

Chapter 1

Introduction and Goals

The literature on residential segregation is one of the oldest empirical research traditions in sociology and has long been a core topic in the study of social stratification and inequality as well as in the study of the demography of spatial population distribution. This literature is guided by the fundamental assumption that group differences in neighborhood residential outcomes are closely associated with social position and life chances. Accordingly, indices measuring segregation, especially the dimension of uneven distribution, are viewed as important summary indicators of overall group standing and scores for segregation indices have been a mainstay of research documenting levels, patterns, and trends in the residential segregation of racial and ethnic groups. Given the extensive attention social scientists have directed to the study of residential segregation, one might assume that the relationship between residential segregation and group differences in neighborhood residential outcomes is well understood. Surprisingly, this is not the case. The issue has received little attention in the literature on segregation measurement. Consequently, researchers are not able to offer precise conclusions about group differences in residential outcomes based on scores for popular and widely used indices of uneven distribution.

In this monograph I address this deficiency in the literature by outlining a new approach to measuring uneven distribution. My goal is not to replace familiar, widely-used indices with new ones. Instead, I wish to place popular indices in a new alternative framework that clarifies the implications they carry for group differences in individual-level residential outcomes. My motivation for doing this rests on two convictions. One is that understanding how segregation is related to individual residential outcomes is desirable for its own sake and brings valuable new options for interpreting segregation index scores and understanding differences between them. The other is that casting segregation indices in terms of group differences in individual residential outcomes brings benefits for segregation measurement and analysis including, as two primary examples, the ability to directly link segregation at the aggregate or macro level to micro-level processes of residential attainment and the

ability to develop versions of the indices that are free of the troublesome problem of inherent upward bias.

Moving from generalities to specifics, my goal in this monograph is to set forth the “difference of means” framework, a new framework for segregation measurement wherein popular indices of uneven distribution are cast as simple differences of group means on residential outcomes that register group contact and exposure based on area racial composition. In accomplishing this goal I establish that all widely used segregation indices including the Gini Index (G), the Delta or Dissimilarity Index (D), the Hutchens Square Root Index (R) – an index with close similarities to the Atkinson Index (A), the Theil Entropy Index (H), and the Separation Index (S) – also known as the variance ratio and a variety of other names, can be expressed as a difference of group means on individual- or household-level residential outcomes (y) that are scored on the basis of index-specific scaling of group contact based on area group proportions.

The indices just listed are all well-known and all have been reviewed in detail in many previous methodological studies (e.g., Duncan and Duncan 1955; Zoloth 1976; James and Taeuber 1985; Stearns and Logan 1986; White 1986; Massey and Denton 1988; Hutchens 2001, 2004; Reardon and Firebaugh 2002). The contribution I seek to make is to clarify a characteristic of these indices that currently is not well understood; namely, the particular way each one relates to and ultimately quantitatively registers group differences in neighborhood residential outcomes. The sociological relevance of segregation index scores rests on the presumption that they carry important implications for group differences in social position and life chances that are associated with area of residence. Segregation researchers and consumers of segregation research thus generally assume that variation in segregation index scores tends to correlate with variation in a broad range of group disparities associated with neighborhood residential outcomes.

It is definitely plausible to assume that summary index scores may serve as proxies for valuable, but usually unavailable information about residentially-based group inequality and disparity. But it is important to recognize that, in the final analysis, the calculations that yield segregation index scores revolve around a simple and very particular aspect of neighborhood residential outcomes – “pairwise” group proportions.¹ This residential outcome can be understood in multiple ways from the point of view of individuals and households. For example, it can be understood as registering levels of contact or exposure based on co-residence with members of the two groups in the comparison. Alternatively, it can be understood as registering exposure to deviations or departures from the racial composition of the city as a whole. One of my goals is to clarify how different indices register group differences on individual residential outcomes relating to area racial mix and group proportions.

