# The *Cartography* of Opportunity

Elijah Knaap

July 19, 2016

**Abstract**

As evidence on the contextual effects of place upon individual outcomes has become increasingly solid over time, so too have urban policies and programs designed to connect underserved people with access to spatial opportunity. To this end, many attempts have been made to quantify the geography of opportunity and quite literally plot it on a map by combining evidence from studies on neighborhood effects with rapidly expanding spatial data resources and GIS technology. Recently, these opportunity maps have not only become increasingly common but their preparation has been encouraged and facilitated by the US Department of Housing and Urban Development (HUD). On one hand, the increasing prominence of opportunity mapping is a useful and important step forward for equity planning. Maps are powerful means of displaying the concept of opportunity and its variation across space. On the other hand, the institutionalization of opportunity mapping portends a need to examine critically the foundations that underlie the construction of opportunity metrics, the display of maps, and the application of these techniques in public policy. A closer look at the conceptual foundations and analytical methods that underlie these exercises offers important lessons not just for the practice of opportunity mapping but also for the practice of equity planning in general.

In the following essay, I examine the practice of opportunity mapping from both theoretical and methodological perspectives, highlighting several weaknesses of the common methods. Following, I outline an improved theoretical framework based on Galster's (2012) categorization of the mechanisms of neighborhood effects. Using data from the Baltimore metropolitan region, I then use confirmatory factor analysis to specify a measurement model that verifies the construct validity of the proposed theoretical framework. The model provides estimates of four latent variables that may be conceived as the essential dimensions of spatial opportunity: Social-Interactive, Environmental, Geographic, and Institutional. Finally, I develop a neighborhood typology by applying an unsupervised machine learning algorithm to the four dimensions of opportunity. The results suggest that the practice of opportunity mapping can be improved substantially through (1) a better connection to the empirical literature on neighborhood effects, (2) a multivariate statistical framework, and (3) more direct relevance to public policy interventions.

## Introduction

Two decades ago Galster and Killen coined the term "geography of opportuntity" in a seminal article published in this journal (George C. Galster and Killen 1995). Since then, the spatial pattern of factors that shape the structure of opportunity in metropolitan regions has been an important topic of public policy and lively subject of research. Conceptually, the notion of spatial opportunity is both simple and intuitive: neighborhoods, as unique packages of resources, institutions, and socializing agents, are likely to have a powerful influence on the welfare and life chances of their residents. In sociology, there is a long tradition of scholarship on spatial opportunity structures, the impacts of which are called neighborhood effects. In the last few decades, however, interest in these topics has burgeoned, and the geography of opportunity has attracted the interest of researchers from across the social and behavioral sciences. Today, the term geography of opportunity is widely used even in popular media, and closing the spatial opportunity gap is a well understood goal of sustainable urban development. What is less well understood, however, is how "opportunity" should be defined, how it should be measured, and how those measures can help close the opportunity gap through spatial policy interventions. Despite these ambiguities, achieving spatial equality of opportunity has become an acutely important agenda, particularly in the fields of housing policy, community development, and equity planning, and it has motivated calls for strategic neighborhood investments, residential mobility programs, or both.

To this end, many attempts have been made to visualize the geography of opportunity and quite literally plot it on a map. Recently, these *opportunity maps* have not only become increasingly common, but their preparation has been encouraged and facilitated by the US Department of Housing and Urban Development (HUD). In many respects, this is overwhelmingly positive, and the increasing prominence of opportunity mapping represents a useful and important step forward for equity planning. Maps are powerful vehicles for revealing spatial patterns in social structure. Like any data visualization technique, however, maps can be misleading or misinterpreted if their underlying assumptions are not met or critically examined. For these reasons, the institutionalization of opportunity mapping portends a need to examine critically the foundations that underlie the construction of opportunity maps, and their application in planning and public policy. A closer look at the conceptual foundations and analytical methods that underlie these exercises offers important lessons not just for the practice of opportunity mapping but also for the implementation of fair housing policy, regional planning, and equity planning in general.

In the following essay, I examine the practice of opportunity mapping from both theoretical and methodological perspectives, highlighting several weaknesses of the common methods. Following, I outline an improved theoretical framework based on Galster's (2012) categorization of the mechanisms of neighborhood effects. Using data from the Baltimore metropolitan region, I then use confirmatory factor analysis to specify a measurement model that verifies the construct validity of the proposed theoretical framework. The model provides estimates of four latent variables that may be conceived as the essential dimensions of spatial opportunity: Social-Interactive, Environmental, Geographic, and Institutional. Finally, I develop a typology of neighborhood opportunity by applying an unsupervised machine learning algorithm to the four dimensions of opportunity. The results suggest that the practice of opportunity mapping can be improved substantially through (1) a better connection to the empirical literature on neighborhood effects, (2) a multivariate statistical framework, and (3) more direct relevance to public policy interventions.

I begin with the premise that opportunity maps are intended to display the spatial variation in structures that influence a variety of socioeconomic outcomes, regardless of personal circumstance. This assertion has two important implications. First, opportunity mapping is distinct from vulnerability mapping. The former is concerned with identifying social, physical, and environmental attributes that affect a hypothetical household that resides in a particular neighborhood. The latter is concerned with identifying particular subpopulations who are at increased risk to external shocks due to their own precarious circumstance. A map of poverty concentration could serve either of these purposes: an impoverished family is at a greater risk of homelessness in the event of a major recession; an impoverished neighborhood imposes a negative externality on all neighborhood residents, whether they are personally below the poverty line or not. A map showing the concentration of homeowners holding high-cost loans is an example of a vulnerability map, *not* an opportunity map. Ceteris paribus, there is no reason to assume that a family living in a neighborhood will be affected by the share of its neighbors paying high interest rates on their mortgages. In practice, these two concepts are often conflated, leading to an erratic and unjustified collection of indicators intended to measure “opportunity”.

The second implication flows from the first: there should be some plausible connection that explains why the spatial structures measured affect socioeconomic outcomes. For example, a neighborhood with high rates of unemployment has less opportunity because the population in that neighborhood is less connected to the labor market and may have relatively little information about potential employment opportunities that may be passed on to job seekers in the neighborhood. For this reason, there is a natural connection between the practice of opportunity mapping and the literature on neighborhood effects–the body of research designed to uncover the causal effects of spatial structure on key socioeconomic outcomes.

Following this premise, I argue that traditional opportunity mapping exercises are often flawed in their theoretical and methodological conceptions. These flaws make opportunity maps difficult to interpret, and limit their utility in the formulation of housing policy. To address these issues, I outline an alternative methodology based on structural equation modeling, and I present an empirical example demonstrating the conceptual superiority over traditional methods.

Finally, I use a simple machine learning algorithm to show how neighborhoods may be classified into a typology useful for planners and policymakers, skirting the typical issues that arise in the presentation and display of opportunity maps.

## The Geography of Opportunity & The Mechanisms of Neighborhood Effects

Among social scientists it is now widely accepted that neighborhoods influence a wide variety of socioeconomic outcomes. In a recent and comprehensive review of the literature, Galster and Sharkey (forthcoming) identify compelling evidence of neighborhood effects on outcomes that include cognitive and behavioral development, educational performance and attainment, teen fertility, physical and mental health, labor force participation and earnings, and crime. Additional effects have been found related to outcomes as diverse as obesity (Diez Roux AV et al. 2006), violence (Sampson, Morenoff, and Raudenbush 2005), crime (J. R. Kling, Ludwig, and Katz 2005), high school graduation (Wodtke, Harding, and Elwert 2011), children's test scores (Burdick-Will et al. 2011), college attendance rates (Chetty, Hendren, and Katz 2015), earnings (Chetty, Hendren, and Katz 2015), intergenerational mobility (Chetty and Hendren 2015), mental health (Sanbonmatsu et al. 2011; Ludwig et al. 2012), physical activity (Kaczynski and Henderson 2007), cognition (Sharkey and Elwert 2011; Sampson, Sharkey, and Raudenbush 2008), infant health (Yang and Chou 2015), mortality (Anderson 2015), employment (G. Galster et al. 2015), and many others.

Some of these studies use similar explanatory variables, but many examine the effect of a particular key variable such as poverty, walkability, or ambient pollution. The challenge for researchers seeking to construct opportunity indices is to synthesize these results into a common formula that represents the most important neighborhood influences on the most important outcomes of interest. Overcoming such a challenge requires first the specification of a particular outcome (e.g. economic mobility), second the identification of appropriate causal neighborhood mechanisms that contribute to the outcome, and third the application of appropriate weights to each mechanism. In formal terms, this is equivalent to the equation given by G. Galster (2008), in which an outcome of interest $O$ observed at time $t$ for individual $i$ residing in neighborhood $j$ in metropolitan area $k$ can be expressed:

(1)

$$O_{it} = \alpha + \beta[P_{it}] + \gamma[P_i] + \varphi[UP_{it}] + \delta[UP_i] + \theta[N_{jt}] + \mu[M_{kt}] + \varepsilon$$

where

$O_{it}$ = employment status or income (model dependent) for individual $i$ at time $t$
$[P_t]$ = observed personal characteristics that can vary over time (e.g., marital or fertility status, educational attainment)
$[P]$ = observed personal characteristics that do not vary over time (e.g., year

and country of birth)
$[UP_t]$ = unobserved personal characteristics that can vary over time (e.g., psychological states, interpersonal networks and relationships)
$[N_t]$ = observed characteristics of neighborhood where individual resides during $t$
$[M_t]$ = observed characteristics of metropolitan area in which individual resides during $t$ (e.g., area unemployment rates)
$\varepsilon$ = a random error term
$i$ = individual
$j$ = neighborhood
$k$ = metropolitan area
$t$ = time period (typically a year)

This equation (henceforth Galster's equation) is a useful vehicle for describing the challenges and assumptions underlying opportunity mapping methodology and will be used throughout the paper. The challenge for opportunity mappers is, thus, to identify the vector(s) of spatial attributes (terms $M$ and $N$) that contribute to a variety of outcomes ($O$), and to develop a valid framework for applying weights to these indicators (i.e. the parameters $\theta$ and $\mu$). This is a tall order since there exists no study to date that permits the estimation of Galster's equation (and there is unlikely to ever be one) (G. Galster 2008). Empirically, it is also impossible to determine whether some outcomes ($O$'s) are more important than others (is cognition, for example, more important than employment?) and philosophically, these issues will draw diverse opinions. Furthermore, there may be important path-dependencies among outcomes; cognition, for example, is likely to impact educational performance, which is in turn likely to impact employment prospects.

### Identifying Outcomes and Indicators

With respect to the voluminous literature on neighborhood effects, several authors have offered opinions on which variables might best represent the vectors $N$ and $M$. Chetty and Hendren (2015), for instance, hold that "Low-income children are most likely to succeed in counties that have less **concentrated poverty, less income inequality, better schools, a larger share of two-parent families, and lower crime rates**. Boys' outcomes vary more across areas than girls, and boys have **especially poor outcomes in highly-segregated areas**". These sentiments are largely shared by Massey (2015), who argues that "The social scientific evidence thus yields several firm conclusions. First, the combination of **racial segregation, class segregation, and high rates of minority poverty** mechanically combine to produce neighborhoods of concentrated disadvantage. Second, exposure to concentrated disadvantage reduces human wellbeing along multiple dimensions, with powerful negative effects on health, cognition, education, employment, and earnings" (Massey 2015).

