"Broken Windows" and the Dynamics of Disorder: 3-1-1 Calls in New York City

Nicole Pangborn

pangborn@princeton.edu

Abstract

Disorder is a common feature of the modern urban landscape. While it has been associated with various negative neighborhood outcomes, it has most notably been linked to crime rates. The famous "broken windows" hypothesis of crime claims there is a connection between disorder in appearance ("physical disorder") and disorder in behavior ("social disorder") in neighborhoods over time. Challengers of broken windows posit that this mechanism does not work the same way in all neighborhoods, since "disorder" is not an objectively-defined entity; however, they have lacked the fine-grained spatiotemporal data necessary for the firm confirmation of their claims. This project examines neighborhood disorder on a more dynamic level. Through the use of 3-1-1 calls in New York City, I find that neighborhoods have markedly different trajectories over time in physical-social disorder space, and higher levels of physical disorder at one timepoint do not necessarily lead to higher levels of social disorder at future timepoints. In addition, different disorder trajectory patterns are associated with different social characteristics of neighborhoods.

1 Introduction

Some urban sociologists have theorized that a certain amount of neighborhood disorder is part and parcel of successful urban life. Heterogeneity, vibrancy, and dynamic sidewalk life can help keep a city alive, functioning to help its residents "learn from people who are unlike [them]selves" [Sennett, 1992]. However, other scholars maintain that disorder is associated with various negative life outcomes for city dwellers. It has been linked to stress [Hill et al., 2005, Latkin and Curry, 2003], poor health [Ross and Mirowsky, 2001], and most notably, crime [Wilson and Kelling, 1982, Skogan, 1990, Kelling and Coles, 1996]. In particular, Wilson and Kelling's famous "broken windows" piece, published in *The Atlantic Magazine* in 1982, connected neighborhood disorder to more serious crime, placing it in the foreground of discussion for urban sociologists, crimonologists, and policy-makers.

The essence of the broken windows hypothesis was that "serious street crime flourishes in areas in which disorderly behavior goes unchecked" [Wilson and Kelling, 1982]. Wilson and Kelling based their idea on a demonstration conducted by social psychologist Philip Zimbardo, who left two abandoned vehicles in the streets of very different neighborhoods (one in the Bronx, NY, and one in Palo Alto, CA) to observe how neighborhood residents would react. In a very short amount of time, the car in the Bronx was broken into and its contents stripped and stolen, while the car in Palo Alto remained untouched. Zimbardo observed, however, that when he smashed the window of the Palo Alto car himself, soon thereafter it was looted in the same way that the vehicle in the Bronx was [Zimbardo, 1969].

Wilson and Kelling used this result to popularize the idea that small types of disorder in appearance — for example, a broken window — spawn more disorder, often of the social and criminal variety. Their argument hinged on the idea that signs of physical disorder serve as visual cues that residents have "given up" on the neighborhood, stigmatizing it (and those living within its borders) as apathetic and uncaring. When a marker of disorder is left unchecked, it becomes a fixture of a neighborhood, leading to the breakdown of the type of informal social control that urban scholars such as Jacobs [1961] maintain is necessary for successful city life. Without it, residents withdraw in fear, creating a breeding ground for criminal activity to flourish.

Urban policy-makers and police commissioners jumped on the broken windows idea immediately, hoping that cracking down on smaller disorderly behaviors would cause an overall reduction in crime rates. Most famously, New York City mayor Rudy Giuliani implemented a type of "Quality of Life" initiative, where minor crimes of disorderly behavior were treated more seriously than they were previously, and officers were even rewarded for high numbers of arrests for "nuisance" crimes [Onishi, 1994]. Police in multiple cities across the U.S. began "zero-tolerance" strategies, charging small offenses of disorder —panhandling, public drunkenness, graffiti — as legitimate criminal offenses.

Critiques

The idea behind broken windows necessitates a causal connection over time between physical disorder and social disorder (and, ultimately, crime, as the most extreme form of social disorder) in any given neighborhood. But, recently, that connection has been shown to be questionable.

First, one set of scholars maintains that the link between disorder and crime is correlational rather than causal. Sampson and Raudenbush [1999] test whether or not measures of neighborhood disorder are essential in explaining neighborhood crime in Chicago. They find that the relationship between disorder and crime is a spurious one: when they examine factors such as concentrated disadvantage, land use, and "collective efficacy"[1], any relationship between disorder and crime essentially disappears [Sampson and Raudenbush, 1999, Harcourt, 2009]. Their conclusions point to the possibility that crime and disorder are both based in the same structural characteristics of neighborhoods, rather than causally related [Sampson, 2012]. Similarly, Yang [2010] finds that high physical disorder doesn't always precede high levels of social disorder in crude longitudinal analysis, also suggesting that a causal mechanism between the two is unconfirmed.

Second, the causal mechanism implicit in Wilson and Kelling's argument involves the *social meaning* of disorder: graffiti always *means* breakdown of informal social control, and can thus serve as a universal signal of the sorts of people who live in a particular neighborhood. But this rests on a large assumption: that "disorder" is an objectively measurable entity, seen in the same way by all groups of people [Harcourt, 2009]. Couldn't we imagine a scenario in which

[1] Collective efficacy is defined as the "differential ability of neighborhoods to realize the common values of residents and maintain effective social controls". For more details, see Sampson et al. [1997].

this assumption might not be true? Take, for example, a gentrifying neighborhood in Brooklyn, where graffiti might serve as a marker for a hip, edgy place to live for young people [Zukin, 2009]. Would this graffiti still be the type which leads to gang-related or criminal behavior?

