Social Science & Medicine 203 (2018) 64–73

Contents lists available at ScienceDirect

Social Science & Medicine

journal homepage: www.elsevier.com/locate/socscimed

A multilevel approach to modeling health inequalities at the intersection of multiple social identities

Clare R. Evans[a,*], David R. Williams[b], Jukka-Pekka Onnela[c], S.V. Subramanian[b]

[a] Department of Sociology, University of Oregon, Eugene, OR, United States
[b] Department of Social and Behavioral Sciences, Harvard T.H. Chan School of Public Health, Boston, MA, United States
[c] Department of Biostatistics, Harvard T.H. Chan School of Public Health, Boston, MA, United States

ARTICLE INFO

Keywords:
Health inequalities
Intersectionality
Eco-epidemiology
Multilevel modeling

ABSTRACT

Rationale: Examining interactions between numerous interlocking social identities and the systems of oppression and privilege that shape them is central to health inequalities research. Multilevel models are an alternative and novel approach to examining health inequalities at the intersection of multiple social identities. This approach draws attention to the heterogeneity within and between intersectional social strata by partitioning the total variance across two levels.
Method: Utilizing a familiar empirical example from social epidemiology—body mass index among U.S. adults (N = 32,788)—we compare the application of multilevel models to the conventional fixed effects approach to studying high-dimension interactions. Researchers are often confronted with the need to explore numerous interactions of identities and social processes. We explore the interactions of five dimensions of social identity and position—gender, race/ethnicity, income, education, and age—for a total of 384 unique intersectional social strata.
Results: We find that the multilevel approach provides advantages over conventional models, including scalability for higher dimensions, adjustment for sample size of social strata, model parsimony, and ease of interpretation.
Conclusion: Considerable variation is attributable to the within-strata level, indicating the low discriminatory accuracy of these intersectional identities and the high within-strata heterogeneity of risk that remains unexplained. Multilevel modeling is an innovative and valuable tool for evaluating the intersectionality of health inequalities.

1. Introduction

Intersectionality is a theoretical framework that is increasingly used to study the patterning of health inequalities because of its focus on the multidimensional, multiplicative nature of disadvantage (Bowleg, 2012; Farmer and Ferraro, 2005; Schulz and Mullings, 2006; Veenstra, 2011; Warner and Brown, 2011), which corresponds with discipline-specific theories such as fundamental causes (Link and Phelan, 1995) and ecosocial theory (Krieger, 2011). Intersectionality theorists posit that inequalities are generated by numerous interlocking systems of privilege and oppression such as racism, classism, sexism, and ageism (Collins, 1990; Crenshaw, 1989; McCall, 2005), and push back against the "additive approach," which treats the advantages or disadvantages conferred through simultaneous occupation of multiple social positions as simply accumulated. Care must be taken in the adoption of intersectionality by public health researchers, however, to ensure that it is properly framed within the context of ongoing debates in epidemiology—namely between the so-called "risk factor" epidemiology and "eco-epidemiology" (Susser and Susser, 1996). Conventional approaches to quantitative intersectionality analysis have also presented several methodological limitations, including issues of scalability, model parsimony, small sample size, and interpretability of results.

In this study, we explore an alternative analytic approach (Evans, 2015; Green et al., 2017; Jones et al., 2016) that resolves some of the key theoretical and methodological tensions inherent to this research. This approach involves applying hierarchical, multilevel models to study large numbers of interactions and intersectional identities while partitioning the total variance between two levels—the *between-strata* (or between category) level and the *within-strata* (or within category) level. This analytic approach is a valuable tool for *exploring* multiple interactions simultaneously and the patterning of inequalities across society. We apply and compare this new approach with the

* Corresponding author.
E-mail address: cevans@uoregon.edu (C.R. Evans).

https://doi.org/10.1016/j.socscimed.2017.11.011
Received 26 January 2017; Received in revised form 23 October 2017; Accepted 5 November 2017
Available online 30 November 2017

conventional, fixed effects approach to interaction models. To demonstrate its potential application in health research, we explore an empirical example that will be familiar to many social epidemiologists—body mass index among U.S. adults.

1.1. Risk factor epidemiology and eco-epidemiology—Framing intersectionality

McCall (2005) has identified three distinct orientations within the current intersectional literature—the *intercategorical*, the *intracategorical* and the *anti-categorical*. The anti-categorical approach involves the critique and deconstruction of analytic categories. The intracategorical approach tends to "focus on particular social groups at neglected points of intersection … in order to reveal the complexity of lived experiences within such groups" (p.1774). The approach to intersectionality most often adopted in social epidemiology is the intercategorical approach because of its natural fit with quantitative analyses of inequalities. The intercategorical approach involves the provisional adoption of analytic categories to document inequalities among groups and explore the interactions between different dimensions of identity, position, and social processes. The conventional intercategorical approach to studying interactions, which we will refer to as the *fixed effects* approach, involves fitting a single-level regression model with a full complement of parameters to account for all points of interaction. This can be accomplished either by using a full set of dummy variables (one representing each possible combination of social identity and position, e.g., young college educated high income black woman) or by including main effects and a saturation of interaction terms. Mathematically these approaches are equivalent, and so we refer primarily to the version with interaction parameters.

When using intercategorical intersectionality in population health research it is critical to correctly situate this framework within existing debates in epidemiology. Rose (1992) famously distinguished between *causes of individual cases* (i.e., why did *this* person get sick with *this* illness at *this* point in time) and *causes of population incidence* (i.e., what caused *this* population to have a higher disease incidence than *that* population). Causes at these distinct levels may or may not resemble each other. Susser and Susser (1996) expanded on this to differentiate *risk factor epidemiology* with its focus on identifying causes of cases from *eco-epidemiology,* which takes a multilevel perspective and considers causal pathways ranging from the societal level to the molecular level. As others have noted (Merlo and Wagner, 2012; Merlo, 2014), this distinction is not always appreciated in modern epidemiology.

The risk factor approach in epidemiology involves the identification of risk factors through the comparison of group averages. In practice the use of variables such as gender, race/ethnicity, and socioeconomic status (SES) in quantitative intersectionality research may make it appear that the mission of intersectionality research corresponds with the risk factor approach, and involves identifying ever-narrower and more specific "risky identities" that are particularly burdened by health inequalities. This is, however, diametrically opposed to a central tenet of intersectionality—namely, that intersectionality does not situate the problems associated with particular identities within individuals or the identities themselves, but within the structural power hierarchies, social processes, and social determinants that shape the social experiences of individuals with those intersectional identities. While categorical variables (gender, race, class) may be used in regression models, care should always be taken to recognize that these may be intended as proxies for the interactions of systems of oppression (sexism, racism, classism) and other social processes in producing population-level incidence (Bauer, 2014).

In ecosocial theory, Krieger (2011) theorizes health inequalities between populations as resulting from numerous interacting *pathways of embodiment* across the life course, through which we come to "literally incorporate, biologically … the material and social world in which we live" (p. 214). Ecosocial theory encourages a broad vision for the determinants of health inequalities—including both the interlocking systems of oppression and privilege (sexism, racism, classism) implicated by intersectionality and *other* social processes. For instance, Krieger points to issues of social and economic deprivation, environmental hazards, and the targeted marketing of harmful commodities to low income populations as key pathways of embodiment, which may not readily be classified as forms of intersectional "classism" per se.