With respect to housing policy, the realm in which much of the research on spatial opportunity has been conducted, authors have argued that “both common sense and a growing body of research evidence teach us that living in a racially isolated, high poverty community undermines a family’s well-being and life chances, yet conversely, we know much less about how to define the”opportunity rich" neighborhoods to which we should be helping families move. We suggest that, instead of simple proxies, such as a neighborhood’s racial composition or poverty rate, destination neighborhoods should be targeted on the basis of concrete opportunities, such as **community safety, quality schools, or access to skill-appropriate jobs**" (Briggs and Turner 2006, 28). Others, meanwhile, have argued that “The neighborhood effects literature stresses that residential mobility may affect individuals by giving them access to better **community resources, schools, labor markets, and immediate neighbors, and moving them away from segregated enclaves and the negative influences in their prior neighborhoods**” (Rosenbaum and Zuberi 2010, 31) (emphasis added).

Synthesizing these results, G. Galster (2010) argues that while “the listings of potential mechanisms differ in labeling and categorizations, there is a broad consensus about how the underlying causal paths are thought to operate in theory. Unfortunately, there are few tentative conclusions, let alone any consensus, about which mechanisms demonstrate the strongest empirical support” (p.1). In other words, the empirical record provides guidance on the vectors that comprise $N$ and $M$ (or at least the mechanisms underlying them), but little to no guidance on the appropriate magnitudes of each $\theta$ and $\mu$. Continuing his comprehensive review, Galster finds empirical evidence for 15 independent causal pathways through which neighborhoods affect different socioeconomic outcomes, which he organizes into four categories: social-interactive, environmental, geographic, and institutional. These categories are largely identical to those outlined by Sampson, Morenoff, and Gannon-Rowley (2002), except that Galster includes an environmental category whereas Sampson et al differentiate two types of social-interactive categories; this difference merely reflects the fact that Sampson et al review only the sociological literature while Galster also incorporates epidemiological perspectives. The categories are also similar to those outlined by Ellen and Turner (1997), and Leventhal and Brooks-Gunn (2000). Given the strength with which these categories are represented in the literature, it is natural that they should be reflected well in the opportunity mapping methodology. In practice, however, this is rarely the case.

## The Practice of Opportunity Mapping

The practice of opportunity mapping has several intellectual roots. As a technical exercise, opportunity mapping builds on techniques developed for suitability analysis by Ian McHarg (Collins, Steiner, and Rushman 2001). Ostensibly, opportunity mapping involves the identification of areas well suited to promote social mobility by combining GIS layers of various social, economic, and environmental variables. More conceptually, the practice builds on equity mapping developed by Toulmin (1988) and Truelove (1993) and applied by Emily Talen (1998) and E. Talen and Anselin (1998). This body of research defines equity in terms of proximity or access to various public facilities or neighborhood attributes. The current practice of opportunity mapping, however, draws from john powell's opportunity-based housing model, and was developed in the context of fair housing litigation (john a. Powell 2003). Specifically, in the case of *Thompson v HUD,* john a.powell testified as follows (J. Powell 2005):

> The segregation of African American public housing residents isolates them from the opportunities that are critical to quality of life, health, stability, and social advancement. The safe and stable neighborhoods, successful schools and employment opportunities generally available to Whites in the greater Baltimore region have been denied to African American public housing residents in the City of Baltimore. To remedy this segregation two objectives must be met: 1) the remedy must give African American public housing residents the opportunity to live in racially integrated areas in the Baltimore region and 2) the remedy must affirmatively connect African American public housing residents to high opportunity neighborhoods in the Baltimore region.

powell then introduced opportunity maps that included measures in three categories–economic opportunity and mobility, educational opportunity, and neighborhood health—-aggregated into an overall opportunity index, and showed that minorities and public housing developments were (and are) disproportionately concentrated in low opportunity areas.

powell's testimony was convincing and led to a ruling in favor of the plaintiff class and the creation of a regional housing mobility program designed explicitly to help connect families with high-opportunity neighborhoods. In recent years, opportunity mapping exercises have been conducted in metropolitan areas across the nation including Seattle, Austin, Minneapolis, Chicago, Baltimore, Boston, and many other places, and these maps have moved well beyond the realm of fair housing litigation into much broader usage including the development of regional housing, transportation, and economic development policies. Much of the work has been conducted by the Kirwan Institute, whose process has become somewhat standardized: (1) Select variables that measure the presence or lack of opportunity, (2) Collect data and assign values to common geographic units, (3) Normalize the data and assign to subcategories, (4) Compute a composite opportunity index, (5) Create thematic maps, (6) Overlay with other variables of interest (Reece and Gambhir 2008). This methodology appears logical and straightforward upon first inspection, however, it also involves a number of subjective decisions and computational tasks at each step in the process that significantly influence the results of the analysis yet are neither discussed in the literature nor widely understood.

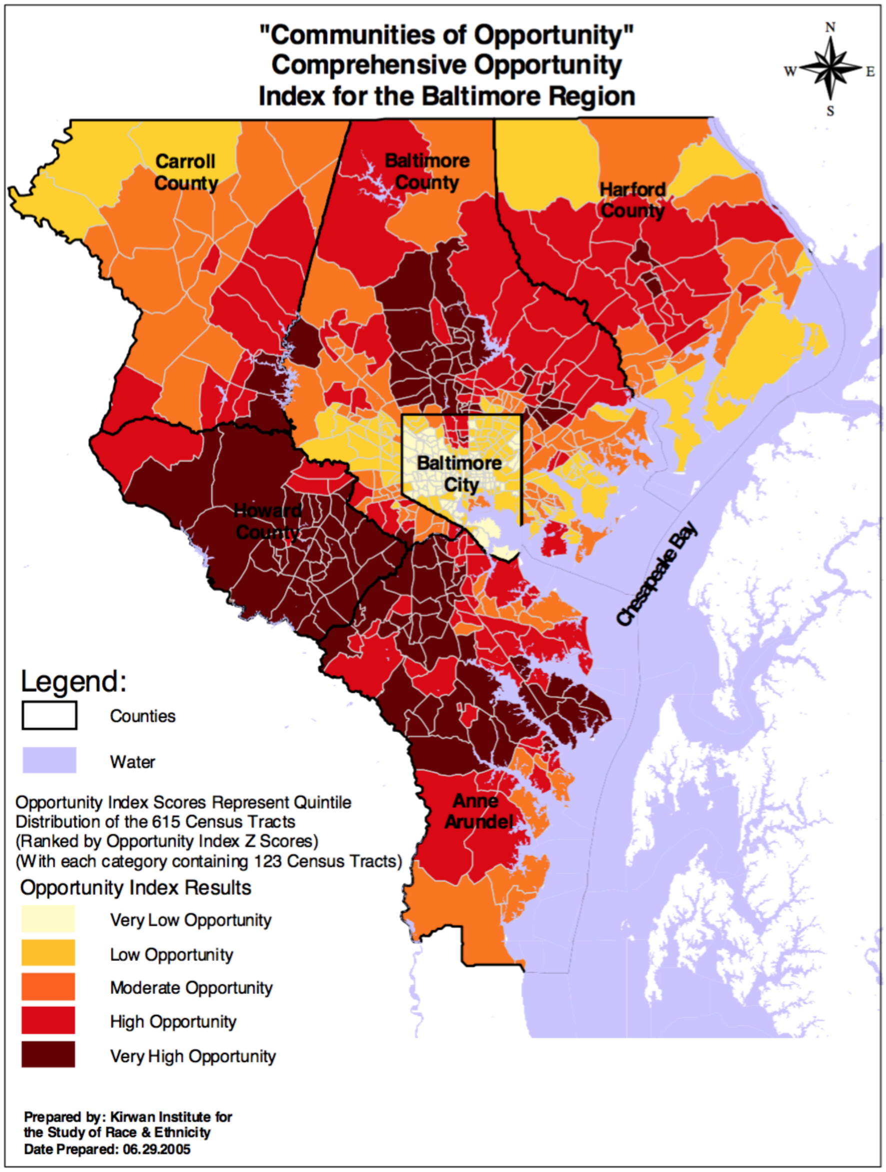


Figure 1: Thompson v. HUD Opportunity Map

## Critique

Despite the widespread and increasingly common use of opportunity maps in the development of urban policy, there exists no published discussion among researchers or practitioners that examines the utility and soundness of the technique. This is an important omission in the literature. According to Giovannini et al. (2005, 14), "composite indicators are much like mathematical or computational models. As such, their construction owes more to the craftsmanship of the modeller than to universally accepted scientific rules for encoding. With regard to models, the justification for a composite indicator lies in its fitness for the intended purpose and in peer acceptance". Absent the discussion and peer acceptance described by Giovannini et al. (2005), the opportunity metrics in common use should be treated with skepticism. Indeed, upon review of the data and methods used to develop most opportunity maps, it is clear that a number of critical flaws exist that limit the utility of opportunity mapping in its current form.

### Indicator Selection and Categorization

The list below represents the indicators and categories used most commonly in these exercises (Kirwan Institute for the Study of Race and Ethnicity 2013):

I. Education Indicators

- Adult education attainment
- Promotion Rates
- Graduation rates
- School Proficiency Index
- Student poverty rates
- Student teacher ratio
- High Quality Teachers

II. Economic Indicators

- Economic climate (change in number of jobs)
- Employment competition (ratio of jobs to labor force within a certain miles)
- Proximity to employment
- Job growth trends
- Population on public assistance
- Unemployment rate

III. Housing and Neighborhood Indicators

- Affordable housing
- Foreclosure rate
- High-cost loan rate
- Housing cost burden

- Home ownership
- Housing vacancy
- Mortgage denials
- Population change 1990-2000
- Poverty rates
- Property appreciation and tax base
- Property values
- Sub-prime loans
- Subsidized housing

IV. Transportation and Mobility Indicators

- Access to automobile
- Mean commute time
- Public transit access
- Transit cost
- Transit dependency
- Transportation cost
- Walkability

V. Health and Environmental Indicators

- Amount of toxic waste release
- Crime index
- Grocery stores
- Parks and open space
- Proximity to toxic waste release sites

Each of the variables identified above is useful for understanding the spatial distribution of inequality. Their categorization and aggregation, however, calls into question the construct validity of each purported subindex: are education, economy, housing and neighborhood, transportation and mobility, and health and environment truly the subdimensions of opportunity? Are the indicators grouped into those categories valid measures of those subdimensions? And is it justified to assume each indicator and each category contributes equally to the geography of opportunity? The answer to all of these questions is "probably not".

What is striking about the variables and categories outline above is the complete omission of a category pertaining to social structure. Indeed, this is even more egregious as several authors have argued that "if neighborhood effects on child outcomes exist, presumably they are constituted from social processes that involve collective aspects of community life" (Sampson, Morenoff, and Earls 1999, 634; Mayer and Jencks 1989). The empirical record appears to confirm this view, and there is strong evidence that at least *some* aspect of social interactive mechanisms contributes to a variety of socioeconomic outcomes (George C Galster 2012). In some cases, social variables are misappropriated into other categories; adult educational attainment, for instance, is often categorized into an "educational" group, which ostensibly measures the educational opportunities in a neighborhood (Reece and Gambhir 2008; Kirwan Institute for the Study of Race and Ethnicity 2013). Conceptually, however, educational opportunities are a dimension of institutions–not people–and adult educational attainment does not measure institutional capacity. Even as a proxy, educational attainment is a poor indicator of school quality; in Baltimore, for example, there are several neighborhoods near the City's inner harbor in which the population is young, affluent, and well educated–and these neighborhoods are served by some of the lowest-performing schools in the state of Maryland.

