

Can a Twitter-Based Measure of Racial Heterogeneity Explain Neighborhood Crime Differences? Proposing an Ambient Population Measure of Social Disorganization

Riley Tucker^{1,2,3} Daniel T. O'Brien^{1,2} Forrest Hangen^{1,2}

¹Northeastern University, ²Boston Area Research Initiative, ³University of Chicago

Published on: Feb 07, 2024

DOI: <https://doi.org/10.21428/cb6ab371.8efcc1b9>

License: [Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License \(CC-BY-NC-ND 4.0\)](https://creativecommons.org/licenses/by-nc-nd/4.0/)

ABSTRACT

In the early 20th century, Chicago School scholars proposed social disorganization theory, which posited that some neighborhoods experience more crime due to disadvantage, residential instability, and racial heterogeneity. However, contemporary studies using annual demographic data have found mixed evidence about the role of racial heterogeneity. We consider that, instead, racial heterogeneity among people in public spaces may act in real time to impact people's willingness to exercise guardianship over crime. To test this possibility, a machine learning strategy was used to train and test a race prediction algorithm and estimate the race of 7,069 Twitter users who sent at least one geotagged 'tweet' from Boston in 2018. By measuring the locations of users' tweets, we create a census tract-level measure representing ambient racial heterogeneity in the average week. Validation exercises indicated that inferred race of individual Twitter users was 90.7% accurate and produced a sample with comparable demographics to the Boston residential population. Results from negative binomial regression analyses suggested that ambient heterogeneity is a strong positive correlate of multiple indicators of violent and non-violent crime. Future research should develop theory and methodological strategies to further investigate how individual characteristics aggregate within ambient populations to explain crime differences across places.

INTRODUCTION

In the early 20th century, scholars from the Chicago School of Sociology asserted that racial heterogeneity among residents is a key driver of crime rates across neighborhoods. Conceptualizing their framework as *social disorganization theory*, they argued places will experience more crime when they have a demographic makeup that impedes trust among residents and thus limits their ability to engage in collective action that can prevent crime (Shaw and McKay 1969). While contemporary research has provided strong evidence that neighborhood-level poverty and residential instability break down social trust in ways that allow crime problems to manifest (Morenoff, Sampson and Raudenbush 2001; Sampson, Raudenbush and Earls 1997; Sampson 2006), analytic results suggest racial heterogeneity among residents has a less salient effect on social disorganization and crime than originally theorized (Letki 2008; Taylor, Twigg and Mohan 2010; Twigg, Taylor and Mohan 2010).

Meanwhile, social scientists across disciplines have repeatedly demonstrated that racial heterogeneity strongly predicts trust among people. At the macro-level, scholars from political scientists (Putnam 2007) to economists (Charles and Kline 2006) to sociologists (Dinesen and Sønderskov 2015; Hipp and Perrin 2009) have demonstrated that individuals living in heterogeneous residential contexts tend to have less social trust and ties than others. At the micro-level, research also shows that interpersonal interactions between individuals are altered when those individuals belong to different racial groups (Eckel and Grossman 2001; Enos 2014; Glaeser et al. 2000). Theoretically, the ability of racial heterogeneity to depress social cohesion, even at extremely small geographic scales (see: Dinesen and Sønderskov 2015), should impede the willingness of

neighborhood residents to engage in social control behavior. As such, it is surprising that contemporary research has not been able to identify a relationship between residential racial heterogeneity and crime.

We propose that scientists have found limited evidence for the relationship between racial heterogeneity and crime because they have focused on residential demographic measures. Given that urban places tend to be composed of both residents and visitors at any given time, we may consider that someone's behavior in a given moment is more directly shaped by the demographics of all people who are around them rather than just their neighbors. In that case, phenomena associated with racial heterogeneity may also operate within *ambient populations* that represent the total collection of people who occupy a physical space together. By more precisely measuring the racial demographics of those who are proximate in public places, we believe we can more strongly model the likelihood that some of those people will be motivated to engage in preventative behavior when a potential crime threat emerges.

To consider this possibility, we apply social disorganization theory to the study of ambient populations to evaluate whether neighborhoods with racially heterogeneous ambient populations tend to experience more crime than others. If heterogeneity is a fundamental social characteristic that leads to social breakdown, people may be less likely to engage in pro-social behaviors in contexts where the ambient population tends to be more heterogeneous. Indeed, results from a series of psychological experiments suggest that individuals are less likely to respond to deviant behavior if there is heterogeneity present between those interacting (Chekroun 2008; Chekroun and Nugier 2011; Moisuc and Brauer 2019). To consider whether this phenomenon scales to the neighborhood level, we introduce a methodological strategy for measuring racial demographics within ambient populations. Utilizing a database with information from over 1.3 million Twitter users who have had their account linked to the U.S. voter registry, we show that it is possible to train a machine learning model that can accurately estimate the racial demographics of Twitter users based on their names and home locations. By using this trained model to estimate the racial demographics of users from an independent sample of Twitter users who made at least one geotagged 'tweet' from Boston in 2018, we can aggregate the racial demographics of users who spent time in the same neighborhood to assess whether ambient racial demographics shape crime beyond residential demographics. To proceed, we begin by reviewing theory and empirical work that support the relationship between racial heterogeneity and crime.

LITERATURE REVIEW

Social Disorganization and the Criminogenic Nature of Racial Heterogeneity

Research from the social disorganization tradition has highlighted residential demographic characteristics help explain why some neighborhoods experience more crime than others. In their seminal study, Shaw and McKay (1969) argued that aggregate demographic features of neighborhood residents explained why certain parts of Chicago had an enduring crime problem. By mapping features such as racial heterogeneity, poverty, and

residential instability, they were able to demonstrate that neighborhoods have entrenched sociodemographic dynamics that shape behavior among the individuals who live there. Labeling their framework *social disorganization theory*, they proposed that when community members have different cultural or social backgrounds it hampers communication and connectivity in a manner that impedes the ability of that community to respond to shared problems such as crime.

Contemporary findings in sociology and political science support the argument that social heterogeneity within a collective group limits trust and cohesion. In one oft-cited example, Robert Putnam (2007) used a national survey to show that Americans living in racially heterogeneous neighborhoods tend to be less trusting of their neighbors, including those from the same racial group. Findings from Hipp and Perrin (2009) also show that social heterogeneity in terms of social demographics and wealth reduces the ability of social ties to form among neighborhood residents. Most recently, Dinesen and Sønderskov (2015) found that feelings of trust among individuals are shaped by the heterogeneity among neighbors living within just 80 meters of their home. This attenuated trust among individuals within racially heterogeneous contexts leads to changed behaviors. Putnam (2007) found that residents of heterogeneous neighborhoods are less likely to vote or participate in community projects and tend to have fewer friends, while a second study showed that residents of heterogeneous neighborhoods are less likely to participate in carpooling (Charles and Kline 2006).

According to the social disorganization perspective, heterogeneity also reduces the likelihood that individuals will step in to take actions that can reduce crime, which has been supported by research on its contemporary extension, *collective efficacy* (see: Sampson, Raudenbush and Earls 1997). While the collective efficacy perspective has been influential across the social sciences, empirical evidence for the role of racial heterogeneity is murky because heterogeneity is strongly correlated to poverty, a second neighborhood contextual feature shown to drive crime. As such, several studies have argued that the linkage between racial heterogeneity and crime dissipates once economic disadvantage is properly modeled (Letki 2008; Taylor, Twigg and Mohan 2010; Twigg, Taylor and Mohan 2010). Furthermore, many of the key studies on collective efficacy neglect to measure racial heterogeneity (Browning and Cagney 2002; Browning, Dietz and Feinberg 2004; Sampson, Morenoff and Earls 1999), leaving an open question about whether racial heterogeneity truly shapes propensity for social control above and beyond the effects of poverty and residential instability.

Empirical evidence of the criminogenic effect of racial heterogeneity may be limited because heterogeneity has generally been conceptualized as a stable characteristic of residential spaces. In other words, racial heterogeneity is not generally treated as a condition that fluctuates across situations or times. However, a series of studies has shown that heterogeneity has contextual effects on behavior that operate in real time. For example, economic studies using experimental games have shown that individuals are less trusting when negotiating or collaborating with individuals from different racial groups (Eckel and Grossman 2001; Glaeser et al. 2000). Most directly relevant to social disorganization theory, a series of experiments from the field of psychology also show that strangers are more likely to take action in response to observed deviance when the

offending actor is a member of the same social group, suggesting that situational heterogeneity seems to be linked to reduced guardianship (Chekroun 2008; Chekroun and Nugier 2011; Moisuc and Brauer 2019). Under this lens, in the event of a crime attempt the likelihood of someone stepping in would be shaped by the racial demographics of those present in the surrounding area. Specifically, if heterogeneity is a central feature limiting trust and prosocial engagement then we would expect places with heterogeneous ambient populations to have heightened levels of crime.

Despite strong theoretical and empirical evidence suggesting that the willingness of people to take preventative action toward crime may be shaped by the racial demographics of those who are around them, researchers have done little to measure racial heterogeneity within ambient population units. In the first study to pursue this line of inquiry, Gu et al. (2023) used information about between-neighborhood mobility flows to estimate racial heterogeneity across census block groups in Cincinnati. Results suggested that ambient racial heterogeneity, rather than the classic residential measure, was correlated with rates of street robbery. While that study represents strong preliminary evidence for the importance of ambient heterogeneity, it faces several limitations. First, the study inferred the ambient population characteristics of a given neighborhood from the residential demographics of the communities of visitors (e.g., assuming that the aggregate racial makeup of individuals who travel from Neighborhood A to Neighborhood B is equivalent to the residential racial makeup of Neighborhood A). These assumptions may then inflate the similarities between ambient and residential population characteristics. Second, because the study only analyzed a single crime type, robbery, it is unclear whether the effects of ambient racial heterogeneity generalize to a larger set of crimes.

As such, many questions about ambient heterogeneity and its relationship to crime remain unanswered. First, it is unclear if ambient population demographics even deviate from residential demographics. Considering past research has shown that between-neighborhood mobility patterns are highly segregated along racial lines (Phillips et al. 2019; Wang et al. 2018), it is possible that the demographics of neighborhood visitors simply reflect residential demographics, in which case there would be little variation between ambient and residential measures of race. In the event that racial demographics of residential and ambient populations do not align it would be unclear which is more criminogenic in nature. Does someone choose to engage in guardianship behavior based on who is around them when a need arises, or are guardianship preferences conditioned by the norms of the residential community someone belongs to? To provide clarification to this point, we need to construct parallel measures of ambient and residential racial demographics to compare their relationships to each other and to crime. In contextualizing this effort, we now move to review past research about ambient population effects on crime.

Ambient Populations and Crime

Rather than utilizing social disorganization theory, criminologists studying ambient populations tend to draw from a second perspective, routine activities theory, to explain crime differences across places. Routine activities theory is built on the simple premise that crime occurs when a motivated offender meets a suitable

target in space and there are no capable guardians present to prevent that crime from occurring (Cohen and Felson 1979; Felson and Cohen 1980). At its root, this suggests that the likelihood of crime happening at a given place and time is a product of the proportional representation of these three groups to one another. If the actors who make up a large crowd are predominantly motivated to act as guardians, they should have high efficacy to prevent crime in the presence of a smaller number of motivated offenders. Conversely, if a crowd is mostly composed of suitable targets and motivated offenders, a small number of guardians will have low capability to respond to all potential crime threats, increasing the likelihood of crime occurring.

