Article:

Leckie, G., Bell, A. orcid.org/0000-0002-8268-5853, Merlo, J. et al. (2 more authors) (Accepted: 2025) The statistical advantages of MAIHDA for estimating intersectional inequalities. Sociological Methods & Research. ISSN: 0049-1241 (In Press)

The Statistical Advantages of MAIHDA for Estimating Intersectional Inequalities

George Leckie[1], Andy Bell[2], Juan Merlo[3], SV Subramanian[4], Clare Evans[5]

[1]Centre for Multilevel Modelling and School of Education, University of Bristol, UK.

[2]Sheffield Methods Institute, School of Education, University of Sheffield, Sheffield, UK

[3]Research Unit of Social Epidemiology, Faculty of Medicine, University of Lund, Sweden

[4]Department of Social and Behavioral Sciences, Harvard T.H. Chan School of Public Health, Boston, MA, USA

[5]Department of Sociology, University of Oregon, Eugene, OR, USA

Contact information for the corresponding author

George Leckie

Centre for Multilevel Modelling and School of Education

University of Bristol

35 Berkeley Square

Bristol BS8 1JA

United Kingdom

g.leckie@bristol.ac.uk

Acknowledgments

We thank the editor and reviewers for their comments.

STATEMENTS AND DECLARATIONS

Ethical considerations

Not applicable.

Consent to participate

Not applicable.

Consent for publication

Not applicable.

Declaration of Conflicting Interest

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

This work was funded by a UK Economic and Social Research Council (ESRC) grants ES/W000555/1 and grant ES/X011313/1.

Data Availability

The Keller et al. (2023) intersection-level predictions are provided in the Supplemental Material.

SUPPLEMENTAL MATERIAL

The Keller et al. (2023) intersection-level predictions are provided in the Supplemental Material.

AUTHOR BIOGRAPHIES

George Leckie is a Professor of Social Statistics at the Centre for Multilevel Modelling and School of Education, University of Bristol. His research interests are in the development, application, and dissemination of multilevel models to analyze educational and other social science data. His recent work has appeared in the American Sociological Review, Journal of Educational and Behavioral Statistics, Journal of the Royal Statistical Society: Series A, Psychological Methods, and Sociological Methodology.

Andrew Bell is an Associate Professor at the Sheffield Methods Institute, University of Sheffield, UK. His current substantive research spans a diverse range of subject areas, including social epidemiology, geography, sociology, political science, public health, and economics. Methodologically, his interests are in the development and application of multilevel models, with work focusing on age-period-cohort analysis and fixed and random effects models.

Jaun Merlo is a Professor of Social Medicine at Lund University, Sweden. His research investigates the role of society as a determinant of individual health. He adopts a multilevel perspective to develop innovative methodologies while bringing together a multidisciplinary team including expertise in the fields of medicine, biology, epidemiology, statistics, sociology, and anthropology. His work has appeared in the American Journal of Epidemiology, Social Science and Medicine, Health and Place, and BMJ group.

SV Subramanian is Professor of population health and geography in the Department of Social and Behavioual Sciences at the Harvard T.H. Chan School of Public Health, USA. His research focuses on understanding the role of geographic, spatial and institutional contexts in influencing population health.

Clare Evans is an Associate Professor in the Department of Sociology and Co-Director of the Global Health program at the University of Oregon, USA. She is a medical sociologist and social epidemiologist with expertise in advanced quantitative analysis. Her research interests focus on the intersectional social determinants of population health inequities, and quantitative methods for evaluating intersectional inequities and environmental injustices. Her work has appeared in Health and Place, Social Science and Medicine, and the American Journal of Epidemiology.

The Statistical Advantages of MAIHDA for Estimating Intersectional Inequalities

ABSTRACT

Multilevel Analysis of Individual Heterogeneity and Discriminatory Accuracy (MAIHDA) is a multilevel regression modeling approach rooted in intersectionality theory. It is used to examine inequalities across intersections of multiple social identities (e.g., gender, ethnicity, social class). Proponents argue that MAIHDA provides more accurate predictions of population means for these groups than the conventional approach of calculating the simple arithmetic mean for each intersectional group—either directly or via a linear regression model that includes all possible interactions between the social identities. In this study, we aim to evaluate and demonstrate the validity of this claim to the community of quantitative intersectionality researchers. We do so by analytically comparing the variance and correlation of simple and MAIHDA-predicted means with population means. Additionally, we assess the bias, variance, and mean squared error in predicting the mean of a given intersection. Our findings show that MAIHDA-based means outperform simple means, particularly when using the MAIHDA means that decompose intersectional means into additive and non-additive effects of social identities. However, the relative advantage of MAIHDA means depends on the nature of the intersectional inequalities and the sizes of the intersectional groups being studied. MAIHDA's benefits are most pronounced when inequalities are subtle or when data on certain intersections, such as those for marginalized groups, are sparse. These conditions are common in practice, highlighting the practical significance of our findings.

Keywords: MAIHDA, intersectionality, inequalities, multilevel models, predicted means, empirical Bayes, posterior means

INTRODUCTION

Multilevel Analysis of Individual Heterogeneity and Discriminatory Accuracy (MAIHDA) is a recently-developed multilevel regression modeling approach designed to explore complex social inequalities in individual outcomes (Evans, Williams, Onnela, and Subramanian 2018). MAIHDA is motivated by intersectionality theory (Collins and Bilge 2020; Crenshaw 1989), which observes that individuals' lived experiences and outcomes are shaped by their positionalities within complex and interlocking systems of oppression, including sexism, racism, and socioeconomic inequality. MAIHDA quantifies inequalities across intersections of multiple social identities and positionalities (e.g., gender, ethnicity, and social class), rather than focusing on one axis of inequality at a time. The growing adoption of MAIHDA reflects widespread interest in quantitative methods that align with intersectionality's demands for expansive consideration of diversity (Bauer et al. 2021; McCall 2005; Merlo 2018).

While early applications were primarily in social epidemiology, MAIHDA is increasingly being used across the social sciences, with applications in criminology (Pina-Sánchez and Tura 2024; Tura et al. 2024), education (Giaconi et al. 2024; Keller et al. 2023; Prior and Leckie 2024; Prior et al. 2022; van Dusen et al. 2024), environmental justice (Alvarez et al. 2022), gender studies (Ivert et al. 2020; Silva and Evans 2020), organizational studies (Humbert 2024), psychiatry (Forrest et al. 2023), and social work (Lister, Hewitt, and Dickerson 2024; Pomeroy and Fiori 2025). For example, Keller et al. (2023)—an application we will return to later—applied MAIHDA to study intersectional inequalities in 15-year-old students' reading scores across four social identities: gender, immigrant status, parental education, and parental occupational status.

Evans et al. (2024) present an introduction and tutorial on MAIHDA for those completely new to the approach. Here, we summarize the key points. MAIHDA was developed in response to perceived weaknesses of the conventional approach to studying intersectionality, which typically involves estimating linear regression models on social identities (Evans et al., 2018). These weaknesses include the assumption that the effects of social identities are additive, which may overlook how systems of oppression interact in more complex, multiplicative ways. When interaction terms are considered, often just a single two-way interaction is included. Including all possible interactions, however, quickly leads to many regression coefficients, resulting in overfitting and challenges with interpretation. Interpretation, if not overfitting, is eased by predicting and comparing the mean outcome for each intersectional group. We refer to these as simple means as they are equal to the arithmetic means obtained by calculating the mean separately for each group.

The MAIHDA approach, in contrast, is grounded in the multilevel modelling framework (Raudenbush and Byrk, 2002; Snijders and Bosker, 2012). It involves fitting a sequence of two multilevel regression models where individuals (level 1) are nested in intersectional social strata (level 2), henceforth referred to as intersections. Thus, Intersection 1 might refer to native female students with low parental education and low parental occupation. Intersection 2 might then be native female students with low parental education and low-to-middle parental occupation, and so on. For simplicity, we focus on MAIHDA models for continuous outcomes, though MAIHDA models can be applied to all outcome types.

The first multilevel model, henceforth referred to as MAIHDA Model 1, is a two-level model without any covariates. The model estimates the overall magnitude of intersectional inequalities in the data and predicts the mean outcome for each intersection. This facilitates the identification of social identity combinations associated with the most and least favorable mean outcomes.

The second multilevel model, henceforth referred to as MAIHDA Model 2, is a two-level model in which the social identity variables used to construct the intersections are included as main-effect covariates. The model examines the extent to which intersectional inequalities deviate from the simplest additive patterns of social identities—for example, whether the gender-based mean outcome difference remains constant across different values of immigrant status, parental education, and parental occupational status. By assessing these deviations, Model 2 can reveal hidden social processes that emerge only for specific social identity combinations, as well as quantify consistent or typical patterns (e.g., women tend to experience worse outcomes than men).

A central argument made by proponents of MAIHDA is that its intersectional means provide more accurate predictions of the population means than do simple means (Evans et al. 2018; Evans et al. 2024b). This claim is based on earlier findings from the statistical literature on multilevel models (Raudenbush and Bryk, 2002; Snijders and Bosker, 2012). However, these findings and their implications are less well understood by applied researchers especially within the context of MAIHDA. Fundamentally, how much more accurate are MAIHDA means than simple means? What does their relative accuracy depend on, and how does this vary across a wide range of possible scenarios? Most importantly, is the choice between simple and MAIHDA means likely to affect conclusions in real-world research? These questions and their answers matter. If the community of intersectional researchers using MAIHDA unknowingly make incorrect choices—leading to results that would have differed if a preferred approach had been used—then they risk mischaracterizing inequalities, inefficiently targeting marginalized groups, and misallocating resources, all of which can have harmful consequences for individuals and society.

The arguments in favor of MAIHDA predicted means over simple means are typically based on two key points from the multilevel literature. First, simple means exhibit high sampling variability when the number of individuals per intersection is low. Second, MAIHDA means address this issue as they are defined as conditional expectations of the population given the data, which shrink their predictions from the simple means toward model-implied means—that is, the means predicted by the intercept and the main effect covariates if included) (Raudenbush and Bryk, 2002; Snijders and Bosker, 2012). Greater shrinkage is applied to the smallest intersections. Shrinkage is viewed as beneficial because it protects against overinterpreting extreme predictions that may have arisen due to chance (sampling variation). When the models are estimated by frequentist methods (MLE or REML) these expectations are calculated using empirical Bayes prediction. When Bayesian methods (MCMC) are used, the posterior distributions of the intersection means are estimated and are then summarized by their posterior means.

Importantly, the means predicted by MAIHDA Model 1 and Model 2 also differ from one another. This variation stems from the difference in their model-implied means—or, in other words, the values toward which the predictions are shrunk. In Model 1, the model-implied means are simply the overall or grand mean, so final predictions are shrunk toward this single value. Thus, the final predictions are informed by both the data from that intersection and the overall data. This shrinkage is therefore sometimes referred to as partial pooling. In Model 2, the model-implied means are the means implied by the estimated additive effects of the social identities used to define the intersections. The predictions in Model 2 are therefore shrunk toward these intersection-specific values. To the extent that the MAIHDA-predicted means are preferable to the simple means, the Model 2 means are expected to be preferable to the Model 1 means, as they will lie closer to the true population means for each intersection.

Several simulation studies have started to examine the predictive accuracy of the MAIHDA means. Bell, Holman, and Jones (2019) simulated data from linear regression models, primarily without interaction terms, and compared Type I error rates among simple means, MAIHDA Model 1, and Model 2 means. Their findings suggest that MAIHDA Model 2 means result in a lower Type I error rate than both Model 1 and simple means.

Mahendran, Lizotte, and Bauer (2022a, 2022b) expanded on this by simulating data from linear and logistic regression models with various interaction terms. They compared MAIHDA Model 2 means to simple means. For both continuous and binary outcomes, they concluded that MAIHDA Model 2 means offer greater accuracy than simple means, particularly for smaller intersection sizes.

Van Dusen et al. (2024) simulated data from a MAIHDA Model 2 and found that Model 2 means outperform simple means, with the performance gap widening as intersection sizes decrease.

While these studies consistently show that MAIHDA Model 2 means outperform simple means, they provide little insight into why this occurs, beyond broadly attributing it to shrinkage. They also offer minimal exploration of MAIHDA Model 1 means and do not examine how the relative merits of all three means might vary based on the nature of the intersectional inequalities being studied.

In this study, we aim to evaluate—and clearly demonstrate to the community of quantitative intersectionality researchers—the claim that MAIHDA means are more accurate than simple means. We seek to explain when and why these different means diverge and to provide guidance on which to report in practice. Specifically, we present and analyze analytical expressions that describe how these means vary across intersections relative to the true variance, and how they correlate with the true intersection means. We then assess the statistical properties of each approach for a given intersection of interest, deriving and analyzing expressions for the bias, variance, and mean squared error (MSE) of the three means based on random samples of individuals within that intersection. While these expressions are not themselves new—since MAIHDA models are multilevel models—their interpretation in the context of MAIHDA, and thus their relevance to intersectional research, is novel. Across both sets of analyses, we examine how these properties vary as a function of the overall magnitude of intersectional inequalities, the extent to which those inequalities follow an additive pattern, and key data characteristics, including the mean and variability of intersection sizes.

THE TWO MAIHDA MODELS

In this section, we provide a brief review of the two MAIHDA models.

Model 1: Empty, Null, or Unadjusted Model

Model 1 is a two-level model without any covariates. Let y_{ij} denote the outcome for individual i ($i = 1, \dots, n_j$) at intersection j ($j = 1, \dots, J$). The model can then be written as:

$$y_{ij} = \mu_j + e_{ij} \tag{1}$$

where μ_j denotes the population or true mean outcome at intersection j—that is, the mean outcome if the entire population of individuals at that intersection were observed, rather than just a sample. The term e_{ij} is an individual residual measuring how each individual's outcome deviates from the true mean for their intersection. The true mean is then specified as:

$$\mu_j = \beta_0 + u_j \tag{2}$$

where β_0 is fixed-effect intercept measuring the overall mean or grand mean across all intersections, and u_j is a random effect measuring how each intersection's true mean deviates from the overall mean. Substituting (2) into (1) and rearranging gives the combined equation:

$$y_{ij} = \beta_0 + u_j + e_{ij}. \tag{3}$$

The intersection random effects and individual residuals are each assumed normally distributed with constant variances σ_u^2 and σ_e^2. Thus, σ_u^2 is also the variance of the true means about the overall mean.

The overall magnitude of intersectional inequalities is measured by the Variance Partition Coefficient (VPC), which quantifies the proportion of outcome variance that lies between the intersection means:

$$\text{VPC} = \frac{\sigma_u^2}{\sigma_u^2 + \sigma_e^2}. \tag{4}$$

The VPC can range from 0 to 1, with higher values indicating greater intersectional inequalities. Most applications find VPCs ranging from 0.01 to 0.20, suggesting that the studied social stratifications account for 1-20% of the outcome variation (Evans et al. 2024a). Conversely, 80-99% of the variation reflects other unmodelled individual characteristics.

When the models are estimated using frequentist methods (MLE or REML), empirical Bayes prediction is applied post-estimation to assign values to the random effects u_j, which are then used to predict the intersection means μ_j. In contrast, Bayesian methods (MCMC) estimate the posterior distributions of the intersection means simultaneously with the model parameters, typically summarizing them by their posterior means. Analyzing these intersection means enables the identification of social identity combinations linked to the most and least favorable outcomes. Henceforth, for simplicity, we refer to predicting the intersection means, without distinguishing between prediction and estimation as defined under frequentist or Bayesian frameworks.

Model 2: Full, Main Effects, or Adjusted Model

Model 2 is a two-level model in which the social identity variables used to construct the intersections are included as main-effect covariates. The model can be written as:

$$y_{ij} = \mu_j + e_{ij} \tag{5}$$

$$\mu_j = \beta_0 + \beta_1 x_{1j} + \cdots + \beta_p x_{pj} + u_j \tag{6}$$

where $x_{1j}, \ldots, x_{pj}$ denotes the social identity variables (entered as dummy variable covariates) used to define the intersections, $\beta_1, \ldots, \beta_p$ are the associated fixed-effect regression coefficients capturing the additive patterns in the intersection inequalities. The term u_j is the intersection random effect, measuring how each intersection's true mean deviates from that implied by the additive effects; in other words, it captures all two-way and higher-order interaction variability. Thus, σ_u^2 now measures the variance associated with this non-additivity.

Substituting (6) into (5) gives the combined equation:

$$y_{ij} = \beta_0 + \beta_1 x_{1j} + \cdots + \beta_p x_{pj} + u_j + e_{ij}. \tag{7}$$

The extent to which the intersectional inequalities are additively patterned is assessed via the Proportion Change in Variance (PCV) statistic. The PCV measures the reduction in intersection variance when moving from Model 1 to Model 2. This is given by:

$$\text{PCV} = \frac{\sigma_{u1}^2 - \sigma_{u2}^2}{\sigma_{u1}^2} \tag{8}$$

where the 1 and 2 subscripts distinguish between the Model 1 and Model 2 intersection random effect variances. Thus, $1 - \text{PCV}$ measures the degree to which intersectional inequalities deviate from those implied by additivity and, therefore, the extent to which hidden social processes may be at play, occurring only for certain social identity combinations. Put differently, the PCV accounts for how much of the overall inequality quantified by Model 1 follows a consistent pattern (e.g., female students consistently score higher than male students, regardless of other characteristics), while $1 - \text{PCV}$ accounts for departures from this pattern (e.g., gender inequality is larger among immigrants).