Other times, social variables such as poverty are grouped into an ambiguous "neighborhood quality" category which also contains variables related to the housing stock, such as vacancy rates and home values (Reece and Gambhir 2008). Because these variables measure multiple unrelated phenomena, it is wholly unclear what this category measures, how it should be interpreted, and how it provides useful information to policymakers. This does not suggest that any of the indicators in this category are necessarily misguided in their own right, but that their combination into a single metric has no theoretical justification and no direct utility.

Instead of focusing on social indicators, opportunity mapping exercises seem to place undue importance on the location of jobs–an ode to the pervasive notion of spatial mismatch (John F. Kain 1968; John F Kain 1992). This is despite the fact that "there is considerable statistical evidence that this spatial mismatch is of less importance to economic outcomes than the social-interactive dimensions of neighborhoods" (G. Galster 2010, 14). The employment category also reveals another problem in that indicators are often grouped into their topical domain rather than their underlying data-generating process. This strategy leads to unemployment rates and access to jobs grouped into the same category, which is problematic because the *pathways* through which unemployment and access to jobs affect socioeconomic mobility are entirely distinct. Furthermore, in many cases, unemployment and job accessibility are almost perfectly correlated in the opposite direction: central cities have large concentrations of unemployed people *and* the best access to jobs (Chapple 2014). Averaging these two indicators together results only in noise.

### Normalization, Weighting, and Aggregation

Returning to Galster's equation, it is natural to frame opportunity mapping through the lens of a linear regression equation; first an outcome variable (opportunity) is defined, then explanatory variables are identified and combined to yield an estimate of the outcome. This regression metaphor is useful for discussing the implicit assumptions that underlie the current practice of opportunity mapping and the potential drawbacks they introduce. I argue that the most important assumptions that should be treated with caution are the identification of coefficients, possible non-linearity of effects, and interactions among variables.

While there appears to be some agreement among researchers about the types of neighborhood structures that contribute to spatial inequality, there is little guidance from the literature regarding *how much* each structure contributes to the pattern. Since the mechanisms that drive neighborhood effects remain very much elusive, it is extremely difficult to quantify the relative contribution of each potential source (G. Galster 2010; G. Galster 2008). For this reason, explanatory variables in an opportunity index are treated as equal, independent, and linear in effect. In the absence of any alternative theoretical framework, these are reasonable simplifying assumptions. In many cases, however, these assumptions may be untrue, which could lead to misinterpretation of opportunity metrics and faulty policy prescriptions.

Indeed, in some cases, there is already evidence that these assumptions are violated. G. Galster, Quercia, and Cortes (2000), for instance, have argued that "when a neighborhood reaches a critical value of a certain indicator, it may trigger more rapid changes in that neighborhood's environment," and that neighborhood poverty rate is one such indicator that has a non-linear relationship with other quality of life indicators such as unemployment and vacancy rate. In the face of this evidence, it is questionable whether poverty (which follows a log-linear distribution) should be treated as a linear variable in opportunity indices or whether its inclusion warrants some other transformation. This consideration is likely to apply for other variables as well. In a similar vain, it is reasonable to assume that each of the opportunity indicators may have a more nuanced relationship with a particular outcome. It would seem unlikely that each of the variables identified above have the exact same significance and magnitude associated with a particular outcome. In other words, in the context of Galster's equation, it is unreasonable to assume that each $\theta$ and $\mu$ are equal and invariant across opportunity indicators.

One possible solution is to develop weights for each indicator using alternative methods that could include surveys/crowdsourcing or an alternative statistical model. With respect to the former, researchers must rely on stated preferences about the types of opportunities that matter and the types of resources people believe contribute to that particular type of opportunity. This method introduces bias relative to the composition of the sample, their knowledge, and their desires. With respect to the latter, researchers must assume that an alternative model carries the appropriate information they wish to distill. One potentially valuable method is the use of decomposition techniques like factor analysis and principal components analysis, which project collections of correlated variables onto a smaller subset of representative factors or components. This technique helps solve the problem of weighting individual indicators within categories but does not overcome the issue of weighting different dimensions of opportunity relative to one another. In other words, factor analysis can help estimate a reduced set of variables that represent meaningful sub-dimensions of opportunity, but does not provide a framework for understanding how those estimated variables should be aggregated into a single univariate metric.

Finally, the research discussed above highlights at least three important heterogeneities with respect to neighborhood effects on particular subpopulations: race, gender, and age (Sharkey and Faber 2014). Younger people and minorities are more prone to neighborhood effects in general; girls experience a larger effect on their mental health, and boys experience a larger effect on their education, employment, and criminality (Chetty and Hendren 2015; G. Galster et al. 2015; Sanbonmatsu et al. 2011). Returning to Galster's equation, these findings imply that the $P$ terms are not only significant, but that depending on the values they take, also change the values of $\theta$ and $\mu$. In the face of these findings, the design of a universal index of opportunity may be problematic. If there is heterogeneity in the way that neighborhood effects are experienced, then there may need to be heterogeneity in the types of metrics that are collected and combined. The typology approach described in later sections addresses this issue by introducing flexibility in the way that opportunity metrics are combined and displayed.

### Geographic Assignment

Apart from the difficulty of conceptualizing measures of opportunity, a number of technical challenges remain concerning the quantification and representation of space, particularly when different data sources refer to different underlying data models (e.g. point, line, polygon, raster). When computing a spatial opportunity index an analyst must combine multiple, unrelated data, nearly all of which are measured along different scales and units. To overcome this issue, data must be standardized according to a common geographical unit that is representative of a neighborhood, and converted to a consistent measure. In practice, this nearly always involves collecting data by census tract, and standardizing the data to z-scores. Census tracts are usually chosen as the geographic unit of analysis because they offer the finest geographic precision for which data are widely available.

The drawbacks associated with using census tracts as a proxy for neighborhoods are many and have been well articulated by others (Sampson, Morenoff, and Gannon-Rowley 2002; Walter and Wang 2016). Census tracts vary in size and may or may not correspond well with a resident's notion of neighborhood. Second, using only the data contained within a single census tract ignores the possibility of spatial spillover and the influence of adjacent or outlying neighborhoods (Sampson, Morenoff, and Earls 1999; Dietz 2002; Sampson 2012; L. Anselin 2013). Third, data that are important but unavailable at the census tract level, like those defined by an alternative geography, (e.g. zip-codes), or disaggregate microdata (i.e. point locations) require additional transformation, introducing an additional source of error. With respect to polygons-to-polygon conversions, this error represents the well known "modifiable areal unit problem" (MAUP). With respect to point-to-polygon conversions, there are many possible sources of error depending on how the conversion is performed.

The simplest way to deal with point data, is to simply aggregate all the points within a single tract, and assume that locations with multiple locations have better access. Indeed, this is a common technique in opportunity mapping (Kirwan Institute for the Study of Race and Ethnicity 2013) This is an undesirable method, however, for a number of reasons. First, many tracts would not have a score at all simply because they did not contain any points. This is problematic because tracts could still have good access to one or more points even if those points do not fall within the tract itself. Second, this method fails to account for the difference in tract size, and will provide a bias toward larger tracts.

Another technique for transforming point locations into tract scores is to calculate a measure of kernel density. Kernel density estimators are common among statistical and GIS and software packages, and are typically applied when performing cluster analyses (e.g. the study of crime). In the context of spatial analysis, kernel estimators split a study region into a raster grid, then for each grid cell, search for any points that fall within a specified distance of the originating cell. Points within the search radius are then applied to a distance decay function so that nearby points have a larger effect than those further away. The result is a crude estimate of spatial exposure; since the score given to each cell in the raster is a distance-weighted sum to a given resource, it can be conceptualized as Euclidian-based accessibility measure. Grid scores can then be averaged for each census tract to provide an overall measure. This method is better than simply aggregating points to tracts because it incorporates the influence of points that lie outside the tract boundaries. It is also preferable to simple buffer-based methods (e.g. aggregating all points that lie within a 1-mile radius of the tract) because it treats space as continuous rather than discrete, and does not assume a constant effect within the buffer. Despite these strength, kernel density methods also suffers some drawbacks. For one, both the size of the grid cells and the search radius must be specified by the analyst, which requires some theoretical foundation for decision making. Second, although a kernel density calculation can be used to estimate access or exposure to a resource, it ignores the impacts of infrastructure connectivity and travel times. Thus, it may be a good choice for estimating the impact of resources that are unrestricted by transport networks (e.g. pollution radiating from a smokestack) but is less desirable for measuring access to amenities like jobs.

A third option is to incorporate locational amenities into a transportation model. For many point-based indicators this may be the most desirable method because it provides the most reliable measure of accessibility, accounting for transportation infrastructure and commute times. This method also provides the potential benefit of disaggregating among modes of travel, for instance, to compute measures of accessibility by transit, which can add an additional layer of sophistication. Since many opportunity analyses are conducted at the metropolitan region, and many metropolitan planning organizations (MPOs) develop travel models, incorporating travel models into opportunity analyses seems like a natural fit. For many spatial variables of interest, this method should be chosen above others when there is an available travel model, though it is not without some disadvantages. One drawback worth noting is that travel models frequently rely on transportation analysis zones (TAZs) or other alternative geographic units like statewide modeling zones (SMZs). Rarely do these types of models incorporate census tracts, so the conversion of model zones to census tracts can introduce error (i.e. the modifiable areal unit problem). Furthermore, accessibility-based measures still require theoretical specification on the analyst's part with respect to the search radius and the travel cost function used to parameterize the measure.

What is clear from this discussion is not that some methods are inherently more preferable than others, but that a sound rationale is required for a series of subjective decisions and that developing indicators of opportunity is less a formulaic exercise than an iterative one in which many alternatives are tested with respect to their theoretical soundness and analytical performance.

### Visualization

In the common practice of opportunity mapping, each spatial indicator is z-transformed, then aggregated into a scale, which typically ranges from around -3 to 3. For the purposes of visualization, this scale is then collapsed into a five-level index with each quintile representing a different echelon of opportunity. From a visual perspective, this is a convenient transformation; maps based on quantiles produce equal shares of each color and typically produce identifiable patterns and appealing aesthetics. From an analytical perspective, however, maps based on quantiles of any number can be misleading because they intentionally break data into a uniform distribution. Each quantile contains the same number of observations as every other, regardless of the shape of the underlying data's distribution. This means that extreme outliers may be classified into the same quantile as observations that are relatively common in the data but happen to fall near the top or bottom of the distribution. Furthermore, when data are highly clustered, quantile breaks may create artificial differentiation among observations that have similar or even the exact same values. This creates the visual distortion that meaningful differences exist in the data when, in truth, none exist at all. Conventionally, opportunity maps are displayed using five quintiles, which comports nicely to a familiar likert-like interpretation, but it is not clear why the data should be divided into five quintiles rather than three or seven or ten divisions.

Again, this is an area with few hard guidelines. As Christopher Ingraham explained in a short article for the Washington Post's Wonkblog, "Visualizing data is as much an art as a science. And seemingly tiny design decisions–where to set a color threshold, how many thresholds to set, etc.–can radically alter how numbers are displayed and perceived by readers" (Ingraham 2016). Rather than force these choices upon an analyst, one alternative is to use a clustering algorithm to divide neighborhoods into discrete categories based on the structure of their underlying data. These techniques help remove subjectivity from the analysis, using a data-driven approach that defines clusters of neighborhoods by maximizing the variation between clusters and minimizing the variation within them. Such an approach is discussed in later sections.