While criminologists have extensively theorized about how crime is produced through the routine activities of people who spend time in a place, it has historically been a challenge to accurately measure the people who are in a given place at a given time. Short of ethnographic and qualitative works observing small geographies over the course of years, scientists have generally struggled to comprehensively measure the characteristics of people who occupy a space in real time. However, advances in geospatial and mobile technology have provided new strategies to measure the presence of people who occupy urban spaces. The earliest studies assessing the relationship between ambient population size and crime relied on LandScan, a satellite-based measurement system that estimates how many people occupied a square-kilometer area during a 24-hour window. Studies utilizing this source of data have shown that ambient population size helps explain crime differences across municipalities (Andresen 2010) and neighborhoods (Andresen 2006).

More recently, new technologies for measuring individual movement have allowed researchers to measure ambient population size in more robust ways. As people use cell phones and social media in their daily lives, they ping cell towers and satellites, creating a record of where they were in space at a given time. By aggregating individual mobility points to count everyone who occupied a shared space, researchers can measure ambient population size at more specific spatiotemporal scales. This strategy has been applied to gain a more granular understanding of how ambient populations generate crime. For example, Song et al. (2018) found that mobile phone-derived measures of ambient population reliably measure ambient populations for 3-hour windows and effectively represent the population at-risk for theft victimization. Studies have also generated ambient population measures from geotagged Twitter posts, with results indicating these metrics help explain crime variation at hourly scales (Hipp et al. 2019; Malleson and Andresen 2015).

However, few studies to date have attempted to measure characteristics of the individuals that compose ambient population units. This is somewhat surprising given that even the first article proposing routine activities theory suggested that differential mobility patterns based on a specific individual characteristic, gender, help explain the geographic distribution of suitable targets (Cohen and Felson 1979). According to routine activities theory, we should expect high crime places to have ambient populations that have high presences of motivated offenders and suitable targets while low presence of capable guardians. Measuring characteristics of the individuals composing ambient populations may help inform us about whether members of ambient populations are more likely to act as offenders, targets, or guardians. As such, measuring the

characteristics of individuals within ambient populations could help explain spatial variation through a distinct mechanism that is not captured when analyzing ambient population size alone.

To our knowledge, efforts to measure characteristics within ambient populations have generally been limited to studies that have analyzed the mobility patterns of individuals to categorize them as locals or visitors in a given ambient population. Using geotagged Twitter posts, Tucker et al. (2021) estimated the home locations of users to create measures of commuter and tourist presence within block-level ambient population units, with results suggesting these measures help explain the frequency of both public violence and private conflict. Similarly, Wo et al. (2022) tagged Twitter user's home block group based on where they sent the most tweets and found that block groups with higher proportions of local users tended to have lower counts of violent crime and property crime. Using mobile phone data, Song et al. (2021) estimated where people live and work to create ambient population measures for residents, employees, and visitors, with their results also suggesting that ambient populations composed of visitors are more criminogenic in their nature. As noted, the one case we know of that sought to study the racial composition of the ambient population measured individual-level race as an extension of the neighborhood's racial composition without drawing additional information from the user's activity (Gu et al 2022).

We propose that there are substantially more individual characteristics beyond home location that can be measured from oft-used ambient population datasets. With the exception of one recent study that analyzed mobility data with individual-level estimates of income to demonstrate that theft occurrence is correlated with the ambient presence of certain income brackets (Song et al. 2023), this possibility has been unexplored. This is surprising given that computer and data scientists have developed methods for estimating the age (Morgan-Lopez et al. 2017), income-level (Preoțiuc-Pietro et al. 2015), mental health status (Coppersmith, Dredze and Harman 2014), and indeed, race (Preoțiuc-Pietro and Ungar 2018) of social media users who can also be studied through the geographic indicators they produce. While much of this work is conducted by computer scientists, we offer that these technologies are also accessible to social scientists and crime scholars, specifically. To provide an example and use case, we use this study to show that basic information available in mobility data can be used to measure variables such as race that are central to understanding social phenomena. To complete this task, we follow the lead of computer scientists who have developed methodologies to estimate the racial demographics of users in an oft-used ambient population data source: Twitter.

CURRENT STUDY

Measuring racial heterogeneity within ambient populations requires a strategy for measuring the features of the individuals composing mobility datasets. Within mobility research, the scholarly convention is to place the highest value on data sources capturing the highest proportion of the population. As such, researchers have gravitated towards mobility datasets generated from cell phone records, including the Safegraph dataset which captures movement data from 47 million cell phone devices and has become a staple of the mobility literature since it was openly published during the COVID-19 pandemic (Coston et al. 2021; Zhang et al. 2022). While

these data and other similar products offer highly accurate mobility information due to the magnitude of their samples, researchers often receive these data in an aggregate format that reports the number of trips to or from places (e.g., the number of people who traveled from neighborhood A to neighborhood B or the number of people who visited Business Establishment X). While this data structure facilitates research about mobility patterns, the obfuscation of individual-level information limits the utility of the data to answer questions about how individual features aggregate within ambient populations. Though some studies have inferred ambient population demographic measures based on mobility patterns and residential demographics (see: Gu et al. 2023), an approach directly measuring the individual-level features of ambient population members should theoretically be more precise. To consider this methodological framework, it is necessary to turn to mobility datasets with rich individual-level information even if it comes at the cost of sample size.

Computer and data scientists have demonstrated that there are many effective strategies for estimating the demographic features of social media users. Specifically, there is a large literature on estimating the characteristics of Twitter users, a data source that is also studied in ambient population research. Studies assessing the demographic characteristics of Twitter users generally utilize a machine learning approach that trains a model on a set of data with known user attributes. For example, Preoțiuc-Pietro and Ungar (2018) conducted an online survey that asked respondents to report both their race and Twitter username. They were then able to identify patterns in word usage and surnames that were predictive of the race of a user. Another study searched Twitter for birthday announcements, which allowed them to label user ages and then train an age prediction model including features of language usage and profile characteristics (Morgan-Lopez et al. 2017). Because users in these studies are associated with validated characteristics, it is possible to implement the prediction algorithm on a testing sub-sample and compare estimated features to the verified attributes, providing a measure of model accuracy. Using such strategies, researchers have shown that it is possible to estimate Twitter-user demographics such as race (Preoțiuc-Pietro and Ungar 2018), gender (Miller, Dickinson and Hu 2012), and age (Miller, Dickinson and Hu 2012) with high accuracy.

Notably, many of these studies using Twitter focus on analyzing user profile characteristics and textual post data, an approach requiring technical skills that are out of reach for many scholars. However, studies have also seen success using more basic types of information, such as name. By calculating the proportion of individuals with a given name who belong to demographic groups of interest, it is possible to create a set of racial probabilities associated with various common names. For example, Mislove et al. (2011) estimated the race of Twitter users using a file released as part of the decennial census that reports the racial distribution for all last names that belong to at least 100 Americans. Preoțiuc-Pietro and Ungar (2018) also found that surname-based probabilities increased their ability to accurately predict the race of Twitter users.

We build on these past approaches by introducing an additional set of features that are predictive of an individual's race: the demographics of their home neighborhood. We propose that by constructing a machine learning model that compares a Twitter user to the demographics of Americans with their last names and to the

residential demographics of the census block group where they live, we can estimate the race of Twitter users with a high enough accuracy to reliably measure racial demographics within census tract-level ambient population units. In return, this will then allow us to compare ambient population demographics to residential demographics and crime.

To train our model, we draw from a dataset of over a million Twitter users who have been linked to the voter registry on the basis of their name and home location. Because the voter registry is an official source of demographic data, it represents an ideal validated source to train a prediction model on. To demonstrate the utility of this methodology, we then proceed in three steps. First, we begin by establishing the reliability and validity of our race estimation strategy by assessing model accuracy on a testing sub-sample and comparing the predicted demographics of Boston residents to the true demographic distribution of the city. If our model is effective, it will estimate user race with a high level of accuracy and, resultingly, the Boston users in our sample will have a racial makeup comparable to the true residential racial makeup. After establishing the validity of our measurement approach, we will compare our ambient racial heterogeneity measure to the traditional residential heterogeneity measure. We hypothesize that ambient and residential racial heterogeneity will be correlated but distinct measures as the visitors to a neighborhood often have a demographic composition that diverges from that of residents. Finally, we test the possibility that ambient racial heterogeneity is correlated to crime. We hypothesize that ambient racial heterogeneity will be positively correlated with crime frequency even when controlling for residential demographic features.

DATA & METHODS TO ESTIMATE THE RACE OF TWITTER USERS

To measure the racial characteristics of individuals who have spent time in Boston, this study utilizes information from two sources of Twitter data. To create Twitter-based ambient population measures, we draw from a collection of tweets posted in 2018 that were geolocated to Boston. To shed light on the racial characteristics of the users who created those tweets, we draw from a second source of Twitter data provided by TargetSmart^[1] that has linked over a million twitter users to the voter registry, providing validated demographic and home location information about each user.

Together, these two datasets provide enough information to develop a modeling strategy to leverage the information about Twitter users with known demographics in the TargetSmart data to estimate the demographics of users in the Boston ambient populations Twitter sample. With estimated racial information for users in the Boston ambient population Twitter sample, we can aggregate these measures within ambient population units to generate demographic metrics. Below, we describe the contents of the two Twitter samples, discuss strategies for developing features for the machine learning model, and outline the methodology for aggregating user characteristics and locations to ambient population metrics.

Twitter Datasets

Boston Ambient Population Twitter Sample

To measure ambient populations in Boston, we draw from a collection of 232,874 geotagged tweets collection through Twitter's Streaming API, a tool that allows users to download information on up to 1% of tweets globally. These data include information about tweets, the users who posted tweets, and the geographic location from which tweets were sent. Importantly, these data were only collected for users who opted in to releasing their information through the Streaming API. To create an analytic sample for the current study, the sample of global tweets collected through the API was limited to tweets that were geolocated in Boston and posted in 2018.

TargetSmart Demographically Identified Twitter Sample

To develop an algorithm for predicting the demographics of users in the Boston ambient population Twitter sample, we draw from a rich sample of information about 1,315,430 Twitter users that has been linked to the US Voter Registry. These data are provided by TargetSmart, who has linked Twitter accounts to the Voter Registry based on the last name and location associated with each account.

Generating Model Features to Estimate Race

Last Name Probabilities

To measure the likelihood that a Twitter user is in a given racial group based on their last name, we used a Census dataset of 162,253 last names that were reported 100 or more times in the 2010 decennial census. For each name, the Census reports the percentage of Americans with that name who identify as belonging to six racial groups of interest^[2]. For both the TargetSmart and Boston ambient population Twitter users, we then merge in the following probabilities based on the last name each user listed on Twitter: likelihood of being white, likelihood of being Black, likelihood of being Asian, and likelihood of being Hispanic.