Such questions have led to more recent focuses on the *perception* of disorder. Sampson and Raudenbush [2004] in particular discuss the social construction of disorder: what does it mean to perceive a broken window, and does that perception vary by place? They combine neighborhood surverys, police records, census data, and systematic social observation methods to examine the social context of perceived disorder in Chicago. They hypothesize that neighborhood social composition (race, class, ethnicity) predicts perceptions of neighborhood disorder *net of systematically observed disorder*. They find that this is indeed the case, concluding that "social structure" — in this case, racial, ethnic, and socioeconomic composition — "proved a more powerful predictor of perceived disorder than did carefully observed disorder" [Sampson and Raudenbush, 2004, p. 336]. Their results suggest that residents supplement any knowledge from observed appearance with prior beliefs about neighborhood repute.

Given these findings, we could say that objectively observed "disorder" might not hold universal meaning. In the words of Sampson [2011], "ugliness more than beauty, we might say, is 'in the eye of the beholder'". Acting on appearance alone might have heterogeneous consequences, dependent on local meanings and social context: if people are different, the meaning of disorder will be different. So, if "broken windows" themselves are socially constructed, the mechanism behind the broken windows hypothesis — that disorder in appearance causes disorder in behavior — cannot universally hold when neighborhoods are socially heterogeneous.

Main questions

The neighborhood heterogeneity described above has been demonstrated only in cross-sectional or crude two-point longitudinal contexts. But with the advent of new digital data collection, we can examine this heterogeneity more closely, shedding light on differences in the connection between physical and social disorder over time. This project uses a new dataset — the full set of phone calls made to 3-1-1 in New York City from 2010 to 2015 — to answer the following questions:

1. What do the fine-grained dynamics of disorder say about neighborhood heterogeneity? and

2. What does such dynamic heterogeneity tell us about the broken windows hypothesis?

2 Data: 3-1-1 calls

3-1-1 is a telephone number for non-emergency municipal services. In 2003, New York City mayor Michael Bloomberg consolidated the thousands of telephone numbers for government assistance into one easy-to-remember 3-digit code [Accenture, 2003]. Residents use this number for everything from personal information assistance (e.g., requesting the next ferry time) to filing complaints about potholes on their streets. The complaints of interest for this particular project were those involving physical or social disorder in a neighborhood.

To acquire these complaints of interest, I first collected the full log of 3-1-1 calls from New

York's `OpenData`[2] website from January 2010 to January 2015. The city's `OpenData` program maintains data sets gathered from a wide range of municipal services and organizations, and makes all datasets available to the public on their webpage. The initial dataset I downloaded consisted of 8.6 million calls, each one an observation, tagged with a timestamp (`%Y-%m-%d %H:%M:%S`), precise location (latitude/longitude), and complaint type (244 unique types). The dataset's size and nature allow for fine-grained spatiotemporal analyses — of both "disorder" and other qualities — on an unprecedented scale.

Classification

I separated out the calls related to "disorder" in a neighborhood as defined by the literature [Skogan, 1990, Sampson and Raudenbush, 1999, 2004, Yang, 2010, Sampson, 2012, Hwang and Sampson, 2014]. Table 1 presents the classification scheme of 30 complaints into "Physical" (relating to neighborhood appearance) or "Social" (relating to the behavior of a person). 2.25 million of the initial 8.6 million calls were classified as dissorder calls. All remaining calls — most of which were complaints about personal or informational issues — were classified as "other" for use in the normalization techniques defined below.

Normalization and main measures

To quantify neighborhood disorder from call data, it is necessary to recognize that 3-1-1 calls are a marker of a few things simultaneously:

1. *Perceived* disorder (not objective)
2. Knowledge of 3-1-1 and municipal services
3. Trust in 3-1-1 and municipal services

For the purposes of this study, measures must reflect only (1): *perceived* disorder. So, capturing differential levels of (2) and (3) — knowledge of and trust in 3-1-1 and municipal services by neighborhood — becomes a main concern.

In order to find a faithful measurement of perceived physical and social disorder to compare neighborhoods over time, I use the relative frequencies of physical disorder calls or social disorder calls, per timepoint (month), per geographic unit (Zip Code Tabulation Area, or ZCTA). ZCTAs used in analyses were restricted to those in which residents actually lived (i.e., I excluded all zip codes which contained only parks, malls, airports, or no residents[3]). To normalize the frequencies of each, I use the number of "other" (or personal, "informational" calls). Per location l and timepoint t, I define physical disorder ($d_{p,l,t}$) and social disorder ($d_{s,l,t}$) as:

$$d_{p,l,t} = \frac{\text{no. of physical disorder calls}}{\text{no. of other calls}} \tag{1}$$

$$d_{s,l,t} = \frac{\text{no. of social disorder calls}}{\text{no. of other calls}} \tag{2}$$

There are two reasons for using "other" calls in the denominator. First, I wish to minimize the dependence between the two measures of physical or social disorder; if I normalized instead by the total call volume per timepoint per geographic unit, physical disorder would be dependent on social disorder — as one increases, the

[2] `https://nycopendata.socrata.com/`

[3] The following zip codes were excluded based on these criteria: 10069, 10282, 11109, 10004, 10006, 10464, 10470, 11693, 11239, 11360, 11430

other would decrease. Second, using "other" call volume (rather than, say, a more traditional per capita measure) allows us to compare neighborhood measures while capturing differential levels in knowledge and trust in municipal services by neighborhood, assuming these can be reflected by rates of calling 3-1-1 about other concerns.