The PCV can range from 0 to 1, with higher values indicating a greater additive structure. When the PCV equals 0, there is no additive structure at all, and Model 2 simplifies to Model 1. However, in practice, we typically see a PCV that suggests a mix of additive (consistent) and interaction (unique departures) inequality patterns. Most empirical applications find PCVs ranging from 0.60 to 0.95, suggesting that 60-95% of the variation in mean outcomes across intersections follows an additive pattern (Evans et al. 2024a). Conversely, 5-40% of the variation reflects more complex interaction effects.

The individual residual variance e_{ij}, in contrast, does not change when moving from Model 1 to Model 2, as the social identities are intersection-level covariates and therefore do not explain outcome heterogeneity within intersections.

As with Model 1, we can predict the intersection means μ_j. Additionally, we can decompose each predicted mean outcome into its additive and non-additive effect, allowing identification of specific intersections where non-additivity occurs. To do this, a 95% confidence interval (for frequentist estimation) or a 95% credible interval (for Bayesian estimation) is typically constructed around the non-additive effect u_j to determine which intersections significantly deviate from what would be expected under an additive model. However, we do not explore this further here.

Illustrative Application: Keller et al. (2023)

Keller et al. (2023) presented the first application of MAIHDA in educational research. They applied MAIHDA to 5,451 student reading performance scores from the German sample of the Programme for International Student Assessment (PISA) 2018. They considered four social identities: gender (male, female), immigrant status (native, immigrant), parental education (low, high, as measured by university entrance certificate), and parental occupational status (low, low-middle, middle, middle-high, high). Combining these categories resulted in 40 intersections ($= 2 \times 2 \times 2 \times 5$). Their descriptive statistics showed that the mean and standard deviation (SD) of reading performance were 498 and 106, respectively (versus the OECD average mean and SD of 485 and 105). Female students scored higher than male students, native students outperformed immigrant students, students whose parents held a university entrance certificate scored higher than those whose parents did not, and students from families with higher occupational status achieved higher scores (see their Table 1).

Table 1 presents their results from MAIHDA Model 1 and Model 2 estimated using Bayesian Markov chain Monte Carlo (MCMC) methods (see their Table 3). The Model 1 VPC statistic reveals substantial intersectional inequalities: 16% of the variance in student achievement lies between the intersection means. The Model 2 regression coefficients estimate the additive structure in the intersectional means, and these results align with their descriptive statistics. The PCV statistic shows that the additive structure accounts for 91% of the variation in the intersection means, meaning that 9% of the variation reflects deviations—both positive and negative—from the model-implied additive patterns of inequality. In other words, variation associated with two-way and higher-order interactions captured by the intersection random effect. Indeed, six of their intersection means deviate by 10 or more points from additivity (approximately 0.1 SD or more), but only one of these departures is statistically significant (see their Figure 4). Intersection 40, Female, native students with university entrance certificate parents and high occupational status—already the highest-scoring intersection in terms of additive effects—scored around 15 points higher (0.15 SD) than what additivity would suggest.

Figure 1 shows the simple means and the MAIHDA Model 1 and Model 2 means for all 40 intersections, and where we have highlighted the six intersections with 20 or fewer individuals. For most intersections, the means are very similar across the three methods. However, as expected, the means for smaller intersections (highlighted) notably vary across methods (see also Table 2). Specifically, compared to the simple means, the MAIHDA Model 1 means are shrunk toward the overall average. Relative to the MAIHDA Model 1 means, the MAIHDA Model 2 means generally increase, except for Intersection 35, which decreases. In the most extreme case, Intersection 39 (male, immigrant, low parental education, high occupational status; $n = 3$), the three means are 362, 438, and 458, with a range of 96 points (approximately 0.96 SD). These results clearly demonstrate that the different prediction methods can yield substantively different results, especially when intersection sizes are small.

PREDICTED MEANS

The predicted intersection means are a central output of a MAIHDA analysis, as they are typically used to identify the social identity combinations associated with the most and least favorable mean outcomes (Evans et al., 2024b). In this section, we describe the three methods for predicting the true intersection means: simple means, the MAIHDA Model 1 means, and the MAIHDA Model 2 means. A key purpose of this section is to demonstrate how shrinkage leads the MAIHDA Model 1 and Model 2 means to differ from the simple means and from each other. It is worth reiterating that both MAIHDA Model 1 and Model 2 are standard multilevel models, and so the equations for their predicted means follow the formulations reported in the multilevel modeling literature (Raudenbush & Bryk, 2002; Snijders & Bosker, 2012). Note that we use subscripts 1 and 2 to distinguish between the parameters and terms associated with Model 1 and Model 2, respectively.

Simple Means

Averaging across n_j individuals within intersection j yields the simple mean:

$$\bar{y}_{.j} = \frac{1}{n_j}\sum_{i=1}^{n_j} y_{ij} \tag{9}$$

This is sometimes referred to as the cross-classification method. The simple mean can also be obtained by estimating a linear regression of y_{ij} on an intercept and $J-1$ dummy variables (or no intercept and J dummy variables) and predicting the fitted values, sometimes referred to as a fixed effects dummy variable model. Alternatively, they can be derived by fitting a linear regression of y_{ij} on the social identity variables that define the intersections, fully interacting these variables. This is sometimes referred to as a saturated fixed effects model, and was until MAIHDA the cutting edge for quantitative intersectional analyses. For example, in the Keller et al. (2023) study, this would involve entering the main effects of the four social identity variables, and then all two-, three-, and four-way interactions terms.

MAIHDA Model 1 Means

The MAIHDA Model 1 mean for intersection j is defined as the conditional expectation of the true intersection mean, given the observed simple mean for that intersection:

$$\tilde{\mu}_{1j} \equiv \mathrm{E}\left(\mu_j \middle| \bar{y}_{.j}\right) = R_{1j}\bar{y}_{.j} + \left(1 - R_{1j}\right)\beta_{1,0} \quad (10)$$

where we use the tilde notation to help distinguish the predicted mean $\tilde{\mu}_{1j}$ from the true mean μ_j. Under frequentist estimation, we plug in the parameter estimates, resulting in empirical Bayes predictions. Under Bayesian estimation, we instead estimate the posterior distribution of μ_j and summarize it using the posterior mean.

The term R_{1j} is the reliability of the simple mean $\bar{y}_{.j}$ as an estimate of the true mean μ_j. These predictions are therefore reliability weighted averages of the simple means $\bar{y}_{.j}$ and the model-implied or overall mean $\beta_{1,0}$. The more reliable $\bar{y}_{.j}$ is as an estimate of μ_j, the more weight is given to $\bar{y}_{.j}$ and the less to $\beta_{1,0}$. The Model 1 means can therefore be viewed as a prediction that starts with the simple mean and is then shrunk toward the overall mean. The degree of shrinkage is given by:

$$\tilde{\mu}_{1j} - \bar{y}_{.j} = \left(R_{1j} - 1\right)\left(\bar{y}_{.j} - \beta_{1,0}\right) \quad (11)$$

Therefore, shrinkage decreases as reliability increases, and it also decreases as the difference between the simple mean and the overall mean becomes smaller.

Reliability is calculated as the ratio of the true mean variance to the observed mean variance, and thus varies from 0 to 1:

$$R_{1j} = \frac{\mathrm{Var}(\mu_j)}{\mathrm{Var}(\bar{y}_{.j})} = \frac{\mathrm{Var}(\beta_{1,0}+u_{1j})}{\mathrm{Var}(\beta_{1,0}+u_{1j}+\bar{e}_{.j})} = \frac{\sigma_{u1}^2}{\sigma_{u1}^2+\frac{\sigma_e^2}{n_j}} \quad (12)$$

where $\bar{e}_{.j}$ is the mean of the individual residuals:

$$\bar{e}_{.j} = \frac{1}{n_j}\sum_{i=1}^{n_j} e_{ij} \quad (13)$$

with variance σ_e^2/n_j. Note that we do not add a 1 subscript to $\bar{e}_{.j}$ or σ_e^2, as these terms do not change when moving from Model 1 to Model 2. Reliability can also be interpreted as the squared correlation between $\bar{y}_{.j}$ and μ_j, that is, as the R-squared or coefficient of determination for predicting μ_j from $\bar{y}_{.j}$. From this expression, we see that reliability increases with the Model 1 VPC, the magnitude of intersectional inequalities, and n_j, the number of individuals at an intersection. This suggests that as VPC or n_j increase, both $\bar{y}_{.j}$ and $\tilde{\mu}_j$ will converge toward μ_j.

MAIHDA Model 2 Means

The MAIHDA Model 2 mean for intersection j is given by the conditional expectation of the true mean, given the simple mean and the covariates for that intersection:

$$\tilde{\mu}_{2j} \equiv \mathrm{E}\left(\mu_j \middle| \bar{y}_{.j}, x_{1j}, \dots, x_{pj}\right) = R_{2j}\bar{y}_{.j} + \left(1 - R_{2j}\right)\left(\beta_{2,0} + \beta_{2,1}x_{1j} + \cdots + \beta_{2,p}x_{pj}\right) \quad (14)$$

where R_{2j} is the reliability of the simple mean $\bar{y}_{.j}$ as an estimate of the true mean μ_j, conditional on the additive effects of the multiple social identities. Thus, the Model 2 means are also weighted averages of the simple means $\bar{y}_{.j}$ and the model-implied means, but where the latter are now given by the fixed-part of the model $\beta_{2,0} + \beta_{2,1}x_{1j} + \cdots + \beta_{2,p}x_{pj}$. Because the additive model implies a different mean for every intersection, the shrinkage is now towards a different point for every intersection rather than the single common point $\beta_{1,0}$. The degree of shrinkage is now given by:

$$\tilde{\mu}_{2j} - \bar{y}_{.j} = \left(R_{2j} - 1\right)\left\{\bar{y}_{.j} - \left(\beta_{2,0} + \beta_{2,1}x_{1j} + \cdots + \beta_{2,p}x_{pj}\right)\right\} \quad (15)$$

and, therefore, again decreases as the now conditional reliability increases, and again increases as the difference between the simple mean and the model-implied mean increases.

The expression for the conditional reliability takes the same form as before:

$$R_{2j} = \frac{\mathrm{Var}\left(\mu_j \middle| x_{1j}, \dots, x_{pj}\right)}{\mathrm{Var}\left(\bar{y}_{.j} \middle| x_{1j}, \dots, x_{pj}\right)} = \frac{\mathrm{Var}\left(\beta_{2,0} + \beta_{2,1}x_{1j} + \dots + \beta_{2,p}x_{pj} \middle| x_{1j}, \dots, x_{pj}\right)}{\mathrm{Var}\left(\beta_{2,0} + \beta_{2,1}x_{1j} + \dots + \beta_{2,p}x_{pj} \middle| x_{1j}, \dots, x_{pj}\right)} = \frac{\sigma_{u2}^2}{\sigma_{u2}^2 + \frac{\sigma_e^2}{n_j}} \quad (16)$$

However, the conditional reliabilities will have lower values than the unconditional reliabilities as $\sigma_{u2}^2 < \sigma_{u1}^2$. Specifically, as the explanatory power of the additive effects of the social identity variables increases, σ_{u2}^2 decreases, resulting in lower values of R_{2j}. Consequently, when computing the MAIHDA means, the shrinkage factor in Model 2 is stronger than in Model 1, $\left|R_{2j} - 1\right| > \left|R_{1j} - 1\right|$. However, because Model 2 typically shrinks the simple means toward more realistic values than Model 1—since the model-implied means in Model 2, based on additive effects, are generally more accurate for most intersections than the overall mean used in Model 1—the Model 2 shrinkage factors are typically applied to smaller differences between the simple means and model-implied means than in Model 1. Whether the resulting degree of shrinkage is generally larger or smaller in Model 1 or Model 2 depends on the specific application. In the presence of a very strong two-way interaction, the model-implied means from a more complex model that includes this interaction at a fixed-effect covariate will be even more realistic. However, even when Model 2 is misspecified by omitting this interaction, its model-implied means will generally provide more realistic values to shrink toward than simply shrinking to the overall intercept from Model 1.

Illustrative Application: Keller et al. (2023)

Figure 2 focuses on the simple mean and the MAIHDA means from Models 1 and 2 for Intersection 37 (male, native, low parental education, high occupational status; n = 14). The plot illustrates how the MAIHDA means for Models 1 and 2 (triangles) are reliability-weighted averages of the simple mean (circles) and their respective model-implied means (squares). For

Model 1, the model-implied mean is the overall mean, whereas for Model 2, it is the mean implied by the estimated additive effects of being male, native, low parental education and high occupational status.

In Model 1, the reliability of the simple mean is low (0.73) as it is based on only 14 individuals, and the simple mean of 463 is some 15 points lower than the model-implied mean of 478. As a result, the simple mean is shrunk upwards by 4 points towards the model-implied mean, resulting in a Model 1 MAIHDA mean of 468.

In Model 2, the conditional reliability of the simple mean is very low (0.18), and the simple mean of 463 is some 42 points lower than the model-implied mean of 505. As a result, the prediction is shrunk upwards by a very large 35 points towards the model-implied mean, resulting in a Model 2 MAIHDA mean of 497.

Reconsider the different predicted means for the six intersections plotted in Figure 1 and listed in Table 2. In Model 1, the MAIHDA means represent the simple means shrunk toward the overall model-implied mean of 478. The reliabilities and therefore the degree of multiplicative shrinkage vary as a function of intersection size from 0.36 to 0.73. In Model 2, the MAIHDA means are instead shrunk toward intersection-specific model-implied means derived from the additive model specification. The conditional reliabilities vary from 0.05 to 0.18. As expected, in both models, the MAIHDA means fall between the simple means and the model-implied means.

ANALYTIC EXPRESSIONS

Thus far, we have defined the Simple Means as well as the MAIHDA Model 1 and Model 2 means, shown that they yield different results, and demonstrated that these differences can be substantial—particularly for small intersectional groups—in a real-world setting. The remainder of the paper aims to evaluate and clarify the statistical properties of these three approaches to predicting intersectional group means. In doing so, we address our motivating questions: How much more accurate are MAIHDA means compared to simple means? What factors influence their relative accuracy? And, most importantly, is the choice between simple and MAIHDA means likely to affect conclusions in real-world research?

We use analytic expressions rather than simulations because they allow us to precisely identify the sources of divergence between the approaches and to assess their performance efficiently across a wide range of scenarios, in a way that prior simulation-based studies have struggled to do so. As with simulation studies, however, analytic approaches require specification of a “true” model or data-generating process (DGP) to serve as a benchmark. It is therefore important that the assumed DGP be a plausible reflection of reality.

MAIHDA Model 1 would be a poor choice, as it assumes no consistent (additive) patterns in intersectional inequalities. The Simple Means model is even less plausible, as it treats each intersectional group as entirely independent—an “island”—offering no information about any other group. In contrast, MAIHDA Model 2 provides a far more realistic basis, combining additive structure in intersectional inequalities with intersection specific deviations. Empirical studies often find that 60–95% of the variation in intersectional inequalities is accounted for by additive effects, supporting the assumptions embedded in Model 2. Moreover, via shrinkage, MAIHDA Model 2 leverages information across groups, reflecting the idea that every intersectional group shares information with others. We therefore assume that the true DGP follows MAIHDA Model 2.

Of course, the true DGP is unknown in any applied study and will vary from case to case. It will also inevitably differ from MAIHDA Model 2 in some respects. The key point is that

Model 2 will more closely resemble most real-world DGPs than either Model 1 or the Simple Means model.

STATISTICAL PROPERTIES OF THE DISTRIBUTION OF PREDICTED INTERSECTION MEANS

We begin our comparison of the statistical properties of three approaches to predicting intersection means by focusing on the big picture—how the approaches vary in capturing the distribution of intersection means.

We start by presenting analytical expressions for the variance of the intersection means across the distribution of intersections. We then graphically demonstrate how these variances change with intersection size and how these relationships are influenced by the Model 1 VPC, the Model 2 PCV, and the variability of intersection sizes. We operationalize the latter in terms of the Coefficient of Variation (CV) of intersection sizes, defined as the ratio of the standard deviation of intersection sizes to the mean. A CV of 0 means all intersections are of equal size. But most applications have unequal intersection sizes and so CVs greater than 0. Finally, we provide analytical expressions for the correlation between each set of means and the true means, and illustrate these correlations graphically in a parallel manner. We do not present analytical expressions for the expectation of each set of means as in all cases the expectation is equal to the mean of the intersection means.

The analytic expressions we present follow from results in the statistical literature on multilevel models. To derive all expressions, we specify the true model as Model 2 and assume that the regression coefficients, intersection random effect variance, and individual residual variance are known, with only the true mean for each intersection being unknown. Thus, all derivations are done in terms of the true parameter values of Model 2, rather than in terms of frequentist or Bayesian estimators. Additionally, we assume many intersections of equal size. We present full derivations in the Supplemental Material. We use simulation to address the case of varying intersection sizes. For plotting purposes, we consider the case where the mean and SD of the outcome are 0 and 1, respectively.

