense filter (murder, robbery, aggravated assault via NIBRS crime type) but returns substantially fewer records. The gap is not explained by differences in filtering logic — the same offense strings return fewer matching records from the endpoint today than they did when the data was originally exported.

Possible explanations include dataset re-publication with different deduplication rules (incident-level vs. offense-level counting), retroactive reclassification of offense types, or replacement of the underlying dataset. We were unable to determine the precise cause, as Socrata does not maintain a public change log for dataset revisions.

For the Dallas analysis in this paper, we use data provided directly by the Dallas Police Department, which matches the earlier (higher) totals and is consistent with DPD's own reporting.

I.3 Implications for the Multi-City Analysis

The concentration and churn metrics central to this paper (Jaccard indices, tier classifications, persistent-cell analysis) measure geographic overlap rather than absolute crime counts. These metrics are robust to uniform undercounting — if every part of a city is undercounted by the same proportion, cells rank in the same order and the same cells appear in the same tiers. The Jaccard values and churn patterns reported in Appendix D are therefore unlikely to be materially affected by the Dallas data issue, provided the undercount is geographically uniform (which is expected, since the filtering applies citywide without a spatial component).

However, absolute crime counts and year-over-year percentage changes at the city level are directly affected. Any citywide totals for Dallas reported in this paper should be understood as drawn from DPD-provided data, not from the current state of the open data portal.

I.4 Recommendations for Reproducibility

Based on this experience, we recommend the following practices for analyses that depend on municipal open data:

1. Archive raw data pulls. Save timestamped copies of raw API responses or CSV exports at the time of analysis. Do not assume the same query will return the same results later.
2. Validate against authoritative sources. Where possible, cross-check open data totals against official department reports, UCR submissions, or direct agency correspondence.
3. Document query parameters. Record the exact API endpoint, filters, and date of access for each data pull.
4. Pin to a data version. If the portal supports dataset versioning or snapshots, reference a specific version rather than the live endpoint.

The raw data files and query parameters used in this analysis are available in the project repository.