¹ The expression “pairwise” ethnic mix or group proportion signifies that the calculations involved use only the population counts for the two groups in the comparison. It can be contrasted with “overall” ethnic mix or group proportion based on the full population including groups *not* in the segregation comparison. Significantly, the presence and distributions of groups other than the two in the comparison has no bearing on scores for indices of uneven distribution.

In doing so I hope to help researchers better understand what indices specifically measure in this regard when they are interpreting index scores and evaluating their relevance as proxies for group position.

Indices of uneven distribution provide quantitative summaries of how groups are differentially distributed across neighborhoods that vary on “pairwise” racial mix. This obviously has direct implications for individual residential outcomes relating to racial mix and indices can be cast in two ways that reflect this fact. One option is to cast indices as simple, overall population averages on individual residential outcomes scored on the basis of area racial mix. I review this option briefly, but I give it limited attention because it is not especially novel and it is not useful for my main goals. The second option is to cast indices of uneven distribution as group differences of means on segregation-relevant neighborhood residential outcomes scored from pairwise racial mix. This approach is the primary focus of my attention because it resonates with substantive interests that motivate much of the research on segregation – namely, concerns about group disadvantage and inequality rooted in differential residential distribution. Additionally, the difference of means approach brings several practical advantages for segregation measurement and analysis.

I offer the difference of means framework for computing indices of uneven distribution in hopes that it will be a useful alternative to prevailing approaches to computing index scores. However, I stress from the outset that I intend this new framework to be an enhancement of and supplement to traditional approaches to segregation measurement, not a wholesale replacement. The difference of means framework does not yield different values for index scores. Instead, it yields identical index scores but draws on new, mathematically equivalent index formulations to gain new understandings of segregation and new options for measurement, interpretation, and analysis. In current practice indices of uneven distribution are formulated and interpreted in ways that focus attention on aggregate-level patterns for spatial units (i.e., areas or neighborhoods). The formulas used generally feature calculations that register the extent to which the racial mix of areas (neighborhoods) within a city depart from the racial composition of the city as a whole. These widely used formulas are tried and true and they are useful and convenient for many purposes. That said, it also is important to recognize what the most widely used computing formulas neglect and obscure. Traditional approaches to measuring uneven distribution do not clarify the how segregation is connected to group differences in neighborhood residential outcomes for individuals. It is obvious that neighborhood departures from city racial composition necessarily carry implications for group differences in residential outcomes. But the specific nature of these implications is not well understood because it is not revealed in prevailing approaches to formulating, computing, and interpreting segregation indices.

The “difference of means” framework for calculating and interpreting popular segregation indices I introduce here addresses this gap in the literature on segregation measurement. The framework highlights something that currently is not widely appreciated – that differences between indices can be understood as arising from a single factor, the particular way each index registers segregation-relevant residential outcomes for individuals as scored from area racial composition. On reflection this

probably should not be surprising. All indices are calculated from the same underlying distribution of residential outcomes on pairwise racial proportions. Consequently, index scores obtained from group differences of means on residential outcomes can differ only by registering these very specific residential outcomes in different ways. These cross-index differences in “scoring” area racial mix provide a new basis for comparing and evaluating indices of uneven distribution.

The difference of means formulation of indices of uneven distribution brings additional practical benefits beyond clarifying how index scores are related to group differences in residential outcomes. One example is that the approach makes it possible to join the study of aggregate segregation with the study of individual-level residential attainment in a seamless way. This becomes possible because segregation index scores now can be viewed as arising from the simple additive aggregation of segregation-relevant, neighborhood residential outcomes for individuals. As a result, segregation index scores can be equated with the effect of race in micro-level regression models predicting the residential attainments of individuals and households that additively determine segregation at the aggregate-level.² These micro-level attainment models can be extended to include multiple individual and household characteristics as predictors in the attainment equation. This then enables researchers to assess segregation – now equated to the effect of race on residential attainments – in multivariate specifications that control for non-racial factors (e.g., income, nativity, language ability, etc.) that also may affect the residential attainments that ultimately determine segregation. The new ability to model the individual-level residential attainments that directly and additively give rise to segregation makes it possible to undertake quantitative standardization and decomposition analyses to assess the extent to which group differences on factors other than race contribute to overall segregation based on their impact on residential outcomes that determine aggregate segregation. Finally, city-specific, individual-level models of residential attainments can be extended to multi-level specifications that can be used to investigate variation in segregation over time and across cities in new ways that previously were not feasible.