Eco-epidemiologists have argued strongly against the "tyranny" (Merlo and Wagner, 2012) of comparing group averages, both because it risks framing inequalities as individual-level issues resolvable by individuals (resulting in "blaming the victim") and because such averages obscure the relatively low predictive power of these labels to distinguish between cases and non-cases. In other words, risk factors typically are unable to discriminate between individuals who will become sick and those who will not (Merlo, 2014), which should caution all of us to frame intercategorical intersectionality research in the health inequalities domain as explicitly eco-epidemiologic.

Paradoxically, as Merlo (2014) has noted, many existing eco-epidemiologic studies continue to utilize a framework reliant on comparing group averages—though admittedly these new risk factors are situated at higher contextual levels, such as comparing neighborhood averages. Eco-epidemiologic approaches should balance consideration of group averages with what Merlo (2014) has called a "multilevel analysis of individual heterogeneity"—or a multilevel examination of variation within and between groups. The approach presented here is explicitly framed with this intention and allows for consideration of both group averages and multilevel variation within and between groups.

As a brief aside, we will henceforth refer to these points of intersection as "social strata" rather than as categories or groups. The intersectionality literature has encouraged us all to become more skeptical of the reification of categorical labels, and therefore we feel that the term "strata"—which alludes to stratified analyses—implies *provisional* acceptance of labels for the purposes of studying inequalities, while remaining aware of the inherent danger in treating social labels as monolithic, unchanging, and inflexible. Similarly, we sometimes use the word "identity" as a shorthand to refer to dimensions of identity, position, and resources. We do not mean to imply that income, for instance, is best understood as a social identity.

1.2. Theoretical and empirical motivation

The conventional fixed effects approach to operationalize intercategorical intersectionality is open to two related theoretical criticisms and poses additional empirical concerns. First, including interaction terms encourages us to only study the intersectionality of *marginalization*. For instance, in a comparative quantitative intersectional analysis where white males are the reference group, we might include main effects for "black" and "female" and an interaction term for "black *and* female." Following current standards, finding this interaction term to be statistically significant would be interpreted as support for the interaction of racism and sexism. However, this setup enables us to only evaluate the interaction effect experienced by black women, while those experiencing multiple privileges (e.g., white men) or mixes of privilege and disadvantage (e.g., white women and black men) are treated as having no observable interaction effect. Theorists have voiced this concern and called for consideration of the points of intersection that mix privilege and marginalization (Bauer, 2014; Choo and Ferree, 2010; Hancock, 2007; Nash, 2008). While studying intersectionality of privilege could be accomplished by switching the reference group to "low SES black females" or by constructing alternative post hoc analyses, ideally, we would be able to determine simultaneously whether *all* intersectional identities exhibit evidence of an interaction (or intersectional) effect above and beyond the contributions of the additive main effects.

In other words, to examine whether a given social stratum shows evidence of an intersectional interaction effect, we would want to compare what is observed for that stratum with what *might have been expected for it based on the additive contributions of the main effects*. Establishing the magnitude and direction of this unique "interaction effect" for each social stratum is also of interest because a simplistic reading of intersectionality might imply that this interaction effect will in some way reflect the number of marginalized or privileged identities. For instance, a naïve reading of foundational works of intersectional scholarship, with their focus on race- and gender-based discrimination (Collins, 1990; Crenshaw, 1989), might lead some to conclude that intersectionality implies that possessing *more* marginalized identities will necessarily result in a *more* harmful interaction effect. The intersectionality of privilege might be assumed to work in the opposite direction—enjoying multiple interlocking privileges will result in an interaction effect that is more beneficial. This interpretation of intersectional thought has been refuted by most scholars since the early days of the approach (e.g., King, 1988). However, researchers such as Bauer (2014), who have called for greater attention to identities that mix privilege and marginalization, have highlighted our continued uncertainty about the effects of possessing a combination of marginalized and privileged identities and social positions when it comes to the social patterning of outcomes such as health inequalities.

It is important to acknowledge that our framing of quantitative, intercategorical intersectionality falls into what some intersectionality scholars have called "intersectionality as testable explanation" (Hancock, 2013), in that it involves an assessment of whether statistically significant interaction effects are detectable. Though becoming more common, this framing of intersectionality remains contentious, with some scholars arguing that intersectionality should be considered as more akin to an analytic tool to be utilized rather than a hypothesis that can be tested. The approach outlined here is also useful as a tool for *exploratory* analysis of inequalities, and does not necessarily require this hypothesis testing of interaction terms.

The second theoretical criticism of the fixed effects approach is that consistently comparing the multiply marginalized to the multiply privileged runs the risk of reinforcing the notion that the dominant, privileged group often used as the reference category (e.g., high SES white males) is the standard against which all other groups *ought* to be compared (Choo and Ferree, 2010). In other words, using the multiply privileged as the yardstick against which marginalized groups are measured reinforces the social primacy of the privileged as the "default" category. While to some extent comparisons to more privileged groups are inherent to the project of documenting inequalities, we would ideally be able to make multiple comparisons simultaneously—rather than relying only on comparisons to one privileged identity.

The fixed effects approach has yielded powerful insights into the patterning of health inequalities across society. However, the ever-increasing demands to examine interactions between dimensions of identity beyond just race and gender (McCall, 2005; Nash, 2008) has pushed researchers up against some of the *methodological limitations* of this approach. Namely, the fixed effects approach to interactions struggles with issues of scalability, model parsimony, reduced sample size in some intersectional strata, and occasionally, issues of interpretability.

From an empirical standpoint, including many dimensions of social identity in interaction models creates a set of additional modeling and interpretation challenges. As the number of dimensions of social identity considered increases, the number of parameters required in a fixed effects model increases *geometrically* to allow for all combinations of first-order, second-order, and higher-order interaction terms. While this poses minute problems when fewer dimensions of identity are interacted, model parsimony and fit does become a concern at higher dimensions. The geometric increase of parameters to interpret can also complicate the examination of results. This is especially true given the non-ideal comparisons built into the model setup described above.

Additionally, as a given sample size is parsed across more intersectional identities the issue of insufficient sample size in many social strata becomes a concern. To illustrate most clearly, imagine a model using the dummy variable fixed effects approach wherein each intersectional identity has its own indicator variable. In a linear regression with a model fully saturated in this way, we might obtain the mean expected outcome for each identity, although this mean is calculated for some groups based on only a handful of respondents. It is then up to the researcher to weed out (often by hand) those intersectional groups of insufficient sample size.

In the recent, seminal work by Jones et al. (2016) they demonstrate the novel use of multilevel models to study high-dimensional interactions in multivariate models. Applying a hierarchical random effects model, individuals are nested within what they term multivariate "contingency tables" but what we might call intersectional social strata, and apply this to study voting (and abstaining) patterns in the 2015 UK general election. A similar and explicitly intersectional approach was proposed by Evans (2015) and Green et al. (2017). These early examples of this approach build on a growing interest in using multilevel models to study variation within and between social groups (Merlo, 2014; Merlo et al., 2016). We propose applying this novel approach to interactions to study intersectional social identities in the domain of health inequalities, and offer an illustration of this alternative method and the ways in which it addresses the criticisms and resolves the tensions we mentioned above.