To examine neighborhood trajectories in physical and social disorder over the five year time period, I first adjusted for the effects of seasonality using a method based on Loess, which is "a filtering procedure for decomposing a seasonal time series into three components: trend, seasonal, and remainder" [Cleveland et al., 1990]. This essentially involves a weighted local regression, weighting points in localized subsets based on their distance to the point of estimation. Here, monthly points were given weights based on their "typical" values to correct for seasonality. The `decompose()` function in R base stats handles this task well. (See Figure 1 for an example of trend, seasonal, and remainder decomposition.) After decomposing each time series in each zip code into its seasonal component, I subtract the seasonal component from the original time series data to arrive at the adjusted data [Coghlan, 2014].

Finally, I smoothed each neighborhood's seasonally adjusted data with an exponential moving average, using a period of 12 for monthly timepoints. The effects of these techniques are to remove the types of "noise" found in messy time-dependent data, which are common across all ZCTAs, in order to arrive at the underlying overall trends in social and physical disorder over time in any given neighborhood.

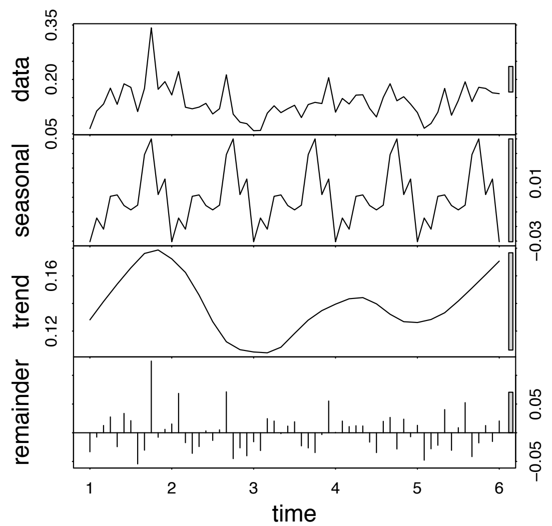

Figure 1: Time Series Decomposition

An example decomposition of time series data in d_s for zip code 10001 in Manhattan. The seasonal component, in panel 2, is subtracted from the data component, top panel, to arrive at the overall seasonally-corrected trend in panel 3. This type of seasonal decomposition was done for each neighborhood.

Demographic data

All demographic data were obtained from the Census's American Fact Finder[4] website.

3 Results

The results below are presented in three subsections. First, I use the relative frequencies of physical and social disorder to present neighborhood trajectories in physical-social disorder space. Second, I use k-means clustering to sep-

[4]`http://factfinder.census.gov/faces/nav/jsf/pages/index.xhtml`

arate the neighborhood trajectories into groups, demonstrating heterogeneity in the dynamics of neighborhood disorder. Finally, I examine the social characteristics of the emergent clusters of neighborhoods to better understand what such dynamic heterogeneity might mean for the social processes behind the broken windows theory.

Trajectories

For each of the ZCTAs in New York City ("neighborhoods"), I calculated the smoothed and seasonally-adjusted trend of relative frequencies of physical and social disorder per month. Since these adjustments discarded the first 11 datapoints in each trend, I was left with 50 timepoints per neighborhood. I plotted each of these points in physical-social disorder space $[(x, y) = (d_p, d_s)]$, drawing an arrow along the segments between timepoints. The goal here was to create a path of movement — similar to a vector field — between values of physical and social disorder over time. I refer to these plots as neighborhood trajectories.

Figure 2 presents an example trajectory for zip code 10001, an area in midtown Manhattan. As we can see, both physical and social disorder decrease overall during the timeframe from 2010 - 2015, and the more recent timepoints (marked by the shift from yellow to red coloring) stay relatively stable in the lower left hand corner of the plot.

We can contrast Figure 2 with Figure 3, which presents the same type of trajectory for zip code 11102, an area in Astoria, Queens. This trajectory resembles somewhat of a U-shape, beginning with high relative physical and social disorder, decreasing in both, and then increasing in d_s as we approach 2015.

After examining all of the trajectories for all

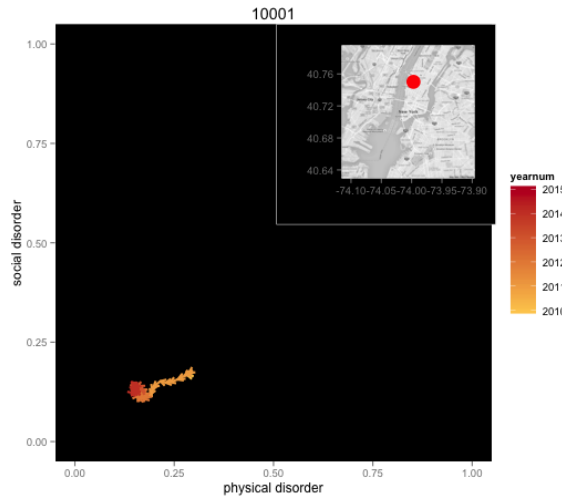

Figure 2: Trajectory for zip code 10001 (midtown Manhattan)