Census Demographics

Additionally, race was estimated as a function of the demographics of the census block group where each user lives. TargetSmart users are linked to their home address via the Voter Registry. Addresses were geocoded and then address coordinates were spatially joined to a shapefile of US census block groups to identify the block group where each user lived. We then merged information from the 2014-2018 American Community Survey to link each user to the following demographics of their home neighborhood: percentage of non-Hispanic white residents, percentage of non-Hispanic Black residents, percentage of non-Hispanic Asian residents, and percentage of Hispanic residents.

The Boston ambient population Twitter sample was not collected from a source reporting home address. Thus, to link each user in this sample to the demographics of their home neighborhood we needed to first estimate

their home location based on the locations of their tweets. To estimate home locations we follow the approach of past research by using the DBSCAN* cluster algorithm to identify areas of about 30 squared-meters from which a user sent three or more tweets between the hours of 9 P.M. and 12 A.M during weekday nights, the period when someone is most likely to be in their home (Tucker et al. 2021; Wang et al. 2018). The centroid of each user's home cluster was then spatial joined to the shapefile of census block groups and racial demographics from the 2014-2018 ACS were merged at the user-level using the same strategy that was applied to the TargetSmart users.

Final Analytic Sample

The machine learning approach utilized in this study necessitates that all features are present in both the training and predicted datasets. Thus, if a user lists a last name that is not present in the census data (22.67% of TargetSmart users, 25.74% of Boston ambient population users with home location data), they are not analyzable. Additionally, users in the Boston ambient mobility sample can only be included if they have produced enough tweets for the DBSCAN* algorithm to estimate their home location (47.2% of users) and 4.86% of the geocoded home locations for TargetSmart users did not spatial join to the census shapefile. Of Boston mobility sample users with estimated home location, 22% of them had estimated home locations outside of the US and thus could not be analyzed because there is no way to link them to the demographic characteristics of their home neighborhood.

Following these exclusion criteria, the final sample contains 760,335 TargetSmart users and 5,170 users in the Boston ambient population sample for whom we have enough data points to proceed with our modeling strategy. We then use a machine learning approach to train a model on the TargetSmart users that is then used to estimate racial characteristics for each of the 5,170 users.

Algorithmic Strategy

We begin our analyses by using a machine learning strategy to estimate the racial demographics of Twitter users in the Boston mobility sample. Specifically, we train a Gaussian Naïve Bayes model on the probabilities associated with the first names, last names, and home locations of TargetSmart users. This model works by applying Bayes' theorem to a set of predictive features that are assumed to be independent. While this assumption is rarely true with real-world data, studies suggest that models based on it perform surprisingly well on many classification tests (Domingos and Pazzani 1997; Zhang 2004). Because this strategy assumes that feature probabilities follow a Gaussian distribution, feature likelihood is estimated using the following calculation where x is the likelihood of a feature given y group:

$$P(x_i|y) = \frac{1}{\sqrt{2\pi\sigma_y^2}} \exp\left(\frac{-(x_i - \mu_y)^2}{2\sigma_y^2}\right)$$

The likelihoods for features based on name and demographics of home neighborhood are then applied to Bayes theorem to calculate the likelihood that a user is a member of group y given x features:

$$P(y|x_1, \dots, x_n) = \frac{P(y)P(x_1, \dots, x_n|y)}{P(x_1, \dots, x_n)}$$

The above calculation is conducted for all races to evaluate the likelihood that each user belongs to each racial group.

Ethics

In using Twitter data to estimate the demographic characteristics and home locations of users, it is important to consider the ethics of using these data. While all respondents in the Boston mobility sample opted-in to having their tweets geotagged, it is unlikely they anticipated these data would be used to uncover their personal information. Similarly, while the users in the TargetSmart sample chose to list identifying characteristics on their public Twitter profile, they probably did not expect this information would be used to link their identity to official data. If data provided in the manuscript or the methodology described within allowed members of the public to identify the home locations or personal characteristics of those within either sample, this study would be putting the subjects at risk of crime victimization (stalking, vandalism, assault, etc.) or personal attacks without gaining their consent. As such, a central aspect of designing this study has been to ensure that subjects are not personally identifiable.

To limit the identifiability of study subjects, we aggregate neighborhood-visitation across the entire year. This means that it would not be possible for a reader to identify a study subject based on their known mobility in 2018. To further reduce the potential for identity breaches, we do not report any statistical results about individual-level mobility patterns. That said, the possibility remains that a member of the public may leverage this study methodology for unethical purposes. However, in 2019, Twitter announced it would begin deprecating its geotagging feature (Twitter 2019). As such, tweets will no longer be geotagged, and it will be impossible to estimate the home locations of users. While there may still be opportunities to utilize our methodology for nefarious purposes using older Twitter data, over time much of the information embedded in tweets will be less representative of true user characteristics, further limiting possibilities for these methods to be used for invasive purposes[3].


Evaluating Accuracy and Implications of Race Estimation Model

How Accurate is Model Prediction?

To assess the efficacy of this approach, we start by using the standard methodology of splitting the TargetSmart data into a training sample composing 80% of the analyzable users ($n = 608,268$) and a testing sample composing 20% of analyzable users ($n = 152,067$). For validation purposes, each user was assigned to the most likely race given their name and location features^[4]. By predicting the race within the testing sample that includes verified racial demographics, we can conduct analyses comparing the estimated values to the true values to evaluate how accurate the model is in predicting race.

Overall, the model accurately predicted the race of 90.7% of users. Figure 1 reports prediction accuracy by race for the 152,067 randomly selected TargetSmart users. Results indicate that the model was most accurate for estimating the race of white (92% accuracy) and Hispanic users (91% accuracy). The model was moderately less accurate for the other two groups of interest, with an accuracy rate of 81% for Black users and 80% for Asian users. For Black users, erroneous predictions predominantly improperly identified them as white (17% of all estimation attempts), while error was more equitably distributed for other groups. While the model tended to be accurate in predicting race overall, it did have erroneous tendencies to label Black users white.

Figure 1. Crosstab of Predicted Race by True Race of Twitter Users (n=152,067)



Sample-Level Validity Assessment

In addition to assessing whether the model accurately estimated the race of individuals, we assess how similar the demographics of users in our final analytic mobility sample are to the residential demographics of Boston. If our measurement strategy is accurately capturing the Boston ambient population, the demographic profile of Twitter users who live in Boston should be very similar to the true demographic profile of residents captured by the census. As such, we analyzed a sub-sample of users in the Boston mobility sample who had an estimated or known home location in the Boston metropolitan region^[5] (n = 2,203) and assessed the percentage of users that fall in each racial group to compare the demographics within our mobility sample to the metropolitan racial demographics reported by the Census. Table 1 reports the percentage of sampled Twitter users and residents who belong to each of the analyzed racial groups.

Table 1. Comparing Racial Demographics of Boston Metro Residents in Twitter Mobility Sample to Boston Metro Residents

	Twitter Users (n = 2,203 users)	Residents
Percent White	77.91	77.23
Percent Black	7.18	6.83
Percent Hispanic	6.54	9.30
Percent Asian	8.37	6.63
Residential demographics provided by the 2010 Decennial Census		

Overall, the demographics of Boston residents in the mobility sample had a comparable demographic profile to the true demographic profile of residents. The most substantial differences between our mobility sample and the true demographics were overrepresentations of Asian individuals (1.74%) and underrepresentation of Hispanic individuals (2.76%). Overall, we consider our sample to be highly representative in terms of racial demographics.

MEASURING AMBIENT RACIAL HETEROGENEITY

Using the machine learning strategy to estimate the racial characteristics of users in the Boston ambient population sample provides estimated probabilities that a given user belongs to each of the follow four racial demographic categories: white, Black, Hispanic, or Asian. In addition to those 5,170 users with estimated racial demographic information, we used a unique identifier provided by Twitter to identify an additional 1,899 users in the Boston ambient population sample that are also captured by the TargetSmart data, allowing us to merge in their verified racial demographic information. In total, this gives us a sample of 7,069 Twitter users with measured racial characteristics who sent one or more tweet from Boston in 2018.

Calculating Ambient Racial Heterogeneity

The 7,069 users analyzed here generated a body of 64,005 geocoded tweets in Boston in 2018 (9.05 tweets per user). We then created tract-level ambient population units based on tweet locations. Data was reduced so each user only had one tweet per week per tract. The racial probabilities associated with each user^[6] were then averaged to estimate weekly percentages of users in each tract who were white, Black, Hispanic, and Asian. To account for week-to-week differences in visitation patterns, we then averaged these variables across weeks to create study metrics that represent a tract's ambient population makeup across the typical week.

Following the standard strategy for calculating demographics of residential populations, these racial demographic variables were used to calculate tract ambient racial heterogeneity scores using the Herfindahl index (see: Rhoades 1993):

$$heterogeneity = 1 - [p(white)^2 + p(Black)^2 + p(Hispanic)^2 + p(Asian)^2]$$

This strategy produced a value for each tract that can potentially range from 0 – 1 and represents the likelihood that two members of the ambient population selected at random would have different racial categories.

Do Ambient Demographics Deviate from Residential Demographics?

As people engage in their routine activities, city residents will leave their home neighborhood while individuals residing in the greater metropolis and beyond simultaneously visit. Because this population movement has the potential to shift the demographics profile of those who are present together in an ambient population, we investigate whether ambient population racial demographic measures diverge from residential measures.

To evaluate consistency between ambient and residential racial demographics, we measured residential socioeconomic indicators at the tract-level using the 2014-2018 American Community Survey. Notably, this provides measures indicating the percentage of residents in each tract who are white, Black, Hispanic, and Asian. These values were used to calculate a residential racial heterogeneity score using the same formula as for ambient racial heterogeneity, allowing direct comparison between the two measures.

Are Ambient Population Demographics Impacted by Sample Bias?

It was also important to consider whether demographic biases across our mobility sample impact our tract-level ambient population measurements. To do so, we created population weights for each race by dividing the percentage of Boston metro residents belonging to a given race by the percentage of users in our Twitter mobility sample belonging to that race. These calculations produced weights^[7] that have been applied to each user in the mobility sample for each tract they visited, adjusting the ambient population demographics of each tract so that they appear as if they were produced by a sample that was more comparable to the true residential population. Bivariate analyses suggested that weighted measures of ambient racial proportions and heterogeneity are all correlated to their non-weighted counterparts at .99, indicating that tract-level ambient demographic measures are not strongly biased by demographic representation or identification issues. In taking the fullest efforts to minimize the impacts of bias on further analyses, we proceed using these weighted ambient population measures.

Comparing Ambient and Racial Demographics

Table 2 reports descriptive statistics for ambient racial demographics and all other measures analyzed in this study. Results from these descriptive analyses already began to hint at a divergence between ambient and residential racial demographic measures. For example, the average tract had an ambient population that is 66%


white, compared to an average residential population of 48%. We also observed differences in terms of racial heterogeneity, where the average tract had a residential heterogeneity score of 0.52 compared to an average ambient heterogeneity score of 0.19, suggesting that ambient population units tend to be substantially more homogeneous.