Variance of the Predicted Means Across the Distribution of Intersections

Let σ_μ^2 denote the variance of the true means. The variance of each set of predicted means can then be written as:

Simple: $$\text{Var}(\bar{y}_{.j}) = \sigma_\mu^2 + \frac{\sigma_e^2}{n} \quad (17)$$

Model 1: $$\text{Var}(\tilde{\mu}_{1j}) = R_1\sigma_\mu^2 \quad (18)$$

Model 2: $$\text{Var}(\tilde{\mu}_{2j}) = \sigma_\mu^2 - (1 - R_2)\sigma_{u2}^2 \quad (19)$$

where we continue to use 1 and 2 subscripts to distinguish between the Model 1 and Model 2 parameters, with the exception for σ_e^2, as $\sigma_{e1}^2 = \sigma_{e2}^2$. Since we assume balanced data, we omit the j subscript from n, R_1 and R_2.

Figure 3 shows the true variance and the variances of the Simple, Model 1, and Model 2 predicted means, plotted against intersection size. The figure is based on a Model 1 VPC of 0.15, a Model 2 PCV of 0.90, and a CV of intersection sizes of 0, indicating equal intersection sizes. These VPC and PCV values are similar to those reported in Keller et al. (VPC = 0.16 and PCV = 0.91). The CV value differs from that reported in Keller et al., where the CV was 0.98. We plot the four variances against intersection size only up to $n = 100$ because the differences between the variances diminish as intersection size increases. In Keller et al., the sizes of the 40 intersections ranged from 3 to 512 students, with a median of 93. Six intersections had fewer than 20 students, and three had fewer than 10.

Given the Model 1 VPC of 0.15 and an outcome variance of 1, it follows that the variance of the true means is also 0.15. This variance remains constant regardless of intersection size. In contrast, the variance of the simple means is initially much larger but decreases as intersection size increases, eventually converging to the true variance. This follows from the expression: $\mathrm{Var}(\bar{y}_{.j}) = \sigma_\mu^2 + \sigma_e^2/n$. As $n \to \infty$, the term $\sigma_e^2/n \to 0$, leading to $\mathrm{Var}(\bar{y}_{.j}) \to \sigma_\mu^2$. In words, the variance of the simple means is equal to the sum of the true variance and the sampling variance in the simple mean. The latter tends to 0 as the number of individuals per intersection increases. Thus, with sufficiently large intersection sizes, the variance of the simple means approaches that of the true means. However, when intersection sizes are smaller, the variance of the simple means overestimates the true variance.

The variance of the Model 1 means is smaller than the true variance when n is small, but it increases and converges to the true variance as intersection size increases. This follows from the expression: $\mathrm{Var}(\tilde{\mu}_{1j}) = R_1\sigma_\mu^2$. As $n \to \infty$, the term $R_1 \to 1$, leading to $\mathrm{Var}(\tilde{\mu}_{1j}) \to \sigma_\mu^2$. Thus, with sufficiently large intersection sizes, the reliabilities approach 1, and so the variance of the Model 1 means also closely approximates the variance of the true means. The fact that the variance of the Model 1 means is smaller than the true variance reflects the conservative nature of shrinkage—shrinkage leads to an underestimation of the true variance, particularly when intersection size is small. Importantly, for any given intersection size, while the variance of the Model 1 means is biased downward, it remains closer to the true variance than the variance of the simple means, which is biased upward.

The variance of the Model 2 means is slightly smaller than the true variance when n is small, but it increases and converges to the true variance as intersection size increases. Compared to Model 1, the variance of the Model 2 means remains close to the true variance even at the smallest intersection sizes. This is because the PCV is 0.90, indicating that 90% of the variation in the Model 2 means is captured by additive social identity effects, which are identified using data pooled across all intersections, and therefore reliably. Only 10% of the variation stems from intersection-specific deviations, which are identified separately for each intersection. As a result, relatively little data per intersection is required for the variance of the Model 2 means to approximate the true variance.

In summary, the variance of the simple means is biased upwards, whereas the variances of the Model 1 and Model 2 means are biased downwards. The magnitude of all three biases decreases as intersection size increases. However, the variance of the Model 2 means is always closest to the variance of the true means (i.e., shows the smallest bias). Thus, regardless of intersection size, the Model 2 means are the preferred means to report. These results have implications for the reporting of tables of intersection means.

Figure 3 was plotted using specific values for the Model 1 VPC, Model 2 PCV, and CV of intersection size. Figures S1, S2, and S3 in the Supplemental Material explore how this plot changes as each of these factors varies. The key finding is that the relative ranking of the three methods of prediction remains unchanged. However, the advantage of the Model 2 mean over the simple mean is most pronounced when intersectional inequalities are less pronounced (low VPC), when inequalities follow a largely additive pattern (high PCV), and when intersections vary in size (high CV).

Correlation Between the Predicted and True means Across the Distribution of Intersections

The correlation of each set of predicted means with the true means is given by:

Simple: $$\mathrm{Corr}(\bar{y}_{.j}, \mu_j) = \sqrt{R_1} \quad (20)$$

Model 1: $$\mathrm{Corr}(\tilde{\mu}_{1j}, \mu_j) = \sqrt{R_1} \quad (21)$$

Model 2: $$\mathrm{Corr}(\tilde{\mu}_{2j}, \mu_j) = \sqrt{\mathrm{PCV} + R_2(1 - \mathrm{PCV})} \quad (22)$$

The higher the correlation, the more the set of predicted means reflect the distribution of true means.

Figure 4 shows the correlation between each set of predicted means and the true means, plotted against intersection size. The figure is based on a Model 1 VPC of 0.15, a Model 2 PCV of 0.90, and a CV of intersection size of 0, indicating equal intersection sizes.

The correlation between the simple means and the true means is initially much smaller than 1 but increases as intersection size grows, eventually converging to a correlation of 1. The low correlation at smaller intersection sizes reflects the low reliability of the simple means as a predictor of the true means when intersection sizes are small. As $n \to \infty$, $R_1 \to 1$, leading to $\mathrm{Corr}(\bar{y}_{.j}, \mu_j) \to 1$. Thus, with sufficiently large intersection sizes, the correlation between the simple means and the true means approaches 1.

Figure 4 shows the correlation between the Model 1 means and the true means is identical to the correlation between the simple means and the true means (their respective line plots lie on top of one another). This is because, in the current case of equal intersection sizes, the Model 1 means are just the simple means shrunk towards the overall mean by a common reliability statistic. The two correlations will diverge when we allow the intersection sizes to vary (as seen in Supplemental Figure S6 and discussed below).

The correlation between the Model 2 means and the true means also converges to a correlation of 1 as intersection sizes increase. Compared to Model 1, the correlation between the Model 2 means and the true means remains close to 1 even at the smallest intersection sizes. This is due to the PCV being 0.90, which means that 90% of the variation in the Model 2 means is explained by additive social identity effects, identified using data pooled across all intersections, and therefore reliably. Only 10% of the variation stems from intersection-specific deviations, which are identified separately for each intersection. As a result, relatively little data per intersection is needed for the correlation between the Model 2 means and the true means to be close to 1.

In summary, all three correlations are biased downwards but approach 1 as intersection size increases. The correlation between the Model 2 means and the true means consistently shows the smallest bias. Therefore, the Model 2 means are the preferred predictor to report, regardless of intersection size.

Figure 4 was plotted using specific values for the Model 1 VPC, Model 2 PCV, and CV of intersection size. Figures S4, S5, and S6 in the Supplemental Material explore how this plot changes as each of these factors varies. The key finding is that the Model 2 means remain the preferred predictor. However, similar to the variance of the predicted means, the advantage of this predictor is most pronounced when intersectional inequalities are less pronounced (low VPC), when inequalities follow a largely additive pattern (high PCV), and when intersections vary in size (high CV).

STATISTICAL PROPERTIES OF PREDICTED INTERSECTION MEAN FOR A GIVEN INTERSECTION

In this section, we continue to address the same two key questions: Which set of predicted means – Simple, Model 1, or Model 2 – is most appropriate to report, and does this choice depend on the nature of intersectional inequalities – magnitude and additivity – being studied and the size of the data? However, our goal here is to examine the statistical properties of the predicted mean *for a given intersection* across random samples of individuals from that intersection. This contrasts the previous section where we focused on the statistical properties of the distribution of predicted intersection means across the population of intersections.

We focus on the bias, variance, and mean squared error (MSE) of each predictor for a given intersection. *Bias* measures the difference between the average predicted mean for a given intersection and its true value. A high bias, whether positive or negative, means the predictor consistently over- or under-predicts. *Variance* here reflects how much the predicted intersection mean varies across different samples of individuals. In other words, if we were to repeat a sampling process multiple times from the same intersection population, how varied would our predictions be? High variance indicates that the model is overly sensitive to the specific individuals randomly selected for the sample. Importantly, when comparing predictors there is typically a *bias-variance tradeoff* whereby the only way to reduce the variance is by introducing bias. The *MSE* therefore combines both bias and variance to assess overall prediction accuracy, with a lower MSE indicating better predictive performance.

We begin by presenting analytic expressions for the bias, variance, and MSE. As in the previous section, the analytic expressions we present follow from results in the statistical literature on multilevel models. We again present full derivations in the Supplemental Material. After presenting analytic expressions for bias, variance, and MSE, we graphically illustrate how these properties change as a function of the difference between the true intersection mean and the model-implied mean and how these relationships are influenced by intersection size, the Model 1 VPC, and the Model 2 PCV.

To derive all expressions, we continue to specify the true model as Model 2 and assume that the regression coefficients, intersection random effect variance, and individual residual variance are known, with only the true mean for each intersection being unknown. For plotting purposes, we again consider the case where the mean and SD of the outcome are 0 and 1, respectively.

Bias

For each predictor, the bias for intersection j measures the difference between the expected value for the predicted mean and the true mean for that intersection:

Simple: $$\text{bias}\left(\bar{y}_{.j} \middle| \mu_j\right) = 0 \tag{23}$$

Model 1: $$\text{bias}\left(\tilde{\mu}_{1j} \middle| \mu_j\right) = -\left(1 - R_{1j}\right)\left(\mu_j - \beta_{1,0}\right) \tag{24}$$

Model 2: $$\text{bias}\left(\tilde{\mu}_{2j} \middle| \mu_j\right) = -\left(1 - R_{2j}\right)\left\{\mu_j - \left(\beta_{2,0} + \beta_{2,1}x_{1j} + \cdots + \beta_{2,p}x_{pj}\right)\right\} \tag{25}$$

where the 1 and 2 subscripts continue to distinguish between the Model 1 and Model 2 parameters.

Figure 5 presents the bias of each predictor as a function of the difference between the true mean μ_j and the model-implied mean for that intersection: $\beta_{1,0}$ in Model 1 or $\beta_{2,0} + \beta_{2,1}x_{1j} + \cdots + \beta_{2,p}x_{pj}$ in Model 2. The figure assumes an intersection size of 10 individuals, a Model 1 VPC of 0.15, and a Model 2 PCV of 0.90. It follows that the Model 1 and 2 reliabilities are 0.64 and 0.15. Notably, in the Keller et al. study, six of the 40 intersections had fewer than 20 students, and three had fewer than 10 students. The x-axis range of the plots for the Simple and Model 1 bias are limited to the expected middle 95% of differences from Model 1, given by

$\pm\Phi^{-1}(0.975)\sqrt{\sigma_{u1}^2}$. Similarly, the x-axis range for the plot of the Model 2 bias is restricted to the expected middle 95% of differences from Model 2, given by $\pm\Phi^{-1}(0.975)\sqrt{\sigma_{u2}^2}$. We apply these limits to avoid disproportionately long bias lines (and variance and MSE lines in Figures 6 and 7) at the extremes, which could distract the reader from the main patterns within the central range of the data.

The simple mean is an unbiased predictor of the true mean, meaning that, averaging across repeated samples, the simple mean equals the true mean.

The Model 1 mean is a biased predictor. The bias is a negative linear function of the difference between the true mean μ_j and the model-implied mean $\beta_{1,0}$, where the latter is simply the overall mean. Specifically, this difference is scaled by $-(1 - R_{1j})$, which is negative and, in this illustration takes the value -0.36 (with a possible range from -1 to 0). As a result, the sign of the difference is reversed, and its magnitude is reduced. Consequently, when the true mean exceeds the model-implied mean, the Model 1 mean is biased downward, whereas when the true mean is lower than the model-implied mean, the bias is upward. Thus, the Model 1 is biased toward the model-implied mean. This once again reflects the conservative nature of shrinkage.

The Model 2 mean is also a biased predictor of the true mean, with the magnitude of its bias again following a negative linear relationship with the difference between the true mean μ_j and the model-implied mean $\beta_{2,0} + \beta_{2,1}x_{1j} + \cdots + \beta_{2,p}x_{pj}$. However, unlike in Model 1, the model-implied mean is now an intersection-specific mean, determined by the additive effects of the relevant social identity characteristics. As a result, these model-implied means are typically closer to the true means than in Model 1. This increases the first term in absolute magnitude and decreases the second term in the expression for the bias. Specifically, the first term, $-(1 - R_{2j})$, is now -0.85. It is now closer to -1 because the conditional reliability R_{2j} is always lower than the unconditional reliability R_{1j}. This results in a larger per-unit multiplicative bias, making the slope of the linear bias relationship steeper than before. However, the decrease in the second term—the deviation between the true and model-implied mean—means this increased multiplier is applied to smaller differences than in Model 1, making the range over which the linear bias relationship operates narrower (note the narrower range of the line plot for Model 2 in Figure 5). The net effect is that the Model 2 mean is expected to be a less biased predictor of the true mean for a given intersection than the Model 1 mean. Nevertheless, the Model 2 mean remains a biased predictor, and one that is biased towards additivity. That is, the predicted intersection means will conform more to additivity than they do in actuality.

In summary, the simple mean is an unbiased predictor of the true intersection mean, whereas the Model 1 and Model 2 means are both biased towards their respective model-implied means with the Model 1 means, on average, exhibiting greater bias.

Figures S7, S8, and S9 in the Supplemental Material explore how the results shown in Figure 5 change as a function of intersection size, Model 1 VPC, and Model 2 PCV. The key finding is that the simple mean remains the preferred predictor as it is always unbiased (when averaged across repeated samples), whereas the Model 1 and Model 2 means exhibit bias in all cases except when the true and model-implied means are equal.

Variance

For each predictor, the variance for intersection j measures the variability of the prediction (their "consistency") across random samples of individuals from that intersection:

Simple: $$\text{Var}\left(\bar{y}_{.j} \middle| \mu_j\right) = \frac{\sigma_e^2}{n_j} \quad (26)$$

Model 1: $$\text{Var}\left(\tilde{\mu}_{1j} \middle| \mu_j\right) = R_{1j}^2 \frac{\sigma_e^2}{n_j} \quad (27)$$

Model 2: $$\mathrm{Var}\left(\tilde{\mu}_{2j} \middle| \mu_j\right) = R_{2j}^2 \frac{\sigma_e^2}{n_j} \tag{28}$$

Figure 6 presents the variance of each predictor for the true mean of intersection j as a function of the difference between the true mean and the model-implied mean for that intersection. The plot's range is again restricted to the expected middle 95% of differences for each model. As with the bias figure, the analysis assumes an intersection size of 10 individuals, a Model 1 VPC of 0.15, and a Model 2 PCV of 0.90.

The variance of each predictor is constant. For the simple mean, the variance is given by the sampling variance across repeated samples, $\frac{\sigma_e^2}{n_j}$. The variance of the Model 1 mean is obtained by multiplying this sampling variance by the squared reliability, R_{1j}^2 which in this illustration is 0.38. As $R_{1j} < 1$, the Model 1 mean exhibits less variation across repeated samples than the simple mean —specifically, 62% less. The variance of the Model 2 mean follows the same general form as the Model 1 mean. However, because the conditional reliability is lower than the unconditional reliability, $R_{2j} < R_{1j}$, the variance of the Model 2 mean is even smaller than that of Model 1. In this illustration, $R_{2j}^2 = 0.0225$, meaning the variance of the Model 2 mean is 97.5% smaller than that of the simple mean.

In summary, the Model 2 mean has the lowest variance (the most consistency) across repeated samples, followed by the Model 1 mean, with the simple mean exhibiting the highest variance (least consistency). Thus, the Model 2 means to a much greater extent replicate themselves across repeated samples. The simples means do not.

Figures S10, S11, and S12 explore how the results shown in Figure 6 change as a function of the intersection size, Model 1 VPC, and Model 2 PCV. The key finding is that the Model 2 mean remains the preferred predictor in terms of having the smallest variance.

However, its advantage over the simple mean is most pronounced when the intersection size is small, intersectional inequalities are less pronounced (low VPC), and when these inequalities follow a largely additive structure (high PCV).