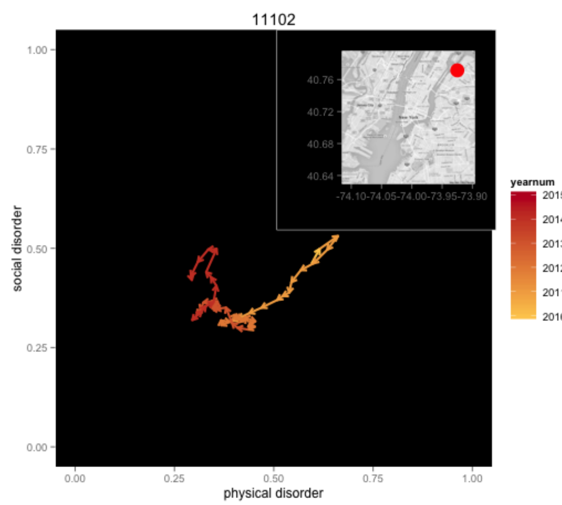

Figure 3: Trajectory for zip code 11102 (Astoria, Queens)

zip codes, many appeared to be quite similar to either Figure 2 or Figure 3. The next step became finding similar trajectories and differentiating the classes of trajectories from one another. Finding such differences in neighborhoods (in the relationship between physical and social disorder over time) would bolster previous critiques of broken windows theory—which emphasize neighborhood heterogeneity—in a more dynamic setting. I use k-means clustering to accomplish this task.

k-means clustering

K-means clustering is a data mining technique which separates n individual cases into k clusters based on "closeness" to a set of designated means.[5] The goal of k-means is to minimize the sum of squares within each cluster (i.e., the sum of the squares of distances between each point within each cluster and the mean of that cluster).

Traditional k-means clustering involves calculating distances within a matrix of independent attributes for each of n individual cases. But for neighborhood trajectories, each "case" is a trajectory — of time-dependent movement in two variables, d_p and d_s. Since time-dependent data cannot be considered as a set of independent variables by each timepoint[6], the baseline k-means implementation must be modified.

In order to take into account the co-evolution of *two* time series — trends in d_p and d_s — I use the `kml3d`[7] package in `R`, which works jointly on more than one variable over time [Genolini et al., 2013]. In this version of k-means, the "means" are the mean trajectories for each cluster of trajectories. In this case, the distance metric must essentially measure the distance between all pairs of disorder trajectory matrices, D_i and D_j, for t timepoints:

$$D_i = \begin{pmatrix} d_{p_{i,1}} & d_{p_{i,2}} & \dots & d_{p_{i,t}} \\ d_{s_{i,1}} & d_{s_{i,2}} & \dots & d_{s_{i,t}} \end{pmatrix}$$

$$D_j = \begin{pmatrix} d_{p_{j,1}} & d_{p_{j,2}} & \dots & d_{p_{j,t}} \\ d_{s_{j,1}} & d_{s_{j,2}} & \dots & d_{s_{j,t}} \end{pmatrix}$$

By default, `kml3d` uses the Euclidean distance between these two matrices:

$$Dist(D_i, D_j) = \sqrt{\sum_t (d_{p_{i,t}} - d_{p_{j,t}})^2 + (d_{s_{i,t}} - d_{s_{j,t}})^2}$$

where t is the set of all timepoints (here, $t = 50$ timepoints for December 2010 - January 2015).

One advantage of using the `kml3d` package is that it allows for seamless normalization of time series trajectory data. The normalization occurs not just at each timepoint (as would be the case if using the typical implementation of k-means), but over the entirety of the time series [Genolini et al., 2013]. So, trajectory data is centered at zero before clusters are computed. I chose this option because this project focuses on how neighborhoods are different from each other in the ways in which they *change*, independent of starting conditions (starting d_p, d_s). In Figure

[5] Steps of the algorithm are as follows. First, means can be initially set randomly within the n observations. The first set of k clusters are created by minimizing the distance between each observation and the initial means. Once the clusters have been set, the centroid of each cluster becomes the new mean for the cluster, and the iterative process begins again. Once no individual cases switch clusters at the new cluster mean stage, the clusters are considered stable and the process is complete [Hartigan and Wong, 1979].

[6] They are essentially repeated measures, where each value is dependent on the previous value.

[7] `http://cran.r-project.org/web/packages/kml3d/index.html`

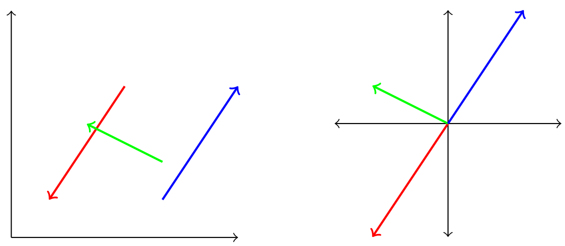

Figure 4: Normalizing vectors

4, we see a simple example of normalizing the vectors on the left panel to center around zero, leaving information only about the *direction* and relative *magnitude* of the vector (no information about its starting values).

I identified the optimal number of clusters for the joint time series data by maximizing the Calinski and Harabatz criterion, $c(k)$:

$$c(k) = \frac{\text{Trace}(B)}{\text{Trace}(W)} \frac{n-k}{k-1}$$

where B is the matrix of variance between, W is the matrix of variance within, n is the number of cases, and k the number of clusters [Caliski and Harabasz, 1974, Genolini et al., 2013].

I ran the `kml3d` version of k-means on the time series data for d_p and d_s in the set of New York City zip codes, setting the range for possible clusters as 2:9[8]. The algorithm identified three clusters of neighborhood trajectories: A (41.6% of neighborhood cases), B (39.8%), and C (18.7%).

Figure 5 presents the neighborhood disorder trajectories by cluster, separated out by single variables d_p and d_s. We can see that the trajectories appear to be more clearly differentiated from each other in the d_p single variable case. Importantly, `kml3d` takes into account joint movement in *both* variables when minimizing the distance between each neighborhood case and cluster means.