2. Method

2.1. Data

The National Epidemiologic Survey on Alcohol and Related Conditions (NESARC) was a longitudinal study launched in 2001 by the National Institute on Alcohol Abuse and Alcoholism. It was designed to include a representative sample of the U.S. non-incarcerated civilian population, including citizens and non-citizens, aged 18 years and older who are residing in the U.S. In this study, we used data from Wave 2, collected between 2004 and 2005 (Grant and Kaplan, 2005). The large sample size and intentional oversampling of young adults, Hispanics, and non-Hispanic blacks means that the data are sufficiently diverse to provide a reasonable sample size within most strata.

2.2. Outcome: body mass index

Body mass index, calculated as body weight in kilograms divided by height in meters squared, is commonly used to classify individuals as underweight, normal weight, overweight, and obese. In NESARC the respondents' height and weight were elicited through self-report. BMI is a useful epidemiologic, population-level indicator of weight status and its social patterning reflects the embodiment of a broad range of social conditions and inequalities (Berkman and Kawachi, 2000; Krieger, 2011; Link and Phelan, 1995).

BMI was the empirical application of choice to showcase this approach for three reasons. First, BMI is a continuous measure with a roughly Gaussian distribution, enabling the use of linear models without transformations of the dependent variable. Second, BMI is well studied in the social epidemiologic literature, enabling the results of these models—though presented in a less familiar format—to be recognizable to researchers. Third, the so-called obesity epidemic has disproportionately affected members of society along a range of social identities, with women, racial/ethnic minorities, older, and low SES populations tending to be disproportionately affected (Clarke et al., 2009; Ogden et al., 2012). Thus, we can illustrate the application of this approach when many social strata are included in the analysis.

2.3. Dimensions of social identity and social strata

We used five dimensions of social identity to construct the intersectional social strata in this study—*gender, race/ethnicity, education, income*, and *age*. These identities were self-reported. Missing demographic data were imputed by the Census Bureau, and are included in the publicly available data set. For a detailed description of imputation methods used, see Grant et al. (2003).

Gender. Interviewers who conducted the survey were instructed to ask the respondent what their sex was, if the sex of the respondent was "*not apparent*." Respondents were given the option of Male or Female. In both measurement and description throughout the NESARC documentation, sex and gender appear to have been conflated—a common and problematic issue (Oakley, 1972). We use the term *gender* under the assumption that gender was more likely to have been observed and reported.

Race/Ethnicity. Respondents were asked, "*Are you of Hispanic or Latino origin*?" and were instructed to "*select 1 or more categories to describe your race*" with the options: American Indian or Alaska Native, Asian, Black or African American, Native Hawaiian or Other Pacific Islander, and White. When a respondent reported more than one race or ethnicity a Census Bureau algorithm to code a single racial/ethnic category was used. Only three racial/ethnic categories are included in this analysis—(1) white, not Hispanic or Latino, (2) black, not Hispanic or Latino, and (3) Hispanic or Latino—because other categories had insufficient sample size for this analysis. For simplicity, we refer to these categories as white, black, and Hispanic.

Education. Education was measured in Wave 1 with fourteen response categories that we condensed into four: (1) *less than high school* (≤11 years), (2) *completed high school* (or equivalency), (3) *some college no degree*, (4) *college degree or more*. These categories were selected for two reasons: first, they reflect socially significant distinctions in educational achievement, and second, it partitioned the full sample into roughly equivalent groups.

Income. Total pre-tax annual family income was measured for the previous 12 months. We created a per capita income measure by dividing annual family income by the number of related persons in the household (including the respondent and children). For simplicity, we refer to this variable simply as *income*. Twenty-one income categories were condensed to four based on percent of poverty threshold in 2000 (U.S. Census Bureau, 2012)—low-income (below 100%), low-middle-income (100%–199%), high-middle-income (200%–399%), and high-income (400% or more). Cutoff values were derived based on estimates for a single person under age 65.

Age. Age was self-reported and recorded as ranging from 18 to 90 or older. Though measured continuously, we partition age into four categories— (1) 18–29 years, (2) 30–44 years, (3) 45–59 years, (4) 60 years and older. These categories were determined based on two factors—a need to distinguish socially meaningful age categories balanced against the benefits of distributing the sample roughly equally between the four categories.

2.4. The multilevel approach

Multilevel models partition the residual variation in a model into within-group (level 1) and between-group (level 2) variation (Raudenbush and Bryk, 2002). Typically, multilevel approaches are used to model the clustering of respondents by some observable context (e.g., children clustered by neighborhood, school, and peer network; Evans et al., 2016), to model panel data, or when we wish to compensate for artifacts of the data collection process (e.g., cluster-based sampling). However, multilevel models are capable of handling clustering from more abstract sources. Clustering occurs when individuals share something that creates *similarity* between them and ignoring this clustering would violate the regression assumption of *independence*. While the clustered individuals may share something concrete—like a neighborhood—they may also share something abstract, like a common set of social exposures associated with their intersectional social identities. In other words, they occupy the same social stratum as defined by their gender, race/ethnicity, income, education, and age.

In this analysis, we employed a two-level hierarchical random intercepts model with individuals (level 1) clustered within social strata (level 2). Each social stratum is assigned a unique identifying number and a separate stratum is defined for every combination of social identities and positions considered. To control for the 'additive' aspects of the social identities, main effect predictors such as 'female' are included in the model as fixed effects. A brief side note on notation: Contrary to most multilevel models where social categories such as race and income are individual-level covariates, here such covariates are properties of the strata level (Supplemental Technical Note 1). *Critically, no interaction terms between these independent variables are included.* The stratum-level residual for each stratum encompass the entirety of the interaction effect, and it is therefore possible to determine the extent to which the mean BMI differs from what was 'expected' based on the contributions of the additive main effects for each stratum.

In its most general form, the linear model is specified as:

$$y_{ij} = \beta\gamma_j + \mu_{0j} + e_{0ij}$$

$$[\mu_{0j}] \sim N(0, \sigma^2_{strata}) \qquad \text{Level 2}$$

$$[e_{0ij}] \sim N(0, \sigma^2_{e0}) \qquad \text{Level 1}$$

Where y_{ij} is the value of the outcome for respondent i in stratum j. γ_j is a vector of the intercept and main effect predictors for stratum j (e.g., female) and β is a row vector of associated parameter values. The difference between the average value of the outcome in stratum j and the expected value of y based on the main effects is given by the strata-level residual μ_{0j}, which is normally distributed with mean 0 and variance σ^2_{strata}. The difference between the value of the outcome for respondent i in stratum j and his or her stratum's average is given by e_{0ij} which is also normally distributed with mean 0 and variance σ^2_{e0}. After adjustment for the main effects, and assuming no omitted variable bias, the stratum-level residuals (μ_{0j}) capture the unique interaction effect for each social stratum. The assumption about omitted variables is an important one: we caution that the strata-level residual estimates may not be a *direct* measure of "intersectionality" because this unexplained residual at the strata-level may encompass both interaction effects and other sources of variation, such as the effects of unmeasured dimensions of social identity.

A useful feature of multilevel models is that estimates obtained for each social stratum ultimately account for the sample size observed in that stratum. In a case where the sample size for a stratum is very small the fixed effect estimate (e.g., mean) for that stratum would become less stable (i.e., easily influenced by outlier observations) and should be less trusted. The estimate for stratum j in the multilevel model (β_{0j}) is obtained by a weighted combination of the estimate for stratum j from a fixed effects model (β^*_{0j}) and the grand multilevel mean (β_0):

$$\beta_{0j} = w_j\beta^*_{0j} + (1 - w_j)\beta_0$$

Thus, as the weight (w_j) decreases, the multilevel estimate for stratum j (β_{0j}) more closely resembles the multilevel grand mean (across all strata) and moves away from the fixed effect estimate for that stratum. The weight w_j is determined by:

$$w_j = \frac{\sigma^2_{strata}}{\sigma^2_{strata} + \left(\sigma^2_{e0}/n_j\right)}$$

where n_j is the sample size observed in stratum j. Thus, as the sample size for stratum j decreases, the weight w_j also decreases and the multilevel estimate for stratum j relies more on what is known about the sample overall (i.e., the grand mean) and less on the sparse and possibly biased observations for that stratum. While similar weighting can be accomplished in a fixed effects model this process is automatically included in the multilevel estimation procedure.