Table 2. Descriptive Statistics Analyzed Variables (n=163)		
	<i>N</i> (%) / Mean (SD)	Range
Ambient White	0.66 (0.23)	0.00 – 1.00
Ambient Black	0.16 (0.19)	0.00 – 1.00
Ambient Hispanic	0.11 (0.16)	0.00 – 0.87
Ambient Asian	0.08 (0.11)	0.00 – 1.00
Ambient Heterogeneity	0.19 (0.12)	0.00 – 0.47
Ambient Population Size	11.88 (45.36)	0.06 – 509.40
Resident White	0.48 (0.29)	0.01 – 0.98
Resident Black	0.20 (0.23)	0.01 – 0.82
Resident Hispanic	0.19 (0.15)	0.00 – 0.70
Resident Asian	0.09 (0.10)	0.00 – 0.56
Resident Heterogeneity	0.52 (0.16)	0.04 – 0.76
Resident Poverty	0.14 (0.12)	0.00 – 0.58
Resident Renting	0.66 (0.19)	0.09 – 1.00
Vacant Units	0.08 (0.05)	0.01 – 0.25
Downtown Tract	11 (6.83%)	
Residential Tract	116 (72.05%)	
Violence	23.49 (23.37)	1.58 – 171.74

Robbery	1.69 (2.14)	0.00 – 15.24
Social Disorder	4.72 (6.98)	0.00 – 48.20

To assess how ambient and residential measures of racial demographics align, Figure 2 visualizes a correlation matrix representing a series of bivariate correlation analyses between ambient and residential racial measures. Results indicated that correlations between ambient and residential demographics were strongest for proportion Black ($r = 0.63, p < .001$) and white ($r = 0.46, p < .001$) but weaker for proportion Hispanic ($r = 0.25, p < .01$) and quite low for proportion Asian ($r = 0.11, p > .05$). As a product of these differences between ambient and residential measures of racial presence, the residential and ambient heterogeneity measures only had a correlation of 0.06 ($p > .05$), suggesting that the residential and ambient ethnic heterogeneity of a neighborhood often differ substantially.

Figure 2. Bivariate Correlations between Residential and Ambient Demographic Measures (n=163)



IS AMBIENT RACIAL HETEROGENEITY CORRELATED WITH CRIME?

Analytic Strategy

Given that ambient and residential measures of racial demographics appear to be distinct from one another, we use a regression strategy to determine whether residential demographics, ambient demographics, or potentially both drive differences in crime rates across Boston census tracts.

Measuring Crime

We integrated 911 dispatch data provided by Boston Police Department (BPD) to evaluate whether ambient racial heterogeneity is associated with crime. BPD provides dispatch data that identifies the type of criminal event that prompted a dispatch and provides coordinates to geolocate the location of each criminal event. These locations have been spatial joined to a shapefile of Boston census tracts in order to create tract-level crime counts.

The current study analyzed four types of crime: violence, robbery, larceny, and social disorder. We have selected these crimes because they represent a diverse set of social processes. Specifically, we conceptualize that violence and robbery represent a type of violent behavior that would require substantial action from potential guardians to successfully intervene against. In contrast, larceny and social disorder constitute less serious, non-violent offenses where bystanders may have increased efficacy to take meaningful anti-criminogenic action. For all crimes, incidents were measured as a rate per 1,000 neighborhood residents.

Controlling for Land Use Differences

Several additional neighborhood measures representing the percentage of tract residents renting their home, the percentage of vacant units in a tract, whether a tract was downtown, and whether the tract was generally non-residential are drawn from the Boston Area Research Initiative's census geographies database (Zoorob et al. 2021). To ensure that ambient heterogeneity did not spuriously represent ambient population size, analyses also control for the average number of unique Twitter users posting from the tract per week.

Regressing Crime on Ambient Racial Heterogeneity

Violent Crime

Finally, we assessed the relationship between ambient racial demographic measures and crime. To begin, Table 3 presents results from a series of negative binomial models regressing rates of violence and robbery onto ambient demographics, residential demographics, and land use features. For each crime, we started by only including the residential demographic and land use variables to set a baseline, then introduced ambient population features into the models to assess how they relate to crime when accounting for residential and land use characteristics.

Table 3. Negative Binomial Models Examining the Tract-level Covariates of Violent Crime (n = 163)

	Violence IRR (Std. Error)			Robbery IRR (Std. Error)		
Intercept	5.37*** (1.65)	4.33 (1.69)		0.26 (1.99)	0.17 (2.07)	
Residential Pop.						
Res. Hetero.	2.60 (2.20)	2.66 (2.25)		2.20 (2.74)	2.88 (2.88)	
Res. Black	2.81** (1.62)	3.16** (1.80)		4.55** (1.81)	5.78** (2.08)	
Res. Hispanic	1.31 (2.30)	1.32 (2.39)		2.49 (2.89)	2.41 (3.05)	
Res. Asian	0.84 (4.29)	0.55 (4.72)		1.32 (5.90)	0.43 (6.89)	
Poverty	3.30 (2.56)	3.28 (2.63)		0.49 (3.34)	0.43 (3.58)	
Renters	0.87 (1.82)	0.87 (1.83)		1.82 (2.22)	1.93 (2.31)	
Land Usage						
Downtown	1.52 (1.45)	1.19 (1.50)		1.72 (1.56)	1.52 (1.65)	
Residential	0.85 (1.26)	0.95 (1.27)		0.86 (1.34)	1.01 (1.35)	
Vacant Units	1.08*** (1.02)	1.05** (1.02)		1.08*** (1.02)	1.04* (1.03)	
Ambient Pop.						
Amb. Hetero.		6.79** (2.29)			24.22*** (2.93)	
Amb. Black		0.90 (1.80)			0.63 (2.14)	
Amb. Hispanic		0.78 (1.77)			0.39 (2.31)	
Amb. Asian		0.59 (2.22)			0.53 (3.29)	

Amb. Pop. Size		1.00 (1.00)			1.00 (1.00)	
Deviance	68.00	59.65		85.98	75.04	

Results from the full models showed that ambient heterogeneity was a strong positive indicator of both violence (IRR = 6.79, $p < .01$) and robbery (IRR = 24.22, $p < .001$). Conversely, residential heterogeneity did not have a statistically significant association with either form of violent crime (IRR = 0.95, $p > .05$ for violence; IRR = 1.01, $p > .05$ for robbery). Additionally, several control variables representing proportions of Black residents (IRR = 3.16, $p < .01$ for violence; IRR = 5.78, $p < .01$ for robbery) and vacant units (IRR = 1.05, $p < .01$ for violence; IRR = 1.04, $p < .05$ for robbery) had statistically significant associations with crime that were robust to the inclusion of ambient population measures.

Do Results Generalize to Non-Violent Crime?

We then replicated these models for the non-violent crimes of larceny and social disorder to evaluate if these relationships were generalizable to a wider variety of crimes (See Table 4). As was the case for violent crimes, ambient racial heterogeneity was strongly correlated with larceny (IRR = 8.27, $p < .01$) and social disorder (IRR = 12.47, $p < .001$). Parameters for control variables indicated that neighborhoods with higher proportions of vacant units have more non-violent crime on average (IRR = 1.07, $p < .001$ for larceny; IRR = 1.08, $p < .001$ for social disorder) and neighborhoods with higher proportions of renters experience more social disorder (IRR = 3.74, $p < .05$). Including ambient population variables in the model did not substantively change any parameters observed in the baseline models.

Table 4. Negative Binomial Models Examining the Tract-level Covariates of Non-Violent Crime (n = 163)

	Larceny IRR (Std. Error)			Social Disorder IRR (Std. Error)		
Intercept	4.55*** (1.65)	3.40** (1.70)		0.59 (1.77)	0.39 (1.83)	
Residential Pop.						
Res. Hetero.	2.21 (2.20)	2.43 (2.26)		2.32 (2.40)	3.31 (2.49)	
Res. Black	1.27 (1.62)	1.44 (1.80)		0.81 (1.70)	0.75 (1.93)	
Res. Hispanic	0.40 (2.32)	0.40 (2.41)		1.14 (2.48)	0.91 (2.59)	

Res. Asian	0.45 (4.31)	0.29 (4.74)		0.32 (4.86)	0.09 (5.44)	
Poverty	1.48 (2.57)	1.31 (2.63)		1.02 (2.80)	1.22 (2.89)	
Renters	1.75 (1.82)	1.90 (1.83)		3.68* (1.96)	3.74* (2.00)	
Land Usage						
Downtown	1.27 (1.45)	1.04 (1.50)		1.35 (1.48)	1.28 (1.53)	
Residential	0.99 (1.26)	1.15 (1.27)		0.88 (1.28)	0.99 (1.30)	
Vacant Units	1.10*** (1.02)	1.07*** (1.02)		1.11*** (1.02)	1.08*** (1.02)	
Ambient Pop.						
Amb. Hetero.		8.27** (2.29)			12.47*** (2.50)	
Amb. Black		0.89 (1.81)			0.90 (1.96)	
Amb. Hispanic		0.56 (1.78)			0.65 (1.93)	
Amb. Asian		0.54 (2.23)			1.08 (2.48)	
Amb. Pop. Size		1.00 (1.00)			1.00 (1.00)	
Deviance	63.81	53.86		112.19	103.30	

Supplementary Analyses

Do home locations of the ambient populace matter?

We concluded our analyses with several supplementary analyses that consider potential alternative explanations for the observed association between ambient heterogeneity and crime. First, social disorganization theory posits that behaviors are learned through neighborhood social norms. As such, the guardianship proclivity of ambient population members may systematically vary dependent on the demographic context of where they live. In this event, ambient heterogeneity could be spuriously capturing collectives of individuals from certain types of neighborhoods, particularly if residential heterogeneity is a mechanism driving guardianship behavior.

To evaluate this possibility, we matched all individuals in the ambient population sample to the racial demographics of their home census block group. The strategy that was utilized to measure ambient racial demographics was then applied to generate variables representing the average residential context of members in each ambient population unit: percent Black residents, percent Hispanic residents, percent Hispanic residents, and racial heterogeneity. Regression results suggested that the demographic features of ambient population member's home neighborhoods were not associated with violent or non-violent crime when residential demographics, ambient population demographics, and land use were controlled for ($p > .05$ for all variables and models) (See Table 5). Moreover, ambient heterogeneity remained statistically associated with crime across crime types, and the associations of other variables generally remained robust to the inclusion of these additional variables.

Table 5. Negative Binomial Models Examining the Tract-level Covariates of Crime in Residential Neighborhoods (n = 117)					
Variable	Violence	Robbery		Larceny	Social Disorder
Variable					
Intercept	4.18*** (1.64)	0.11*** (2.12)		4.35*** (1.64)	0.33* (1.82)
Residential Pop.					
Res.Hetero.	1.63 (2.51)	1.73 (3.45)		1.29 (2.52)	2.00 (2.89)
Res. Black	3.68* (2.02)	6.05** (2.50)		2.58 (2.03)	0.91 (2.24)
Res. Hispanic	1.83 (2.88)	3.01 (3.90)		1.24 (2.90)	1.48 (3.22)
Res. Asian	1.63 (9.43)	1.92 (16.95)		1.34 (9.52)	1.09 (12.00)
Poverty	6.46 (3.82)	0.82 (5.74)		1.40 (3.84)	2.61 (4.48)
Renters	0.77 (2.11)	3.32 (2.86)		1.40 (2.11)	3.32 (2.40)
Land Usage					
Downtown	1.00 (1.00)	1.00 (1.00)		1.00 (1.00)	1.00 (1.00)
Vacant Units	1.04 (1.03)	1.04 (1.04)		1.05 (1.03)	1.07** (1.03)
Ambient Pop.					