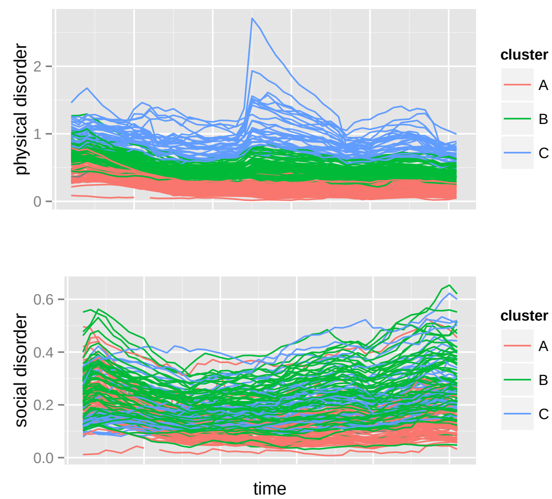

Figure 5: kml3d clustering results

Neighborhood disorder trajectories by cluster, separated out by single variables d_p and d_s. In order to show as much differentiation as possible, Y-axis scales are not the same between the two panels.

Figure 6 (next page) displays the three average neighborhood trajectories by cluster.

Cluster A, which decreases in both values steadily and stabilizes at a roughly equal value of d_p and d_s, akin to Figure 2, is seen on the top; in the center, cluster B decreases in d_p similarly to cluster A, but with an increase in d_s in more recent timepoints, resembling Figure 3; and cluster C, on the bottom, displays more variable d_p and overall increase in d_s. Importantly, in *none*

[8]This range was chosen following the "rule of thumb" of k-means clustering [Hartigan, 1975], which states that a good estimation of the number of k clusters in n observations is $k \approx \sqrt{n/2}$, while allowing for a potentially smaller number of clusters.

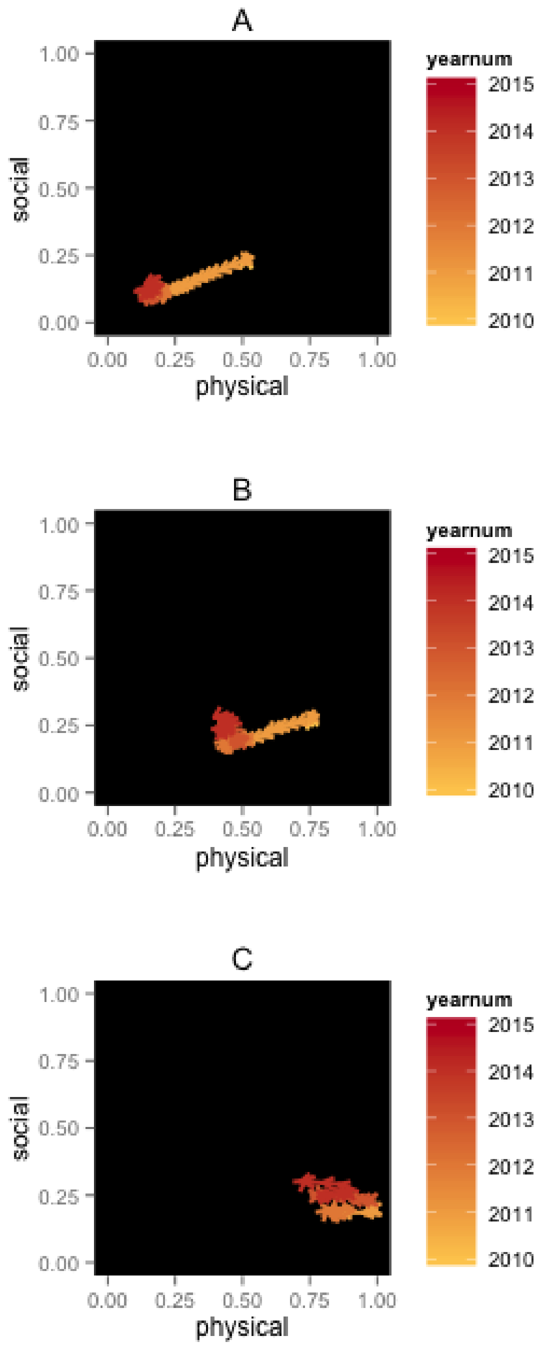

Figure 6: Average neighborhood disorder trajectory, by cluster

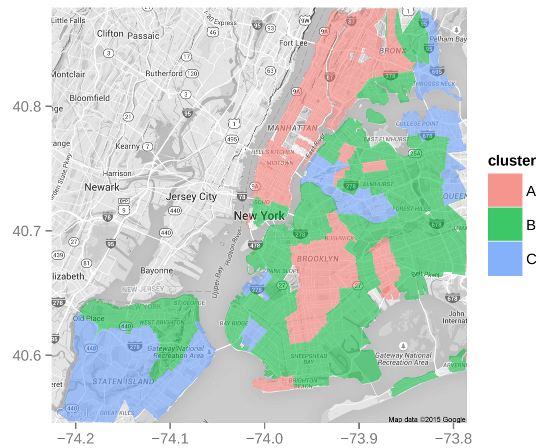

Figure 7: ZCTAs by trajectory cluster

Each ZCTA is colored by its assigned trajectory cluster. Neighborhoods in the same cluster have "similar" trajectories in d_p and d_s over time. Interpretation?

of these clusters do high levels of physical disorder consistently directly precede higher levels of social disorder, which is the essence behind the causal mechanism of the broken windows hypothesis.