## A Measurement Model of Opportunity Structure

Given the conceptual and technical measurement issues outlined above, it is clear that opportunity is a difficult concept to operationalize, let alone measure and visualize. One way to address this problem is to treat the quantification of opportunity as a measurement error problem. Through a liberal interpretation, this may be viewed as an extension of ecometrics, a methodology concerned with developing measures of neighborhood social ecology (Raudenbush and Sampson 1999; Mujahid et al. 2007; O'Brien, Sampson, and Winship 2013). In this framework, opportunity and its subdimensions are viewed as latent variables that cannot be measured directly, but can be estimated by modeling the covariation among the indicators through which they manifest. As with any measurement model, however, opportunity metrics require a sound theoretical framework for organizing and specifying relationships among variables. As described above, a major weakness of opportunity analyses to date has been the lack of a sound framework for organizing indicators into categories of metrics. To address this issue, I argue that the literature on neighborhood effects offers a sound organizing framework for classifying subdimensions of opportunity. Specifically, I propose that neighborhood indicators should be categorized according to the four mechanisms of neighborhood effects outlined by George C Galster (2012): social-interactive, environmental, geographic, and institutional. These categories are well supported by the empirical literature, and are theoretically grounded in causal processes that generate socioeconomic outcomes. Following the outline of this framework, I use confirmatory factor analysis to verify the construct validity of the proposed approach by showing that indicators load on the theoretically-defined factors in the expected patterns and that the measurement model is valid with respect to overall model fit indices. This gives both theoretical and empirical justification to the notion that the selected indicators measure what they purport to measure, and that four resulting metrics are reasonable estimates of each theorized dimension.

Both exploratory factor analysis (EFA) and principal components analysis (PCA) are becoming increasingly common in urban research although neither approach is truly widespread. Ewing et al. (2003), for example, uses PCA to develop a measure of urban sprawl, and Nosoohi and Zeinal-Hamadani (2011) use EFA to study measures of welfare and development in Iran. In the context of opportunity mapping, other researchers have advocated for similar data-reduction techniques in the computation of opportunity metrics. Walter and Wang (2016), for example, suggest the use of geographically weighted principal components analysis (GWPCA). They do not, however describe their theoretical framework for categorizing indicators into groups, nor do they provide any argument as to why their measurements are causally related to important socioeconomic outcomes. Indeed, their only rationale for choosing indicators is "accounting for and considering the variables used in all previous studies" (Walter and Wang 2016). Thus, although they provide suggestions for enhancing the methodological rigor of opportunity analyses, they do not demonstrate that their approach has any greater level of construct validity than traditional approaches. In other words, they do not report whether their indicators load on components in the expected fashion and whether those components capture enough variation to be viewed as valid composite indicators. In many cases, this seems unlikely, at least when applied to a diverse range of metropolitan areas. Access to primary care physicians and the percentage of area coverage by parks and green space, for example, which are classified under the "Healthy Environment" category are unlikely to load on a single factor, with physician access likely biased towards urban areas and green space likely biased toward suburban and exurban areas. When these indicators do not load strongly on a single component, it is unclear how the results of GWPCA should be interpreted and is unlikely that the first component may serve as a composite metric for the category. Furthermore, Walter and Wang (2016) fall victim to the same categorization issues that plague traditional approaches by conflating institutional measures like school proficiency and early childhood neighborhood participation alongside social-interactive measures like poverty and labor market engagement. Although these indicators may be highly correlated and may well load on a single component, they are quite distinct from a conceptual perspective.

By contrast, confirmatory factor analysis (CFA) approaches based on structural equation modeling are comparatively rare in urban research. One notable exception is Bagley and Mokhtarian (2001), who use confirmatory factor analysis to study the effects of neighborhood conditions on travel behavior. CFA differs significantly from EFA and PCA approaches because it requires the specification of an a priori theoretical model. For this reason, CFA is often used for the purposes of construct validity because it demonstrates that observed data conform to the particular theory imposed by the researcher. In the context of opportunity mapping, this makes CFA a particularly attractive approach because it facilitates the evaluation of the particular theoretical frame used to devise opportunity metrics. Another benefit of the CFA approach is that indicators can be specified to load on distinct but correlated latent variables (e.g. assigning educational attainment to an institutional dimension rather than a social dimension) rather than simply mining the data for common covariance, as is the case with EFA. The preceding section describes the theoretical framework and the data sources uses to construct such a model.

## Data

### Institutional

Institutional variables are designed to capture "actions by those typically not residing in the given neighborhood who control important institutional resources located there and/or points of interface between neighborhood residents and vital markets" (George C Galster 2012). Although institutional data comprises more than educational systems, schools are a primary vehicle through which institutional capital and institutional opportunity is transmitted. Furthermore, schools are perhaps the only type of institution that facilitates evaluation of *quality* in some form. For this reason, variables in the institutional category are designed to capture multiple dimensions of school quality and include the share of highly qualified teachers[1], the share of Students achieving a passing grade on state administered exams, performance on Advanced Placement exams, SAT scores, and high school dropout rates. These data are collected from the Maryland State Department of Education which provides annual statistics for each school in the state of Maryland.

Maryland State Assessments (MSA), administered every year to students in grades 3 through 8, and High School Assessments (HSAs) are administered in grades 9-11. For this analysis, subject scores for each school are averaged into an overall measure of students who score passing grades on the exams. Individual school measures are then assigned to census tracts by collecting catchment areas from each jurisdiction in the Baltimore region, matching each school with its catchment area, then geocoding each tract to the applicable catchment area. The best fitting model is achieved by incorporating a nested structure in which separate factors are estimated for high school and elementary school. These two factors are combined with two measures of middle school quality (highly qualified teachers and standardized test achievement) to yield the overall institutional factor[2].

### Geographic

Geographic variables are designed to capture aspects of neighborhoods that affect residents' life-courses "purely because of the neighborhood's location relative to larger-scale political and economic forces," and encompass the subdimensions of spatial mismatch and public services (George C Galster 2012). Operationally, these are dimensions of the built environment, particularly accessibility to necessary goods and services. For the purpose of this analysis, geographic variables include jobs accessible by walk and transit, access to healthcare facilities, access to public institutions, and access to social organizations.

Job location data is collected from the U.S. Census Longitudinal Employment-Household Dynamics (LEHD) via the LEHD Origin-Destination Employment Statistics (LODES) database. The LODES database records the location of jobs,

[1] To be classified as a "highly qualified" teacher, MSDE requires that instructors in core academic subject areas must: Hold at least a bachelor's degree from a regionally accredited institution of higher education (IHE); hold a valid Standard Professional Certificate or Advanced Professional Certificate or Resident Teacher Certificate in the subject area they are teaching; and satisfy additional requirements associated with specific teaching levels and experience. For additional information, see http://www.marylandpublicschools.org/msde/programs/esea/docs/TQ_Regulations/general_definition.htm

[2] a middle school factor is not estimated because it only includes two distinct data points

annually, at the census block level. Although the current analysis is limited to the Baltimore Metropolitan region, it is necessary to collect and analyze LODES data for the entire state of Maryland, as well as portions of Washington D.C., Virginia, West Virginia, Pennsylvania, and Delaware to reflect the fact that, while families must reside in the Baltimore region, they still have access to jobs in nearby counties and states. I use total jobs accessible rather than low/mid-skill jobs because, in Maryland, there is very little spatial differentiation among jobs of differing skill levels.

Healthcare facility data is collected from the 2013 Maryland Quarterly Census of Employment and Wages (QCEW), a database of employment records updated quarterly by the Maryland Department of Labor, Licensing, and Regulation (DLLR). The QCEW database contains information on the location, employment levels, and industrial classification for each employer in the state. Social organizations (e.g. philanthropies and non-profits), public institutions (like schools and universities) and applicable healthcare facilities were identified from the QCEW data. Healthcare facilities were selected using the following NAICS codes:

- 6211: Offices of Physicians
- 6214: Outpatient Care Centers
- 6219: Other Ambulatory Healthcare Services
- 6221: General Medical and Surgical Hospitals

Because the key variables in the geographic category are *destinations*, (i.e. locations that require travel to and from a household's residence) computing access to these destinations requires data on transportation infrastructure that facilitates such travel (i.e. pedestrian, transit, and automobile networks). Data for walk networks is collected from OpenStreetMap (OSM), a worldwide, open source repository for spatial data. In many urban areas, OSM is the most comprehensive and up-to-date source of roadway, bikeway, and pedestrian infrastructure. Data for transit networks is provided by the Maryland Transit Administration and the Central Maryland Regional Transit Agency in the form of General Transit Feed Specifications (GTFS). GTFS data is published by these agencies on a regular basis and includes information on bus and rail transit stops, frequencies, and schedules.

Accessibility (sometimes called "cumulative opportunity") measures the ease with which a person can consume a resource located in space. Typically, accessibility is used to measure how many jobs are available to a person living in a particular residential area, though it can also be applied to any type of origin or destination. There are many measures of accessibility, most of which are variations on the seminal work of Hansen (1959). In simple terms, accessibility for a given location can be calculated as the weighted sum of all activities that can be reached within some specified cost. Formally, accessibility $A_i$ can be expressed by

(2)

$$A_i(C) = \sum_j a_j f(c_{ij})$$

Where:
$a_j$ is the quantity of some resource at location $j$ obtainable within a generalized cost parameter $C$
$c_{ij}$ is the generalized cost of travel between origin $i$ and destination $j$
$f(c_{ij})$ is an impedance function that quantifies the disutility of the travel

In regional science and urban economics literature, $f(c_{ij})$ is typically a linear or exponential decay function, yielding what is often called a gravity model. Gravity equations of this type are commonly used in econometric location choice models to help explain why certain sets of amenities (like jobs) "pull" households into locating in a certain area.

In this analysis, the cost parameter $C$ is fixed at a threshold of 60 minutes, meaning that only activities that can be reached within a 60 minute commute will be included in the measure. The cost of travel $c_{ij}$ (i.e. time) varies by mode, with automobiles often covering longer distances than walk/transit in shorter periods of time, except in cases of extreme congestion. For job accessibility, the impedance function is fixed at 1, meaning that the accessibility measure represents simply the total sum of jobs that can be reached within 60 minutes. For healthcare and occupational training a linear decay function is applied that weights nearer activities higher those which are further away. Automobile accessibility is computed via the MSTM; walk accessibility is computed via the Pandana software library for the python programming language[3]; transit accessibility is computed using the TransportAnalyst software platform[4].

### Environmental

Environmental variables are designed to capture the "natural and human-made attributes of the local space that may affect directly the mental and/or physical health of residents without affecting their behaviors" (George C Galster 2012). Unfortunately, there are relatively few data sources that can provide such variables at the necessary scale, so this category includes only three: a crime victimization index (which captures the social environment), proximity to designated toxic release sites (which captures the ambient physical environment), and the share of vacant housing units (a measure of physical disorder). Additional variables that could be incorporated into future analyses might include lead contamination in the water, access to parks and open space, or signs of physical disorder collected through systematic social observation (Bader et al. 2015; O'Brien, Sampson, and Winship 2013).

In lieu of actual crime statistics, which are unavailable for the region (except Baltimore City) a Crime risk index is used. The crime risk index is developed by Applied Geographic Solutions[5], and uses data from the FBI Uniform Crime Reports (available at the county level) combined with data from local jurisdictions

[3] http://udst.github.io/pandana/
[4] https://github.com/conveyal/analyst-server
[5] http://www.appliedgeographic.com/MethodologyStatements_2015/CrimeRisk2015A.pdf

to model crime risk down to the census tract level. To be sure, this is a weakness of the current analysis. Reported crime would be a better measure of a negative environmental externality, and modeled data are likely to be confounded by the sociodemographic data with which it is estimated.