Mean Squared Error

For each predictor, the mean squared error (MSE) for intersection j reflects both its variance and bias. As a result, the MSE provides a measure of the predictor's overall accuracy. A lower MSE indicates a more accurate predictor. Specifically, The MSE is defined as the sum of the variance and the bias squared:

$$\text{Simple:} \quad \text{MSE}(\bar{y}_{.j}|\mu_j) = \frac{\sigma_e^2}{n_j} \tag{29}$$

$$\text{Model 1:} \quad \text{MSE}(\tilde{\mu}_{1j}|\mu_j) = R_{1j}^2 \frac{\sigma_e^2}{n_j} + (1 - R_{1j})^2 (\mu_j - \beta_{1,0})^2 \tag{30}$$

$$\text{Model 2:} \quad \text{MSE}(\tilde{\mu}_{2j}|\mu_j) = R_{2j}^2 \frac{\sigma_e^2}{n_j} + (1 - R_{2j})^2 \{\mu_j - (\beta_{2,0} + \beta_{2,1}x_{1j} + \cdots + \beta_{2,p}x_{pj})\}^2 \tag{31}$$

Figure 7 presents the MSE of each predictor for the true mean of intersection j as a function of the difference between the true mean and the model-implied mean for that intersection. The plot's range is again restricted to the expected middle 95% of differences for each model, and the analysis continues to assume an intersection size of 10 individuals, a Model 1 VPC of 0.15, and a Model 2 PCV of 0.90.

The MSE for the Simple Mean is equal to its variance, as its bias is zero.

The MSE for the Model 1 mean follows a quadratic function of the difference between the true mean and the model-implied mean, $\mu_j - \beta_{1,0}$, where the model-implied mean is simply the overall mean. The MSE reaches its minimum when the true mean matches the model-implied mean, at which point it equals the variance of the Model 1 mean, $R_{1j}^2 \frac{\sigma_e^2}{n_j}$. As the true mean deviates from the model-implied mean, the MSE increases.

Similarly, the MSE for the Model 2 mean is also a quadratic function of the difference between the true mean and the model-implied mean. However, in this case, the model-implied mean represents the intersection-specific mean, determined by the additive effects of the relevant social identity characteristics, $\beta_{2,0} + \beta_{2,1}x_{1j} + \cdots + \beta_{2,p}x_{pj}$. As with the Model 1 mean, the MSE is minimized when the true mean matches the model-implied mean, at which point it equals the variance of the Model 2 mean, $R_{2j}^2 \frac{\sigma_e^2}{n_j}$. Since this variance is lower than that of the Model 1 mean, the MSE for Model 2 at this point is also lower.

While the MSE of the Model 2 mean increases more rapidly than that of the Model 1 mean as the true mean increasingly deviates from the model-implied mean, this effect is mitigated by the fact that the model-implied means in Model 2 are generally much closer to the true means than those in Model 1. As a result, the MSE for the Model 2 mean will almost always be smaller than that for the Model 1 mean for a given intersection.

In summary, when the model-implied mean is closer to the true mean, both the Model 1 and, especially, Model 2 predictors exhibit lower MSE than the simple mean. As the model-implied mean deviates further from the true mean, the MSE for both models increases. However, the ranking of prediction method remains consistent, except in cases of extreme deviations. For instance, consider an intersection with a very low true mean of -0.7. Across repeated samples of 10 individuals, the MSE for the Model 1 mean at that intersection (0.099) would actually exceed that of the simple mean (0.085). However, the MSE for the Model 2 mean would be significantly lower.

Figures S13, S14, and S15 explore how the results shown in Figure 7 change as a function of the intersection size, Model 1 VPC, and Model 2 PCV. The key finding is that the Model 2 mean remains the preferred predictor in terms of the lowest MSE. However, its advantage over the simple mean is most pronounced when intersection sizes are small, intersectional inequalities are less pronounced (low VPC), and the inequalities follow an additive structure (high PCV). These results align with the variance patterns observed earlier.

DISCUSSION

This article examines the claim that the predicted intersection means derived from a MAIHDA analysis provide better predictions of population intersection means than the simple arithmetic means. This claim is made in many applications of MAIHDA (Evans et al. 2024b) based on findings from the statistical literature on multilevel models, but until now this claim has not been formally explored in the context of MAIHDA analyses.

Our conclusion is that the predicted means from MAIHDA Model 2 are statistically more accurate than those from MAIHDA Model 1, and both outperform the simple means. This conclusion assumes that Model 2—which treats the additive main effects as fixed effects and all remaining two-way and higher-order interaction variability as an intersection-specific random effect—adequately represents the true data-generating process (DGP), a point we elaborate on below. Specifically, the Model 2 means exhibit the closest variance to, and the highest correlation with, the true means. While the Model 2 mean for a given intersection is biased (unlike the simple mean), it displays much lower variance across repeated samples, resulting in a lower MSE and, therefore, greater predictive accuracy. In essence, we accept a slight bias toward additivity in exchange for a substantial reduction in statistical noise. The improved statistical accuracy of the Model 2 means over the simple means arises from their ability to optimally combine the imprecise, intersection-specific simple means with precise, model-implied means that draw on information from all intersections.

This ranking of prediction methods holds across several key factors, including intersection size, the magnitude of these inequalities as measured by the variance partitioning coefficient (VPC), the extent to which inequalities are driven by the additive effects of the social identities that form the intersections, as indicated by the proportion change in variance (PCV), and the variation in intersection sizes, measured by the coefficient of variation (CV).

Importantly, when all intersection sizes are large, the differences between the three methods diminish, making even the simple means reasonably accurate. However, when at least some intersection sizes are small, the difference in predicted intersection means become pronounced, particularly when the VPC is low and the PCV is high. Crucially, small intersection sizes, low VPCs, and high PCVs are very common in empirical MAIHDA applications (Evans et al. 2024a). Indeed, the intersections of very greatest interest—those representing multiply marginalized groups—tend to be the very smallest, and so most sensitive to choice of prediction method.

Our analytical expressions are derived under the assumption that Model 2 is the data-generating process (DGP). In any real-world application, however, Model 2 can only approximate the unknown true DGP. Consequently, the greater the divergence of the real-world DGP from Model 2, the more likely it is that a more complex model could yield statistically more accurate estimates of intersectional means. For instance, in a particular application where a specific two-way interaction—such as that between gender and race—is substantial and supported by sufficient data, the true DGP might be better approximated by modeling this interaction as a fixed effect regression coefficient, rather than implicitly treating it as random, as in our current formulation.

However, the increase in predictive accuracy from specifying such a two-way interaction as fixed is likely to be relatively small compared to that achieved with the MAIHDA Model 2 means—and certainly smaller than the gains observed when moving from Simple Means to MAIHDA Model 1 means, and then to Model 2. The intuition is that, in most applications, the additive main effects in MAIHDA Model 2 account for approximately 60–95% of the variation in intersectional inequalities—a substantial to very large majority. This leaves limited room for further improvements in predictive accuracy by explicitly modelling specific interactions rather than absorbing them into the random effect. We illustrate this argument using a specific fixed two-way interaction DGP in the Supplemental Materials.

Nonetheless, it is important to further explore these and other arguments. In particular, it will likely be most informative to examine how the increase in predictive accuracy for a given intersection varies depending on how that intersection is involved in the specific two-way interaction. A key challenge in such explorations is that, as the true model becomes more complex, analytical derivations become increasingly difficult. In these cases, simulation studies will likely be necessary to assess the statistical properties of the resulting approaches.

As our analytic expressions relate to these true parameters, we have largely abstracted from choice of estimation method. However, for the purpose of inference, this choice is important, as different estimation methods—frequentist approaches such as maximum likelihood (MLE) and restricted maximum likelihood (REML) followed by empirical Bayes prediction, versus Bayesian methods such as MCMC—differ in how they propagate uncertainty in the estimated parameters into the predicted intersection effects. ML does not propagate uncertainty in the estimation of any model parameters. REML propagates uncertainty in the regression coefficients but not in the variance components. In contrast, Bayesian methods propagate uncertainty in all parameters. As a result, in any applied analysis, REML—and especially MLE—tend to understate uncertainty around the predicted intersection means. This would be expected to lead to Type I errors of inference—declaring particular intersections to show statistically significant deviations from additivity when they don't. In this regard, Bayesian methods are preferable for MAIHDA analyses. Here too, simulation studies could be used to explore and quantify these differences and their implications.

While we focused on MAIHDA models for continuous outcomes, we see no reason why our findings will not apply more generally to MAIHDA models for binary (Evans et al. 2024b; Mahendran et al. 2022b) or other categorical, count, or survival outcomes. Similarly, we see no reason why they will not also apply to longitudinal MAIHDA models which fit separate mean trajectories for each intersection (Bell et al. 2024), or other random slope models. Future studies should confirm these intuitions. Our findings are also applicable when MAIHDA is used beyond intersectionality research, in the broader study of high-dimensional interactions among multiple categorical variables. Such applications have been termed multicategorical MAIHDA, to distinguish them from those focused specifically on intersectionality (Evans, 2024; Rodriguez-Lopez, Leckie, Kaufman, & Merlo, 2023).

Viewed more broadly, MAIHDA applies shrinkage to simple means through multilevel modeling. When recast as a linear regression of the outcome on a full set of intersection dummy variables, this process shrinks the estimated coefficients toward the overall mean, effectively functioning as a form of regularized regression. This perspective highlights the potential of alternative approaches—such as LASSO, Ridge regression, or elastic net—to achieve similar goals whether for continuous or other outcome types (Hastie, Tibshirani, and Wainwright, 2015). By fitting these models and using their predicted values, one will typically again obtain improved estimates of intersection means over the simple means. Future work might therefore also explore how these regularization methods compare to MAIHDA in terms of predictive accuracy and interpretability.

We have highlighted the benefits of shrinkage for predicting intersection means. However, shrinkage introduces a potential theoretical tension. Intersectionality emphasizes that unique combinations of social identities give rise to distinct lived experiences and, consequently, different outcomes. Yet, shrinkage pulls estimates for specific intersections toward the overall pattern across all intersections, which may seem to undermine the uniqueness of those experiences. This is the bias–variance tradeoff. To achieve greater predictive accuracy, we must accept some bias in exchange for substantially reduced variance. In our analysis, we have used minimization of the MSE as the criterion for balancing bias and variance. However, we acknowledge that some intersectionality researchers may wish to prioritize unbiasedness over predictive accuracy. In such cases, simple means may be preferable.

In sum, for optimal prediction of intersection means, we recommend using MAIHDA Model 2 over both Model 1 and simple means, especially when inequalities are subtle—whether in magnitude or due to hidden processes affecting specific combinations of social identities—and when data for certain intersections, such as those representing multiply marginalized groups, is limited.

REFERENCES

Alvarez, Camila H., Anna Calasanti, Clare Rosenfeld Evans, and Kerry Ard. 2022. "Intersectional inequalities in industrial air toxics exposure in the United States." *Health & Place*, 77, 102886.

Bauer, Greta R., Siobhan M. Churchill, Mayuri Mahendran, Chantel Walwyn, Daniel Lizotte, and Alma Angelica Villa-Rueda. 2021. "Intersectionality in quantitative research: A systematic review of its emergence and applications of theory and methods." *SSM-Population Health 14*: 100798.

Bell, Andrew, Clare Evans, Daniel Holman, and George Leckie, G. 2024. “Extending intersectional multilevel analysis of individual heterogeneity and discriminatory accuracy (MAIHDA) to study individual longitudinal trajectories, with application to mental health in the UK.” *Social Science & Medicine* 351, 116955.

Bell, Andrew, Daniel Holman, and Kelvyn Jones. 2019. "Using shrinkage in multilevel models to understand intersectionality." *Methodology 15(2)*: 88–96.

Collins, Patricia Hill, and Sirma Bilge. 2020. *Intersectionality*. John Wiley & Sons.

Crenshaw, Kimberlé. 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*, pp. 139-167.

Evans, Clare R. 2024. “Overcoming Combination Fatigue: Addressing High-Dimensional Effect Measure Modification and Interaction in Clinical, Biomedical, and Epidemiologic Research Using Multilevel Analysis of Individual Heterogeneity and Discriminatory Accuracy (MAIHDA).” *Social Science & Medicine* 340:116493.

Evans, Clare R., Luisa N. Borrell, Andrew Bell, Daniel Holman, S. V. Subramanian, and George Leckie. 2024a. "Clarifications on the intersectional MAIHDA approach: A conceptual guide and response to Wilkes and Karimi (2024)." *Social Science & Medicine* 350: 116898.

Evans, Clare R., George Leckie, S. V. Subramanian, Andrew Bell, and Juan Merlo. 2024b. "A tutorial for conducting intersectional multilevel analysis of individual heterogeneity and discriminatory accuracy (MAIHDA)." *SSM-population health*: 101664.

Evans, Clare R., David R. Williams, Jukka-Pekka Onnela, and S. V. Subramanian. 2018. "A multilevel approach to modeling health inequalities at the intersection of multiple social identities." *Social Science & Medicine* 203: 64-73.

Forrest, Lauren N., et al. 2023. "Intersectional prevalence of suicide ideation, plan, and attempt based on gender, sexual orientation, race and ethnicity, and rurality." *JAMA psychiatry* 80(10): 1037-1046.

Giaconi, Valentina, Manuela Mendoza, Ngaire Honey, and Enoc Lemus-Cortez. 2024. "The intersection of gender, rurality, income and school track in the university access mathematics tests scores." *Frontiers in Education*, vol. 9, p. 1306426. Frontiers Media SA,.

Hastie, Trevor, Robert Tibshirani, and Martin Wainwright. 2015. *Statistical Learning with Sparsity: The Lasso and Generalizations*. Boca Raton, FL: CRC Press.

Humbert, Anne Laure. 2024. "Methodological insights into multilevel analysis of individual heterogeneity and discriminatory accuracy: An empirical examination of the effects of strata configurations on between-stratum variance and of fixed effects across hierarchical levels." *PLoS One* 19(3): e0297561.

Ivert, Anna-Karin, Enrique Gracia, Marisol Lila, Maria Wemrell, and Juan Merlo. 2020. "Does country-level gender equality explain individual risk of intimate partner violence against women? A multilevel analysis of individual heterogeneity and discriminatory accuracy (MAIHDA) in the European Union." *European Journal of Public Health* 30(2): 293-299.

Keller, Lena, Oliver Lüdtke, Franzis Preckel, and Martin Brunner. 2023. "Educational Inequalities at the Intersection of Multiple Social Categories: A n Introduction and Systematic Review of the Multilevel Analysis of Individual Heterogeneity and Discriminatory Accuracy (MAIHDA) Approach." *Educational Psychology Review* 35(1)1-31.

Lister, Jennie, Catherine Hewitt, and Josie Dickerson. 2024. "Using I-MAIHDA to extend understanding of engagement in early years interventions: an example using the Born in Bradford's Better Start (BiBBS) birth cohort data." *Social Sciences & Humanities Open 10*: 100935.

Mahendran, Mayuri, Daniel Lizotte, and Greta R. Bauer. 2022a. "Describing intersectional health outcomes: an evaluation of data analysis methods." *Epidemiology 33, no. 3*: 395-405.

Mahendran, Mayuri, Daniel Lizotte, and Greta R. Bauer. 2022b. "Quantitative methods for descriptive intersectional analysis with binary health outcomes." *SSM-population health 17*: 101032.

McCall, Leslie. 2005. "The complexity of intersectionality." *Signs: Journal of Women in Culture and Society 30.3*: 1771-1800.

Merlo, Juan. 2018. "Multilevel analysis of individual heterogeneity and discriminatory accuracy (MAIHDA) within an intersectional framework." *Social Science & Medicine 203*: 74-80.

Pina-Sánchez, Jose, and Ferhat Tura. 2024. "Intersectional disparities in sentencing: An application of MAIHDA to the US Federal Courts." *CrimRxiv*.

Pomeroy, Edward, and Francesca Fiori. 2025. "Competing demands on adult children: How do they shape their provision of informal care?." *SSM-Population Health* 29: 101754.

Prior, Lucy, Clare Evans, Juan Merlo, and George Leckie. 2022. "Sociodemographic inequalities in student achievement: An intersectional multilevel analysis of individual heterogeneity and discriminatory accuracy (MAIHDA)." *Sociology of Race and Ethnicity* 23326492241267251.

Prior, Lucy, and George Leckie. 2024. "Student intersectional sociodemographic and school variation in GCSE final grades in England following Covid-19 examination cancellations." *Oxford Review of Education* 1-22.

Raudenbush, Stephen W., and Anthony S. Bryk. 2002. *Hierarchical Linear Models: Applications and Data Analysis Methods*. 2nd ed. Thousand Oaks, CA: Sage.

Rodriguez-Lopez, Merida, George Leckie, Jay S. Kaufman, and Juan Merlo. 2023. "Multilevel Modelling for Measuring Interaction of Effects Between Multiple Categorical Variables: An Illustrative Application Using Risk Factors for Preeclampsia." Paediatric and *Perinatal Epidemiology* 37(2):154–164.

Snijders, Tom A. B., and Roel J. Bosker. 2012. *Multilevel Analysis: An Introduction to Basic and Advanced Multilevel Modeling*. 2nd ed. London: SAGE Publications.

Silva, Tony, and Clare R. Evans. 2020. "Sexual identification in the United States at the intersections of gender, race/ethnicity, immigration, and education." *Sex Roles* 83(11): 722-738.