Clustered characteristics

Figure 7 presents a map of New York City ZCTAs, colored by their trajectory clusters. Many of the zip codes in Manhattan and the Bronx are grouped in cluster A, which means their trajectories of physical and social disorder tend to decrease and stabilize. By contrast, with the curious exception of a group of zip codes in the center of the region, Brooklyn and Queens are mostly colored green, indicating that its zip codes have disorder trajectories which increase in social disorder in recent years. Cluster C, with more volatile disorder trajectories, forms Staten Island

and some of the outer zip codes in Queens.

After examining these average trajectories, I used census data to understand the social characteristics of the neighborhoods within each cluster. Figure 8 (next page) presents boxplots of neighborhood social characteristics by cluster. We can make a few observations here:

- Cluster A is formed by zip codes whose residents are relatively young, mobile (have not lived in their current residence for more than 1 year), and do not own their residences;

- Cluster B is formed by zip codes whose residents are also relatively young, but less mobile and more likely to be owners of their residences; and

- Cluster C is formed by zip codes whose residents are relatively old, the least mobile, and are very likely owners of their residences.

Despite the above varying population characteristics, we can see that all three clusters maintain about the same average rent change from 2011 to 2013.

4 Conclusions

This project examines the connection between physical and social disorder in urban neighborhoods through use of a new "big data" source. It puts forth both methodological and substantive contributions to the fields of urban sociology and criminology.

First, it presents a new way to measure the change in physical and social disorder over time. New administrative sources for "big data" (e.g., 3-1-1) can provide more finely-grained spatiotemporal analyses of neighborhood life than were previously possible. Gleaned from this source, neighborhood trajectories in physical-social disorder space can illuminate the connection between a neighborhood's appearance and the behavior of its residents over time, in their own eyes.

Second, these trajectories in physical-social disorder space are heterogeneous enough to warrant more support for critiques of the broken windows hypothesis. Different neighborhoods do not display the same relationship between physical and social disorder over time, and higher levels of physical disorder at one timepoint do not necessarily lead to higher levels of social disorder at future timepoints.

In the third set of results above, an implementation of the k-means clustering algorithm separates neighborhood trajectories into clusters. If the broken windows mechanism is indeed a product of social construction [Sampson and Raudenbush, 2004], examining the social characteristics within each cluster of trajectories might aid in our understanding of how different groups of people perceive and construct the notion of neighborhood disorder.

It seems that zip codes of residents who are older, less mobile, and own their own homes tend to display trajectories with the most volatile relationship between physical and social disorder as marked by 3-1-1 calls (cluster C). In this cluster, more upward movement in d_s relative to d_p is observed over the entire timespan, with both displaying higher values relative to other clusters. It is possible that these residents might feel more "ownership" over their neighborhoods, and thus might be more compelled to fix any problems in appearance or behavior that they might observe. This would be in accordance with the