Toxic release sites are collected from the Environmental Protection Agency's (EPA) Toxic Release Inventory Program[6], which publishes an annual database of toxic chemical and pollution emitting locations. To quantify the toxic effect of environmental pollution, it is necessary to estimate exposure to such pollutants. Although data provides the point sources of pollution emitting locations, the myriad factors that may influence particulate dispersal (e.g. wind speeds, weather patterns, atmospheric pressure, etc) estimating the precise level of exposure to pollutants is challenging. Following some simple assumptions, however, it is possible to estimate a reasonable approximation using standard spatial analysis procedures.

Kernel Density Estimation (KDE) is a statistical technique used to produce a smooth density surface of point events over space (Xie and Yan 2008). In spatial analyses, kernel density estimators are used commonly to identify "hot spots" of point occurrences, like crimes or traffic accidents. The estimators work by splitting a study area into a regular grid, then specifying a search radius over which the density kernel will be calculated. In the context of toxic release sites, the use of KDE helps incorporate the effects of multiple, overlapping sites in close proximity. For this analysis, a search radius of five miles is specified, which suggests that each toxic release site could have harmful effects of to five miles away, and that the effects decrease with distance according to a quadratic decay function. Tracts are assigned the average value of the grid cells that fall inside them. KDE is performed using ArcGIS (version 10.2) via the Kernel Density tool.

### Social-Interactive

Social-Interactive variables measure the attributes of people living in each neighborhood, and are designed to capture the "social processes endogenous to neighborhoods," which may include social capital, collective efficacy, collective socialization, and parental mediation among others (George C Galster 2012). These include a number of variables commonly used to measure concentrated affluence and concentrated disadvantage including owner occupancy rate, income, share of residents with a high school diploma or greater, poverty rate, unemployment rate, and welfare receipt (Sampson, Morenoff, and Earls 1999; DiPasquale and Glaeser 1999; Sampson, Sharkey, and Raudenbush 2008; Hedman et al. 2013). Although these variables are unable to tap directly the social processes that affect socioeconomic outcomes, they have been shown to correlate strongly with these social structures and thus can be viewed as important proxy measures (Sampson, Morenoff, and Earls 1999). This is consistent with the notion of a

[6] http://www2.epa.gov/toxics-release-inventory-tri-program

measurement model in that each variable is viewed as an imperfect measurement of the underlying construct.

Naturally, some of these variables will be positively related to social opportunity (e.g. educational attainment) while others will be negatively related (e.g unemployment). Models were tested that included a nested structure that included sub-dimensions of concentrated affluence and concentrated disadvantage, however the non-nested model yielded a better fit. All variables are collected from the 2010 Census American Community Survey (ACS) via the Neighborhood Change Database (NCDB) provided by Geolytics Inc. As with other categories, the variables selected here represent only a small portion of those that might be included under ideal circumstances. Better data would tap resident perceptions of social cohesion and control through direct survey measures, and might also include information about membership in community organizations, voluntary associations and neighborhood activism (Sampson, Morenoff, and Earls 1999). The omission of these data sources does not imply their lack of importance, but rather their difficulty in collection (Walter and Wang 2016).

### Results

To identify the four dimensions of opportunity outlined above, I construct a confirmatory factor analysis, occasionally referred to as a measurement model in the structural equation modeling literature. The model is estimated using the lavaan package in the R statistical language (Rosseel 2012; R Development Core Team 2011). The results confirm the emergence of the hypothesized factors, and the indicators load strongly in the expected fashion.[7]

Table 1: Factor Loadings

| Factor | Indicator | Loading | Std. Err | Std. Loading |
|---|---|---|---|---|
| social | income | 1.000 | 0.000 | 0.884 |
| social | edu_diploma | 0.915 | 0.034 | 0.808 |
| social | owner_occupied_housing | 0.840 | 0.035 | 0.743 |
| social | poverty | -0.994 | 0.031 | -0.879 |
| social | unemployment | -0.798 | 0.037 | -0.705 |
| social | welfare | -0.764 | 0.038 | -0.675 |
| geographic | walk_score | 1.000 | 0.000 | 0.870 |
| geographic | density_public_institutions | 1.113 | 0.028 | 0.968 |
| geographic | density_social_orgs | 1.103 | 0.028 | 0.960 |
| geographic | jobs_transit | 0.907 | 0.034 | 0.789 |
| geographic | access_healthcare | 1.011 | 0.031 | 0.879 |

[7] Note that because all three variables in the environmental dimension are *negative*, it should be assumed that this represents the inverse of opportunity. In other words, crime, toxic exposure, and vacancy are allowed to load positively on the environmental factor, and the inverse of the environmental factor is taken to be a measure of opportunity.

| Factor | Indicator | Loading | Std. Err | Std. Loading |
|---|---|---|---|---|
| HS | hs_performance | 1.000 | 0.000 | 0.977 |
| HS | ap_scores | 0.957 | 0.016 | 0.936 |
| HS | sat_score | 1.023 | 0.009 | 1.000 |
| HS | teachers_high | 0.771 | 0.027 | 0.753 |
| ES | reading_3rd | 1.000 | 0.000 | 0.928 |
| ES | math_3rd | 0.958 | 0.025 | 0.890 |
| ES | reading_5th | 1.003 | 0.024 | 0.931 |
| ES | math_5th | 0.977 | 0.025 | 0.907 |
| institutional | HS | 1.000 | 0.000 | 0.914 |
| institutional | ES | 0.939 | 0.032 | 0.902 |
| institutional | ms_performance | 0.990 | 0.030 | 0.884 |
| institutional | teachers_middle | 0.883 | 0.033 | 0.788 |
| environmental | toxic | 1.000 | 0.000 | 0.570 |
| environmental | crime | 1.266 | 0.089 | 0.721 |
| environmental | vacancy | 1.275 | 0.089 | 0.727 |

The relationship between the four latent variables, measured quantitatively by their correlation (Table 2) and visually by their maps (Figures 3-6) is substantial. Given the high degree of correlation among the latent factors, a better model might be obtained by allowing cross loadings or specifiing only one or two factors. While such a strategy might facilitate a better fitting model, however, each of the factors measures a distinctly different construct and are estimated using data from different sources they are treated as distinct.

Table 2: Model Fit Indices

| Fit Index | Value |
|---|---|
| SRMR | 0.048 |
| IFI | 0.903 |
| CFI | 0.903 |
| NFI | 0.892 |
| TLI | 0.890 |
| RMSEA | 0.107 |
| Chi Square | 2093.331 |

There are no hard guidelines for determining whether CFA represents an adequate model fit. Most authors suggest that a model with an incremental fit index (e.g. IFI, CFI or TLI) higher than 0.9 represents an adequate fit while indices greater than 0.95 represent a good fit. Absolute fit indices, such as the Square Root Mean Residual (SRMR) should be below 0.8 and ideally below 0.5, whereas the Root Mean Square Error of Approximation (RMSEA) should ideally be below .10. These criteria are merely guidelines, however. To illustrate, in a field in which previous models generate CFI values of .70 only, a CFI value of .85 represents progress and thus should be acceptable (Bollen 1989). Indeed, Marsh, Hau, and Wen (2004) argue that even in psychometrics in which factor models are common and measurement items are relatively standardized, “there is some evidence to suggest that even the old cutoff values (e.g., RNI and TLI > .90) are overly demanding in relation to a normative criterion of appropriateness based on the best existing psychological instruments. Hence, the new, more demanding cutoff values proposed by Hu and Bentler (1998, 1999) appear to be largely unobtainable in appropriate practice” (Marsh, Hau, and Wen 2004, 326). Following this advice, the model presented here appears to be adequate, though not perfect. Although many of the incremental indices are modest, they meet the suggested minimum criteria and there exist no similar models in the literature with which to compare them. The SRMR reports a good fit and the RMSEA sits on the edge.

The four variables estimated by the model are presented as maps in figures 3-6 below. Following the standard convention, maps are shown in quintiles, with green colors representing higher levels of opportunity, purple colors representing lower levels, and white representing the average.