Amb. Hetero.	3.06 (2.98)	9.90 (4.18)		1.59 (2.99)	4.66 (3.41)
Amb. Black	0.91 (1.84)	0.63 (2.26)		0.85 (1.85)	0.95 (2.03)
Amb. Hispanic	1.02 (1.88)	0.57 (2.23)		0.74 (1.89)	0.88 (2.05)
Amb. Asian	1.02 (2.30)	1.01 (3.46)		0.99 (2.30)	1.70 (2.59)
Amb. Pop. Size	1.04* (1.02)	1.04 (1.03)		1.07*** (1.02)	1.05** (1.02)
Deviance	29.26	43.81		21.33	59.24

Does the Twitter-based measure of ambient heterogeneity capture effects of young people?

Past research has shown that crime is over-proportionally committed by young people (Farrington 1986; Rocque, Posick and Hoyle 2015). Given critiques that young people are over-represented on Twitter (Oktay, Firat and Ertem 2014; Sloan et al. 2015), it is possible the ambient heterogeneity metric used in this study was also capturing differences in the presence of young people across neighborhoods. In this event, the relationship between ambient heterogeneity and crime could be spuriously explained by the presence of young people. To address this possibility, several indicators were constructed to measure differences in the presence of young people across neighborhoods. To measure the presence of young residents, we utilize two variables from the ACS: percentage of residents under the age of 18 and percentage of residents between 18 and 34. To explain variation in the presence of young people across ambient population units, place of interest data from Foursquare was spatially joined to tax assessment records to measure the percentage of buildings in each neighborhood that feature land uses such as bars or college buildings that systematically attract young people (see Appendix A).

We estimated regression models including the three measures intended to capture the presence of young people (see Table 6). While one age indicator, the percentage of residents between the ages of 18 and 34, had a statistically significant negative association with all four analyzed crime types, the coefficients for ambient heterogeneity remained statistically significant across models. These results suggest the relationship between ambient heterogeneity and crime is not simply explained by differences in the presence of young people across neighborhoods.

Table 6. Negative Binomial Models Examining the Tract-level Covariates of Crime Including Indicators of Young People Presence (n = 163)

Variable	Violence	Robbery		Larceny	Social Disorder
----------	----------	---------	--	---------	-----------------

Variable					
Intercept	11.63*** (2.27)	0.46 (2.94)		13.85 (2.28)	2.41 (2.50)
Residential Pop.					
Res.Hetero.	1.86 (2.30)	2.11 (2.98)		1.47 (2.31)	1.77 (2.55)
Res. Black	2.75 (2.02)	5.16* (2.50)		1.61 (2.04)	0.99 (2.23)
Res. Hispanic	1.05 (2.79)	1.90 (3.78)		0.46 (2.81)	1.42 (3.11)
Res. Asian	0.88 (4.80)	0.71 (7.26)		0.51 (4.82)	0.19 (5.60)
Poverty	2.85 (3.19)	0.47 (4.58)		1.62 (3.20)	2.21 (3.59)
Renters	3.01 (2.38)	7.52* (3.22)		5.10* (2.39)	8.98** (2.67)
Land Usage					
Downtown	0.90 (1.53)	1.16 (1.70)		0.77 (1.53)	0.92 (1.58)
Residential	0.88 (1.28)	0.89 (1.38)		1.06 (1.29)	0.88 (1.32)
Vacant Units	1.05** (1.02)	1.04 (1.03)		1.07** (1.02)	1.08*** (1.02)
Young People					
Res. Under 18	0.08 (13.09)	0.05 (27.92)		0.02 (13.21)	0.00** (17.44)
Res. 18-34	0.09** (2.97)	0.07* (4.04)		0.07** (2.97)	0.04*** (3.26)
Youth Land Uses	0.97 (1.06)	0.94 (1.08)		0.98 (1.06)	0.97 (1.06)
Ambient Pop.					
Amb. Hetero.	4.64* (2.32)	17.55*** (3.03)		5.41** (2.33)	8.32** (2.56)
Amb. Black	0.84 (1.82)	0.55 (2.20)		0.81 (1.83)	0.74 (1.99)

Amb. Hispanic	0.70 (1.77)	0.37 (2.29)		0.52 (1.77)	0.59 (1.93)
Amb. Asian	0.63 (2.21)	0.64 (3.22)		0.56 (2.22)	1.15 (2.46)
Amb. Pop. Size	1.00 (1.00)	1.00 (1.00)		1.00 (1.00)	1.00 (1.00)
Deviance	54.26	70.40		47.13	95.83

Does ambient heterogeneity drive crime in residential neighborhoods?

We also considered whether ambient heterogeneity may operate differently across residential and non-residential neighborhoods. Descriptively, we observed that non-residential neighborhoods ($\mu = 0.27$) tend to have higher levels of ambient heterogeneity than residential neighborhoods ($\mu = 0.16$; $t = 5.06$, $p < .001$). If there is a tipping point where a certain degree of ambient heterogeneity needs to proliferate for crime to emerge, ambient heterogeneity may play a smaller role in explaining crime differences across residential places.

To focus our inquiry to residential contexts, we re-conducted our main regression analyses on a subset of neighborhoods that only included tracts labeled residential by the land use data ($n = 117$; see Table 7). Interestingly, after limiting the data to only residential neighborhoods, ambient heterogeneity no longer had a statistically significant association with any form of crime. Instead, we observed a new relationship where ambient population size was positively associated with violence (IRR = 1.04, $p < .05$), larceny (IRR = 1.07, $p < .001$), and social disorder (IRR = 1.05, $p < .01$). When limiting analyses to residential places, which tend to have more homogeneous ambient populations, it appears that ambient population size supersedes ambient heterogeneity as the mechanism linking ambient population units to crime frequency.

Table 7. Negative Binomial Models Examining the Tract-level Covariates of Crime Including Demographics of Home Neighborhoods ($n = 163$)

Variable	Violence	Robbery		Larceny	Social Disorder
Variable					
Intercept	4.11** (1.82)	0.18** (2.29)		2.76* (1.82)	0.51 (1.97)
Residential Pop.					
Res.Hetero.	2.37 (2.65)	4.10 (3.54)		1.55 (2.66)	3.58 (2.98)
Res. Black	3.47* (1.96)	4.74* (2.32)		1.63 (1.97)	0.94 (2.12)
Res. Hispanic	1.24 (3.05)	1.95 (4.10)		0.62 (3.08)	0.85 (3.43)

Res. Asian	0.61 (5.23)	0.28 (7.83)		0.46 (5.25)	0.09 (6.11)
Poverty	3.66 (2.73)	0.39 (3.75)		1.12 (2.74)	1.65 (3.00)
Renters	0.86 (1.85)	1.94 (2.35)		1.98 (1.85)	4.03** (2.02)
Land Usage					
Downtown	1.19 (1.51)	1.60 (1.66)		0.98 (1.51)	1.37 (1.55)
Residential	0.96 (1.27)	1.00 (1.36)		1.14 (1.27)	1.06 (1.30)
Vacant Units	1.05** (1.02)	1.04 (1.03)		1.07*** (1.02)	1.08*** (1.02)
Ambient Pop.					
Amb. Hetero.	6.15** (2.48)	27.92*** (3.38)		6.66** (2.49)	17.15*** (2.77)
Amb. Black	0.94 (2.41)	0.63 (3.24)		0.82 (2.43)	1.52 (2.72)
Amb. Hispanic	0.80 (1.86)	0.33 (2.46)		0.65 (1.87)	0.69 (2.07)
Amb. Asian	0.62 (2.45)	0.44 (3.63)		0.53 (2.46)	1.74 (2.80)
Amb. Pop. Size	1.00 (1.00)	1.00 (1.00)		1.00 (1.00)	1.00 (1.00)
Home Pop.					
Home Hetero.	1.56 (5.07)	0.31 (9.00)		3.91 (5.14)	0.68 (6.40)
Home Black	0.67 (6.21)	1.73 (9.88)		0.86 (6.30)	0.19 (7.82)
Home Hispanic	1.04 (7.61)	2.26 (13.70)		0.27 (7.81)	0.91 (9.70)
Home Asian	0.54 (10.54)	6.55 (19.15)		0.33 (10.70)	0.09 (14.22)
Deviance	59.50	74.65		53.17	101.23

Why is ambient heterogeneity only associated with crime in non-residential places?

Given that the criminogenic effects of ambient heterogeneity appear to be limited to non-residential areas, we then investigated whether the relationship between ambient heterogeneity and crime is mediated by exogenous features that systematically vary across residential and non-residential contexts, focusing on the fact that non-residential areas tend to have more varied human activity. First, we considered the role of neighborhood visitors because outside visitors should be more concentrated in non-residential areas and they are likely to be a direct source of ambient heterogeneity. Furthermore, past research has shown that geographies with higher proportions of visitors in the ambient population tend to experience more crime (Song et al. 2021; Tucker et al. 2021; Wo et al. 2022), so if ambient heterogeneity and visitation are highly correlated it is possible the association between ambient heterogeneity and crime is spurious.

Table 8. Negative Binomial Models Examining the Tract-level Covariates of Crime Including Proportion of Visitors in Ambient Population (n = 163)					
Variable	Violence	Robbery		Larceny	Social Disorder
Variable					
Intercept	3.84** (1.72)	0.16** (2.13)		3.26** (1.72)	0.35* (1.87)
Residential Pop.					
Res.Hetero.	3.00 (2.29)	2.91 (2.94)		2.55 (2.29)	3.68 (2.53)
Res. Black	3.27** (1.80)	5.80** (2.09)		1.45 (1.81)	0.77 (1.93)
Res. Hispanic	1.17 (2.41)	2.38 (3.10)		0.38 (2.43)	0.82 (2.61)
Res. Asian	0.47 (4.77)	0.41 (7.00)		0.27 (4.79)	0.08 (5.51)
Poverty	3.45 (2.63)	0.42 (3.58)		1.32 (2.64)	1.27 (2.90)
Renters	0.90 (1.83)	1.94 (2.31)		1.92 (1.83)	3.90* (2.00)
Land Usage					
Downtown	1.22 (1.50)	1.53 (1.65)		1.05 (1.50)	1.29 (1.54)
Residential	0.95 (1.27)	1.02 (1.35)		1.15 (1.27)	0.98 (1.30)

Vacant Units	1.05** (1.02)	1.04 (1.03)		1.07*** (1.02)	1.08*** (1.02)
Ambient Pop.					
Amb. Hetero.	6.66** (2.30)	24.18*** (2.94)		8.16** (2.30)	12.24** (2.51)
Amb. Black	0.87 (1.80)	0.62 (2.15)		0.87 (1.81)	0.86 (1.97)
Amb. Hispanic	0.75 (1.79)	0.38 (2.33)		0.55 (1.80)	0.60 (1.95)
Amb. Asian	0.63 (2.22)	0.54 (3.29)		0.55 (2.23)	1.15 (2.48)
Amb. Pop. Size	1.00 (1.00)	1.00 (1.00)		1.00 (1.00)	1.00 (1.00)
Local Pop.	1.49 (1.70)	1.04 (2.08)		1.15 (1.71)	1.43 (1.83)
Deviance	58.64	75.04		53.80	102.99

To estimate the visiting population of each neighborhood, we calculated the percentage of Twitter users in each tract who had an estimated home location within that tract. To ensure the measure was consistent with other ambient population metrics in this study, visiting population was measured as the proportion of visitors in each tract during the average week in 2018. We introduced this variable into our city-wide regressions of crime and found that visiting population did not have a statistically significant effect to violent or non-violent crime types when controlling for other study variables (See Table 8). The effect sizes and p-values of other modeled variables remain stable after including this variable, providing no evidence of a mediating effect.