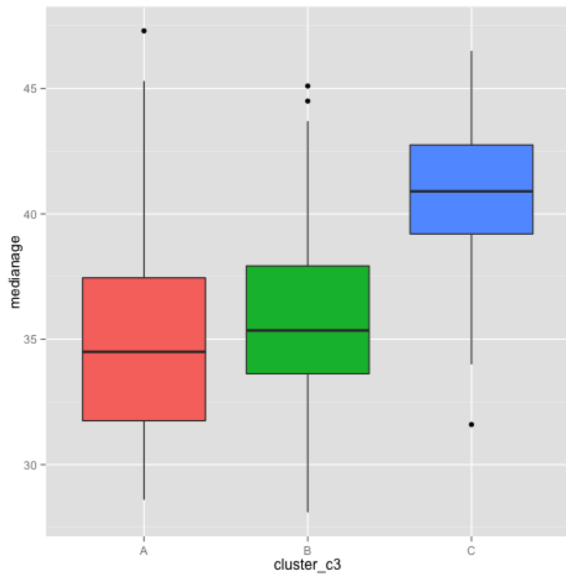

(a) Median age

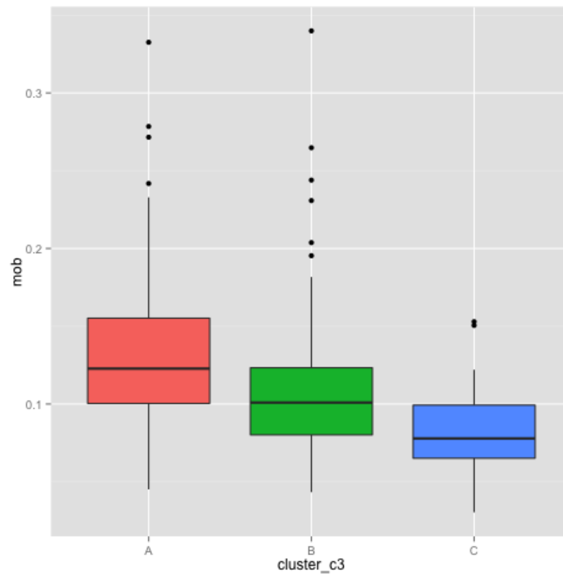

(b) Residential mobility

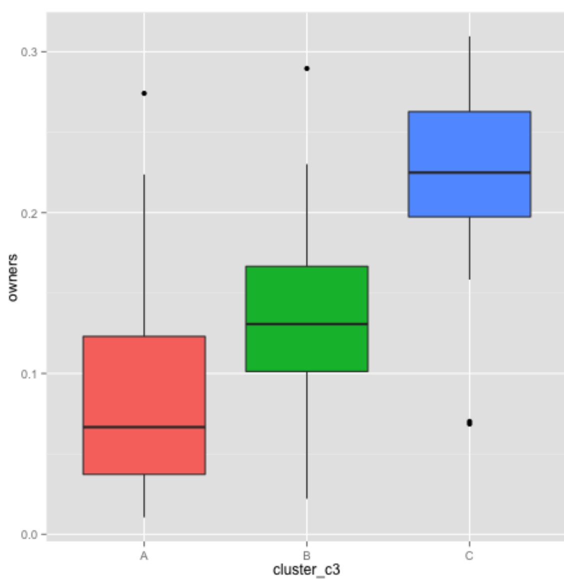

(c) Percent owners

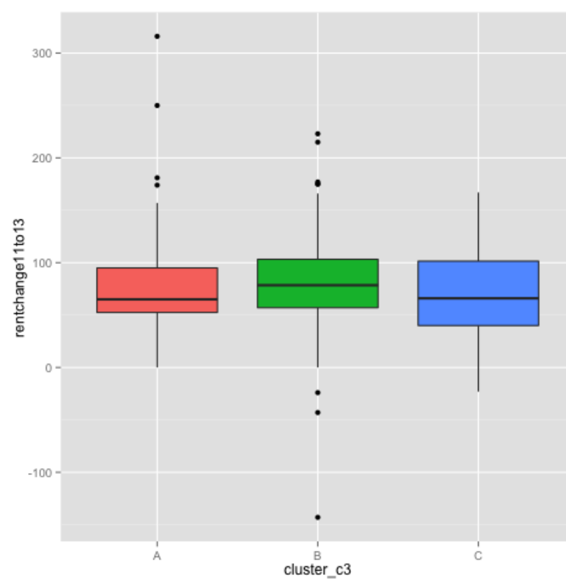

(d) Rent change ('11 - '13)

Figure 8: Social characteristics of neighborhoods, by trajectory cluster

civic participation literature [McCabe, 2013].

By contrast, zip codes of residents who are young, very mobile, and most likely renters display trajectories with steadily decreasing levels of both physical and social disorder, with values even stabilizing in recent years (cluster A). This perhaps presents the opposite case of cluster C; it is possible that these residents represent a more transitory population, who might not "see" disorder as a problem waiting to be fixed given their temporary (or *new*) status in their neighborhoods.

Perhaps the most interesting cluster of trajectories, however, is cluster B, which displays a marked increase in social disorder in more recent years, but an overall decrease in physical disorder over the entire timespan. Many (but not all) of these areas colored in green in Figure 7 are zip codes which are recently gentrifying. Rent change is typically used as an indicator for gentrification, but, as we can see in the boxplots in Figure 8d, there are no substantial differences in rent change between clusters. This could be due to my data limitations: unfortunately, the most recent data I could gather from the American Community Survey about rents were from 2013, which is before the timepoint after which the trajectories for cluster A and B become most different (B sees an uptick in social disorder post-2013, whereas A does not).

It is possible that access to more recent rent data would show gentrification more clearly, and gentrifying areas might be an interesting group in which to study different meanings of social and physical disorder: often, young gentrifyers look for the *appearance* of a "gritty" neighborhood [Zukin, 2009]; but what they are used to *socially* might be quite another story [Hwang and Sampson, 2014].

This study has several limitations. First, by the nature of the anonymized administrative dataset, I have no information about the individuals actually making the calls. So, at most, to avoid the classic ecological fallacy, I can examine neighborhood characteristics and discuss only *possibilities* of "who sees disorder" given aggregate rates by zip code. Second, previous literature theorizes that the "broken windows" process might occur at the block level [Sampson, 2012], but this project examines disorder levels by zip code only. Third, I do not have any details about potential changes made over time to 3-1-1 call options, and any "information" option changes would change the denominators in my measures. Finally, my analysis covers only New York City, and these processes might work in different ways in different cities, given varied physical and social landscapes.

Despite these limitations, this study suggests that the mechanism of physical-to-social disorder underlying the broken windows hypothesis indeed fails to hold up when examined on a dynamic level: not only are neighborhoods different from each other by perceived disorder at static snapshots, as critics have suggested, but they are also different from each other in the ways in which they *change* over time. Since the broken windows theory involves an evolution from physical to social disorder, and different groups of people define each of these in different ways, one must take into account how neighborhood populations themselves are composed and change over time. Such examinations can aid in our understanding of the social construction of disorder more generally. Future work might examine the dynamics of disorder in other cities as well.

Acknowledgements

I would like to thank the following colleagues and friends for their ideas and feedback over the course of this project: Robert Attorri, Clark Bernier, Sharon Cornelissen, Mitchell Duneier, Andrés Lajous, Matthew Salganik, David Schwartz, Tessa Stewart, Janet Vertesi, Andreas Wimmer, and Han Zhang.

All code for this project was written in `R` (`http://cran.r-project.org/`) with the use of the following packages: `data.table`, `dplyr`, `ggmap`, `ggplot2`, `grid`, `kml3d`, `maptools`, `RColorBrewer`, `reshape2`, `rgl`, `TTR`.

References

Accenture. Accenture Transforms Citizen Services in New York City. Technical report, 2003.

Tadeusz Caliski and Jerzy Harabasz. A dendrite method for cluster analysis. *Communications in Statistics-theory and Methods*, 3(1): 1–27, 1974.