The kinds of analysis options just described have been available and used on a routine basis for decades in the broader literature investigating racial differences in most domains of socioeconomic attainment (e.g., education, income, occupation, etc.). Until now, however, they have not been available in segregation research. The reason for this is that segregation, in contrast to racial inequality in other socioeconomic attainments such as education, occupation, and income, has not been explicitly formulated in terms of group differences on individual attainments. Placing indices of uneven distribution in the difference of means framework thus puts segregation analysis on similar conceptual footing with research traditions that analyze other aspects of racial socioeconomic disparity and inequality.

²Specifically, the segregation index score is equal to the value of the unstandardized regression coefficient for race (coded as 0 or 1) in an individual-level regression predicting residential outcomes.

The difference of means formulation of indices of uneven distribution brings other benefits as well. One conceptual benefit is to introduce a new basis for evaluating and choosing among familiar indices; namely, whether and to what degree the individual-level residential outcomes registered by a given index are relevant for theories of segregation dynamics and racial socioeconomic stratification. Another practical benefit is that the approach makes it easy to implement spatial versions of popular segregation indices.

Last but not least, the difference of means formulation of segregation indices provides a basis for gaining a better understanding the source of index bias – a well-known and vexing problem that can make scores of standard versions of indices of uneven distribution untrustworthy and potentially misleading. This new understanding then makes it possible to develop unbiased versions of popular indices based on implementing surprisingly simple refinements to index formulas that eliminate this problematic behavior of index scores.

In the chapters that follow I introduce the difference of means formulations of widely used segregation indices and provide more detailed reviews of the new options for measurement, interpretation, and analysis just mentioned. In Chaps. 2, 3, 4, and 5 I introduce the difference of means framework and explore differences between indices as revealed through the lens of this framework. I begin in Chap. 2 by noting that scores of popular indices of uneven distribution can be obtained using a variety of mathematically equivalent formulas and I briefly review selected formulas to highlight how they support different insights about segregation measurement. I conclude the chapter by introducing the difference of means formulas that are used throughout this monograph. In Chap. 3 I provide a general overview of the difference of means framework. I then expand on this in Chap. 4 by offering a more detailed discussion of how individual measures of uneven distribution can be cast as difference of group means on residential outcomes scored from area racial proportions. In Chap. 5 I note a useful insight about uneven distribution that emerges from the difference of means framework; namely, that differences between indices can be seen as arising from a single source – how each index registers individual residential outcomes scored from area group proportions.

In Chaps. 6, 7, and 8 I review the logical and empirical differences among popular measures of uneven distribution and offer suggestions regarding how to understand and interpret these differences. In Chap. 6 I document that, in contrast to findings reported in some previous methodological studies, popular indices of even distribution can and often do yield highly discrepant scores. The analyses I present here establish that the findings of earlier methodological studies – which reported that popular indices tended to be highly correlated in empirical application – are a byproduct of focusing primarily on White-Minority segregation in a small subset of large metropolitan areas where the minority group is a substantial presence in terms of relative group size and where group residential distributions are characterized by a particular pattern of “prototypical” segregation. This is a pattern of uneven distribution in which group displacement from parity involves a high level of group separation and area racial polarization because *both groups* are disproportionately concentrated in homogeneous areas. I refer to uneven distribution with this pattern

of “concentrated displacement” as “prototypical segregation” because this signature pattern – in which all popular measures of uneven distribution take high scores – is always present in crafted examples used to illustrate high segregation in didactic discussions of segregation measurement. Similarly, it also is invariably present in empirical cases used to illustrate high levels of segregation. So it is easy to understand that many would not be aware that popular indices can take substantially discrepant scores.