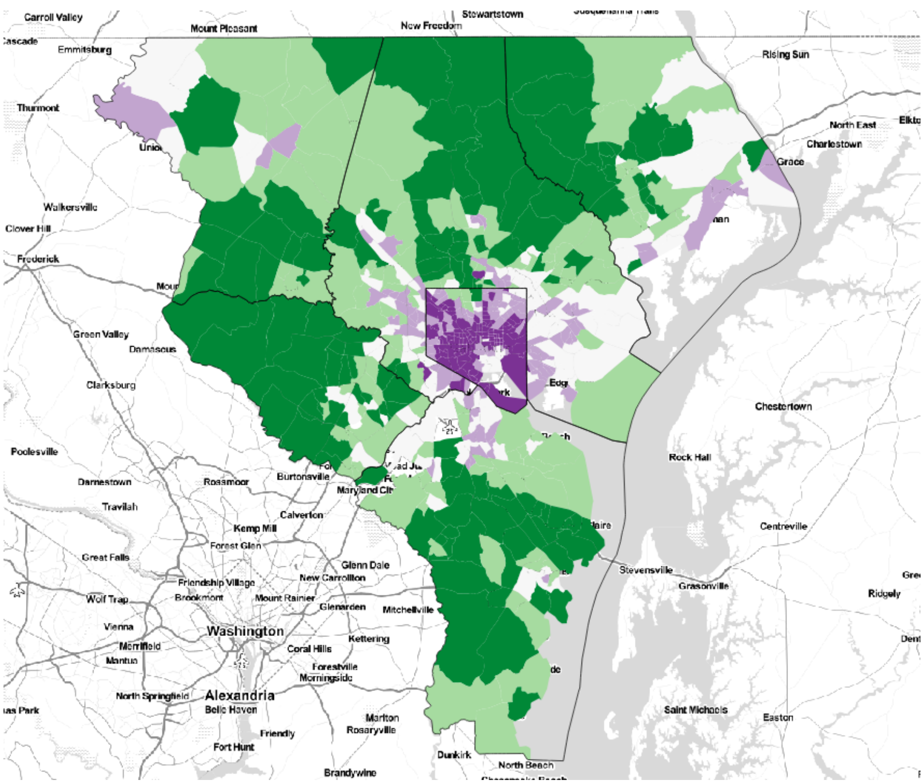


Figure 2: Social-Interactive Map

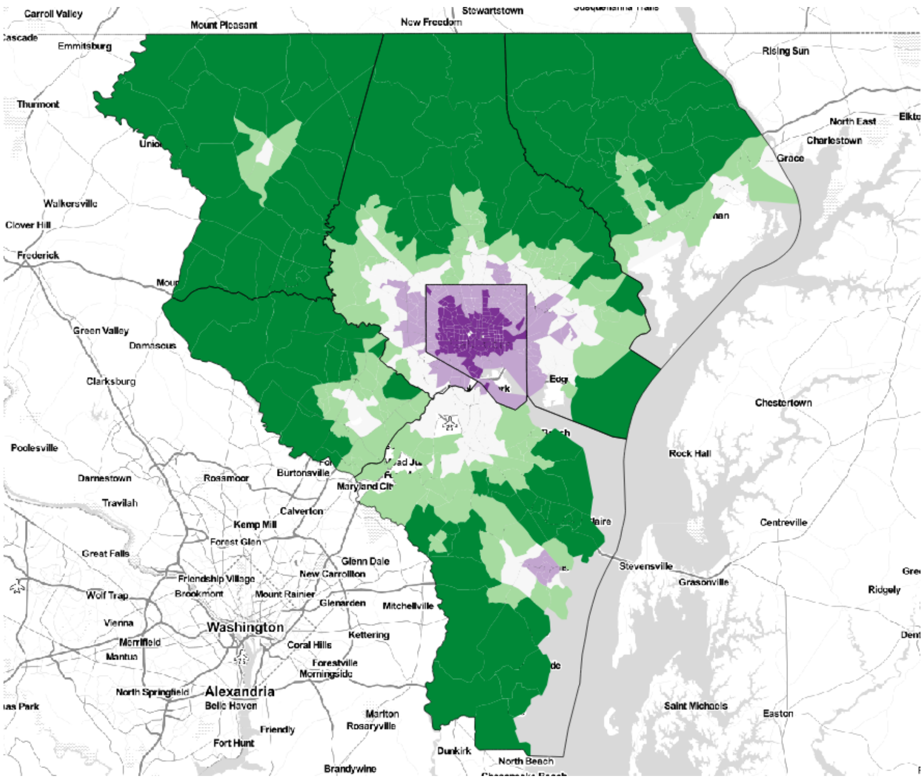


Figure 3: Environmental Map

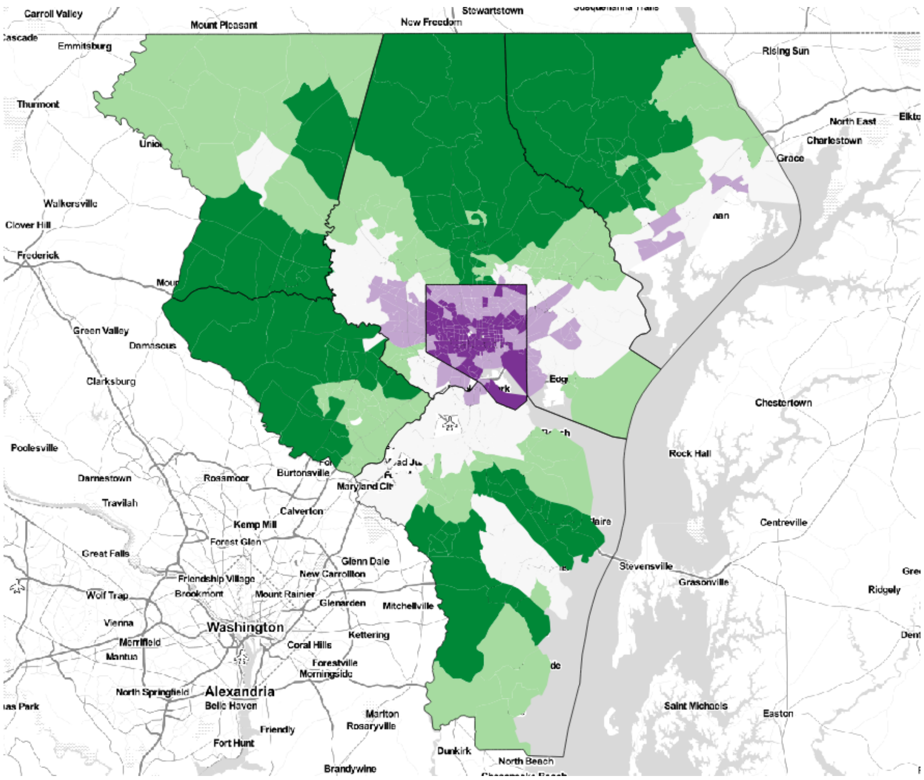


Figure 4: Institutional Map

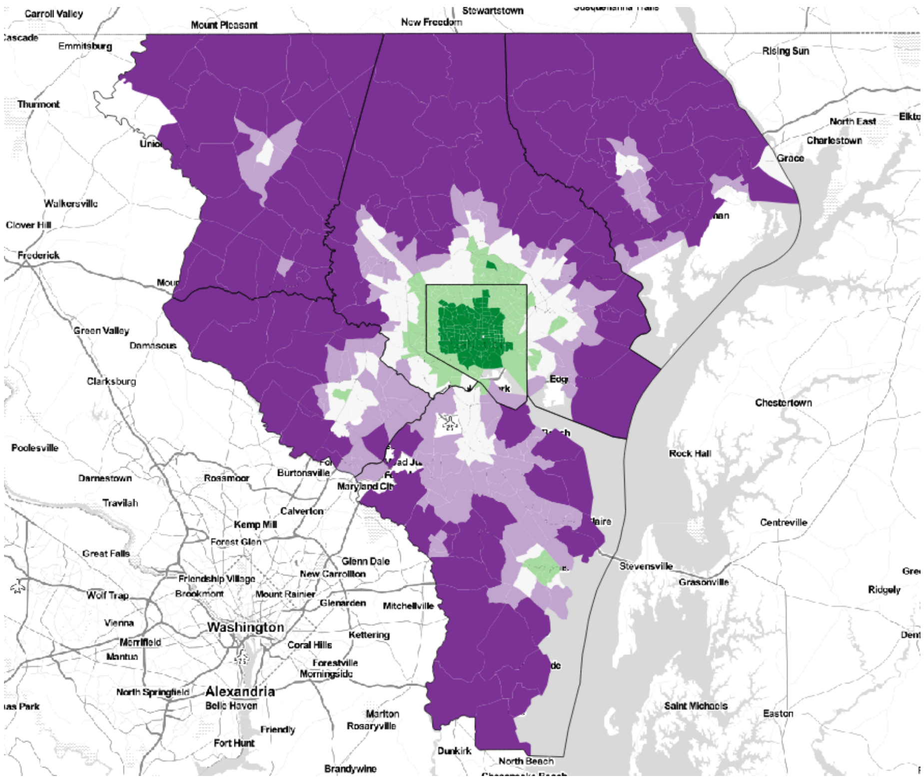


Figure 5: Geographic Map

What is immediately clear is that the social, environmental, and institutional dimensions of opportunity are all highly related to one another; high opportunity tends to be in outer suburbs with most of the disadvantaged neighborhoods clustered in Baltimore City. The one exception to this trend is the geographic dimension of opportunity, which reflects the fact that agglomeration economies and infrastructure provision still favor the City strongly.

Table 3: Latent variable Correlation

| | social | geographic | institutional | environmental |
|---|---|---|---|---|
| social | 1 | -0.78 | 0.90 | 0.88 |
| geographic | | 1 | -0.82 | -0.95 |
| institutional | | | 1 | 0.95 |
| environmental | | | | 1 |

For opportunity mapping applications that require the construction of a single univariate index, such as the analysis of disparate impact in fair housing, analysts are left with few alternatives than to simply average the four components together. This process yields the composite opportunity map presented below.

Again, the composite map may be useful for certain applications, but it does little to provide guidance for interventions that seek to shape the geography of opportunity because it does not provide an indication of why a particular neighborhood has a particular score. For this reason, it may be useful to create an alternative map (or series thereof) which permits greater flexibility in interpretation.

## A Typology of Spatial Opportunity

As an alternative to the composite opportunity index, it may be useful to develop a typology of neighborhoods based on the empirical clustering of opportunity sub-dimensions. This can be accomplished by applying any of several clustering algorithms common to data science and machine learning. Again, this technique, as applied to opportunity data is not entirely novel. A similar approach is applied by Walter and Wang (2016). In their work, however, they use discriminant analysis (which categorizes neighborhoods horizontally) to develop a typology that categorizes neighborhoods along a vertical continuum (e.g. high to low). In my view, this removes the benefit provided by clustering algorithms, which is that the resulting categories need not be ranked in an ordinal hierarchy. An illustrative example is given by Spielman and Singleton (2015) who use cluster analysis applied to U.S. Census data to develop a geodemographic typology. Their typology is descriptive rather than normative and allows them to move away from the "variables paradigm" which seeks to describe neighborhoods along a singular continuum. Instead, they apply "a contextual mode of analysis, [in

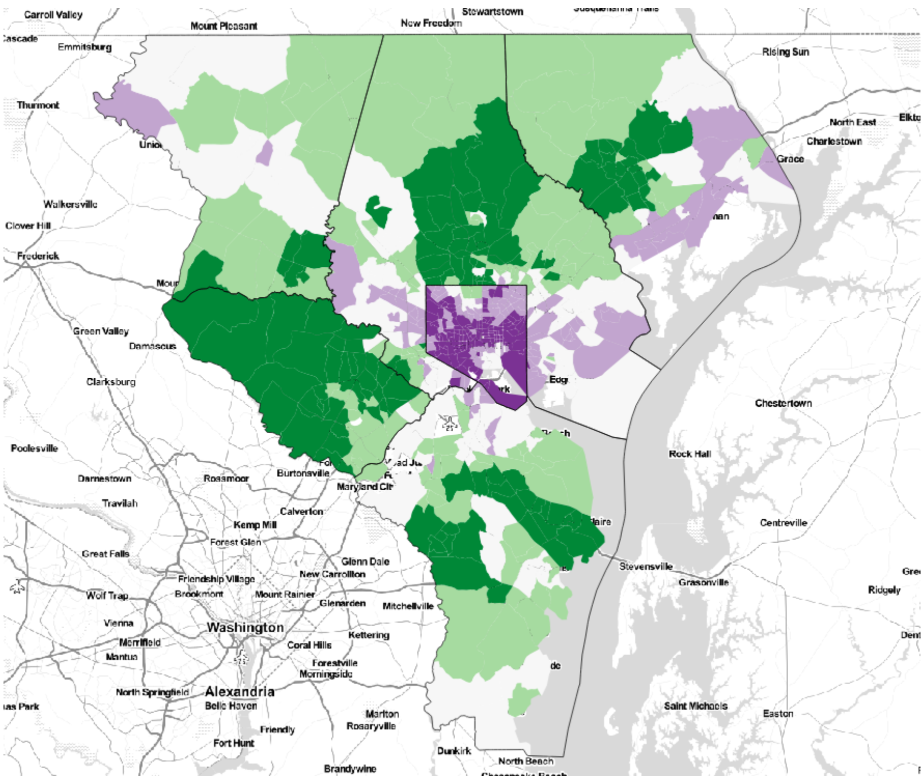


Figure 6: Composite Map

which] neighborhood-to-neighborhood differences are conceptualized as changes of type, not increments to variables" (Spielman and Singleton 2015, 1004).

In the context of opportunity mapping, this is an especially useful paradigm shift for three reasons. First, it recognizes that the *utility* garnered by each dimension of opportunity is heterogeneous across households and aggregation may, therefore, obscure more information than it reveals; as discussed above, young children likely benefit far more from the institutional dimension (e.g. school quality) than the geographic dimension (e.g. access to jobs) (Sharkey and Faber 2014). Second, it removes the need to develop a system for weighting each of the four opportunity dimensions relative to one another. In other words, this removes the necessity of assigning arbitrary values of $\theta$. Third, it provides guidance for policymakers seeking to improve the opportunity level in a given community. A single univariate scale, ranging from high to low can be useful for some policy applications, such as proving disparate impact in the siting of public housing. For other applications, however, such as community development, a univariate scale is less useful. In these cases, policy analysts require information about *why* a particular location has a low opportunity score and what might be done to improve it.

To illustrate this idea, I construct a neighborhood typology by applying a clustering algorithm to the four latent variables estimated in the previous section. One particular benefit of this strategy is that clusters do not contain a pre-defined number of neighborhoods; unlike dividing the composite index into quintiles, each cluster contains a unique number of tracts, defined by their empirical relationship. The cluster analysis identifies five distinct neighborhood types, which vary in their levels of opportunity measures. Formally, this is a gaussian finite mixture model fitted by an expectation-maximization (EM) algorithm performed using the Mclust package for the R statistical language (Fraley et al. 2012). The algorithm tests a variety of different cluster specifications and selects (a) the optimal number of clusters and (b) the assignment of each tract to the optimal cluster based on Bayesian Information Criteria (BIC) (Fraley and Raftery 2002).