Another advantage of the multilevel approach is its parsimony as interacting dimensions of identity are added. While a fixed effects model would grow *geometrically* as dimensions of identity are added, the multilevel model would grow *linearly* because each additional dimension of identity added (e.g., education) will only require adding sufficient main effect parameters to cover the desired categories of education. In a fixed effect model, it would be necessary to add the main effects for education categories as well as all additional first-, second- and higher-order interaction terms. In this analysis, the five dimensions of social identity considered have the following number of categories each: gender (2), race/ethnicity (3), education (4), income (4), and age (4). Thus, the main effects are treated as categorical—there are 12 dichotomous main effect predictors included in the model. The individuals at level 1 are nested within the strata ($N = 384$) at level 2. The analogous fixed effects model would require 384 main effect parameters. This feature of multilevel models explains why this approach scales better.

The fixed effects and multilevel approaches can be compared using Bayesian Information Criterion (BIC) (Schwarz, 1978), a standard criterion based on the likelihood function that takes into account both model fit and parsimony. Holding constant the number of observations in a data set, the more parameters added to the model, the more the model's BIC is penalized. To illustrate this point, we conducted a supplemental simulation analysis comparing the two modeling approaches across a wide range of numbers of strata—from 4 to 384 social strata. See Supplemental Fig. 1 and its accompanying description for details.

In this study's empirical analysis, two multilevel models were fit for each outcome—a *null model* and a *main effects model* (which included the additive parameters). In the null model, the between-strata variance term (σ^2_{strata}) represents the *total* amount of variability *between* strata, including that contributed by the main effects. In the main effects model, the same parameter is what remains of the between-strata variation *after the additive parameters are adjusted for*. In this case, a large between-strata variance would imply that—for most strata—the main effects do an inadequate job of capturing the embodiment of the outcome of interest and that there is a significant interaction between the dimensions of identity. In contrast, a small between-strata variance would imply that main effects do a relatively good job of explaining between-strata differences. Dividing the strata-level variance term from the main effects model by the strata-level variance term from the null model provides an estimate for each outcome of the *proportion of the between-strata variability that is unexplained by the additive main effects*.

2.5. Empirical analysis

First, the strata-specific values of mean BMI were obtained using both the fixed effects and multilevel approaches and then graphically compared to determine whether the multilevel model provides more stable estimates for strata with small sample sizes. All multilevel analyses were conducted in MLwiN version 2.26 (Rasbash et al., 2012) using Bayesian Markov Chain Monte Carlo (MCMC) estimation procedures (Browne, 2009). See Supplemental Technical Note 2 for details. The regression models were first fit using iterative generalized least squares (IGLS) estimation to provide the Bayesian MCMC procedure with initialization values; non-informative priors were used in all analyses. The fixed effects version of the model was fit using IGLS. Then, patterns that emerged in the strata-level residuals (μ_{0j}) and between-strata variance parameters (σ^2_{strata}) were examined, with particular attention to how this approach yields results that are easily interpretable.

3. Results

From the initial sample size of $N = 34{,}653$ in Wave 2, 330 respondents were excluded because of missing responses to the height and weight items, and a further 1535 were excluded because they were not classified in the three race/ethnicity categories considered. Thus, the final sample includes 32,788 respondents. The demographic profile of the sample appears in Table 1. For details on the sample size in each intersectional social stratum, see Supplemental Table 1.

Table 1
Descriptive characteristics of sample.

	Frequency *N*	BMI Mean (Std)	% Obese BMI ≥ 30
TOTAL	32788	27.90 (6.04)	29.49
Gender			
Male	13840	27.96 (5.11)	27.16
Female	18948	27.85 (6.64)	31.20
Race/Ethnicity			
White Non-Hispanic	19955	27.26 (5.78)	25.49
Black Non-Hispanic	6526	29.50 (6.67)	40.58
Hispanic/Latino	6307	28.25 (5.83)	30.70
Education			
Less than high school	5243	28.40 (6.35)	33.09
Completed high school	9038	28.41 (6.29)	33.07
Some college no degree	7043	27.93 (6.07)	29.99
College degree or more	11464	27.24 (5.61)	24.72
Income (% Poverty Threshold in 2000)			
Low income (Below 100%)	7666	28.57 (6.89)	33.89
Low-middle income (100%–199%)	9144	28.01 (6.06)	31.20
High-middle income (200%–399%)	9548	27.71 (5.67)	27.72
High income (400% or more)	6430	27.20 (5.35)	24.45
Age			
18–29 years	4628	26.96 (6.02)	24.68
30–44 years	9975	28.06 (6.18)	30.45
45–59 years	9148	28.63 (6.20)	33.52
60 + years	9037	27.44 (5.63)	26.83
US Region			
South	12485	27.92 (6.08)	29.80
Northeast	5757	27.90 (6.09)	29.29
Midwest	6190	28.05 (6.01)	30.58
West	8356	27.74 (5.97)	28.37

Note: Estimates are unadjusted for disproportionate sampling and therefore represent sample estimates (unadjusted sample proportions), not population estimates.

The number of observations per stratum is important to consider in this analysis. While the multilevel approach adjusts estimates for strata means depending on the sample size, it is still necessary to ensure that a sufficient *proportion* of strata have a reasonable number of observations. If too few strata had sufficient sample sizes then most of the strata considered would have estimates weighted towards the global multilevel mean, which would shrink the magnitude of the between-strata variance parameter (σ^2_{strata}). In this case, the sample size is sufficiently large and well distributed across the strata to prevent significant underestimation of the between-strata variability. Of the 384 strata considered in this analysis, 382 have at least one observation, and most have reasonable sample sizes (Supplemental Table 2); 81% of strata have 20 or more respondents and over 90% have 10 or more respondents. Predictably, those strata with the fewest observations represent particularly unlikely combinations, such as minorities with both low education and high income.

The mean BMI for each stratum was calculated using both modeling approaches and histograms comparing the distribution of these results are presented in Fig. 1. As a useful alternative visual, logistic versions of the models were fit as well for obesity (BMI $\geq$30 kg/m^2) and the proportion of respondents classified as obese in each stratum were calculated and plotted in histograms as well. As is apparent visually, the different handling of strata sample size between the multilevel and fixed effects approaches results in the distribution of values estimated using the fixed effects approach being much wider with more extreme values. For instance, the small sample size in some strata resulted in undoubtedly biased estimates for those strata—including strata with estimates of 0% and 100% obese (Fig. 1). The multilevel approach, however, provides more realistic estimates for all strata—including those with few respondents.

The results of multilevel models are presented in Table 2. It is first important to note that BMI is patterned as expected in the population.