The empirical analyses I present in Chap. 6 document that uneven distribution does not always take the form of prototypical segregation. To the contrary, the analyses instead reveal that broader samples of cities include a large number of cases with a sharply contrasting pattern of “dispersed displacement” wherein uneven distribution involves extensive group displacement from parity but does not involve group separation and area racial polarization. In these situations, index scores can be highly discrepant. Specifically, indices that are sensitive to differential displacement – such as the gini index (G) and the dissimilarity index (D) which Duncan and Duncan (1955) aptly also termed the displacement index – will take high scores while the Theil index (H) and the separation index (S) – which Stearns and Logan (1986) note is sensitive to residential separation and area racial polarization – will take low scores.

In Chap. 7 I review the distinction between concentrated and dispersed displacement in more detail. The chapter makes two important points. One is that the sociological implications of uneven distribution involving “prototypical segregation” and D-S concordance are fundamentally different from the sociological implications of uneven distribution with dispersed displacement and substantial D-S divergence. Simply put, a high level of group separation is obviously substantively compelling and necessarily entails a high level of displacement. But the reverse is not true. Thus, high levels of displacement do not always entail high levels of group separation and this should be noted when it occurs because the literature on segregation measurement provides no clear basis for viewing differential displacement without group separation as sociologically important. The second point I make in this chapter is that the largely unrecognized but empirically common outcome of dispersed displacement is not an artifact of relative group size or deficiencies in indices that are more sensitive to group separation than displacement. To make this point I introduce and exercise simple analytic models to show that when non-trivial displacement from even distribution is present, it can be concentrated or it can be dispersed. Concentrated displacement produces “prototypical segregation” wherein the score of S will approach or even equal the score for D indicating that displacement involves group separation and area polarization. In the case of dispersed displacement, D will be equally high but S will be low signaling that group separation and area polarization are minimal, sometimes to the point of being negligible. I review the principles of transfers and exchanges from segregation measurement theory to establish that D-S discrepancies of this sort arise because D is flawed and does not conform to these accepted principles of segregation measurement.

Chapter 8 supplements the analytic results by discussing the sociological dynamics that are likely to influence whether non-trivial displacement takes the form of

“prototypical segregation” or the substantively less compelling pattern of dispersed displacement. It also reviews case studies of empirical examples of high-D-high S combinations that in communities where the minority group is small in relative size. The discussion here drives home two important points. One is that scores for D and S can be congruent or discrepant in any setting where displacement from uneven distribution is non-trivial. The other is that sociological dynamics, not artifacts of index construction, determine whether in fact D and S are congruent or discrepant in a given community.

In Chap. 9 I show how the difference of means framework creates new options for research by joining micro- and macro-level analysis of segregation. At the simplest level, casting segregation as a difference of group means on residential outcomes leads to the new insight that segregation index scores are exactly mathematically equivalent to the effect of race in bivariate regression analyses predicting segregation-determining residential outcomes for individuals. I then argue that this insight opens the door to the new possibility of using multivariate regression analyses to quantitatively assess how segregation arises from two sources. The first source is group differences on distributions of social and economic characteristics that are salient in residential attainment processes. The second source is group differences in the efficacy of how inputs to residential attainment processes translate into segregation-determining residential outcomes. In this framework, segregation can be analyzed in greater detail and sophistication by using standardization and decomposition analysis in combination with multivariate regression analysis of attainments, methods that are routinely applied to the study of racial inequality in education, occupation, income, health, and other important stratification outcomes. This is a major advance as research on segregation has lagged behind research on group disparities in other domains where aggregate-level outcomes on group disparities have long been routinely analyzed as outgrowths of micro-level attainment processes.