The cluster means are presented in the table below. Clusters one and two (which together comprise about 50% of the census tracts) are strong on environmental, social, and institutional measures, with cluster one performing slightly better. Both clusters are lower than average on the geographic dimension. Cluster three is near the regional average on most measures, with geographic measures slightly better and others slightly worse. Clusters four and five have higher than average geographic measures but lower than average on all others. Cluster five appears to be particularly disadvantaged, with extremely low social, institutional and environmental scores.

Table 4: Cluster Means

| cluster | n | proportion | social | geographic | institutional | environmental |
|---|---|---|---|---|---|---|
| 1 | 152 | 0.224 | 0.825 | -0.933 | 0.924 | 0.624 |
| 2 | 199 | 0.309 | 0.354 | -0.312 | 0.390 | 0.235 |
| 3 | 143 | 0.223 | -0.240 | 0.185 | -0.190 | -0.108 |
| 4 | 115 | 0.174 | -0.760 | 0.950 | -1.040 | -0.684 |
| 5 | 47 | 0.069 | -1.601 | 1.494 | -1.513 | -1.030 |

The map generated by the cluster analysis reveals similar macro-level patterns to the composite index, but the interpretation of each neighborhood is more nuanced. It would be convenient to label clusters one and two as high opportunity, cluster three as moderate opportunity, and clusters four and five as low opportunity, but this normative interpretation ignores the tradeoffs these neighborhoods embody. For a household with young children, clusters one and two offer vital resources for development: good schools, supportive social environments, and safe clean air. If this household is transit dependent, however, clusters one and two may not be viable options at all because they lack important access to services captured by the geographic dimension. For this household, utility may be maximized by a neighborhood somewhere in cluster three (or possibly four); in these neighborhoods, it is possible to find a decent school and still maintain decent access to goods and services served by transit.

From a policy perspective, the cluster approach also yields important entry points for particular policy prescriptions. Cluster five desperately needs investment in its institutional infrastructure, and improving the schools may result in positive externalities in the social dimension. Clusters four and five are important candidates for inclusionary zoning and affordable housing policies to ensure that all people have access to the good schools and safe neighborhoods they offer. Cluster four may be ripe for investments in public transportation to help bridge “last mile” connections between homes and workplaces. Adding more data to the cluster analysis can further improve the policy relevance of the results. The inclusion of housing market conditions, for example could differentiate neighborhoods in cluster three–those with improving markets might focus on preserving affordable housing before displacement becomes a concern whereas those with declining markets would be strong candidates for capital injections like block grants.

## Conclusions

Cartography–the making of maps–is among the oldest and best techniques for data visualization. According to the famed data visualization expert Edward Tufte, “No other method for the display of statistical information is so powerful” (Tufte 1983). But like the statistics that underlie them, data maps can mislead

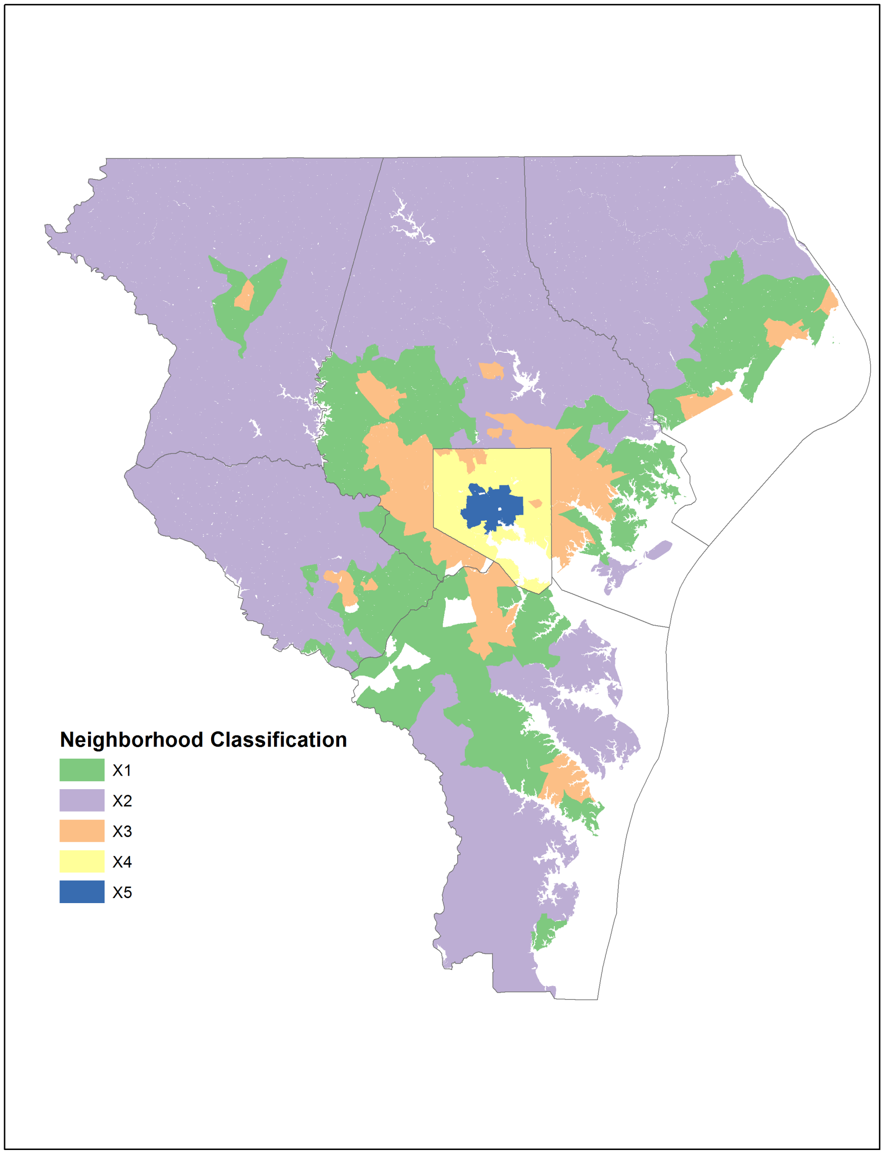


Figure 7: Opportunity Typology

as well as inform. This paper is designed to elucidate the difficulty inherent in designing opportunity maps. While the paper identifies a number of crucial pitfalls common in the current state of practice, it also identifies a number of strategies for overcoming these issues and developing better, more useful metrics and visualizations.

The foundation of this work is a measurement model used to estimate four latent sub-dimensions of opportunity. The model presented here should not be interpreted as the perfect implementation of a spatial opportunity analysis, but rather an example of a general methodology that can and should be extended by additional data and research. This methodology includes a theoretical framework for selecting indicators and dividing them into categories, a measurement strategy that facilitates the evaluation of construct validity, and a visualization strategy that provides more relevance for policymakers.

Despite the solutions proposed in this paper, however, the practice of opportunity mapping still faces a number of serious challenges. As with any quantitative analysis, opportunity mapping is not a purely technical exercise and requires that a series of important subjective decisions be made by an analyst. The only way to validate these decisions is for opportunity analyses to be conducted transparently and vetted by the research community. Beyond the subjectivity of analysts, challenges still remain. Although confirmatory factor analysis and structural equation modeling can help estimate the sub-dimensions of opportunity, more research is necessary to help determine how these dimensions relate to one another and how they combine to produce the socioeconomic outcomes of greatest interest. As the empirical record on neighborhood effects expands, these challenges should wane.

# References

Anderson, Michael L. 2015. "As the Wind Blows: the Effects of Long-Term Exposure To Air Pollution on Mortality." *NBER Working Paper Series* 21578.

Anselin, Luc. 2013. *Spatial econometrics: methods and models*. Vol. 4. Springer Science & Business Media.

Bader, Michael D M, Stephen J Mooney, Yeon Jin Lee, Daniel Sheehan, Kathryn M Neckerman, Andrew G Rundle, and Julien O Teitler. 2015. "Development and deployment of the Computer Assisted Neighborhood Visual Assessment System (CANVAS) to measure health-related neighborhood conditions." *Health & Place* 31: 163–72. doi:10.1016/j.healthplace.2014.10.012.

Bagley, Michael N., and Patricia L. Mokhtarian. 2001. "The impact of residential neighborhood type on travel behavior: A structural equations modeling approach." *The Annals of Regional Science*.

Bollen, K. a. 1989. "A New Incremental Fit Index for General Structural Equation Models." *Sociological Methods & Research* 17: 303–16. doi:10.1177/0049124189017003004.

Briggs, Xavier De Souza, and Margery Austin Turner. 2006. "Assisted housing mobility and the success of low-income minority families: Lessons for policy, practice, and future research." *Nw. JL & Soc. Pol'y* 1 (1). HeinOnline: 25.

Burdick-Will, Julia, Jens Ludwig, Stephen Raudenbush, Robert Sampson, Lisa Sanbonmatsu, and Patrick Sharkey. 2011. "Converging evidence for neighborhood effects on children?'s test scores: An experimental, quasi-experimental, and observational comparison." *Whither Opportunity*, 255–76.

Chapple, Karen. 2014. *Planning sustainable cities and regions: towards more equitable development.* Routledge.

Chetty, Raj, and Nathaniel Hendren. 2015. "The Impacts of Neighborhoods on Intergenerational Mobility: Childhood Exposure Effects and County-Level Estimates and numerous seminar participants for helpful comments and discussions." May.

Chetty, Raj, Nathaniel Hendren, and Lawrence Katz. 2015. "The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment." August. National Bureau of Economic Research. doi:10.3386/w21156.

Collins, MICHAEL G., FREDERICK R. Steiner, and MICHAEL J. Rushman. 2001. "Land-use suitability analysis in the United States: Historical development and promising technological achievements." *Environmental Management* 28 (5): 611–21. doi:10.1007/s002670010247.

Dietz, Robert D. 2002. "The estimation of neighborhood effects in the social sciences: An interdisciplinary approach." *Social Science Research* 31 (4). Elsevier: 539–75. doi:10.1016/S0049-089X(02)00005-4.

Diez Roux AV, Wing S Morland K, Kimberly Morland, Ana V Diez Roux, Steve Wing, and Wing S Morland K Diez Roux AV. 2006. "Supermarkets, other food stores, and obesity: The Atherosclerosis Risk in Communities Study." *American Journal of Preventive Medicine* 30 (4). Elsevier: 333–39. doi:10.1016/j.amepre.2005.11.003.

DiPasquale, Denise, and Edward L. Glaeser. 1999. "Incentives and social capital: are homeowners better citizens?" *Journal of Urban Economics* 45 (2). Elsevier: 354–84. http://www.sciencedirect.com/science/article/pii/S0094119098920988.

Ellen, Ingrid Gould, and Margery Austin Turner. 1997. "Does neighborhood matter? Assessing recent evidence." *Housing Policy Debate* 8 (4). Taylor & Francis: 833–66.

Ewing, Reid, Tom Schmid, Richard Killingsworth, Amy Zlot, and Stephen Raudenbush. 2003. "Relationship Between Urban Sprawl and Physical Activity, Obesity, and Morbidity." *American Journal of Health Promotion, Inc.* 18 (01): 11. doi:10.4278/0890-1171-18.1.47.