One potential explanation for why ambient heterogeneity is not driving crime in residential places is that such communities have far smaller ambient populations than non-residential places. To evaluate if the association between ambient heterogeneity and crime is stronger in neighborhoods with larger ambient populations, we introduced a term into our model representing the interaction between ambient heterogeneity and ambient population size (See Table 9). The interaction term did not reach statistical significance in any crime models, indicating that the localized effect of ambient heterogeneity in non-residential areas cannot be explained by the interaction of ambient heterogeneity and ambient population size.

Table 9. Negative Binomial Models Examining the Tract-level Covariates of Crime Including Interaction Between Ambient Heterogeneity and Ambient Population Size (n = 163)

Variable	Violence	Robbery		Larceny	Social Disorder
----------	----------	---------	--	---------	-----------------

Variable					
Intercept	4.15*** (1.70)	0.16** (2.08)		3.22** (1.70)	0.38 (1.83)
Residential Pop.					
Res.Hetero.	2.45 (2.26)	2.81 (2.90)		2.20 (2.27)	3.16 (2.50)
Res. Black	3.64** (1.82)	6.14** (2.14)		1.77 (1.83)	0.82 (1.96)
Res. Hispanic	1.62 (2.44)	2.64 (3.14)		0.55 (2.46)	1.04 (2.65)
Res. Asian	0.80 (4.85)	0.49 (7.15)		0.49 (4.87)	0.12 (5.60)
Poverty	2.94 (2.67)	0.40 (3.68)		1.05 (2.68)	1.15 (2.94)
Renters	0.82 (1.84)	1.87 (2.31)		1.75 (1.84)	3.61* (2.00)
Land Usage					
Downtown	1.19 (1.50)	1.54 (1.64)		1.07 (1.50)	1.28 (1.53)
Residential	0.95 (1.27)	1.02 (1.36)		1.14 (1.27)	0.98 (1.30)
Vacant Units	1.05** (1.02)	1.04 (1.03)		1.07*** (1.02)	1.08*** (1.02)
Ambient Pop.					
Amb. Hetero.	5.88** (2.30)	23.32*** (2.94)		6.74** (2.30)	11.15** (2.51)
Amb. Black	0.91 (1.80)	0.63 (2.14)		0.89 (1.81)	0.90 (1.96)
Amb. Hispanic	0.82 (1.77)	0.40 (2.31)		0.60 (1.77)	0.67 (1.93)
Amb. Asian	0.67 (2.23)	0.56 (3.30)		0.64 (2.24)	1.18 (2.48)
Amb. Pop. Size	1.02 (1.02)	1.01 (1.02)		1.03 (1.02)	1.01 (1.02)
Hetero*Pop. Size	0.95 (1.05)	0.98 (1.06)		0.93 (1.05)	0.97 (1.05)

Deviance	58.58	74.92		51.74	102.97
----------	-------	-------	--	-------	--------

DISCUSSION

Early in the 20th century, social disorganization theory was constructed to explain why certain urban areas experienced more crime than others. In identifying the place features that lead to social breakdown, Shaw and McKay (1942) and others focused on the features of neighborhood residents. From this perspective, residents are the central actors driving social outcomes in their neighborhoods. However, the ways people live have changed greatly over the last hundred years, with a main difference being that we have much greater spatial mobility in our daily lives; for example, one study found that commute distances in Toronto doubled from 1900 to 1950 (Bloomfield and Harris 1997) and a second study found that in the last century there has been a massive reduction in the proportion of British people walking and biking to work (Pooley and Turnbull 2000). Results from the present study suggest that as people engage in their routine activities across the city, their movements aggregate to shift the geo-racial structure of neighborhoods. This shift appears to have strong implications for urban social phenomena, evidenced by the finding that racial heterogeneity among ambient populations is a key correlate of crime.

Findings from this study suggest that we need an updated framework for thinking about the role collective demographics play in driving crime. At the time of Shaw & McKay's research, residential and ambient demographics may have been quite strongly correlated. According to the concentric zone model, working class people lived in the city center because these neighborhoods offered close proximity to the factories where they worked. In this geospatial structure where people had agency to move close to where they worked, it seems plausible that people could engage in their routine activities without leaving their home neighborhood. However, as processes of urban growth occurred across US cities the urban structure has become decentralized (Palumbo, Sacks and Wasylenko 1990; Schnore 1957). Additionally, the advent of robust public transportation networks and the proliferation of automobiles have facilitated travel across the urban landscape (see: Clark and Kuijpers-Linde 1994). As highlighted by the crime pattern theory literature, this means cities have a pulse where urban dwellers move across neighborhoods engaging in their routine activities (Felson and Boivin 2015)

While the geospatial organization of cities has changed, America's racial hierarchy and the ensuing conflict endure. As such, there is no reason to believe that the connection between racial heterogeneity and social bonding has changed over time. While residential neighborhoods remain racially segregated, it is easier than ever before for people to travel to different urban areas. Incidentally, when people leave their neighborhoods to engage in routine activities across more commercialized areas there is a convergence of individuals from diverse neighborhoods and racial backgrounds. Our findings suggest this convergence has important consequences in terms of rising crime rates. More broadly, results from this study suggest we can gain a fuller understanding of how social dynamics shape these types of places by shifting our perspective from residential populations to ambient populations. This study represents a first step in this agenda by providing a

methodology for measuring social characteristics of ambient populations and by demonstrating the relevance of measures generated through this technique.

Ambient Racial Heterogeneity

By developing a method to measure features of ambient populations, we have gained several important insights. First, we highlight that ambient population demographics do not necessarily align with residential population demographics. As neighborhood residents leave and visitors enter, the racial demographics of those present in the neighborhood do not perfectly mirror residential demographics. Notably, the magnitude of the gap between ambient and residential measures varies across races. Neighborhoods predominantly composed of white or Black residents are moderately associated with ambient populations of the same race. Moreover, white residents are associated with reduced ambient Black populations while Black residents are associated with reduced ambient white populations. These parallel demographic patterns are consistent with research showing mobility patterns are segregated along neighborhood racial demographics (Phillips et al. 2019; Wang et al. 2018). On the other hand, we observed much weaker correlations between the ambient and residential measures of Hispanic and Asian population. This could be because the mobility patterns of these groups are less entrenched in residential segregation patterns or because individuals from these groups are under-represented in the Twitter mobility sample, thereby making them prone to measurement error. Additionally, because Boston's Asian population is comparably small, measures for this group are likely more subject to marked fluctuations.

In calculating across the racial demographic measures, we find that ambient racial heterogeneity is weakly correlated to the traditional measure of residential heterogeneity. This illuminates a new dimension of a long-conceptualized social demographic measure. While Shaw and McKay's (1969) framework about residential heterogeneity has been immensely useful for explaining crime, its emphasis on the enduring ramifications of entrenched residential demographics has left a blind spot where criminologists have done little to consider what happens to the racial heterogeneity of places as residents leave their home neighborhood and engage in routine activities. Our results suggest that as these routine activities occur, population movements across the city shift the demographic profile of places away from their demographic characteristics to a magnitude that ambient racial heterogeneity represents a distinct sociodemographic neighborhood feature.

Ambient Racial Composition and Crime

While Shaw and McKay focused on how racial heterogeneity impacted community-building among residents of the same place, they may have also captured something deeper about the way that personal characteristics facilitate or impede pro-social behavior. Our results suggest that ambient heterogeneity positively explains rates of violent and non-violent crime beyond residential demographics. Ambient populations appear to have their own social dynamics that create or limit criminal opportunity depending on who is spending time in a given place. Guardianship appears to be increased when ambient population members occupy space with

others from a similar background. This is consistent with past research suggesting that heterogeneity impedes trust and cohesion (Charles and Kline 2006; Putnam 2007) and can impact the likelihood of someone taking action when deviant behavior occurs (Chekroun 2008; Chekroun and Nugier 2011; Moisuc and Brauer 2019).

Results from supplementary analyses suggest the positive association between ambient heterogeneity and crime is largely localized to non-residential areas. This finding, and the complementary insight that crime differences across residential neighborhoods can be explained by ambient population size, facilitates a more nuanced perspective of how ambient population dynamics shape crime. Urban mobility flows represent streams that carry inhabitants from across metropolitan areas into common “downtown” areas that are typically characterized by the proliferation of commercial and recreational activity. In these areas where the demographic makeup of those present is vulnerable to variability and rapid change, we see implicit evidence that a high degree of ambient racial heterogeneity can help explain crime differences across urban neighborhoods. On the contrary, the relationship between ambient heterogeneity and crime does not extend to residential places. There are several potential explanations of this finding. On one hand, given that residential neighborhoods tend to have lower levels of ambient heterogeneity on average, it is possible that ambient heterogeneity doesn’t become criminogenic until it reaches a certain threshold. It may be that people don’t consciously observe small amounts of racial heterogeneity in their surroundings or that such observations of this nature simply don’t alter guardianship behavior. While the present study was not able to find evidence of an interaction effect between ambient heterogeneity and ambient population size, if the interactive relationship is non-linear it may not be modeled appropriately by the current set of analyses.

To further understand how ambient population dynamics may encourage or discourage guardianship behavior, criminologists should seek to identify specific anti-crime behaviors that can be measured within ambient populations to evaluate if racial heterogeneity diminishes the presence of such behaviors. For example, a literature on ‘target hardening’ suggests that people can reduce the possibility of victimhood by engaging in practices that make them or their property less suitable targets, such as locking doors or building fences (Bowers, Johnson and Hirschfield 2004; Hirschfield, Newton and Rogerson 2010). Alternatively, psychologists have considered the potential deterrent effect that could arise if someone speaks up or expresses disapproval in response to deviant behavior (Chekroun 2008; Chekroun and Nugier 2011; Moisuc and Brauer 2019). Measuring the frequency of such behaviors across neighborhoods and evaluating their association with ambient heterogeneity would provide more nuanced insight into the mechanisms linking ambient heterogeneity to increased crime rates and allow further understanding about why this relationship appears to be stronger in some contexts than others.

Limitations

Of course, this study is not without its limitations. For one, our approach had some stringent requirements that substantially limited the amount of Twitter users we were able to analyze. Because our approach required users to have first and last names that were independently common enough in the US to be represented in the

external datasets used to establish race-name probabilities and also required users in the Boston mobility sample to provide enough geocoded tweets to estimate their home location, we were only able to analyze 45.06% of TargetSmart users for training and testing purposes and 23.53% of users in the Boston mobility sample for ambient population measurement. This could potentially bias study results if our sampled users have distinct characteristics from the rest of the data population. To address this issue in future work, and we encourage researchers to design more robust estimation strategies that rely on accessible, widely available information. More broadly, since Twitter represents a very small portion of the population, it is possible our sample lacks generalizability.