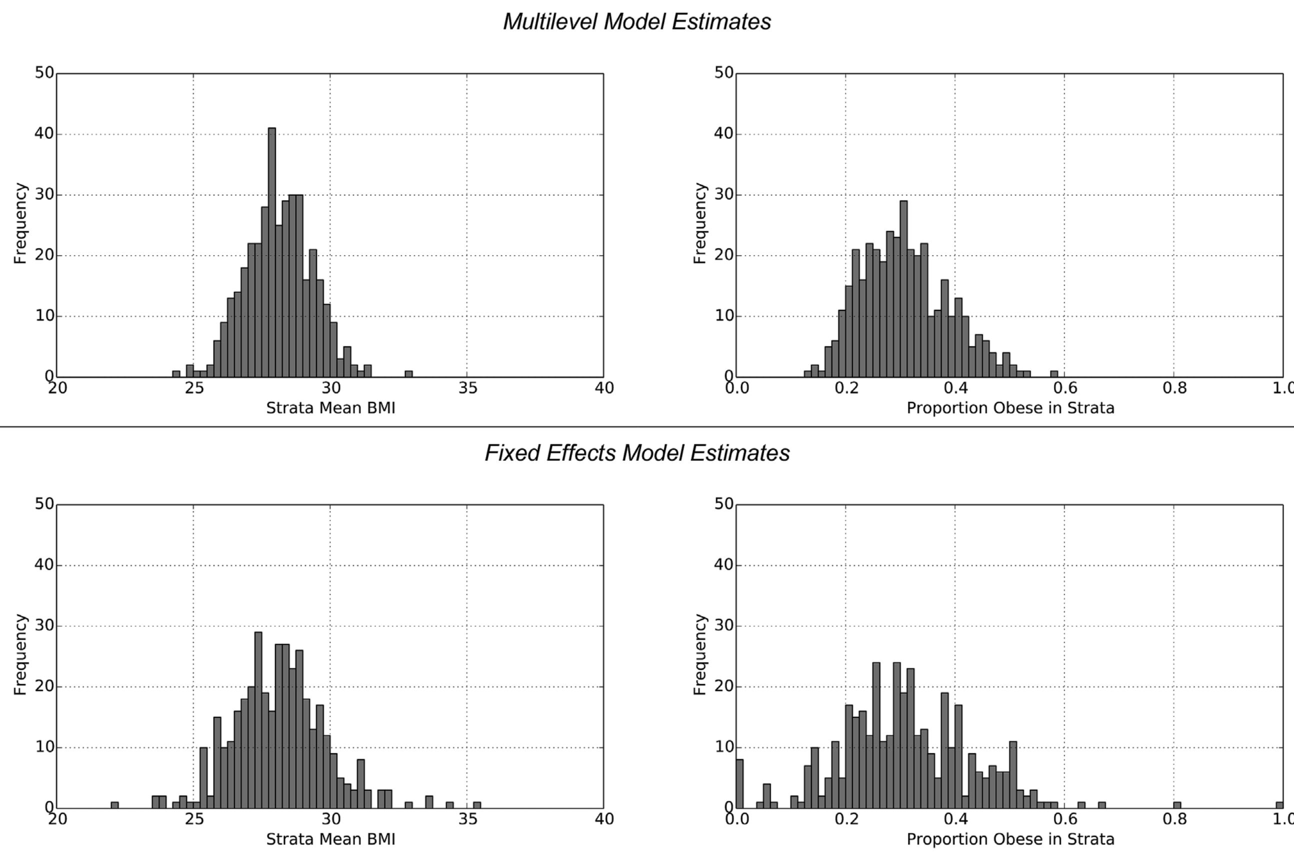

Fig. 1. Strata-Level Predictions of Mean Body Mass Index (BMI) and Proportion Obese obtained using Multilevel approach versus Fixed Effects approach.

Belonging to a social stratum with more marginalized dimensions of identity is associated with a higher mean BMI in absolute terms. See Supplemental Table 1 for stratum-specific mean BMIs and 95% credible intervals predicted by the multilevel approach. This table allows for free comparisons of *expected* mean BMIs (based on the additive main effects) and the *predicted* mean BMIs (inclusive of the additive main effects and strata-specific residuals). While these results speak to the patterning of BMI in *absolute* terms, the question remains whether inequalities are patterned in such a way that significant *interaction* effects exist.

Comparing the magnitude of the strata-level residuals between null and main effects models, we see that approximately 35% of between-strata variability is unexplained by the main effects for BMI. The magnitude of the strata-level residuals in many cases is considerable. For instance, the effect on predicted BMI of belonging to a stratum that is of the highest income level compared with the lowest is -0.584 kg/m^2. Of the 382 strata evaluated, 22.8% experience a strata-level residual of this magnitude or larger. Thus, a substantial variation in mean BMI between strata remains unaccounted for by the additive effects of the five social identity dimensions considered.

We provide an alternative way to visualize the multilevel model results by condensing the five dimensions of identity to three and plotting the strata-level residuals (μ_{0j}) for each combination of gender, race, and high/low SES. In this Fig. 2, a residual of zero indicates that a social stratum experiences precisely the mean BMI predicted by the additive main effects. Positive values represent higher mean BMI than expected based on the additive main effects, while negative residual values indicate lower mean BMI than expected.

While some multiply marginalized social strata do tend to experience higher mean BMI than predicted by the additive main effects (i.e., positive values for strata-level residuals), this is not universally true. Additionally, social strata enjoying multiple privileged identities, such as young white males of high income and education, are not universally experiencing weight-status dividends above and beyond what we might expect for them based on the main effects (i.e., negative values for strata-level residuals). Within the super-category for white males of high income, for instance, markers indicating the strata-level residuals for strata of varying education and age levels show that white males of high income tend to experience *higher* mean BMI than expected based on additive main effects. In other words, while these privileged strata do enjoy lower mean BMI in *absolute* terms, they do not appear to leverage their multiply privileged identities for a beneficial *interaction* effect.

Low income black females, as predicted based on their multiple marginalized identities, do experience higher mean BMI than expected based on additive main effects. Some strata with mixed privileged and marginalized identities—low income black males, high income white females, and high income Hispanic females—tend to experience *lower* mean BMI than expected. These results indicate that the *number* of marginalized identities out of the five considered does not appear to linearly predict magnitude or direction of unexplained strata-level residuals, as the naïve interpretation of intersectional thought might have implied. Furthermore, whatever intersectional *experiences* of privilege the multiply advantaged enjoy, it does not consistently translate into interaction effects that reduce predicted BMI.

Another key result is that in the null model we can explicitly decompose the total variance into within-strata variance (σ^2_{e0}=34.5) and between-strata variance (σ^2_{strata}=1.8). We find that approximately 5% of the total variance is attributable to between-strata differences. This finding supports the argument made by Merlo (2014) and others that while strata may differ from each other in terms of relative risk, there is still remarkably low discriminatory accuracy of these identity labels and exceptionally high within-strata heterogeneity that remains unexplained.

Table 2
MCMC parameter estimates for the two-level hierarchical Bayesian linear regression model of body mass index (kg/m^2).