In Chap. 10 I show how the regression analysis of individual-level residential attainments can subsume comparative analysis of cross-city variation in segregation and create new possibilities for investigating the factors contributing to variation in segregation across cities and over time. The new approach involves extending city-specific analysis of segregation using bivariate and multivariate models of individual-level residential attainment to multi-level specifications that reveal how the process determining segregation varies across cities and over time. I first note that findings from aggregate-level analyses of cross-city variation in segregation can be exactly reproduced using multi-level specifications of segregation-attainment models. I then outline how this specification opens the door for improving the ability of researchers take accurately assess the role that non-racial characteristics such as income may play in shaping cross-city variation in segregation.

Previous research has often tried to assess the impact of group differences on income and other individual-level characteristics by the method of aggregate-level regression analysis. I note that this approach is prone to yield flawed results because it runs afoul of the “aggregate fallacy.” The problem is hidden from view and less obvious when segregation is viewed only as a macro-level outcome. It becomes

clear and more readily evident when segregation is analyzed within the difference of means framework where the outcome of segregation at the aggregate-level is exactly determined by individual-level attainment processes. I demonstrate the importance of the problem by showing that results of analyses that assess the impact of group income differences on segregation using aggregate-level regressions are contradicted by multi-level regression analyses that avoid the aggregate fallacy and properly take account of the effects of income at the individual-level.

In Chaps. 11, 12, and 13 I review topics that benefit from insights and perspectives gained from drawing on the difference of means framework for analyzing segregation. In Chap. 11 I note that the difference of means framework makes it easy for researchers to implement spatial versions of popular segregation indices if they desire to do so. The reason for this is simple; the residential attainments for individuals that determine segregation scores can be computed using mutually-exclusive bounded areas, or using overlapping, spatially-defined areas. The former yields a traditional aspatial index score. The latter yields a “spatial” index score that is affected by how neighborhoods that vary in racial composition are distributed in space.

In Chaps. 12 and 13 I argue that the difference of means framework leads to new perspectives regarding what aspects of residential segregation researchers will view as most compelling on substantive grounds. In Chap. 12 I argue that group separation is a more compelling substantive concern than mere displacement from even distribution. I frame the issue as follows. It is non-controversial to assert that group separation area racial polarization is substantively important because it is a logical prerequisite for group disparity and inequality on neighborhood-based residential outcomes. In contrast, there is no established basis for arguing that displacement from even distribution is substantively important when it does not involve group separation and area racial polarization. The only candidate is the “volume of movement” interpretation of D in which a high value of D does indicate that a large fraction of one group must move to bring about exact even distribution. But it is rendered irrelevant in situations where movement to exact even distribution has no impact on group separation and area racial polarization.

In Chap. 13 I consider how being sensitive to different aspects of uneven distribution makes different indices more or less relevant for theories of segregation. I note that measures rooted in the segregation curve – G and D – are sensitive to rank-order differences on the residential outcome of area racial proportion but are relatively insensitive to the quantitative magnitude of the differences involved. In contrast, the separation index is sensitive to the quantitative magnitude of the differences because it registers the residential outcome of area racial composition in its natural metric. Segregation dynamics such as “tipping,” resulting from group differentials in entries and exits to areas, and discrimination to exclude groups from areas are thought to be triggered by area group proportions. In contrast, theories of segregation dynamics rarely direct attention to rank order position on area racial composition over and above its association with area racial composition itself.