Tura, Ferhat, Jane Healy, Clare R. Evans, and George Leckie. 2024. "An intersectional analysis of stranger, acquaintance, and domestic violence victimisation in England and Wales using MAIHDA." *CrimRxiv*.

Van Dusen, Ben, Heidi Cian, Jayson Nissen, Lucy Arellano, and Adrienne D. Woods. 2024. "Comparing the Efficacy of Fixed-Effects and MAIHDA Models in Predicting Outcomes for Intersectional Social Strata." *Sociology of Education* 00380407241254092.

TABLES

Table 1.

Results of MAIHDA Models 1 and 2 for reading achievement in the Keller et al. (2023) study.

	Model 1		Model 2	
	Est.	SE	Est.	SE
Intercept	478.3	[464.3, 492.3]	468.4	[452.8, 483.3]
Gender				
Female (Ref.)			–	
Male			-25.6	[-36.5, -14.9]
Immigrant background				
Native (Ref.)			–	
Immigrant			-42.5	[-53.7, -30.9]
Highest parental education				
Below uni. entrance certificate (Ref.)			–	
At least uni. entrance certificate			29.8	[18.2, 42.0]
Highest parental occupation status				
Low (Ref.)			–	
Low to middle			10.1	[-5.9, 26.8]
Middle			27.6	[11.4, 44.6]
Middle to high			53.1	[35.2, 70.4]
High			62.6	[40.5, 81.2]
Intersection variance	1698.3		144.4	
Student variance	9011.7		9021.2	

VPC	15.9%	1.6%
PCV	−	91.2%

Note. Table adapted from Table 3 of Keller et al. (2023). The model was estimated using Bayesian Markov chain Monte Carlo (MCMC) methods. Est. = estimate; SE = Standard Error; Uni. = University; Ref. = reference category; VPC = variance partition coefficient; PCV = proportional change in stratum variance.

Table 2.

Predicted intersection means for the six intersections with fewer than n = 20 individuals in the Keller et al. (2023) study.

							MAIHDA Model 1			MAIHDA Model 2		
ID	Gender	Migrant	Education	Occupation	Size	Simple mean	Rel.	Model implied mean	MAIHDA Mean	Rel.	Model implied mean	MAIHDA mean
j					n_j	$\bar{y}_{.j}$	R_j	$\beta_{0,1}$	$\tilde{\mu}_j$	R_j	$\mathbf{x}_j'\boldsymbol{\beta}_2$	$\tilde{\mu}_j$
15	Female	Immigrant	Below UEC	Mid to high	9	446	0.63	478	459	0.13	479	475
17	Female	Native	Below UEC	High	8	536	0.60	478	512	0.11	531	532
19	Female	Immigrant	Below UEC	High	3	412	0.36	478	456	0.05	489	485
35	Male	Immigrant	Below UEC	Mid to high	11	471	0.67	478	473	0.15	453	456
37	Male	Native	Below UEC	High	14	463	0.73	478	468	0.18	505	497
39	Male	Immigrant	Below UEC	High	3	362	0.36	478	438	0.05	463	458

Note. Table adapted from Table S3 of Keller et al. (2023). We calculated R_j for Model 1 and 2 by applying Equations (12) and (16) to the estimates of σ_u^2 and σ_e^2 (Table 1) and n_j. We then calculated $\bar{y}_{.j}$ by rearranging and applying Equations (10) and (14) in terms of

$\bar{y}_{.j}$, using the Model 1 and 2 values of R_j, $\tilde{\mu}_j$, and $\beta_{0,1}$ and $\mathbf{x}_j'\boldsymbol{\beta}_2$, respectively. Rel. – Reliability. Simple mean and reliability statistics are our own calculations; UEC = University Entrance Certificate. $\mathbf{x}_j'\boldsymbol{\beta}_2 = \beta_{2,0} + \beta_{2,1}x_{1j} + \cdots + \beta_{2,p}x_{pj}$.

FIGURES

Figure 1.

Predicted simple means, MAIHDA Model 1 means, and MAIHDA Model 2 means for all 40 intersections in the Keller et al. (2023) study. The six intersections with fewer than n = 20 individuals are emphasized both in the plot and in Table 2. Additionally, Intersection 37 is highlighted in the plot, as it is examined in more detail in Figure 2. The simple means are defined in (9), the MAIHDA Model 1 means in (10), and the MAIHDA Model 2 means in (14).

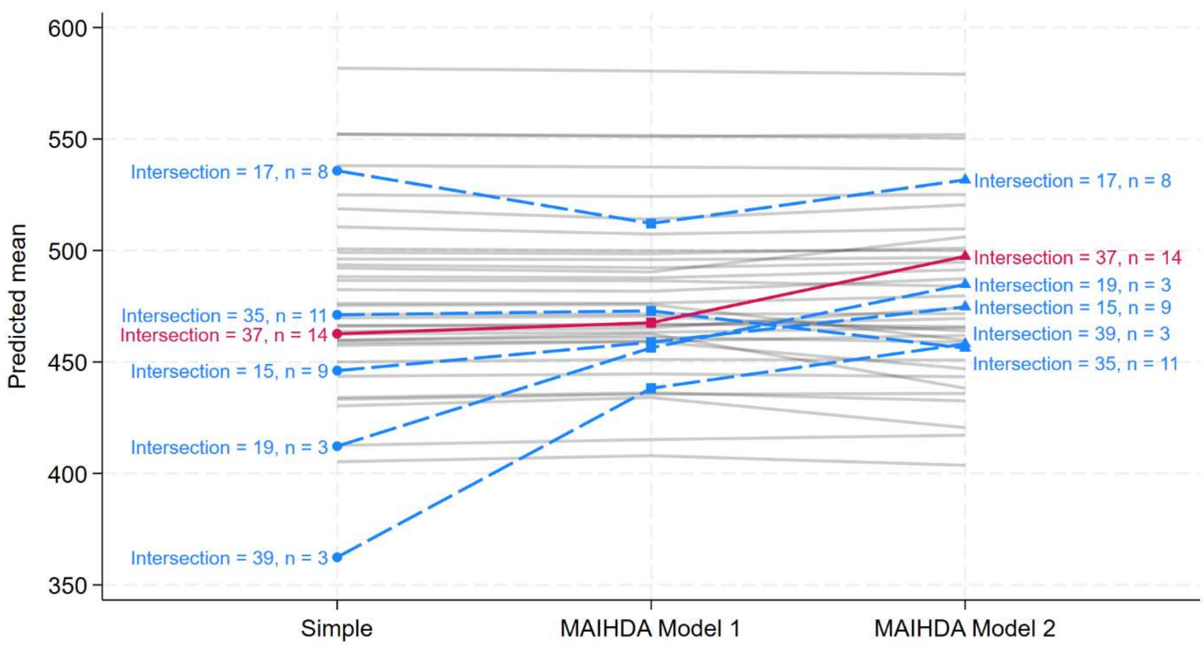

Figure 2.

Predicted simple mean, MAIHDA Model 1 mean, and MAIHDA Model 2 mean for Intersection 37 in the Keller et al. (2023) study. The plot illustrates how the MAIHDA Model 1 and Model 2 means are weighted averages of the simple mean and the corresponding model-implied mean. Intersection 37 includes male, native students with low parental education and high occupational status (n = 14).

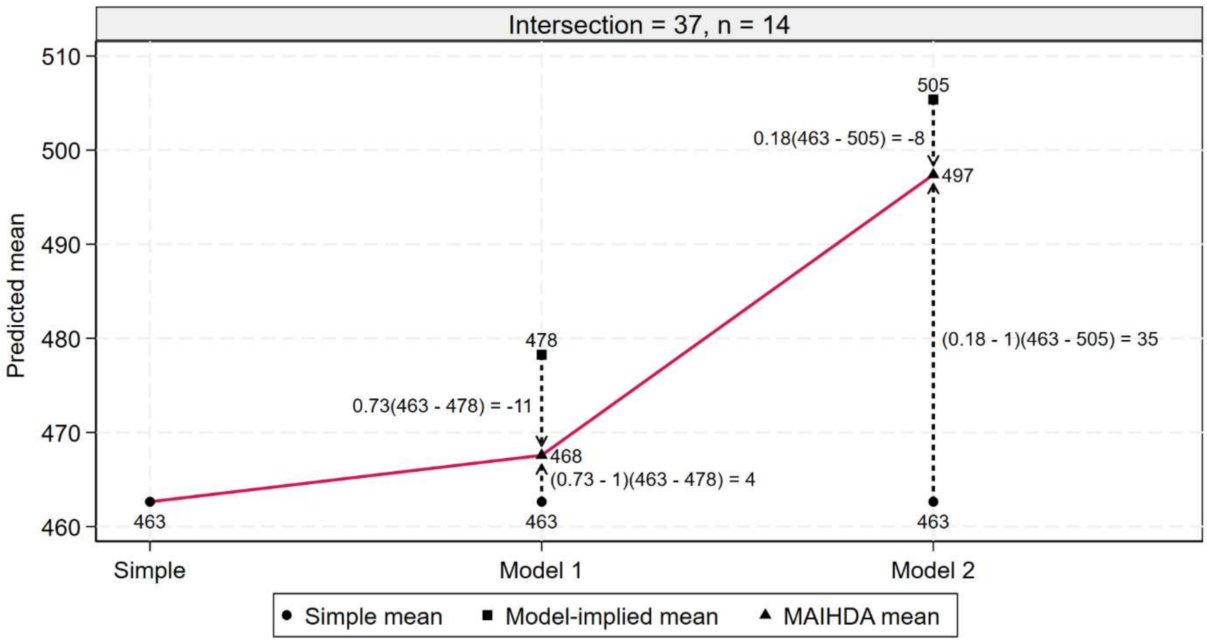

Figure 3.

Variance of the simple means (17), MAIHDA Model 1 means (18), and MAIHDA Model 2 means (19) across the distribution of intersections, plotted against intersection size. The plot assumes a Model 1 VPC of 0.15, a Model 2 PCV of 0.90, and equal intersection sizes.

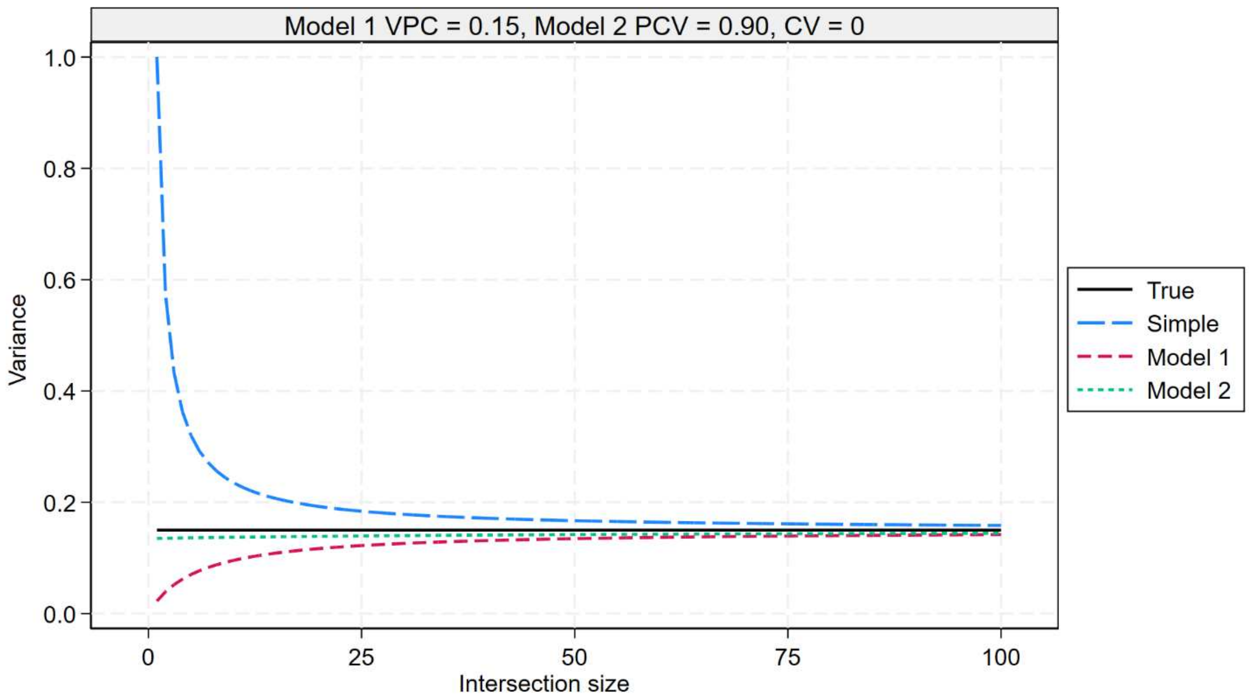

Figure 4.

Correlation between the simple means and the true means (20), MAIHDA Model 1 means and the true means (21), and MAIHDA Model 2 means and the true means (22), across the distribution of intersections, plotted against intersection size. The plot assumes a Model 1 VPC of 0.15, a Model 2 PCV of 0.90, and equal intersection sizes. The line plots for the simple means and the MAHDA Model 1 means lie on top of one another.

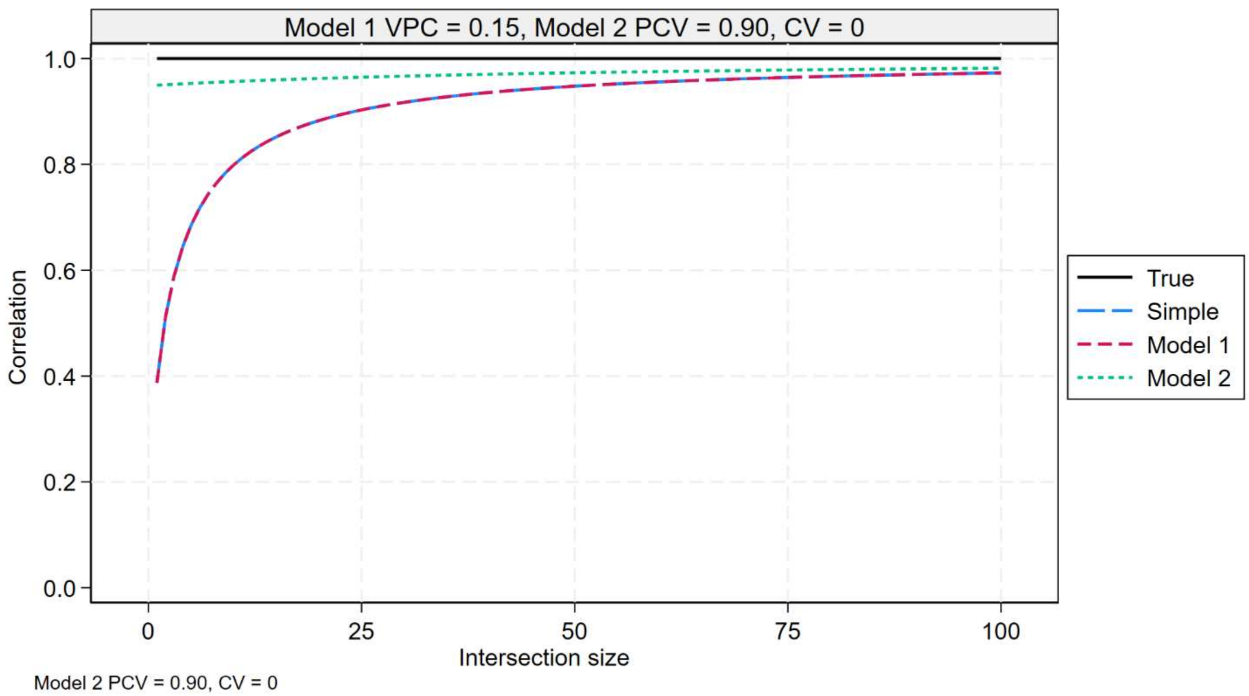

Figure 5.

Bias of the simple mean (23), MAIHDA Model 1 mean (24), and MAIHDA Model 2 mean (25) across repeated samples of individuals for a given intersection, plotted against the difference between the true mean and the model-implied mean for that intersection. The plot assumes an intersection size of 10 individuals, a Model 1 VPC of 0.15, and a Model 2 PCV of 0.90. The Model 1 model-implied mean corresponds to the overall mean, while the Model 2 model-implied mean reflects the additive effects of the social identities defining the intersection. The bias of the MAIHDA Model 2 mean is shown over a narrower range due to smaller differences between the true and model-implied means in this model.