Robert B. Cleveland, William S. Cleveland, Jean E. McRae, and Irma Terpenning. STL: A seasonal-trend decomposition procedure based on loess. *Journal of Official Statistics*, 6(1):3–73, 1990.

Avril Coghlan. A Little Book of R For Time Series. *Wellcome Trust Sanger Institute: Cambridge, UK*, 2014.

C. Genolini, J.B. Pingault, T. Driss, S. Ct, R.E. Tremblay, F. Vitaro, C. Arnaud, and B. Falissard. KmL3d: A non-parametric algorithm for clustering joint trajectories. *Computer Methods and Programs in Biomedicine*, 109(1):104–111, January 2013.

Bernard E. Harcourt. *Illusion of order: The false promise of broken windows policing*. Harvard University Press, 2009.

J. A. Hartigan and M. A. Wong. Algorithm AS 136: A K-Means Clustering Algorithm. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 28(1):100–108, January 1979.

John A. Hartigan. *Clustering algorithms*. Wiley, 1975.

Terrence D. Hill, Catherine E. Ross, and Ronald J. Angel. Neighborhood disorder, psychophysiological distress, and health. *Journal of Health and Social Behavior*, 46(2):170–186, 2005.

Jackelyn Hwang and Robert J. Sampson. Divergent Pathways of Gentrification Racial Inequality and the Social Order of Renewal in Chicago Neighborhoods. *American Sociological Review*, 79(4):726–751, 2014.

Jane Jacobs. *The death and life of great American cities*. Vintage Books, New York, 1961.

George L. Kelling and Catherine M. Coles. *Fixing Broken Windows: Restoring Order And Reducing Crime In Our Communities*. Free Press, New York, January 1996.

Carl A. Latkin and Aaron D. Curry. Stressful Neighborhoods and Depression: A Prospective Study of the Impact of Neighborhood Disorder. *Journal of Health and Social Behavior*, 44(1):34–44, March 2003.

Brian J. McCabe. Are Homeowners Better Citizens? Homeownership and Community Participation in the United States. *Social Forces*, 91(3):929–954, March 2013.

Norimitsu Onishi. Police Announce Crackdown On Quality-of-Life Offenses. *The New York Times*, March 1994.

Catherine E. Ross and John Mirowsky. Neighborhood Disadvantage, Disorder, and Health. *Journal of Health and Social Behavior*, 42(3): 258, September 2001.

Robert J. Sampson. When Things Aren't What They Seem: Context and Cognition in Appearance-Based Regulation. *Harv. L. Rev. F.*, 125:97, 2011.

Robert J. Sampson. *Great American city: Chicago and the enduring neighborhood effect.* University of Chicago Press, 2012.

Robert J. Sampson and Stephen W. Raudenbush. Systematic Social Observation of Public Spaces: A New Look at Disorder in Urban Neighborhoods. *American Journal of Sociology*, 105(3):603–651, November 1999.

Robert J. Sampson and Stephen W. Raudenbush. Seeing disorder: Neighborhood stigma and the social construction of broken windows. *Social psychology quarterly*, 67(4):319–342, 2004.

Robert J. Sampson, Stephen W. Raudenbush, and Felton Earls. Neighborhoods and Violent Crime: A Multilevel Study of Collective Efficacy. *Science*, 277(5328):918–924, August 1997.

Richard Sennett. *The Uses of Disorder: Personal Identity and City Life*. W. W. Norton & Company, New York, August 1992.

Wesley G. Skogan. *Disorder and decline: Crime and the spiral of decay in American neighborhoods.* Univ of California Press, 1990.

James Q. Wilson and George L. Kelling. Broken windows. *Atlantic monthly*, 249(3):29–38, 1982.

Sue-Ming Yang. Assessing the SpatialTemporal Relationship Between Disorder and Violence. *Journal of Quantitative Criminology*, 26(1):139–163, March 2010.

Philip G. Zimbardo. The human choice: Individuation, reason, and order versus deindividuation, impulse, and chaos. In *Nebraska symposium on motivation.* University of Nebraska press, 1969.

Sharon Zukin. *Naked city: The death and life of authentic urban places.* Oxford University Press, 2009.

	Complaint Type	Count	Physical or Social
1	Street Condition	401560	Physical
2	Street Light Condition	274069	Physical
3	Damaged Tree	149312	Physical
4	Traffic Signal Condition	132488	Physical
5	Rodent	111358	Physical
6	Missed Collection (All Materials)	86966	Physical
7	Graffiti	80762	Physical
8	Derelict Vehicle	64663	Physical
9	Derelict Vehicles	54763	Physical
10	Sidewalk Condition	33024	Physical
11	Air Quality	32703	Physical
12	Street Sign - Damaged	31297	Physical
13	Construction	21199	Physical
14	Street Sign - Missing	17951	Physical
15	Vacant Lot	8231	Physical
16	Street Sign - Dangling	7562	Physical
17	Public Toilet	198	Physical
18	Blocked Driveway	293592	Social
19	Illegal Parking	194178	Social
20	Noise - Commercial	125297	Social
21	Noise - Vehicle	67099	Social
22	Traffic	14415	Social
23	Noise - Park	13654	Social
24	Homeless Encampment	11270	Social
25	Smoking	8791	Social
26	Drinking	6128	Social
27	Disorderly Youth	3732	Social
28	Homeless Person Assistance	2479	Social
29	Urinating in Public	1895	Social
30	Violation of Park Rules	1777	Social

Table 1: 3-1-1 call classification scheme