In Chaps. 14, 15, and 16 I give attention to the problem of index bias. All indices of uneven distribution have the undesirable property that their scores are subject to inherent upward bias that can be non-negligible and varies in magnitude across

individual cases. I draw on the difference of means framework to first identify the source of index bias and then identify a solution for obtaining unbiased versions of all popular indices of uneven distribution. I use formal analytic models and empirical exercises to demonstrate that the unbiased versions of G, D, R, H, and S behave as desired in analytic exercises and in empirical applications. Significantly, the difference of means framework is crucial because it provides the vantage point needed to identify both the source of the problem and its solution both of which turn out to be surprisingly simple and intuitive. In this new formulation index scores are calculated as differences of group means on *individual* residential outcomes scored from area racial proportion. The source of bias can be traced to how area racial proportion is assessed from the perspective of individuals. In the “standard” (biased) formulation, the individual in question is *included* in the area counts used to calculate area racial proportion. The value of this residential outcome for the individual thus reflects a combination of two things: the individual’s own contribution to area racial mix and the racial mix of neighbors – the other individuals in the area. Under random assignment the racial mix of neighbors is a random draw and every individual, regardless of group membership has the same expected distribution of outcomes on racial mix of neighbors. In contrast, an individual’s own contribution to area racial mix is fixed and, importantly, differs systematically with group membership. This is the source of index bias. Once seen from this vantage point, the problem of index bias can then be eliminated by assessing area racial proportion for individuals based on *neighbors* instead of *area population*.

The solution to index bias I offer in this monograph is attractive for many reasons. To begin, when working within the difference of means framework for calculating indices of uneven distribution, the solution is simple and intuitive, even “obvious.” Second, the unbiased measures do not require radical changes in research practices. Researchers can continue to use the same measures they have used for decades. But now they can use refined versions of these measures that will yield scores that are free of bias at the level of individual cases in situations where previously researchers could not trust index scores and in other situations the scores of the refined versions will be essentially identical to scores obtained using standard computing formulas. In sum, the new versions will exactly replicate research findings obtained using standard index versions when measurement is non-problematic and will yield superior results when standard calculations cannot be trusted.

In Chap. 17 I offer final comments on the contributions of the monograph overall and reiterate my hope is that the new options for measurement and analysis I introduce here will enable researchers to investigate residential segregation in more detail and depth than has previously been possible. Significantly, the benefits gained from using the new options of measurement and analysis are “cost free”; there are no penalties or sacrifices associated with adopting them. Researchers do not have to put aside familiar measures and replace them with unfamiliar ones. The difference of means framework for measuring segregation permits researchers to exactly replicate results of past studies while at the same time giving them new options for refined measurement, expanded analysis, and attractive substantive interpretations. Thus, researchers can maintain continuity with previous studies of aggregate

segregation while simultaneously having the option of taking advantage of opportunities to analyze segregation in new ways to gain a deeper, more detailed understanding of segregation patterns.

References

- Duncan, O. D., & Duncan, B. (1955). A methodological analysis of segregation indices. *American Sociological Review*, *20*, 210–217.
- Hutchens, R. (2001). Numerical measures of segregation and their desirable properties. *Mathematical Social Sciences*, *42*, 13–29.
- Hutchens, R. (2004). One measure of segregation. *International Economic Review*, *45*, 555–578.
- James, D., & Taeuber, K. (1985). Measures of segregation. *Sociological Methodology*, *13*, 1–32.
- Massey, D. S., & Denton, N. A. (1988). The dimensions of residential segregation. *Social Forces*, *67*, 281–309.
- Reardon, S. F., & Firebaugh, G. (2002). Measures of multi-group segregation. *Sociological Methodology*, *31*, 33–67.
- Stearns, L. B., & Logan, J. (1986). Measuring segregation: Three dimensions, three measures. *Urban Affairs Quarterly*, *22*, 124–150.
- White, M. J. (1986). Segregation and diversity: Measures of population distribution. *Population Index*, *65*, 198–221.
- Zoloth, B. S. (1976). Alternative measures of school segregation. *Land Economics*, *52*, 278–298.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial 2.5 International License (<http://creativecommons.org/licenses/by-nc/2.5/>), which permits any noncommercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

