

METHODOLOGY AND DATA COMPANION

OPEN ACCESS

The Safe-City Myth

Methodology and data companion to the working paper

Gavriel Brown, PhD

A full white paper, with the complete introduction, related literature, data and design, results, discussion, limitations, conclusion, and references, is available on request.

© 2026 Heshbon and the author. This companion explains how the working paper was built and how to read its results. It can be shared and cited with attribution.

At a glance

Research question	Is the city or metro the right unit at which to judge residential safety, or does safety vary so much within a metro that the city label misleads?
Data	A database of 319 Orthodox Jewish communities compiled by the author from public sources, with zip-level national safety percentiles for 264 to 265 of them, plus home values and rents.
Unit of analysis	The individual community, keyed to its most common residential zip code, within its metropolitan area.
Statistical method	A percentile-rank index of zip-code-level national safety, read through within-metro dispersion and Pearson correlation.
Headline finding	Within one metro, communities span nearly the whole national range (New York runs from the 6th to the 99th percentile), and a city-relative grade flatters national standing 86% of the time.
What it shows	Safety varies at the neighborhood scale, so the name of the city carries little decision-relevant information for a household.

The question behind the paper

Families usually begin a move by asking whether a city is safe or dangerous, a question that feels natural and answerable. The paper argues that this is close to the wrong question, because the household actually chooses a specific neighborhood and address, not a metro.

The gap between the two scales is not a rounding error. As the results show, it is frequently the difference between the safest decile of the country and the least safe, which is why the unit at which safety is measured turns out to matter as much as the measurement itself.

Where the data come from

The community database. The primary dataset is a database of 319 identified Orthodox Jewish communities compiled by the author from public sources. Because observant families need walking access to a synagogue, a mikvah (ritual bath), and day schools, these communities are

tightly bounded, well-identified neighborhoods embedded in larger metros, which is the structure needed to separate the city question from the neighborhood question. Each community is keyed to its most common residential zip code.

Two safety instruments. Safety is measured two ways, both built on FBI Uniform Crime Reporting and NIBRS data. A national percentile (CrimeGrade) ranks a community's zip code against all U.S. zip codes, where a higher percentile denotes lower reported-crime density, and it is available for 264 to 265 of the 319 communities. A city-relative grade (DoorProfit) assigns an A to F letter reflecting how a community compares within its own metro.

Prices. Home values from Zillow ZHVI and rents from the Census American Community Survey are joined at the zip level for the price and safety analysis, so that the relationship between what a place costs and how it scores nationally can be examined directly.

The statistical method, named and explained

The paper does not run a single hypothesis test; its analytic engine is a percentile-rank index of national safety, computed at the zip-code level, that scores each community against all U.S. zip codes so that a higher percentile means lower reported-crime density. The core argument is a unit-of-analysis comparison in which the same percentile is read first at the metro level and then at the community level to show how far the two diverge. Three descriptive cuts follow: the range and median of community percentiles within each metro, a comparison of reputationally affluent and reputationally rough communities against their percentiles, and a community-by-community gap between the city-

relative letter grade and the national percentile. For the housing analysis the paper reports Pearson correlation coefficients between price or rent and the safety percentile. The reading throughout is that a low percentile in a dense, affluent enclave reflects the ambient-population denominator, that is, high daytime footfall rather than street-level danger.

The central dispersion finding rests on the full rated set of 264 to 265 communities, which is what gives it weight. The reputation groups and the 36-community mapped subsample are described as illustrative and non-random, and the city-relative grade was confirmable for only two communities, so those pieces are offered as demonstrations rather than as effect-size estimates.

Key results from the paper

Table 1. National safety percentiles within metros containing four or more communities.

Metropolitan area	Communities (n)	Percentile range	Median
New York to Newark to Jersey City	92	6 to 99	58
Miami to Fort Lauderdale	16	15 to 95	60
Washington, DC	7	5 to 96	53
Philadelphia	8	28 to 77	57
Boston	5	26 to 78	40
Chicago	6	22 to 63	42
Los Angeles to Long Beach	12	6 to 42	27

Note. Higher percentile denotes lower reported-crime density. The within-metro spread reaches 90 percentile points in New York, Miami, and Washington, while Los Angeles is the exception in which no community rises above the 42nd percentile.

Table 2. Home value versus national safety, 260 communities split at the medians.

Home value group	Safer half	Riskier half
Below median value (\$613,344)	56	75
Above median value	78	51

Note. The safety split is at the 42nd national percentile. All four quadrants are populated. Home value is essentially uncorrelated with the safety percentile (3-bedroom $r = +0.003$, $n = 260$), and rents show only a mild safety premium (3-bedroom rent $r = +0.258$, $n = 261$).

Figure 1. Affluent, reputationally posh communities cluster in the bottom third nationally.



Note. National safety percentile for the reputationally posh set (median 24th percentile). The reference line marks the median U.S. zip code. Lower percentiles reflect property-crime density in high-footfall areas, not street-level danger.

How to read the results

Table 1 is the heart of the paper. When several communities share a metro they scatter across

nearly the entire national range, so New York's 92 communities run from the 6th to the 99th percentile with a median of 58, and the metro name does not tell a family which case it faces. Figure 1 illustrates the reputation mismatch, since glamorous coastal enclaves such as Beverly Hills and San Francisco sit near the bottom of the national distribution once safety is scored on a property-crime-weighted percentile.

Table 2 completes the picture by showing that price is not a usable proxy for safety, because all four value-and-safety quadrants are populated and the correlation between home value and percentile is essentially zero. The separate finding that a city-relative grade is the rosier reading 86% of the time, by a median of four letter-grades, explains why the reassuring local narrative survives contact with the data: the local grade answers the easier question of whether a place is nice relative to its metro.

What the numbers do not claim

The national percentile is a total-reported-crime measure weighted toward property offenses, so it is not a violent-crime or street-level danger index, and low percentiles for dense enclaves reflect the ambient-population denominator rather than the risk faced by someone living there. Figures are keyed to each community's most common zip code, which can span several miles and does not resolve the block around a specific institution. Both instruments are commercial products whose proprietary weighting is not fully transparent, the city-relative grade was confirmed for only two communities, and the reputation groups and 36-community mapped subsample are illustrative and non-random. None of these qualifications disturbs the central, scale-level finding, which rests on the full rated set.

Reproducing the analysis

The community list was compiled by the author from public sources and keyed to residential zip codes, which is the step most of the work depends on. From there the national percentiles are retrievable by zip from CrimeGrade, the city-relative grades from DoorProfit, and both trace back to the FBI Uniform Crime Reporting and NIBRS program; home values come from Zillow ZHVI and rents from the Census American Community Survey. An analyst could rebuild the results by summarizing within-metro percentile ranges and medians, tabulating the value-and-safety quadrants at the median split, and

computing Pearson correlations between price or rent and the safety percentile.