	Null Model Estimate (95% CI)	Main Effects Model Estimate (95% CI)
Fixed Effects		
Intercept	28.126 (27.965, 28.293)	26.858 (26.433, 27.288)
Gender		
Male (reference)		–
Female		0.081 (−0.149, 0.316)
Race/Ethnicity		
White Non-Hispanic (reference)		–
Black Non-Hispanic		1.791 (1.511, 2.066)
Hispanic/Latino		0.659 (0.383, 0.941)
Education		
Less than high school (reference)		–
Completed high school		0.087 (−0.255, 0.433)
Some college no degree		−0.240 (−0.591, 0.123)
College degree or more		−0.813 (−1.167, −0.460)
Income (% Poverty Threshold in 2000)		
Low income (Below 100%) (reference)		–
Low-middle income (100%–199%)		−0.066 (−0.370, 0.245)
High-middle income (200%–399%)		−0.258 (−0.574, 0.060)
High income (400% or more)		−0.584 (−0.953, −0.210)
Age		
18–29 years (reference)		–
30–44 years		1.282 (0.944, 1.624)
45–59 years		1.814 (1.477, 2.152)
60 + years		0.523 (0.184, 0.862)
Random Effects		
Strata	1.823 (1.503, 2.196)	0.643 (0.488, 0.826)
Individual	34.506 (33.984, 35.035)	34.511 (33.977, 35.047)
Percent of Between-Strata Variation Unexplained by Main Effects		**35.272%**

Note: 95% Credible Intervals in parentheses. *P*-values are associated with frequentist approaches and are not available for Bayesian estimations.

4. Discussion

In this study, we have demonstrated a new approach for modeling health inequalities between multiple intersectional social identities. The multilevel approach has several advantages over the conventional fixed effects approach when the number of interactions becomes large, summarized in Table 3. First, multilevel models present a more parsimonious approach because they grow *linearly* as opposed to *geometrically* as new intersectional identities are added. The scalability of the multilevel approach enables us to answer calls from social epidemiologists and intersectional theorists (Bowleg, 2012; McCall, 2005; Nash, 2008) for greater consideration of how numerous dimensions of social identity interact.

Second, multilevel estimation automatically adjusts strata estimates based on the sample size observed. This enables a more realistic estimate to be obtained for social strata with few respondents. Importantly, however, obtaining estimates of a total "intersectional effect" (i.e., stratum-level residual μ_{0j}) in strata with very small sample sizes remains infeasible using the new approach.

Third, the multilevel approach brings the methods used to study interactions closer into line with what intersectional theorists and social epidemiologists have advocated—the assessment of intersectional effects for all social identities, including those that mix privilege and disadvantage (Bauer, 2014; Nash, 2008).

Fourth, the multilevel approach provides an alternative way to understand and visualize the results of high-dimensional intersectional analyses. While a fixed effects approach is readily able to assess the extent to which interactions are statistically significant, the multilevel approach provides a summary statistic of the magnitude of this "intersectional effect" across social strata—σ^2_{strata}. This between-strata variance parameter, when compared in the null and main effects versions of the multilevel model, enables an assessment of the total between-strata variance that remains *unexplained* by the additive main effects. If the entirety of the between-strata variance were accounted for by the main effects, then this would be reduced to 0%. As it is, we found that 35% of strata-level variation in BMI remains unexplained after adjustment for main effects. This finding confirms the intersectional hypothesis that the unique social experiences, norms, opportunities, and sources of resilience that social strata possess are reflected in the non-additive social patterning of the obesity epidemic.

Furthermore, visualizations of the strata-level residuals such as those in Fig. 2 where social strata are condensed into "super-categories" characterized by gender, race/ethnicity, and SES, illustrates clearly the substantial variation *within* these super-categories. Intersectionality theorists' calls for consideration of within-group variation and further parsing by additional relevant dimensions of social identity are justified when studying the obesity epidemic, and likely for other health outcomes as well. These visualizations also enable the detection of general patterns in the intersectional results. For instance, while estimates of mean BMI in absolute terms follow the general patterns predicted in the literature, with multiply marginalized social identities experiencing higher BMI on average and multiply privileged identities experiencing lower BMI on average, the *interaction effects* are patterned in less obvious ways. Both the multiply privileged high SES white male identities and multiply marginalized low SES black female identities tend to experience higher average BMI *than expected* based on the additive contributions of the main effects. High income white females and Hispanic females and low income black males, on the other hand, tend to experience lower average BMI *than expected*.

Finally, the multilevel approach to intercategorical intersectionality research moves us closer to an explicitly eco-epidemiologic take on health inequalities. The conventional fixed effects approach to intercategorical intersectionality produces only group level averages or odds ratios with no explicit modeling (or recognition) of the residual (and typically very large) variation *within* intersectional identities that remains unexplained. Though the multilevel approach presented here does not explicitly model causal pathways ranging from the societal to the molecular level, it does allow for more explicit consideration of variation at the within-group and between-group levels. Parsing the total variance between levels allows us to evaluate interlocking systems of oppression as social causes of population incidence, while simultaneously acknowledging the low discriminatory accuracy of these identities in predicting individual cases.

There are two important points that interested readers should be aware of if they desire to fit similar, though logistic, versions of these models. First, a logistic version of these models (e.g., where the outcome of interest is dichotomously coded obesity) would not estimate the within-strata variance parameter σ^2_{e0}; in logistic multilevel models, the variance of the binomial distribution is known and therefore level 1 variance is not an estimable parameter. Second, if the multilevel model is being used to examine the strata-level residuals to determine whether statistically significant interaction effects are evident overall and/or in particular social strata, then a logistic version of the model would be inappropriate. This is because much of the quantitative intersectional literature defines "interaction effects" on the additive scale, which is what we obtain in linear models. Logistic models, on the other hand, are fit on a multiplicative scale. On the multiplicative scale, inclusion of the main effects (e.g., gender) may already account for some of the interaction effect detected on an additive scale. On the other hand, if