Fraley, Chris, and Adrian E Raftery. 2002. "Model-Based Clustering, Discriminant Analysis, and Density Estimation." *Journal of the American Statistical Association* 97 (458): 611–31. doi:10.1198/016214502760047131.

Fraley, Chris, Adrian E Raftery, T Brendan Murphy, and Luca Scrucca. 2012. "mclust Version 4 for R: Normal Mixture Modeling for Model-Based Clustering, Classification, and Density Estimation."

Galster, George. 2008. "Quantifying the Effect of Neighbourhood on Individuals: Challenges, Alternative Approaches, and Promising Directions." *Schmollers Jahrbuch* 128 (1): 7–48.

———. 2010. "The Mechanism (s) of Neighborhood Effects Theory, Evidence, and Policy Implications." *ESRC Seminar*, 1–32. doi:10.1007/978-94-007-2309-2.

Galster, George C. 2012. "The mechanism (s) of neighbourhood effects: Theory, evidence, and policy implications." In *Neighbourhood Effects Research: New Perspectives*, 23–56. Springer.

Galster, George C., and Sean P. Killen. 1995. "The geography of metropolitan opportunity: A reconnaissance and conceptual framework." *Housing Policy Debate* 6 (1): 7–43. doi:10.1080/10511482.1995.9521180.

Galster, George, Roberto Quercia, and Alvaro Cortes. 2000. "Identifying neighborhood thresholds: An empirical exploration." *Housing Policy Debate* 11 (3). Taylor & Francis: 701–32. doi:10.1080/10511482.2000.9521383.

Galster, George, Anna Santiago, Jessica Lucero, and Jackie Cutsinger. 2015. "Adolescent neighborhood context and young adult economic outcomes for low-income African Americans and Latinos." *Journal of Economic Geography*. Oxford Univ Press, lbv004.

Giovannini, Enrico, Michela Nardo, Michaela Saisana, Andrea Saltelli, Stefano Tarantola, and Anders Hoffman. 2005. "Handbook on constructing composite indicators: methodology and user guide." OECD Statistics Working Paper, STD/DOC OECD publishing.

Hansen, Walter G. 1959. "Accessibility and Residential Growth." *Massachusetts Institute of Technology*.

Hedman, Lina, David Manley, Maarten van Ham, J. Osth, and John Östh. 2013. "Cumulative exposure to disadvantage and the intergenerational transmission of neighbourhood effects." *Journal of Economic Geography*. Oxford Univ Press, lbt042. http://joeg.oxfordjournals.org/cgi/doi/10.1093/jeg/lbt042.

Ingraham, Christopher. 2016. "The Dirtly Little Secret that Data Journalists Aren't Telling You." Washington DC. https://www.washingtonpost.com/news/

wonk/wp/2016/04/11/the-dirty-little-secret-that-data-journalists-arent-telling-you/.

Kaczynski, Andrew T., and Karla a. Henderson. 2007. "Environmental Correlates of Physical Activity: A Review of Evidence about Parks and Recreation." *Leisure Sciences* 29 (4): 315–54. doi:10.1080/01490400701394865.

Kain, John F. 1992. "The spatial mismatch hypothesis: three decades later." *Housing Policy Debate* 3 (2). Taylor & Francis: 371–460.

Kain, John F. 1968. "Housing Segregation, Negro Employment, and Metropolitan Decentralization." *The Quarterly Journal of Economics* 82 (2): 175–97. doi:10.2307/1885893.

Kirwan Institute for the Study of Race and Ethnicity. 2013. "Opportunity Metrics Resource of Data and Research." Columbus, OH: Kirwan Institute for the Study of Race; Ethnicity. http://kirwaninstitute.osu.edu/docs/FINAL{_}2013 Opp Metrics.pdf.

Kling, Jeffrey R., Jens Ludwig, and Lawrence F. Katz. 2005. "Neighborhood Effects on Crime for Female and Male Youth: Evidence from a Randomized Housing Voucher Experiment." *The Quarterly Journal of Economics* 120 (1): 87–130. doi:10.2307/25098732.

Leventhal, Tama, and Jeanne Brooks-Gunn. 2000. "The neighborhoods they live in: the effects of neighborhood residence on child and adolescent outcomes." *Psychological Bulletin* 126 (2). Center for Children; Families, Teachers College, Columbia University, New York, New York 10027, USA. TL91@columbia.edu; American Psychological Association: 309. doi:10.1037/0033-2909.126.2.309.

Ludwig, Jens, Greg J Duncan, Lisa a Gennetian, Lawrence F Katz, Ronald C Kessler, Jeffrey R Kling, and Lisa Sanbonmatsu. 2012. "Neighborhood effects on the long-term well-being of low-income adults." *Science (New York, N.Y.)* 337 (6101): 1505–10. doi:10.1126/science.1224648.

Marsh, Herbert, Kit-Tai Hau, and Zhonglin Wen. 2004. "In Search of Golden Rules: Comment on Hypothesis-Testing Approaches to Setting Cutoff Values for Fit Indexes and Dangers in Overgeneralizing Hu and Bentler's (1999) Findings." *Structural Equation Modeling: A Multidisciplinary Journal* 11 (3): 320–41. doi:10.1207/s15328007sem1103_2.

Massey, Douglas S. 2015. "The Social Science of Affordable Housing." *Housing Policy Debate* 25 (3). Taylor & Francis: 634–38.

Mayer, S E, and C Jencks. 1989. "Growing up in poor neighborhoods: how much does it matter?" *Science (New York, N.Y.)* 243 (4897): 1441–5. doi:10.1126/science.243.4897.1441.

Mujahid, Mahasin S., Ana V. Diez Roux, Jeffrey D. Morenoff, and Trivellore Raghunathan. 2007. "Assessing the measurement properties of neighborhood scales: from psychometrics to ecometrics." *Am J Epidemiol* 165 (8). United

States: 858–67. doi:10.1093/aje/kwm040.

Nosoohi, Iman, and Ali Zeinal-Hamadani. 2011. "Urban Planning With the Aid of Factor Analysis Approach: The Case of Isfahan Municipality." *Theoretical and Empirical Researches in Urban Management* 6 (1).

O'Brien, Daniel, Robert Sampson, and Christopher Winship. 2013. "Ecometrics in the Age of Big Data: Measuring and Assessing 'Broken Windows' Using Administrative Records." doi:10.1177/0081175015576601.

Powell, John. 2005. "Remedial Phase Expert Report In Thompson v. HUD." Kirwan Institute.

Powell, john a. 2003. "Opportunity-Based Housing." *Journal of Affordable Housing & Community Development Law* 12 (2): 188–228.

R Development Core Team, R. 2011. *R: A Language and Environment for Statistical Computing*. Vol. 1. 2.11.1. doi:10.1007/978-3-540-74686-7.

Raudenbush, Stephen, and Robert Sampson. 1999. "Ecometrics: toward a science of assessing ecological settings, with application to the systematic social observation of neighborhoods." *Sociological Methodology* 29 (1). Wiley Online Library: 1–41.

Reece, Jason, and Samir Gambhir. 2008. "The Geography of Opportunity Review of Opportunity Mapping Research Initiatives." Kirwan Intsitute for the Study of Race; Ethnicity.

Rosenbaum, James E, and Anita Zuberi. 2010. "Comparing residential mobility programs: design elements, neighborhood placements, and outcomes in MTO and Gautreaux." *Housing Policy Debate* 20 (1). Taylor & Francis: 27–41.

Rosseel, Y. 2012. "lavaan: An R Package for Structural Equation Modeling." *Journal of Statistical Software* 48 (2): 1–36. doi:10.18637/jss.v048.i02.

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

Sampson, Robert, Jeffrey Morenoff, and Felton Earls. 1999. "Beyond social capital: Spatial dynamics of collective efficacy for children." *American Sociological Review* 64 (5). JSTOR: 633–60. doi:10.2307/2657367.

Sampson, Robert, Jeffrey Morenoff, and Thomas Gannon-Rowley. 2002. "Assessing 'Neighborhood Effects': Social Processes and New Directions in Research." *Annual Review of Sociology* 28 (1): 443–78. doi:10.1146/annurev.soc.28.110601.141114.

Sampson, Robert, Jeffrey Morenoff, and Stephen Raudenbush. 2005. "Social anatomy of racial and ethnic disparities in violence." *Am J Public Health* 95 (2). United States: 224–32. doi:10.2105/AJPH.2004.037705.

Sampson, Robert, Patrick Sharkey, and Stephen Raudenbush. 2008. "Durable effects of concentrated disadvantage on verbal ability among African-American children." *Proc Natl Acad Sci U S A* 105 (3). United States: 845–52. doi:10.1073/pnas.0710189104.

Sanbonmatsu, Lisa, Jens Ludwig, Lawrence F Katz, Lisa A Gennetian, Greg J Duncan, Ronald C Kessler, Emma Adam, Thomas W McDade, and Stacy Tessler Lindau. 2011. "Moving to Opportunity for Fair Housing Demonstration Program–Final Impacts Evaluation." US Department of Housing & Urban Development, PD&R.

Sharkey, Patrick, and Felix Elwert. 2011. "The Legacy of Disadvantage : Multigenerational Neighborhood Effects on Cognitive Ability." *American Journal of Sociology* 116 (6). United States: 1934–81.

Sharkey, Patrick, and Jacob Faber. 2014. "Where, When, Why, and For Whom Do Residential Contexts Matter? Moving Away from the Dichotomous Understanding of Neighborhood Effects." *Annual Review of Sociology* 40 (1): 559–79. doi:10.1146/annurev-soc-071913-043350.

Spielman, Seth E., and Alex Singleton. 2015. "Studying Neighborhoods Using Uncertain Data from the American Community Survey: A Contextual Approach." *Annals of the Association of American Geographers* 5608 (August): 1–23. doi:10.1080/00045608.2015.1052335.

Talen, E., and L. Anselin. 1998. "Assessing spatial equity: An evaluation of measures of accessibility to public playgrounds." *Environment and Planning A* 30 (4): 595–613. doi:10.1068/a300595.

Talen, Emily. 1998. "Visualizing Fairness: Equity Maps for Planners." *Journal of the American Planning Association* 64 (1): 22–38. doi:10.1080/01944369808975954.

Toulmin, Llewellyn M. 1988. "Equity as a decision rule in determining the distribution of urban public services." *Urban Affairs Review* 23 (3). Sage Publications: 389–413.

Truelove, Marie. 1993. "Measurement of spatial equity." *Environment and Planning C: Government and Policy* 11 (1). SAGE Publications: 19–34.

Tufte, Edward R. 1983. *The visual display of quantitative information.* Vol. 2. 9. Graphics press Cheshire, CT.

Walter, Rebecca J., and Ruoniu Wang. 2016. "Searching for Affordability and Opportunity: A Framework for the Housing Choice Voucher Program." *Housing Policy Debate*, no. May: 1–22. doi:10.1080/10511482.2016.1163276.

Wodtke, Geoffrey T, David J Harding, and Felix Elwert. 2011. "Neighborhood Effects in Temporal Perspective." *Am Sociol Rev* 76 (5): 713–36. doi:10.1177/0003122411420816.

Xie, Zhixiao, and Jun Yan. 2008. "Kernel Density Estimation of traffic accidents in a network space." *Computers, Environment and Urban Systems* 32 (5):

396–406. doi:10.1016/j.compenvurbsys.2008.05.001.

Yang, Muzhe, and Shin-Yi Chou. 2015. “Impacts of Being Downwind of a Coal-Fired Power Plant on Infant Health at Birth: Evidence from the Precedent-Setting Portland Rule.” National Bureau of Economic Research.