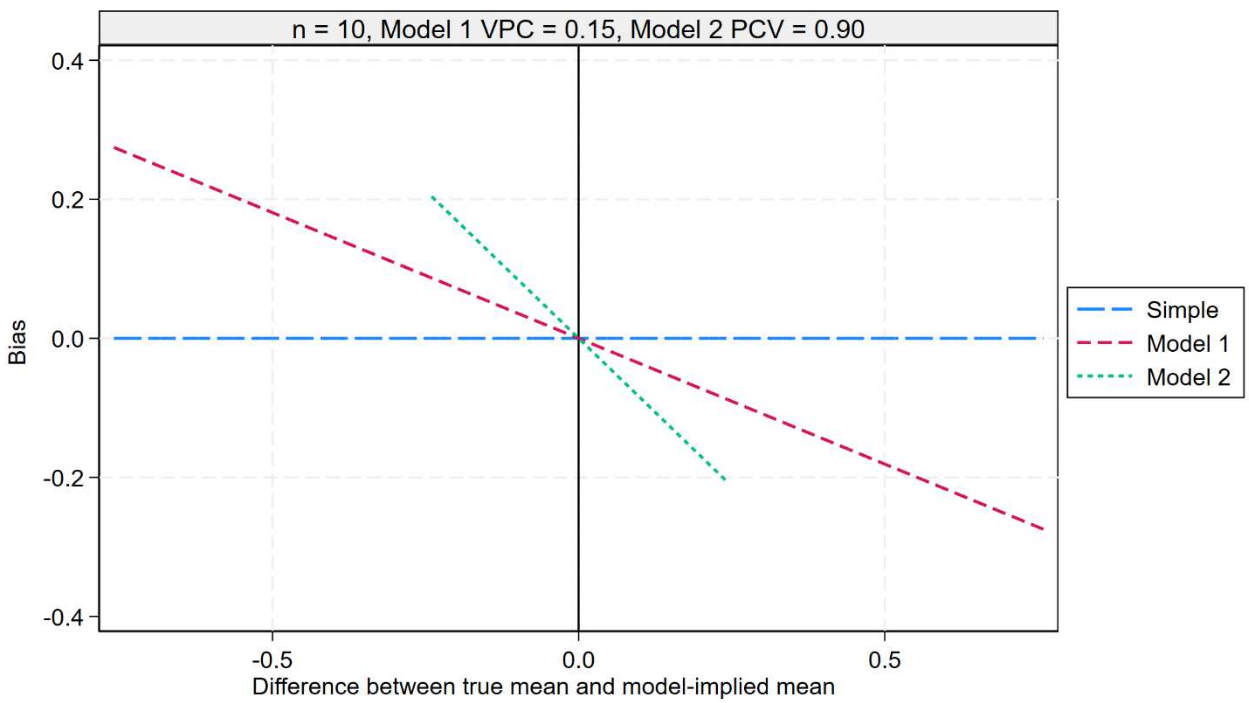

Figure 6.

Variance of the simple mean (26), MAIHDA Model 1 mean (27), and MAIHDA Model 2 mean (28) across repeated samples of individuals for a given intersection, plotted against the difference between the true mean and the model-implied mean for that intersection. The plot assumes an intersection size of 10 individuals, a Model 1 VPC of 0.15, and a Model 2 PCV of 0.90. The Model 1 model-implied mean corresponds to the overall mean, while the Model 2 model-implied mean reflects the additive effects of the social identities defining the intersection. The bias of the MAIHDA Model 2 mean is shown over a narrower range due to smaller differences between the true and model-implied means in this model.

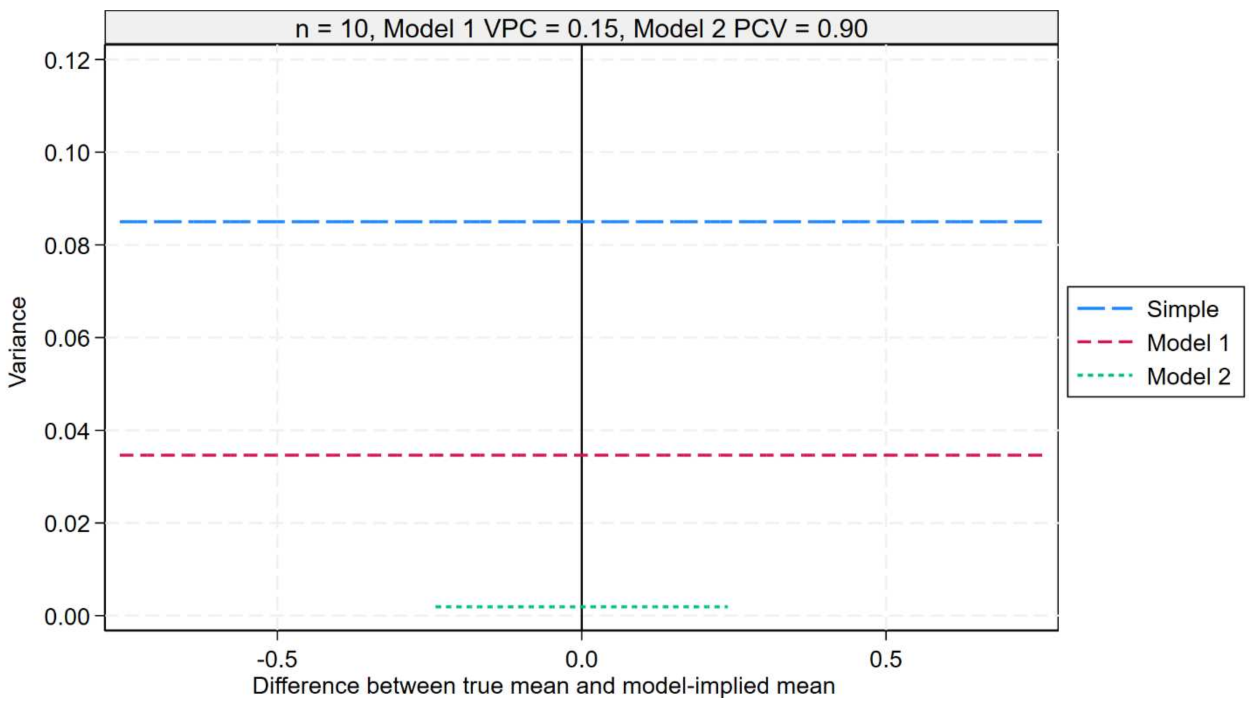

Figure 7.

Mean squared error of the simple mean (29), MAIHDA Model 1 mean (30), and MAIHDA Model 2 mean (31) across repeated samples of individuals for a given intersection, plotted against the difference between the true mean and the model-implied mean for that intersection. The plot assumes an intersection size of 10 individuals, a Model 1 VPC of 0.15, and a Model 2 PCV of 0.90. The Model 1 model-implied mean corresponds to the overall mean, while the Model 2 model-implied mean reflects the additive effects of the social identities defining the intersection. The bias of the MAIHDA Model 2 mean is shown over a narrower range due to smaller differences between the true and model-implied means in this model.

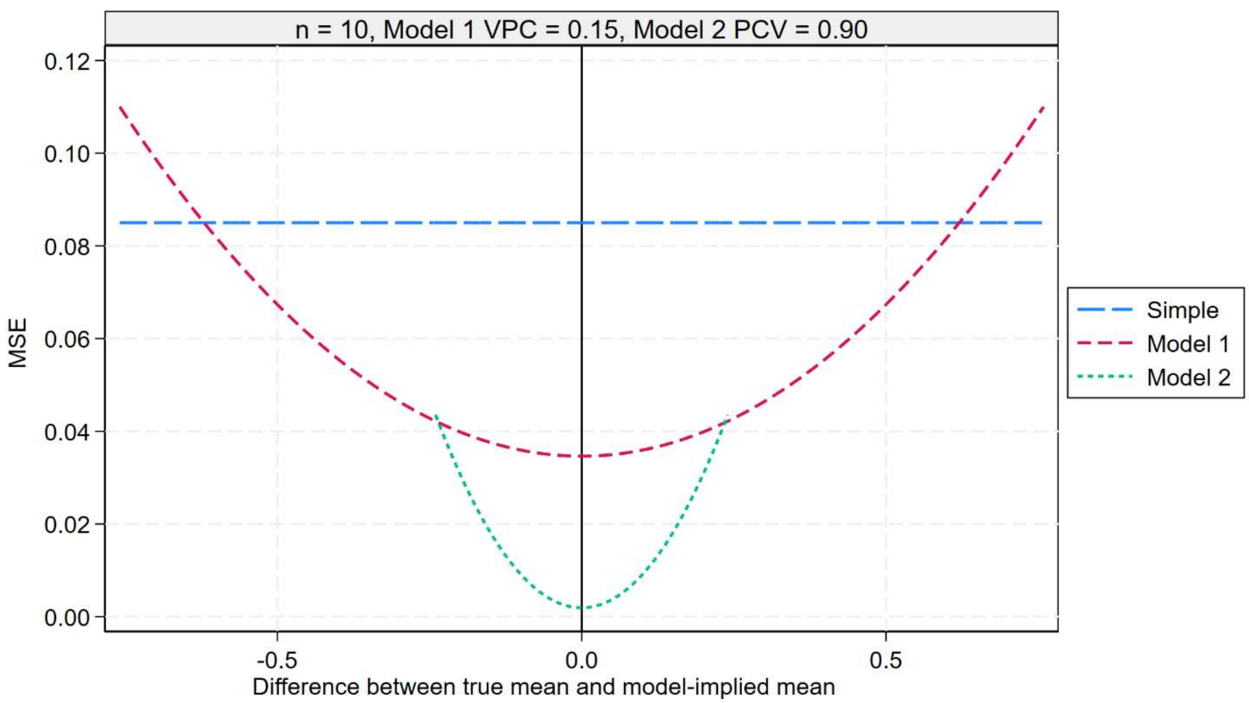

Supplemental Material

These Supplemental Materials are organized into three sections. The first section presents full derivations of all analytic expressions included in the main text. The second section explores how our main results vary as a function of the VPC, PCV, and the coefficient of variation of intersection size. The third section examines how our results are affected when the true data-generating process is more complex than MAIHDA Model 2, specifically by including a fixed two-way interaction between two of the social identities that define the intersections.

S1. Derivations of all analytic expressions

To derive all expressions, we specify the true model as Model 2 and assume that the regression coefficients, random intersection variance, and individual residual variance are known, with only the true mean for each intersection being unknown. See the article for the definition of all terms. All expressions given in this section are also shown for ease of comparison in Table S1.

PREDICTED MEANS

The predicted means are given as follows:

Simple Means

$$\bar{y}_{.j} = \frac{1}{n_j}\sum_{i=1}^{n_j} y_{ij}$$

MAIHDA Model 1 Means

$$\tilde{\mu}_{1j} = R_{1j}\bar{y}_{.j} + \left(1 - R_{1j}\right)\beta_{1,0}$$

See Equation 3.41 in Raudenbush and Bryk (2002, p.46).

MAIHDA Model 2 Means

$$\tilde{\mu}_{2j} = R_{2j}\bar{y}_{.j} + \left(1 - R_{2j}\right)\left(\beta_{2,0} + \beta_{2,1}x_{1j} + \cdots + \beta_{2,p}x_{pj}\right)$$

See Equation 3.48 in Raudenbush and Bryk (2002, p.48).

See Equation 2.56 in McCulloch, Searle, and Neuhaus (2008).

VARIANCE OF THE PREDICTED MEANS ACROSS THE DISTRIBUTION OF INTERSECTIONS

The variance of each set of predicted means can be derived as follows:

Simple.

$$\mathrm{Var}\left(\bar{y}_{.j}\right) = \mathrm{Var}\left(\mu_j + \bar{e}_{.j}\right)$$

$$= \sigma_\mu^2 + \frac{\sigma_e^2}{n}$$

See Equation 3.6 in Raudenbush and Bryk (2002, p.40).

See Equation 4.13 in Snijders and Bosker (2012, p.61).

Model 1.

$$\mathrm{Var}\left(\tilde{\mu}_{1j}\right) = \mathrm{Var}\left\{R_1\bar{y}_{.j} + (1 - R_1)\beta_0\right\}$$

$$= \mathrm{Var}\left(R_1\bar{y}_{.j}\right)$$

$$= R_1^2\mathrm{Var}\left(\bar{y}_{.j}\right)$$

$$= R_1^2\left(\sigma_\mu^2 + \frac{\sigma_e^2}{n}\right)$$

$$= R_1\left(\frac{\sigma_\mu^2}{\sigma_\mu^2 + \frac{\sigma_e^2}{n}}\right)\left(\sigma_\mu^2 + \frac{\sigma_e^2}{n}\right)$$

$$= R_1\sigma_\mu^2$$

See last equation in Rabe-Hesketh and Skondal (2021, p123).

See first equation in Skrondal and Rabe-Hesketh (2004, p.233).

Model 2.

$$
\begin{aligned}
\mathrm{Var}\left(\tilde{\mu}_{2j}\right) &= \mathrm{Var}\left\{R_2\bar{y}_{.j} + (1-R_2)\left(\beta_0 + \beta_1 x_{1j} + \cdots + \beta_p x_{pj}\right)\right\} \\
&= \mathrm{Var}\left\{R_2\bar{y}_{.j} + (1-R_2)\left(\bar{y}_{.j} - u_{2j} - \bar{e}_{.j}\right)\right\} \\
&= \mathrm{Var}\left\{R_2\bar{y}_{.j} + \bar{y}_{.j} - u_{2j} - \bar{e}_{.j} - R_2\bar{y}_{.j} + R_2 u_{2j} + R_2\bar{e}_{.j}\right\} \\
&= \mathrm{Var}\left\{\bar{y}_{.j} - u_{2j} - \bar{e}_{.j} + R_2 u_{2j} + R_2\bar{e}_{.j}\right\} \\
&= \mathrm{Var}\left\{\mu_j - u_{2j} + R_2 u_{2j} + R_2\bar{e}_{.j}\right\} \\
&= \mathrm{Var}\left\{\mu_j + (R_2-1)u_{2j} + R_2\bar{e}_{.j}\right\} \\
&= \mathrm{Var}\left\{\mu_j + (R_2-1)u_{2j}\right\} + R_2^2\mathrm{Var}\left\{\bar{e}_{.j}\right\} \\
&= \mathrm{Var}\left(\mu_j\right) + 2(R_2-1)\mathrm{Cov}\left(\mu_j, u_{2j}\right) + (R_2-1)^2\mathrm{Var}\left(u_{2j}\right) + R_2^2\mathrm{Var}\left\{\bar{e}_{.j}\right\} \\
&= \sigma_\mu^2 + 2(R_2-1)\sigma_u^2 + (R_2-1)^2\sigma_u^2 + R_2^2\frac{\sigma_e^2}{n} \\
&= \sigma_\mu^2 + \{2(R_2-1) + (R_2-1)^2\}\sigma_u^2 + R_2^2\frac{\sigma_e^2}{n} \\
&= \sigma_\mu^2 + \{2R_2 - 2 + (R_2-1)^2\}\sigma_u^2 + R_2^2\frac{\sigma_e^2}{n} \\
&= \sigma_\mu^2 + (2R_2 - 2 + R_2^2 - 2R_2 + 1)\sigma_u^2 + R_2^2\frac{\sigma_e^2}{n} \\
&= \sigma_\mu^2 + (R_2^2 - 1)\sigma_u^2 + R_2^2\frac{\sigma_e^2}{n} \\
&= \sigma_\mu^2 - \sigma_u^2 + R_2^2\sigma_u^2 + R_2^2\frac{\sigma_e^2}{n} \\
&= \sigma_\mu^2 - \sigma_u^2 + R_2^2\left(\sigma_u^2 + \frac{\sigma_e^2}{n}\right) \\
&= \sigma_\mu^2 - \sigma_u^2 + R_2\left(\frac{\sigma_u^2}{\sigma_u^2 + \frac{\sigma_e^2}{n}}\right)\left(\sigma_u^2 + \frac{\sigma_e^2}{n}\right)
\end{aligned}
$$

$$= \sigma_\mu^2 - \sigma_u^2 + R_2\sigma_u^2$$

$$= \sigma_\mu^2 + R_2\sigma_u^2 - \sigma_u^2$$

$$= \sigma_\mu^2 + (R_2 - 1)\sigma_{u2}^2$$

$$= \sigma_\mu^2 - (1 - R_2)\sigma_{u2}^2$$

CORRELATION BETWEEN THE PREDICTED AND TRUE MEANS ACROSS THE DISTRIBUTION OF INTERSECTIONS

The correlation of each set of predicted means with the true means can be derived as follows:

Simple

$$\mathrm{Corr}\left(\bar{y}_{.j}, \mu_j\right) = \frac{\mathrm{Cov}\left(\bar{y}_{.j}, \mu_j\right)}{\sqrt{\mathrm{Var}\left(\bar{y}_{.j}\right)}\sqrt{\mathrm{Var}\left(\mu_j\right)}}$$

$$= \frac{\mathrm{Cov}\left(\bar{y}_{.j}, \mu_j\right)}{\sqrt{\sigma_\mu^2 + \frac{\sigma_e^2}{n}}\sqrt{\sigma_\mu^2}}$$

$$= \frac{\mathrm{Cov}\left(\mu_j + \bar{e}_{.j}, \mu_j\right)}{\sqrt{\sigma_\mu^2 + \frac{\sigma_e^2}{n}}\sqrt{\sigma_\mu^2}}$$

$$= \frac{\mathrm{Cov}\left(\mu_j, \mu_j\right)}{\sqrt{\sigma_\mu^2 + \frac{\sigma_e^2}{n}}\sqrt{\sigma_\mu^2}}$$

$$= \frac{\sigma_\mu^2}{\sqrt{\sigma_\mu^2 + \frac{\sigma_e^2}{n}}\sqrt{\sigma_\mu^2}}$$

$$= \frac{\sqrt{\sigma_\mu^2}}{\sqrt{\sigma_\mu^2 + \frac{\sigma_e^2}{n}}}$$

$$= \sqrt{\frac{\sigma_\mu^2}{\sigma_\mu^2 + \frac{\sigma_e^2}{n}}}$$

$$= \sqrt{R_1}$$

Model 1

$$\begin{aligned}
\mathrm{Corr}(\tilde{\mu}_{1j}, \mu_j) &= \frac{\mathrm{Cov}(\tilde{\mu}_{1j}, \mu_j)}{\sqrt{\mathrm{Var}(\tilde{\mu}_{1j})}\sqrt{\mathrm{Var}(\mu_j)}} \\
&= \frac{\mathrm{Cov}(\tilde{\mu}_{1j}, \mu_j)}{\sqrt{R_1\sigma_\mu^2}\sqrt{\sigma_\mu^2}} \\
&= \frac{\mathrm{Cov}\{R_1\bar{y}_{.j} + (1 - R_1)\beta_0, \mu_j\}}{\sqrt{R_1\sigma_\mu^2}\sqrt{\sigma_\mu^2}} \\
&= \frac{\mathrm{Cov}\{R_1\bar{y}_{.j}, \mu_j\}}{\sqrt{R_1\sigma_\mu^2}\sqrt{\sigma_\mu^2}} \\
&= \frac{R_1\mathrm{Cov}(\bar{y}_{.j}, \mu_j)}{\sqrt{R_1\sigma_\mu^2}\sqrt{\sigma_\mu^2}} \\
&= \frac{R_1\mathrm{Cov}(\mu_j + \bar{e}_{.j}, \mu_j)}{\sqrt{R_1\sigma_\mu^2}\sqrt{\sigma_\mu^2}} \\
&= \frac{R_1\mathrm{Cov}\{\mu_j, \mu_j\}}{\sqrt{R_1\sigma_\mu^2}\sqrt{\sigma_\mu^2}} \\
&= \frac{R_1\sigma_\mu^2}{\sqrt{R_1\sigma_\mu^2}\sqrt{\sigma_\mu^2}} \\
&= \frac{R_1\sigma_\mu^2}{\sqrt{R_1}\sqrt{\sigma_\mu^2}\sqrt{\sigma_\mu^2}} \\
&= \frac{R_1\sigma_\mu^2}{\sqrt{R_1}\sigma_\mu^2} \\
&= \frac{R_1}{\sqrt{R_1}} \\
&= \sqrt{R_1}
\end{aligned}$$

Model 2

$$\mathrm{Corr}(\tilde{\mu}_{2j},\mu_j)=\frac{\mathrm{Cov}(\tilde{\mu}_{2j},\mu_j)}{\sqrt{\mathrm{Var}(\tilde{\mu}_{2j})}\sqrt{\mathrm{Var}(\mu_j)}}$$

$$=\frac{\mathrm{Cov}(\tilde{\mu}_{2j},\mu_j)}{\sqrt{\sigma_\mu^2+(R_2-1)\sigma_{u2}^2}\sqrt{\sigma_\mu^2}}$$

$$=\frac{\mathrm{Cov}\{R_2\bar{y}_{.j}+(1-R_2)(\beta_0+\beta_1x_{1j}+\cdots+\beta_px_{pj}),\mu_j\}}{\sqrt{\sigma_\mu^2+(R_2-1)\sigma_{u2}^2}\sqrt{\sigma_\mu^2}}$$

$$=\frac{\mathrm{Cov}\{R_2(\mu_j-\bar{e}_{.j})+(1-R_2)(\mu_j-u_{2j}),\mu_j\}}{\sqrt{\sigma_\mu^2+(R_2-1)\sigma_{u2}^2}\sqrt{\sigma_\mu^2}}$$

$$=\frac{\mathrm{Cov}\{R_2(\mu_j-\bar{e}_{.j}),\mu_j\}+\mathrm{Cov}\{(1-R_2)(\mu_j-u_{2j}),\mu_j\}}{\sqrt{\sigma_\mu^2+(R_2-1)\sigma_{u2}^2}\sqrt{\sigma_\mu^2}}$$

$$=\frac{\mathrm{Cov}\{R_2\mu_j-R_2\bar{e}_{.j},\mu_j\}+\mathrm{Cov}\{\mu_j-u_{2j}-R_2\mu_j+R_2u_{2j},\mu_j\}}{\sqrt{\sigma_\mu^2+(R_2-1)\sigma_{u2}^2}\sqrt{\sigma_\mu^2}}$$

$$=\frac{\mathrm{Cov}\{R_2\mu_j,\mu_j\}+\mathrm{Cov}\{\mu_j-u_{2j}-R_2\mu_j+R_2u_{2j},\mu_j\}}{\sqrt{\sigma_\mu^2+(R_2-1)\sigma_{u2}^2}\sqrt{\sigma_\mu^2}}$$

$$=\frac{\mathrm{Cov}\{R_2\mu_j,\mu_j\}+\mathrm{Cov}\{\mu_j,\mu_j\}+\mathrm{Cov}\{-u_{2j},\mu_j\}+\mathrm{Cov}\{-R_2\mu_j,\mu_j\}+\mathrm{Cov}\{R_2u_{2j},\mu_j\}}{\sqrt{\sigma_\mu^2+(R_2-1)\sigma_{u2}^2}\sqrt{\sigma_\mu^2}}$$

$$=\frac{R_2\mathrm{Cov}\{\mu_j,\mu_j\}+\mathrm{Cov}\{\mu_j,\mu_j\}-\mathrm{Cov}\{u_{2j},\mu_j\}-R_2\mathrm{Cov}\{\mu_j,\mu_j\}+R_2\mathrm{Cov}\{u_{2j},\mu_j\}}{\sqrt{\sigma_\mu^2+(R_2-1)\sigma_{u2}^2}\sqrt{\sigma_\mu^2}}$$

$$=\frac{\mathrm{Cov}\{\mu_j,\mu_j\}-\mathrm{Cov}\{u_{2j},\mu_j\}+R_2\mathrm{Cov}\{u_{2j},\mu_j\}}{\sqrt{\sigma_\mu^2+(R_2-1)\sigma_{u2}^2}\sqrt{\sigma_\mu^2}}$$

$$= \frac{\sigma_\mu^2 - \sigma_{u2}^2 + R_2\sigma_{u2}^2}{\sqrt{\sigma_\mu^2 + (R_2 - 1)\sigma_{u2}^2}\sqrt{\sigma_\mu^2}}$$

$$= \frac{\sigma_\mu^2 + (R_2 - 1)\sigma_{u2}^2}{\sqrt{\sigma_\mu^2 + (R_2 - 1)\sigma_{u2}^2}\sqrt{\sigma_\mu^2}}$$

$$= \frac{\sqrt{\sigma_\mu^2 + (R_2 - 1)\sigma_{u2}^2}}{\sqrt{\sigma_\mu^2}}$$

$$= \sqrt{\frac{\sigma_\mu^2 + (R_2 - 1)\sigma_{u2}^2}{\sigma_\mu^2}}$$

$$= \sqrt{\frac{\sigma_{u1}^2 + (R_2 - 1)\sigma_{u2}^2}{\sigma_{u1}^2}}$$

$$= \sqrt{\frac{\sigma_{u1}^2 + R_2\sigma_{u2}^2 - \sigma_{u2}^2}{\sigma_{u1}^2}}$$

$$= \sqrt{\frac{\sigma_{u1}^2 - \sigma_{u2}^2}{\sigma_{u1}^2} + \frac{R_2\sigma_{u2}^2}{\sigma_{u1}^2}}$$

$$= \sqrt{\frac{\sigma_{u1}^2 - \sigma_{u2}^2}{\sigma_{u1}^2} + R_2\left(\frac{\sigma_{u2}^2 + \sigma_{u1}^2 - \sigma_{u1}^2}{\sigma_{u1}^2}\right)}$$

$$= \sqrt{\frac{\sigma_{u1}^2 - \sigma_{u2}^2}{\sigma_{u1}^2} + R_2\left(\frac{\sigma_{u1}^2}{\sigma_{u1}^2} + \frac{\sigma_{u2}^2 - \sigma_{u1}^2}{\sigma_{u1}^2}\right)}$$

$$= \sqrt{\frac{\sigma_{u1}^2 - \sigma_{u2}^2}{\sigma_{u1}^2} + R_2\left(\frac{\sigma_{u1}^2}{\sigma_{u1}^2} - \frac{\sigma_{u1}^2 - \sigma_{u2}^2}{\sigma_{u1}^2}\right)}$$

$$= \sqrt{\mathrm{PCV} + R_2(1 - \mathrm{PCV})}$$

BIAS

For each predictor, the bias for intersection j can be derived as follows:

Simple

$$\begin{aligned}\text{bias}\left(\bar{y}_{.j}\middle|\mu_j\right) &= E\left(\bar{y}_{.j}\middle|\mu_j\right) - \mu_j \\ &= E\left(\mu_j + \bar{e}_{.j}\middle|\mu_j\right) - \mu_j \\ &= \mu_j - \mu_j \\ &= 0\end{aligned}$$

Model 1

$$\begin{aligned}\text{bias}\left(\tilde{\mu}_{1j}\middle|\mu_j\right) &= E\left(\tilde{\mu}_{1j}\middle|\mu_j\right) - \mu_j \\ &= \mathrm{E}\left\{R_1\bar{y}_{.j} + (1 - R_1)\beta_0\middle|\mu_j\right\} - \mu_j \\ &= \mathrm{E}\left\{R_1\left(\mu_j + \bar{e}_{.j}\right) + (1 - R_1)\beta_0\middle|\mu_j\right\} - \mu_j \\ &= R_1\mu_j + (1 - R_1)\beta_0 - \mu_j \\ &= (R_1 - 1)(\mu_j - \beta_0) \\ &= -(1 - R_1)(\mu_j - \beta_0)\end{aligned}$$

Model 2

$$\begin{aligned}\text{bias}\left(\tilde{\mu}_{2j}\middle|\mu_j\right) &= E\left(\tilde{\mu}_{2j}\middle|\mu_j\right) - \mu_j \\ &= \mathrm{E}\left\{R_{2j}\bar{y}_{.j} + \left(1 - R_{2j}\right)\left(\beta_0 + \beta_1x_{1j} + \cdots + \beta_px_{pj}\right)\middle|\mu_j\right\} - \mu_j \\ &= \mathrm{E}\left\{R_{2j}\left(\mu_j + \bar{e}_{.j}\right) + \left(1 - R_{2j}\right)\left(\beta_0 + \beta_1x_{1j} + \cdots + \beta_px_{pj}\right)\middle|\mu_j\right\} - \mu_j \\ &= R_{2j}\mu_j + \left(1 - R_{2j}\right)\left(\beta_0 + \beta_1x_{1j} + \cdots + \beta_px_{pj}\right) - \mu_j \\ &= \left(R_{2j} - 1\right)\left\{\mu_j - \left(\beta_0 + \beta_1x_{1j} + \cdots + \beta_px_{pj}\right)\right\} \\ &= -\left(1 - R_{2j}\right)\left\{\mu_j - \left(\beta_0 + \beta_1x_{1j} + \cdots + \beta_px_{pj}\right)\right\}\end{aligned}$$

VARIANCE

The variance we derive here is the sampling variance of the predicted mean for a particular intersection, with true mean μ_j, across repeated samples of individuals from that intersection. For each predictor, the variance for intersection j can be derived as follows:

Simple

$$\mathrm{Var}(\bar{y}_{.j}|\mu_j) = \mathrm{Var}(\mu_j + \bar{e}_{.j}|\mu_j)$$

$$= \frac{\sigma_e^2}{n_j}$$

Model 1

$$\mathrm{Var}(\tilde{\mu}_{1j}|\mu_j) = \mathrm{Var}\{R_{1j}\bar{y}_{.j} + (1 - R_{1j})\beta_0|\mu_j\}$$

$$= \mathrm{Var}(R_{1j}\bar{y}_{.j}|\mu_j)$$

$$= R_{1j}^2\mathrm{Var}(\bar{y}_{.j}|\mu_j)$$

$$= R_{1j}^2\frac{\sigma_e^2}{n_j}$$

See last equation in Skrondal and Rabe-Hesketh (2004, p.233). Some rearranging is required. Note also that this equation is for $\mathrm{Var}(\tilde{u}_{1j}|u_{1j})$, but when all parameters are assumed known (as they are here), $\mathrm{Var}(\tilde{\mu}_{1j}|\mu_j) = \mathrm{Var}(\tilde{u}_{1j}|u_{1j})$.

Model 2

$$\mathrm{Var}(\tilde{\mu}_{2j}|\mu_j) = \mathrm{Var}\{R_{2j}\bar{y}_{.j} + (1 - R_{2j})(\beta_0 + \beta_1 x_{1j} + \cdots + \beta_p x_{pj})|\mu_j\}$$

$$= \text{Var}\left(R_{2j}\bar{y}_{.j} \middle| \mu_j\right)$$

$$= R_{2j}^2 \text{Var}\left(\bar{y}_{.j} \middle| \mu_j\right)$$

$$= R_{2j}^2 \frac{\sigma_e^2}{n_j}$$

MEAN SQUARED ERROR

For each predictor, the mean squared error (MSE) for intersection j can be derived as follows:

Simple

$$\mathrm{MSE}(\bar{y}_{.j}|\mu_j) = \mathrm{Var}(\bar{y}_{.j}|\mu_j) + \{\mathrm{bias}(\bar{y}_{.j}|\mu_j)\}^2$$

$$= \frac{\sigma_e^2}{n_j} + 0^2$$

$$= \frac{\sigma_e^2}{n_j}$$

Model 1

$$\mathrm{MSE}(\tilde{\mu}_{1j}|\mu_j) = \mathrm{Var}(\tilde{\mu}_{1j}|\mu_j) + \{\mathrm{bias}(\tilde{\mu}_{1j}|\mu_j)\}^2$$

$$= R_{1j}^2 \frac{\sigma_e^2}{n_j} + \{-(1 - R_{1j})(\mu_j - \beta_0)\}^2$$

$$= R_{1j}^2 \frac{\sigma_e^2}{n_j} + (1 - R_{1j})^2(\mu_j - \beta_0)^2$$

Model 2

$$\mathrm{MSE}(\tilde{\mu}_{2j}|\mu_j) = \mathrm{Var}(\tilde{\mu}_{2j}|\mu_j) + \{\mathrm{bias}(\tilde{\mu}_{2j}|\mu_j)\}^2$$

$$= R_{2j}^2 \frac{\sigma_e^2}{n_j} + [-(1 - R_{2j})\{\mu_j - (\beta_0 + \beta_1 x_{1j} + \cdots + \beta_p x_{pj})\}]^2$$

$$= R_{2j}^2 \frac{\sigma_e^2}{n_j} + (1 - R_{2j})^2\{\mu_j - (\beta_0 + \beta_1 x_{1j} + \cdots + \beta_p x_{pj})\}^2$$

S2. Further Results

STATISTICAL PROPERTIES OF THE DISTRIBUTION OF PREDICTED INTERSECTION MEANS

Variance of the Predicted Means Across the Distribution of Intersections

Figure 3 was plotted using specific values for the Model 1 VPC, Model 2 PCV, and CV of intersection size: 0.15, 0.90, and 0, respectively. Figures S1, S2, and S3 examine how this plot changes as each of these factors varies. We use simulation to address the case of varying intersection sizes. Here, we state the key findings. In all three cases, the relative ranking of the three methods of prediction remains unchanged, regardless of the VPC, PCV, or CV of intersection size. Specifically, the variance of the Model 2 means consistently lies closest to the true variance, followed by the variance of the Model 1 means, and finally the variance of the Simple means.

1. Effect of VPC: As the VPC increases (from 0.01 to 0.15 to 0.50), the variances of all three sets of predicted means become closer to one another and closer to the true variance (Figure S1). This is because, as the VPC rises, the simple means become a more reliable predictor of the true means.
2. Effect of PCV: As the PCV increases (from 0.50 to 0.90 to 0.99), the variances of the simple and Model 1 means remain constant, while the variance of the Model 2 means gets closer to the true variance (Figure S2). This makes sense, as the PCV reflects the explanatory power of the additive social identity effects, which are only explicitly incorporated into the calculation in the Model 2 means.
3. Effect of CV of intersection sizes: As the CV of intersection sizes increases (from 0 to 1 to 2), the variances of all three predicted means become less similar to each other (Figure

S3). Thus, choice of prediction method is more consequential when intersection sizes become more unequal.

In conclusion, the Model 2 means continue to be the preferred predictor, regardless of changes in the VPC, PCV, or CV of intersection size. However, the relative advantage of this predictor is most pronounced when intersectional inequalities are less pronounced and when intersections vary in size.

Correlation Between the Predicted and True means Across the Distribution of Intersections

Figure 4 was plotted using specific values for the Model 1 VPC, Model 2 PCV, and CV of intersection size: 0.15, 0.90, and 0, respectively. Figures S4, S5, and S6 examine how this plot changes as each of these factors varies. Here, we provide a summary of the key findings. In all three cases, the Model 2 means show the highest correlation with the true means, regardless of the VPC, PCV, or CV of intersection size.