Additionally, our study is fairly limited in terms of its spatiotemporal scope. Due to the fairly small ambient population sample, we only had enough users and location points to reliably measure ambient population demographics on a weekly basis using census tract geographies. This is in contrast to previous research analyzing ambient populations at spatial scales as small as census blocks (Tucker et al. 2021) and time scales as specific as 2-hour windows (Hipp et al. 2019). Because our methodology lacks this granularity, we may not observe micro-bursts of ambient heterogeneity or crime that occur. As such, future research should re-test the role of ambient racial demographics at smaller spatiotemporal scales.

While these limitations of Twitter data are notable, we defend that it represents one of the best available options for measuring the characteristics of individuals within ambient populations. To improve the precision of mobility metrics, some scholars have used products from companies such as Safegraph (see: Zhang et al. 2022) and Cuebiq (Fraiberger et al. 2020; Lucchini et al. 2021) that measure population movement through cell phone records, providing a much more representative sample of Americans. However, in these datasets individual movements are generally aggregated to geographies, meaning they are void of information about individual-level characteristics. As such, to measure individual features within ambient populations scholars need to rely on data outside these oft-used sources of mobility data. Given evidence from past research that suggests Twitter can be used to construct valid mobility measures (Gao et al. 2014; Lenormand et al. 2014; McNeill, Bright and Hale 2017), we believe that it is an ideal source for this purpose despite the tradeoff between mobility accuracy and depth of individual-level information.

Beyond data, our theoretical interpretation is limited in that we focus on guardianship as the mechanism linking ambient heterogeneity to crime. There may be alternative explanations for why heterogeneous ambient populations experience heightened crime. For example, if ambient heterogeneity reduces cohesion, motivated offenders might feel less guilty about targeting individuals from other racial groups. While we focus on guardianship due to its centrality in both the social disorganization and routine activities literatures, future research should identify and measure potential mechanisms to further explain the connection between ambient heterogeneity and crime.

CONCLUSION

In the early 20th century, sociologists such as W.E.B. DuBois, Robert Park, and Ernest Burgess birthed a perspective that we can gain a better understanding of social phenomena by mapping demographic characteristics across space. Due to data limitations, these maps generally characterized the social state of places over the course of a year or even many years. Today we are seeing a gold rush of new types of user-generated data sources that allow us to measure individual features across space at much finer temporal scales. Criminologists should leverage these sources of information to rethink classical frameworks rooted in limited data and to understand how social systems change and operate in real time to drive outcomes. To guide these efforts, as a field we need to be asking as which characteristics of ambient populations impact the likelihood that present individuals will fill the roles of guardians, targets, and offenders. In doing so, we should reach beyond social disorganization theory and crime and place scholarship by drawing from individual-level theories of crime. By measuring theoretically emphasized individual-level features such as self-control (Akers 1991; Gottfredson and Hirschi 1990) and emotional strain (Agnew 1992; Agnew 2001), we can come to a much more robust understanding of how the mobility patterns of individuals aggregate to create ambient population dynamics that promote or limit crime across places.

REFERENCES

- Agnew, Robert. 1992. "Foundation for a General Strain Theory of Crime and Delinquency." *Criminology* 30(1):47-88.
- Agnew, Robert. 2001. "Building on the Foundation of General Strain Theory: Specifying the Types of Strain Most Likely to Lead to Crime and Delinquency." *Journal of Research in Crime and Delinquency* 38(4):319-61.
- Akers, Ronald L. 1991. "Self-Control as a General Theory of Crime." *Journal of Quantitative Criminology* 7(2):201-11.
- Andresen, Martin A. 2006. "Crime Measures and the Spatial Analysis of Criminal Activity." *British Journal of criminology* 46(2):258-85.
- Andresen, Martin A. 2010. "Diurnal Movements and the Ambient Population: An Application to Municipal-Level Crime Rate Calculations." *Canadian Journal of Criminology and Criminal Justice* 52(1):97-109.
- Bloomfield, A Victoria and Richard Harris. 1997. "The Journey to Work: A Historical Methodology." *Historical Methods: A Journal of Quantitative and Interdisciplinary History* 30(2):97-109.
- Bowers, Kate J, Shane D Johnson and Alex Hirschfield. 2004. "The Measurement of Crime Prevention Intensity and Its Impact on Levels of Crime." *British Journal of criminology* 44(3):419-40.

- Brantingham, Patricia and Paul Brantingham. 1995. "Criminality of Place." *European journal on criminal policy and research* 3(3):5-26.
- Browning, Christopher R and Kathleen A Cagney. 2002. "Neighborhood Structural Disadvantage, Collective Efficacy, and Self-Rated Physical Health in an Urban Setting." *Journal of health and social behavior*:383-99.
- Browning, Christopher R, Robert D Dietz and Seth L Feinberg. 2004. "The Paradox of Social Organization: Networks, Collective Efficacy, and Violent Crime in Urban Neighborhoods." *Social Forces* 83(2):503-34.
- Charles, Kerwin Kofi and Patrick Kline. 2006. "Relational Costs and the Production of Social Capital: Evidence from Carpooling." *The Economic Journal* 116(511):581-604.
- Chekroun, Peggy. 2008. "Social Control Behavior: The Effects of Social Situations and Personal Implication on Informal Social Sanctions." *Social and Personality Psychology Compass* 2(6):2141-58.
- Chekroun, Peggy and Armelle Nugier. 2011. "'I'm Ashamed Because of You, So Please, Don't Do That!': Reactions to Deviance as a Protection against a Threat to Social Image." *European Journal of Social Psychology* 41(4):479-88.
- Clark, William AV and Marianne Kuijpers-Linde. 1994. "Commuting in Restructuring Urban Regions." *Urban Studies* 31(3):465-83.
- Cohen, Lawrence E and Marcus Felson. 1979. "Social Change and Crime Rate Trends: A Routine Activity Approach." *American sociological review*:588-608.
- Coppersmith, Glen, Mark Dredze and Craig Harman. 2014. "Quantifying Mental Health Signals in Twitter." Pp. 51-60 in *Proceedings of the workshop on computational linguistics and clinical psychology: From linguistic signal to clinical reality*.
- Coston, Amanda, Neel Guha, Derek Ouyang, Lisa Lu, Alexandra Chouldechova and Daniel E Ho. 2021. "Leveraging Administrative Data for Bias Audits: Assessing Disparate Coverage with Mobility Data for Covid-19 Policy." Pp. 173-84 in *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*.
- Dinesen, Peter Thisted and Kim Mannemar Sønderskov. 2015. "Ethnic Diversity and Social Trust: Evidence from the Micro-Context." *American sociological review* 80(3):550-73.
- Domingos, Pedro and Michael Pazzani. 1997. "On the Optimality of the Simple Bayesian Classifier under Zero-One Loss." *Machine learning* 29(2):103-30.
- Eckel, Catherine C and Philip J Grossman. 2001. "Chivalry and Solidarity in Ultimatum Games." *Economic inquiry* 39(2):171-88.

Enos, Ryan D. 2014. "Causal Effect of Intergroup Contact on Exclusionary Attitudes." *Proceedings of the National Academy of Sciences* 111(10):3699-704.

Farrington, David P. 1986. "Age and Crime." *Crime and justice* 7:189-250.

Felson, Marcus and Lawrence E Cohen. 1980. "Human Ecology and Crime: A Routine Activity Approach." *Human Ecology* 8(4):389-406.

Felson, Marcus and Rémi Boivin. 2015. "Daily Crime Flows within a City." *Crime Science* 4(1):31.

Fraiberger, Samuel P, Pablo Astudillo, Lorenzo Candeago, Alex ChUNET, Nicholas KW Jones, Maham Faisal Khan, Bruno Lepri, Nancy Lozano Gracia, Lorenzo Lucchini and Emanuele Massaro. 2020. "Uncovering Socioeconomic Gaps in Mobility Reduction During the Covid-19 Pandemic Using Location Data." *arXiv preprint arXiv:2006.15195*.

Gao, Song, Jiue-An Yang, Bo Yan, Yingjie Hu, Krzysztof Janowicz and Grant McKenzie. 2014. "Detecting Origin-Destination Mobility Flows from Geotagged Tweets in Greater Los Angeles Area." in *Eighth Intl. Conference on Geographic Information Science (GIScience'14)*.

Glaeser, Edward L, David I Laibson, Jose A Scheinkman and Christine L Soutter. 2000. "Measuring Trust." *The Quarterly journal of economics* 115(3):811-46.

Gottfredson, MR and T Hirschi. 1990. "General Theory of Crime."

Gu, Xin, Lin Liu, Minxuan Lan and Hanlin Zhou. 2023. "Measuring Perceived Racial Heterogeneity and Its Impact on Crime: An Ambient Population-Based Approach." *Cities* 134:104188.

Hipp, John R and Andrew J Perrin. 2009. "The Simultaneous Effect of Social Distance and Physical Distance on the Formation of Neighborhood Ties." *City & Community* 8(1):5-25.

Hipp, John R, Christopher Bates, Moshe Lichman and Padhraic Smyth. 2019. "Using Social Media to Measure Temporal Ambient Population: Does It Help Explain Local Crime Rates?". *Justice Quarterly* 36(4):718-48.

Hirschfield, Alex, Andrew Newton and Michelle Rogerson. 2010. "Linking Burglary and Target Hardening at the Property Level: New Insights into Victimization and Burglary Protection." *Criminal Justice Policy Review* 21(3):319-37.

Lanfear, Charles C. 2022. "Collective Efficacy and the Built Environment." *Criminology* 60(2):370-96.

Lardier, David T, Ijeoma Opara, Yan Lin, Emily Roach, Andriana Herrera, Pauline Garcia-Reid and Robert J Reid. 2021. "A Spatial Analysis of Alcohol Outlet Density Type, Abandoned Properties, and Police Calls on Aggravated Assault Rates in a Northeastern Us City." *Substance use & misuse* 56(10):1527-35.