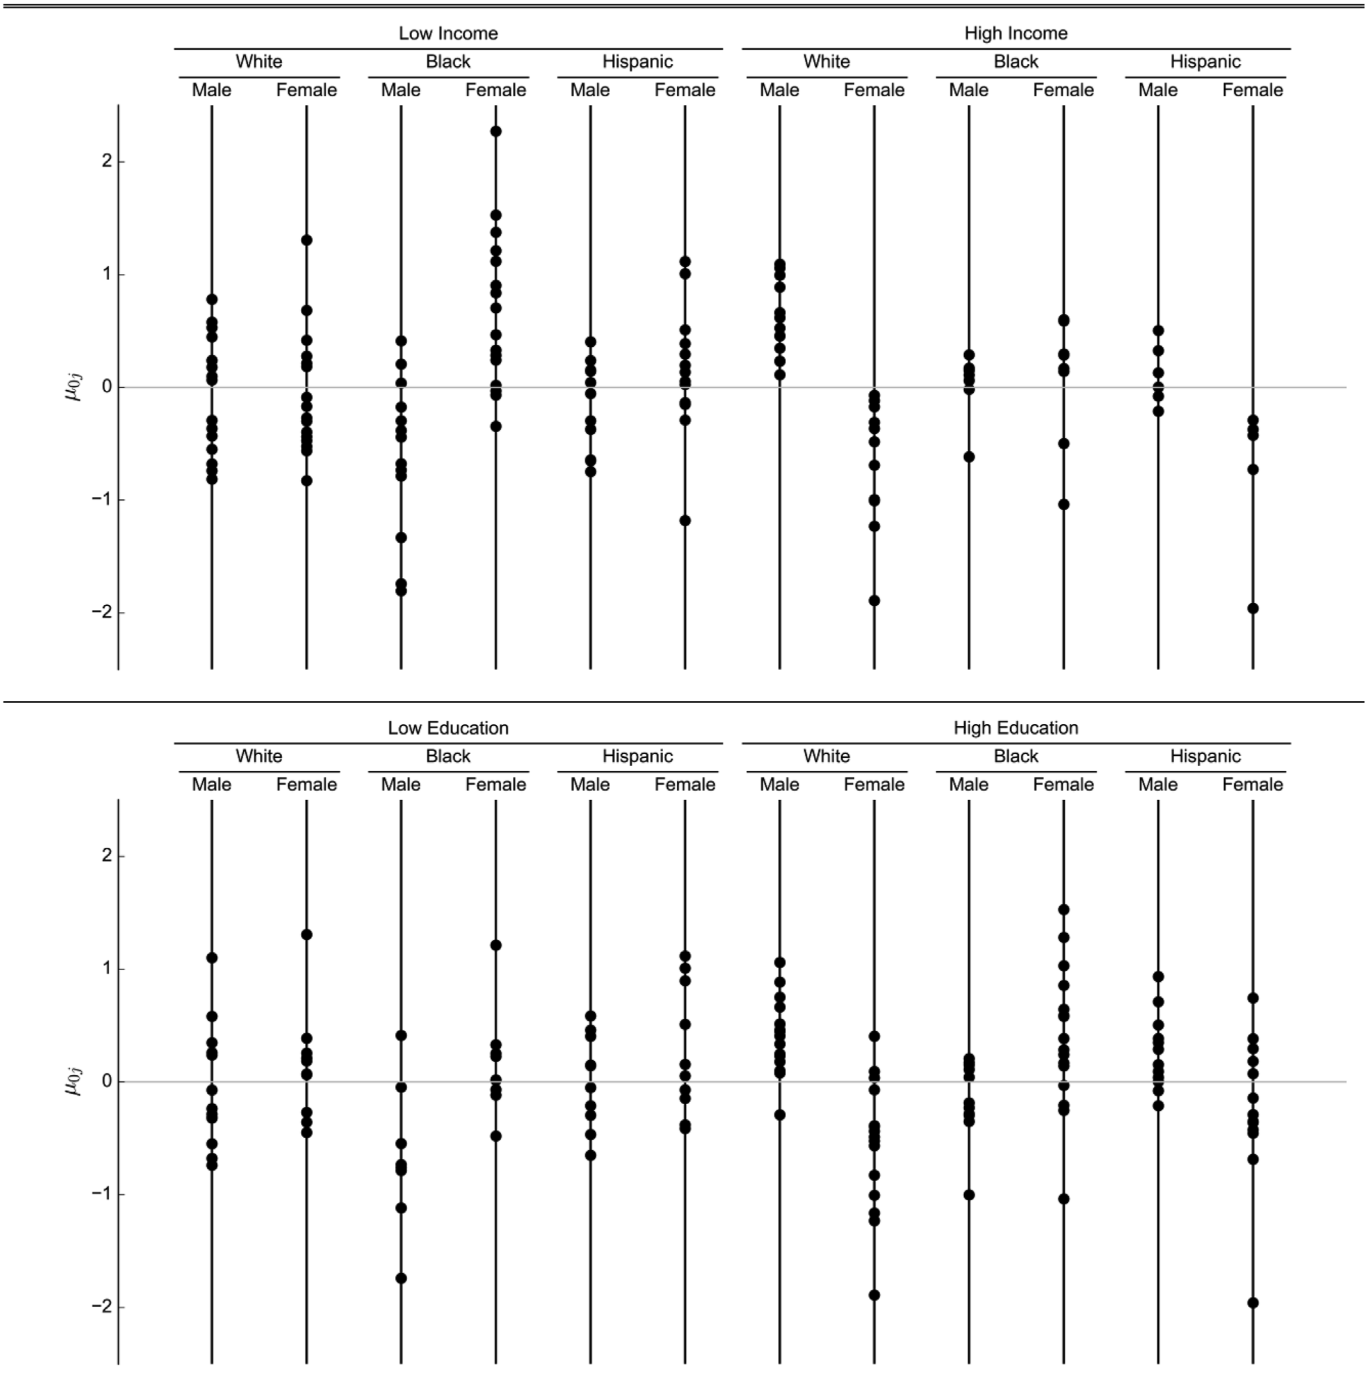

Fig. 2. Strata-level residuals (μ_{0j}) from multilevel models of BMI by gender, race/ethnicity and socioeconomic status.

Note: Each marker indicates one stratum's strata-level residual value. Strata are grouped by gender, race/ethnicity, and socioeconomic status (income and education iteratively). Negative values indicate lower BMI 'than expected' based on the additive contributions of the main effects, while positive values indicate higher BMI 'than expected' for that stratum. Only strata with 20 or more observations are shown.

the purpose of these logistic models is to estimate strata-level effects (as is the case in Jones et al., 2016)) to *explore* the patterning of inequalities across society, and not to examine strata-level residuals, then logistic multilevel intersectionality models would be appropriate.

4.1. Limitations

While the multilevel approach has several apparent advantages over conventional interaction models, it does have some limitations. For instance, the model still requires assumptions about the functional form of the data, such as observations being i.i.d. within strata. This approach also requires the use of large data sets with sufficient diversity in the sample for reliable estimation of effects. Additionally, this approach requires categorization into discrete social strata and therefore some dimensions of identity that could be treated as continuous in other approaches (such as income or age) are not. Finally, it is important to reiterate that the estimates of strata-level variance and strata-level residuals obtained should be interpreted as *unexplained* strata-level variance or residuals, potentially attributable to interaction effects after adjustment for additive effects.

Table 3
Comparison of modeling approaches.

Comparison Criterion	Fixed Effect Model	Multilevel Model
As new dimensions are added to the model, the number of fixed effect parameters increases:	**Geometrically**	**Linearly**
Bayesian Information Criterion (BIC) score increases significantly as new dimensions added to model?	**Yes**	**No**
Are estimates adjusted to account for the sample size observed in a given stratum?	**No**	**Yes**
Comparison made in this approach?	Dummy FE approach: **No formal comparison made**, except in post hoc analysis. Saturated interaction parameters approach: **Strata compared to a single reference group**, and each aspect of interaction is considered separately.	**Each stratum is compared with itself**. The mean/prevalence for a stratum is compared with what was expected for that stratum based on the contributions of the additive main effects.

The fixed effects approach does have some advantages over the multilevel approach that should be taken seriously. First, the fixed effects interaction approach is somewhat simpler to fit and allows for a frequentist approach as opposed to a Bayesian one (Supplemental Technical Note 2). Second, the fixed effects approach may still be preferable when the number of interactions is relatively small (e.g., less than 20). In a multilevel model, it is important to consider the sample size at both levels—the number of individuals within strata and the number of social strata overall. The between-strata variance parameter (σ^2_{strata}) is less meaningful when too few strata are studied and, as illustrated in Supplemental Fig. 1, the advantages of the multilevel approach in terms of BIC are less significant when the number of social strata is small. The sample should also be sufficiently distributed across the strata to allow for reasonable estimates to be obtained for most strata (Supplemental Table 2). We therefore advocate for the continued use of fixed effects models to study interactions, and present the multilevel approach as a valuable *alternative* for some analysis situations.

A third approach to modeling intersectionality worth noting, recently demonstrated by Wemrell et al. (2017), involves combining the fixed effects approach with measures of discriminatory accuracy (the area under the ROC curve—or AUC). The AUC approach provides analogous estimates to the variance partition coefficient (VPC), which can be useful when assessing both total variation in a sample, and variation unexplained by additive effects alone.

5. Conclusion

The present study contributes to our methodological repertoire by broadening our vision of the potential applications of multilevel models in social epidemiology. Multilevel modeling has already demonstrated its value in the public health and social science literature by enabling the study of ecological effects. The approach we have outlined illustrates that the statistical dependencies created by abstract hierarchical social structures within which respondents are embedded can also be modeled.

Funding

No direct funding supported this study.

Acknowledgements and Credits

We wish to thank the reference librarians of the Harvard University Countway Library of Medicine for their assistance in performing a systematic review of the literature.

Appendix A. Supplementary materials

Supplementary materials related to this article can be found at http://dx.doi.org/10.1016/j.socscimed.2017.11.011.

References

Bauer, G.R., 2014. Incorporating intersectionality theory into population health research methodology: challenges and the potential to advance health equity. Soc. Sci. Med. 110, 10–17.
Berkman, L.F., Kawachi, I. (Eds.), 2000. Social Epidemiology. Oxford University Press, Oxford.
Bowleg, L., 2012. The problem with the phrase women and minorities: intersectionality—an important theoretical framework for public health. Am. J. Public Health 102 (7), 1267–1273.
Browne, W.J., 2009. MCMC Estimation in MLwiN v2.1: Centre for Multilevel Modelling. University of Bristol.
Choo, H.Y., Ferree, M.M., 2010. Practicing intersectionality in sociological research: a critical analysis of inclusions, interactions, and institutions in the study of inequalities. Sociol. Theory 28 (2), 129–149.
Clarke, P., O'Malley, P.M., Johnston, L.D., Schulenberg, J.E., 2009. Social disparities in BMI trajectories across adulthood by gender, race/ethnicity and lifetime socio-economic position: 1986-2004. Int. J. Epidemiol. 38 (2), 499–509.
Collins, P.H., 1990. Black Feminist Thought: Knowledge, Consciousness, and the Politics of Empowerment. Routledge, New York.
Crenshaw, K., 1989. Demarginalizing the Intersection of Race and Sex: a Black Feminist Critique of Antidiscrimination Doctrine, Feminist Theory, and Antiracist Politics. University of Chicago Legal Forum.
Evans, C.R., 2015. Innovative Approaches to Investigating Social Determinants of Health - Social Networks, Environmental Effects and Intersectionality. Doctoral dissertation Department of Social and Behavioral Sciences, Harvard T.H. Chan School of Public Health Available at: http://nrs.harvard.edu/.
Evans, C.R., Williams, D.R., Onnela, J.P., Subramanian, S.V., 2016. Multiple contexts and adolescent body mass index: schools, neighborhoods, and social networks. Soc. Sci. Med. 162, 21–31.
Farmer, M.M., Ferraro, K., 2005. Are racial disparities in health conditional on socioeconomic status? Soc. Sci. Med. 60 (1), 191–204.
Grant, B.F., Kaplan, K., Shepard, J., Moore, T., 2003. "Source and Accuracy Statement for Wave 1 of the 2001-2002 National Epidemiologic Survey on Alcohol and Related Conditions.". National Institute on Alcohol Abuse and Alcoholism, Bethesda, MD.
Grant, B.F., Kaplan, K.D., 2005. Source and Accuracy Statement for the Wave 2 National Epidemiologic Survey on Alcohol and Related Conditions (NESARC). National Institute on Alcohol Abuse and Alcoholism, Rockville, Maryland.
Green, M.A., Evans, C.R., Subramanian, S.V., 2017 Apr. Can intersectionality theory enrich population health research? Soc. Sci. Med. 178, 214–216. http://dx.doi.org/10.1016/j.socscimed.2017.02.029. Epub 2017 Feb 22.
Hancock, A., 2007. When multiplication Doesn't equal quick addition: examining intersectionality as a research paradigm. Perspect. Polit. 5 (1), 63–79.
Hancock, A.M., 2013. Empirical intersectionality: a tale of two approaches. UC Irvine Law Rev. 3 (2), 259–296.
Jones, K., Johnston, R., Manley, D., 2016. Uncovering interactions in multivariate contingency tables: a multi-level modelling exploratory approach. Methodol. Innov. 9, 1–17.
King, D.K., 1988. Multiple jeopardy, multiple consciousness: the context of a black feminist ideology. Signs 14 (1), 42–72.
Krieger, N., 2011. Epidemiology and the People's Health - Theory and Context. Oxford University Press Inc, New York, NY.
Link, B.G., Phelan, J.C., 1995. Social conditions as fundamental causes of disease. J. Health Soc. Behav. 35 (Extra Issue), 80–94.
McCall, L., 2005. The complexity of intersectionality. Signs 30 (3), 1771–1800.
Merlo, J., Wagner, P., 2012. The tyranny of the averages and indiscriminate use of risk factors in public health: a call for revolution. Eur. J. Epidemiol. 28, 148.
Merlo, J., Wagner, P., Ghith, N., Leckie, G., 2016. An original stepwise multilevel logistic regression analysis of discriminatory accuracy: the case of neighborhoods and health. PLoS One 11 (4), e0153778.
Merlo, J., 2014. Invited commentary: multilevel analysis of individual heterogeneity—a

fundamental critique of the current probabilistic risk factor epidemiology. Am. J. Epidemiol. 180, 208–212.
Nash, J.C., 2008. "Re-Thinking intersectionality." feminist review. 89, 1–15.
Oakley, A., 1972. Sex, Gender and Society. Maurice Temple Smith, London.
Ogden, C.L., Carroll, M.D., Kit, B.K., Flegal, K.M., 2012. Prevalence of obesity in the United States, 2009-2010. NCHS Data Brief. 2012, 1–8.
Rasbash, J., Browne, W.J., Healy, M., Cameron, B., Charlton, C., 2012. *MLwiN* Version 2.26: Centre for Multilevel Modelling. University of Bristol.
Raudenbush, S.W., Bryk, A.S., 2002. Hierarchical Linear Models: Applications and Data Analysis Methods. Sage Publications, Thousand Oaks, CA.
Rose, G., 1992. The Strategy of Preventive Medicine. Oxford University Press, Oxford England, New York.
Schulz, A.J., Mullings, L. (Eds.), 2006. Gender, Race, Class, and Health: Intersectional Approaches. Jossey-Bass, San Francisco, CA.
Schwarz, G., 1978. Estimating the dimension of a model. Ann. Statistics 6 (2), 461–464.
Susser, M., Susser, E., 1996. Choosing a future for epidemiology: ii. From black box to Chinese boxes and eco-epidemiology. Am. J. Public Health 86 (5), 674–677.
U.S. Census Bureau, 2012. "Poverty thresholds 2000": social, economic, and housing statistics division: poverty. Retrieved 1 April 2014. http://www.census.gov/hhes/www/poverty/data/threshld/thresh00.html.
Veenstra, G., 2011. Race, gender, class, and sexual orientation: intersecting axes of inequality and self-rated health in Canada. Int. J. Equity Health 10 (1), 3.
Warner, D.F., Brown, T.H., 2011. Understanding how race/ethnicity and gender define age-trajectories of disability: an intersectionality approach. Soc. Sci. Med. 72 (8), 1236–1248.
Wemrell, M., Mulinari, S., Merlo, J., 2017. Intersectionality and risk for ischemic heart disease in Sweden: categorical and anti-categorical approaches. Soc. Sci. Med. 177, 213–222.