- Lenormand, Maxime, Miguel Picornell, Oliva G Cantú-Ros, Antonia Tugores, Thomas Louail, Ricardo Herranz, Marc Barthelemy, Enrique Frias-Martinez and José J Ramasco. 2014. "Cross-Checking Different Sources of Mobility Information." *PloS one* 9(8).
- Letki, Natalia. 2008. "Does Diversity Erode Social Cohesion? Social Capital and Race in British Neighbourhoods." *Political studies* 56(1):99-126.
- Linning, Shannon J, Ajima Olaghere and John E Eck. 2022. "Say Nope to Social Disorganization Criminology: The Importance of Creators in Neighborhood Social Control." *Crime Science* 11(1):1-11.
- Lucchini, Lorenzo, Simone Centellegher, Luca Pappalardo, Riccardo Gallotti, Filippo Privitera, Bruno Lepri and Marco De Nadai. 2021. "Living in a Pandemic: Adaptation of Individual Mobility and Social Activity in the Us." *arXiv preprint arXiv:2107.12235*.
- Malleson, Nick and Martin A Andresen. 2015. "Spatio-Temporal Crime Hotspots and the Ambient Population." *Crime Science* 4(1):10.
- McNeill, Graham, Jonathan Bright and Scott A Hale. 2017. "Estimating Local Commuting Patterns from Geolocated Twitter Data." *EPJ Data Science* 6(1):24.
- Miller, Zachary, Brian Dickinson and Wei Hu. 2012. "Gender Prediction on Twitter Using Stream Algorithms with N-Gram Character Features."
- Mislove, Alan, Sune Lehmann, Yong-Yeol Ahn, Jukka-Pekka Onnela and J Niels Rosenquist. 2011. "Understanding the Demographics of Twitter Users." in *Fifth international AAAI conference on weblogs and social media*.
- Moisuc, Alexandrina and Markus Brauer. 2019. "Social Norms Are Enforced by Friends: The Effect of Relationship Closeness on Bystanders' Tendency to Confront Perpetrators of Uncivil, Immoral, and Discriminatory Behaviors." *European Journal of Social Psychology* 49(4):824-30.
- Morenoff, Jeffrey D, Robert J Sampson and Stephen W Raudenbush. 2001. "Neighborhood Inequality, Collective Efficacy, and the Spatial Dynamics of Urban Violence." *Criminology* 39(3):517-58.
- Morgan-Lopez, Antonio A, Annice E Kim, Robert F Chew and Paul Ruddle. 2017. "Predicting Age Groups of Twitter Users Based on Language and Metadata Features." *PloS one* 12(8).
- Oktay, Huseyin, Aykut Firat and Zeynep Ertem. 2014. "Demographic Breakdown of Twitter Users: An Analysis Based on Names." *Academy of Science and Engineering (ASE):A*.
- Palumbo, George, Seymour Sacks and Michael Wasylenko. 1990. "Population Decentralization within Metropolitan Areas: 1970–1980." *Journal of Urban Economics* 27(2):151-67.

Park, Minsu, Chiyong Cha and Meeyoung Cha. 2012. "Depressive Moods of Users Portrayed in Twitter."

Phillips, Nolan E, Brian L Levy, Robert J Sampson, Mario L Small and Ryan Q Wang. 2019. "The Social Integration of American Cities: Network Measures of Connectedness Based on Everyday Mobility across Neighborhoods." *Sociological Methods & Research*:0049124119852386.

Pooley, Colin G and Jean Turnbull. 2000. "Modal Choice and Modal Change: The Journey to Work in Britain since 1890." *Journal of transport geography* 8(1):11-24.

Preoțiu-Pietro, Daniel, Svitlana Volkova, Vasileios Lampos, Yoram Bachrach and Nikolaos Aletras. 2015. "Studying User Income through Language, Behaviour and Affect in Social Media." *PloS one* 10(9).

Preoțiu-Pietro, Daniel and Lyle Ungar. 2018. "User-Level Race and Ethnicity Predictors from Twitter Text." Pp. 1534-45 in *Proceedings of the 27th International Conference on Computational Linguistics*.

Putnam, Robert D. 2007. "E Pluribus Unum: Diversity and Community in the Twenty-First Century the 2006 Johan Skytte Prize Lecture." *Scandinavian political studies* 30(2):137-74.

Rhoades, Stephen A. 1993. "The Herfindahl-Hirschman Index." *Fed. Res. Bull.* 79:188.

Roberts, Kirk, Michael A Roach, Joseph Johnson, Josh Guthrie and Sanda Harabagiu. 2012. "Empatweet: Annotating and Detecting Emotions on Twitter." Pp. 3806-13 in *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*.

Rocque, Michael, Chad Posick and Justin Hoyle. 2015. "Age and Crime." *The encyclopedia of crime and punishment*:1-8.

Sampson, Robert J, Stephen W Raudenbush and Felton Earls. 1997. "Neighborhoods and Violent Crime: A Multilevel Study of Collective Efficacy." *Science* 277(5328):918-24.

Sampson, Robert J, Jeffrey D Morenoff and Felton Earls. 1999. "Beyond Social Capital: Spatial Dynamics of Collective Efficacy for Children." *American sociological review*:633-60.

Sampson, Robert J. 2006. "Collective Efficacy Theory: Lessons Learned and Directions for Future Inquiry." *Taking stock: The status of criminological theory* 15:149-67.

Schnore, Leo F. 1957. "Metropolitan Growth and Decentralization." *American journal of sociology* 63(2):171-80.

Shaw, Clifford R. and Henry D. McKay. 1969. *Juvenile Delinquency and Urban Areas; a Study of Rates of Delinquency in Relation to Differential Characteristics of Local Communities in American Cities*. Chicago,: University of Chicago Press.

Sloan, Luke, Jeffrey Morgan, Pete Burnap and Matthew Williams. 2015. "Who Tweets? Deriving the Demographic Characteristics of Age, Occupation and Social Class from Twitter User Meta-Data." *PloS one* 10(3):e0115545.

Song, Guangwen, Lin Liu, Wim Bernasco, Luzi Xiao, Suhong Zhou and Weiwei Liao. 2018. "Testing Indicators of Risk Populations for Theft from the Person across Space and Time: The Significance of Mobility and Outdoor Activity." *Annals of the American Association of Geographers* 108(5):1370-88.

Song, Guangwen, Yanji Zhang, Wim Bernasco, Liang Cai, Lin Liu, Bo Qin and Peng Chen. 2021. "Residents, Employees and Visitors: Effects of Three Types of Ambient Population on Theft on Weekdays and Weekends in Beijing, China." *Journal of Quantitative Criminology*:1-39.

Song, Guangwen, Liang Cai, Lin Liu, Luzi Xiao, Yuhan Wu and Han Yue. 2023. "Effects of Ambient Population with Different Income Levels on the Spatio-Temporal Pattern of Theft: A Study Based on Mobile Phone Big Data." *Cities* 137:104331.

Stier, Andrew J, Kathryn E Schertz, Nak Won Rim, Carlos Cardenas-Iniguez, Benjamin B Lahey, Luís MA Bettencourt and Marc G Berman. 2021. "Evidence and Theory for Lower Rates of Depression in Larger Us Urban Areas." *Proceedings of the National Academy of Sciences* 118(31):e2022472118.

Taylor, Joanna, Liz Twigg and John Mohan. 2010. "Investigating Perceptions of Antisocial Behaviour and Neighbourhood Ethnic Heterogeneity in the British Crime Survey." *Transactions of the Institute of British Geographers* 35(1):59-75.

Tucker, Riley, Daniel T O'Brien, Alexandra Ciomek, Edgar Castro, Qi Wang and Nolan Edward Phillips. 2021. "Who 'Tweets' Where and When, and How Does It Help Understand Crime Rates at Places? Measuring the Presence of Tourists and Commuters in Ambient Populations." *Journal of Quantitative Criminology* 37(2):333-59.

Twigg, Liz, Joanna Taylor and John Mohan. 2010. "Diversity or Disadvantage? Putnam, Goodhart, Ethnic Heterogeneity, and Collective Efficacy." *Environment and planning A* 42(6):1421-38.

Twitter. 2019. in *Most people don't tag their precise location in Tweets, so we're removing this*

ability to simplify your Tweeting experience. You'll still be able to tag your precise location in Tweets through our updated camera. It's helpful when sharing on-the-ground moments., edited by @TwitterSupport: Twitter.

Wang, Qi, Nolan Edward Phillips, Mario L Small and Robert J Sampson. 2018. "Urban Mobility and Neighborhood Isolation in America's 50 Largest Cities." *Proceedings of the National Academy of Sciences* 115(30):7735-40.

Wang, Wenbo, Lu Chen, Krishnaprasad Thirunarayan and Amit P Sheth. 2012. "Harnessing Twitter" Big Data" for Automatic Emotion Identification." Pp. 587-92 in *2012 International Conference on Privacy, Security, Risk and Trust and 2012 International Confernece on Social Computing*: IEEE.

Wo, James C, Ethan M Rogers, Mark T Berg and Caglar Koylu. 2022. "Recreating Human Mobility Patterns through the Lens of Social Media: Using Twitter to Model the Social Ecology of Crime." *Crime & Delinquency*:00111287221106946.

Zhang, Harry. 2004. "The Optimality of Naïve Bayes." *Aa* 1(2):3

Zhang, Mengxi, Siqin Wang, Tao Hu, Xiaokang Fu, Xiaoyue Wang, Yaxin Hu, Briana Halloran, Zhenlong Li, Yunhe Cui and Haokun Liu. 2022. "Human Mobility and Covid-19 Transmission: A Systematic Review and Future Directions." *Annals of GIS* 28(4):501-14.

Zoorob, Michael, Alina Ristea, Saina Sheini and Daniel T. O'Brien. 2021. "Geographical Infrastructure for the City of Boston (V. 2021)." *Harvard Dataverse: Boston Area Research Initiative*.

APPENDIX A

Place of interest data was collected from Foursquare, a company that hosted a platform allowing users to crowd-source information about local businesses and public spaces. Using an API provided by Foursquare, we downloaded information on 17,241 places of interest in the greater Boston area. For each of place of interest, Foursquare provides geographic coordinates and a tag indicating the natural of land use at each place. To measure the presence of young people attracting places, the data was restricted to places from the following categories: alcohol-focused nightlife, performance arts venues, and buildings associated with college students (see table below for specific land use tags included in each category). To measure the percentage of buildings in each neighborhood with young-attracting land uses, the place of interest locations were spatial joined to land parcel polygons provided by the Tax Assessment Database to count the number of relevant buildings per neighborhood. This count was divided by the total number of parcels in each neighborhood to generate the analyzed measure.

Place of Interest Tags for Land Uses Attracting Young People	
Category	Tags

Alcohol & Nightlife	'Speakeasy', 'Bar', 'Brewery', 'Karaoke Bar', 'Pub', 'Gastropub', 'Nightclub', 'Cocktail Bar', 'Dive Bar', 'Gay Bar', 'Wine Bar', 'Beer Garden', 'Winery', 'Strip Club', 'Sports Bar', 'Lounge', 'Hotel Bar', 'Roof Deck', 'Hookah Bar', 'Other Nightlife'
Performance Arts Venues	'Comedy Club', 'Rock Club', 'Dance Studio', 'Theater', 'Music Venue', 'Concert Hall', 'Event Space', 'Performing Arts Venue', 'Public Art'
College Buildings	'College Academic Building', 'General College & University', 'College Gym', 'College Administrative Building', 'University', 'College Arts Building', 'College Auditorium', 'Auditorium', 'College Classroom', 'College Quad', 'College Residence Hall', 'College Lab', 'College Library', 'College Bookstore', 'College & University', 'Law School' 'College Cafeteria', 'Student Center', 'College Science Building', 'College Rec Center', 'Fraternity House', 'College Engineering Building',

[1] <https://targetsmart.com/>

[2] Missing values were recoded to 0.

[3] To date, geotagged tweets remain downloadable as of February 2023.

[4] Other analyses in this study utilize the raw probabilities to minimize measurement bias.

[5] Of these users, 68.77% (n = 1,515) had their race estimated through the Bayes model and 31.23% (n = 688) had their race identified via linkage to the TargetSmart sample.

[6] Users with verified racial demographics were treated as having a 100% probability of belonging to the racial group assigned by TargetSmart.

[7] 0.99 for white users, 0.79 for Asian users, 0.95 for Black users, and 1.42 for Hispanic users