1. As the VPC increases, the correlations between all three sets of predicted means and the true means approach 1 (Figure S4). This makes sense, as an increase in the VPC reflects a reduction in the sampling variability or error in the simple means, which in turn increases the variability in the Model 1 and 2 means, as reflected in higher reliability statistics. As the variances approach the true variance the correlations approach the true correlation.
2. As the PCV increases, the correlation between the simple and true means and the correlation between the Model 1 and true means remain constant, while the correlation between the Model 2 means and true means moves closer 1 (Figure S5). This makes sense, as the PCV reflects the explanatory power of the additive social identity effects, which are only explicitly incorporated into the calculation the Model 2 means.

3. As the CV of intersection sizes increases, the correlations between each set of means and the true means decrease. Importantly, the correlations between the simple means and Model 1 means and the true means are no longer equal; they begin to diverge, with the gap widening as the CV increases. Specifically, the correlation between the Model 1 means and the true means now exceeds the correlation between the simple means and the true means. Thus, the simple means perform as well as the Model 1 means only in the rather artificial situation of equal intersection sizes. When intersection sizes vary, the Model 1 means are preferred over the simple means. However, the Model 2 means continue to be preferred over both the simple means and the Model 1 means.

In conclusion, the Model 2 means continue to be the preferred predictor, regardless of changes in the VPC, PCV, or CV of intersection sizes. However, the relative advantage of this predictor is most pronounced when intersectional inequalities are less pronounced and when intersections vary in size.

STATISTICAL PROPERTIES OF PREDICTED INTERSECTION MEAN FOR A GIVEN INTERSECTION

Bias

Figure 5 was generated using specific values for intersection size, Model 1 VPC, and Model 2 PCV: 10, 0.15, and 0.90, respectively. Figures S7, S8, and S9 explore how this plot changes when these factors vary. Below, we summarize the key findings.

The simple mean remains unbiased in all cases. The bias of the Model 1 mean as a function of the difference between the true and model-implied mean follows a steeper linear trend than that of the Model 2 mean. However, the bias in the Model 2 mean operates over a narrower range of differences.

1. Effect of Intersection Size: For a given difference between the true and model-implied mean, increasing the intersection size (from 5 to 10 to 50 individuals) rapidly reduces the magnitude of the Model 1 bias, whereas the Model 2 bias declines more gradually (Figure S7).
2. Effect of VPC: As the VPC increases (from 0.01 to 0.15 to 0.50), the bias in the Model 1 mean decreases more rapidly than the bias in the Model 2 mean (Figure S8).
3. Effect of PCV: Increasing the PCV (from 0.50 to 0.90 to 0.99) does not affect the bias of the Model 1 mean (Figure S9). However, for the Model 2 mean, the bias relationship with the difference between the true and model-implied mean strengthens, though this effect is mitigated by a more restricted range of differences.

In conclusion, the simple mean remains the preferred predictor in terms of bias, as it is always unbiased, whereas the Model 1 and Model 2 means exhibit bias in all cases except when the true and model-implied means are equal.

Variance

Figure 6 was generated using specific values for the intersection size, Model 1 VPC, and Model 2 PCV: 10, 0.15, and 0.90, respectively. Figures S10, S11, and S12 explore how this plot changes when these factors vary. Below, we summarize the key findings.

1. Effect of Intersection Size: As the intersection size increases, the variance of the simple mean decreases rapidly (Figure S10). The variance of the Model 1 mean also decreases, but at a slower rate. In contrast, the variance of the Model 2 mean increases, as the squared reliability term R_{2j}^2 grows more rapidly than the decrease in the simple mean variance term σ_e^2/n_j.
2. Effect of VPC: As the VPC increases, the variance of the simple mean decreases, whereas the variance of both the Model 1 and 2 means increases (Figure S11).
3. Effect of PCV: Increasing the PCV does not affect the variance of simple mean or the Model 1 mean, but it decreases the variance of the Model 2 mean (Figure S12).

In conclusion, the Model 2 mean remains the preferred predictor in terms of having the smallest variance, regardless of changes in the intersection size, VPC, or PCV. However, its relative advantage is most pronounced when the intersection size is small, intersectional inequalities are less pronounced, and when these inequalities follow a largely additive structure.

Mean Squared Error

Figure 7 was plotted using specific values for the intersection size, Model 1 VPC, and Model 2 PCV: 10, 0.15, and 0.90, respectively. Figures S13, S14, and S15 explore how this plot changes when these factors vary. Below, we summarize the key findings.

The MSE of the simple mean remains constant and equal to its variance, as the simple mean is always unbiased.

1. Effect of Intersection Size: As the intersection size increases, the MSE of all three predictors decreases and converges toward 0 (Figure S13).
2. Effect of VPC: As the VPC increases, the MSE of the Simple mean decreases, whereas the MSE of the Model 1 mean, and to a lesser extent, the Model 2 mean, increases (Figure S14).
3. Effect of PCV: Increasing the PCV does not affect the MSE of the Simple mean or the Model 1 mean, but it reduces the MSE of the Model 2 mean (Figure S15).

In conclusion, the Model 2 mean remains the preferred predictor in terms of the lowest MSE, regardless of changes in the intersection size, VPC, or PCV. However, its advantage is most pronounced when intersection sizes are small, intersectional inequalities are less pronounced, and the inequalities follow an additive structure. These results align with the variance patterns observed earlier.

S3. A More Complex Data Generating Process with a Two-Way Fixed Interaction

INTRODUCTION

In the main text, we used analytic expressions to compare the statistical properties of Simple Means, MAIHDA Model 1 means, and MAIHDA Model 2 means against the true intersection means. As with simulation studies, this approach required specifying a “true” model or data-generating process (DGP) to serve as a benchmark; we assumed Model 2 to be the true DGP. Our results showed that MAIHDA Model 2 means were the most accurate, followed by Model 1 means, with Simple Means performing worst. We also demonstrated that the differences between these three approaches become particularly pronounced when intersection sizes are small.

A natural question is how our results would change if the true model were more complex than MAIHDA Model 2. In the main text, we considered an example in which the DGP includes a two-way interaction as a fixed effect, rather than capturing it implicitly through the intersection random effect. We argued that while fitting a so-called MAIHDA Model 3—which includes this two-way interaction as a fixed effect—would naturally yield more accurate predictions than MAIHDA Model 2, the improvement would likely be marginal compared to the substantial gains observed when moving first from Simple Means to MAIHDA Model 1 means, and then from MAIHDA Model 1 means to MAIHDA Model 2 means.

In this section, we illustrate the argument with a single simple example. There is clear value in conducting more extensive investigations that examine a wider range of potential deviations from the Model 2 DGP, and we encourage others to pursue this.

MAIHDA MODEL 3

Suppose the true model is more complex than MAIHDA Model 2 in that it additionally includes a fixed two-way interaction between two of the social identity variables that define the intersections. We refer to this expanded model as MAIHDA Model 3. All other social identity variables are included as additive main effects, as in previous models. For simplicity, we consider the case of two binary social identity variables, denoted x_{1j} and x_{2j}. The model can be written as:

$$y_{ij} = \mu_j + e_{ij}$$

$$\mu_j = \beta_0 + \beta_1 x_{1j} + \beta_2 x_{2j} + \beta_3 x_{3j} \cdots + \beta_p x_{pj} + \beta_{p+1} x_{1j} x_{2j} + u_j$$

$$u_j \sim N(0, \sigma_u^2)$$

$$e_{ij} \sim N(0, \sigma_e^2)$$

where $x_{1j}x_{2j}$ represents the additional fixed two-way interaction term and β_{p+1} denotes its coefficient. All other terms are as in MAIHDA Model 2.

THE SOCIAL IDENTITY VARIABLES

For simplicity, we consider the case where x_{1j} and x_{2j} are independent (e.g., gender and ethnicity) and each has a prevalence of 0.5. We also assume x_{1j} and x_{2j} are independent of all remaining social identity variables in the model. For ease of exposition, we represent all additional social identities using a single composite variable z_j^*.

$$z_j^* = \beta_3 x_{3j} \cdots + \beta_p x_{pj}$$

Thus, allows us to rewrite the model more simply as:

$$y_{ij} = \mu_j + e_{ij}$$

$$\mu_j = \beta_0 + \beta_1 x_{1j} + \beta_2 x_{2j} + 1 z_j^* + \beta_{p+1} x_{1j} x_{2j} + u_j$$

$$u_j \sim N(0, \sigma_u^2)$$

$$e_{ij} \sim N(0, \sigma_e^2)$$

Note, the regression coefficient on z_j^* is equal to 1. We choose to rescale z_j^* to have variance 1 and so now specify a regression coefficient β_z.

$$y_{ij} = \mu_j + e_{ij}$$

$$\mu_j = \beta_0 + \beta_1 x_{1j} + \beta_2 x_{2j} + \beta_z z_j + \beta_{p+1} x_{1j} x_{2j} + u_j$$

$$u_j \sim N(0, \sigma_u^2)$$

$$e_{ij} \sim N(0, \sigma_e^2)$$

SPECIFYING THE TRUE PARAMETER VALUES

Overview

We specify MAIHDA Model 3 as the true model and then derive the values implied for MAIHDA Models 1 and 2 from it. The true parameter values for MAIHDA Model 3 are summarized in the table below, along with the variance partition coefficients (VPCs), conditional reliabilities, and proportion change in variances (PCVs). As in the main text, we assume many intersections of equal size. We showed there that the choice of approach matters most when intersection size is small. We therefore focus on the case of 10 individuals per intersection. Corresponding implied values for Models 1 and 2 are also listed. The remainder of this section details how these true values were specified.

Variable	Parameter	Model 1	Model 2	Model 3
1	β_0	0.000	-0.3576	-0.2980
x_{1j}	β_1		0.3576	0.2384
x_{2j}	β_2		0.3576	0.2384

x_{3j}	β_3		0.2598	0.2598
$x_{1j}x_{2j}$	β_4			0.2384
u_j	σ_u^2	0.1500	0.0186	0.0150
e_{ij}	σ_e^2	0.8500	0.8500	0.8500
	VPC	0.1500	0.0214	0.0173
	PCV vs. M1		0.8763	0.9000
	$R = \frac{\sigma_u^2}{\sigma_u^2 + \frac{\sigma_e^2}{10}}$	0.6383	0.1792	.15

Note:

Outcome Mean and Variance, and Model 1 and 3 VPC and PCV

We specify the true parameter values for MAIHDA Model 3 such that the individual outcome variable has a mean of 0 and a variance of 1.

We start by setting the MAIHDA Model 3 intersection variance $\sigma_{u3}^2 = 0.015$ and individual variance $\sigma_{e3}^2 = 0.85$. None of the models include individual-level covariates (the intersection defining social identities operate at the intersection level). It therefore follows that $\sigma_{e1}^2 = \sigma_{e2}^2 = \sigma_{e3}^2 = \sigma_e^2 = 0.85$. Then we set the Model 3 proportional change in variance (PCV) to 0.90. This implies $\sigma_{u1}^2 = 0.15$ and therefore that the Model 1 VPC is 0.15 as is variance of the true means $\sigma_\mu^2 = \sigma_{u1}^2 = 0.15$. Given this, the variance explained by the fixed part of Model 3 is therefore $\sigma_\mu^2 - \sigma_{u3}^2 = 0.15 - 0.015 = 0.135$. We decompose this into two equal parts of 0.0675 each.

The first part corresponds to the variance associated with x_{1j}, x_{2j}, and their interaction $x_{1j}x_{2j}$. The second part corresponds to the variance associated with all remaining social identity variables as captured by the compositive variable x_{3j}. Because we have assumed that x_{1j} and x_{2j} are independent of the remaining social identity variables, there is no covariance between these two components. Thus, half the variance attributed to the fixed effects covariates is due to just two social identity variables and their interaction. That is, we have made the effects of the variables involving the two-way interaction large. Below we will also make the effect of the interaction itself large.

Next we calculate the true values for the regression coefficients of these fixed-effects regression coefficients consistent with this information.

The Variance of the First Component of the Fixed Part of the Model

The variance of the first component of the fixed part of MAIHDA model 3 is given by:

$$Var(\beta_0 + \beta_1 x_{1j} + \beta_2 x_{2j} + \beta_{p+1} x_{1j} x_{2j}) =$$

$$= \beta_1^2 Var(x_{1j}) + \beta_2^2 Var(x_{2j}) + \beta_{p+1}^2 Var(x_{1j} x_{2j}) + 2\beta_1 \beta_{p+1} Cov(x_{1j}, x_{1j} x_{2j})$$

$$+ 2\beta_2 \beta_{p+1} Cov(x_{2j}, x_{1j} x_{2j})$$

Recall that the prevalence of x_{1j} and x_{2j} are $q_1 = q_2 = 0.5$. It follows that:

$$Var(x_{1j}) = q_1 q_1 = 0.5 \times 0.5 = 0.25$$

$$Var(x_{2j}) = q_2 q_2 = 0.5 \times 0.5 = 0.25$$

$$Var(q_{1j} q_{2j}) = q_1 q_2 (1 - q_1 q_2) = 0.5 \times 0.5(1 - 0.5 \times 0.5) = 1.875$$

$$Cov(x_{1j}, x_{1j} x_{2j}) = q_1 q_2 (1 - q_1) = 0.5 \times 0.5(1 - 0.5)$$

$$Cov(x_{2j}, x_{1j} x_{2j}) = q_1 q_2 (1 - q_2) = 0.5 \times 0.5(1 - 0.5)$$

Substituting these results into the expression for the first component of the fixed part gives:

$$Var(\beta_0 + \beta_1 x_{1j} + \beta_2 x_{2j} + \beta_{p+1} x_{1j} x_{2j})$$

$$= \beta_1^2 0.25 + \beta_2^2 0.25 + Var(\beta_2 x_{3j} + \cdots + \beta_p x_{pj}) + \beta_{p+1}^2 0.1875$$

$$+ 2\beta_1 \beta_{p+1} 0.125 + 2\beta_2 \beta_{p+1} 0.125$$

For simplicity, we assume $\beta_1 = \beta_2 = \beta_{p+1}$. We refer to this common coefficient as β. We note that this suggests a large positive interaction effect. The table below presents the expected value of the outcome as a function of the four combinations of x_{1j} and x_{2j}. Thus when either $x_{1j} = 1$ or $x_{2j} = 1$ the expected outcome increases by β, but when both $x_{1j} = 1$ and $x_{2j} = 1$ the expected outcomes increases by 3β rather than 2β.

Combination	x_{1j}	x_{2j}	$x_{1j}x_{2j}$	$E(y_{ij} \mid x_{1j}, x_{2j})$
1	0	0	0	β_0
2	0	1	0	$\beta_0 + \beta$
3	1	0	0	$\beta_0 + \beta$
4	1	1	1	$\beta_0 + 3\beta$

Imposing this common coefficient in the expression for the first component of the MAIHDA Model 3 fixed part gives:

$$Var(\beta_0 + \beta x_{1j} + \beta x_{2j} + \beta x_{1j} x_{2j}) = \beta^2 0.25 + \beta^2 0.25 + \beta^2 0.1875 + 2\beta^2 0.125 + 2\beta^2 0.12$$

$$= \beta^2 0.25 + \beta^2 0.25 + \beta^2 0.1875 + \beta^2 0.250 + \beta^2 0.250$$

$$= \beta^2 (0.25 + 0.25 + 0.1875 + 0.250 + 0.250)$$

Recall, that we set the variance of the fixed-part of the model associated with x_{1j} and x_{2j} to equal 0.0675. Substituting in this value gives:

$$0.0675 = \beta^2(1.1875)$$

Rearranging gives the true value for the common coefficient:

$$\beta = \sqrt{\frac{0.0675}{1.1875}} = .23841582$$

The Variance of the Second Component of the Fixed Part of the Model

The variance of the second component of the fixed part of MAIHDA Model 3 is given by:

$$Var(\beta_z z_j) = \beta_z^2 \sigma_z^2$$

For simplicity we let z_j be a standard normal variate.

$$Var(\beta_z z_j) = \beta_z^2$$

Recall, that the variance of this second-component is also set equal to 0.0675 so that the two components have equal explanatory power. Substituting in this value gives:

$$0.0675 = \beta_z^2$$

Rearranging gives

$$\beta_z = \sqrt{0.0675} = .25980762$$

The Intercept

Recall that we wish the mean outcome to be 0. To derive the true value for the intercept we first note:

$$0 = \mathrm{E}(\beta_0 + \beta x_{1j} + \beta x_{2j} + \beta x_{1j} x_{2j} + \beta_z \mathrm{z_j})$$

$$= \beta_0 + 0.5\beta + 0.5\beta + 0.25\beta + 0\beta_z$$

$$= \beta_0 + 1.25\beta$$

Rearranging gives:

$$\beta_0 = -1.25\beta$$
$$= -1.25 \times 0.23841582$$
$$= -.29801977$$

For simplicity, we assume balanced data.

The Implied Parameter Values for Model 2

We know straight away that the MAIHDA Model 2 residual variance is $\sigma_e^2 = 0.85$. We also know that $\beta_z = .2598076$. The latter is identical to the Model 3 β_z as we have assumed x_{1j} and x_{2j} are independent of all other social identities and therefore $z_j 0$. We just need to find β and then we can calculate the variance associated with the fixed-part of the model which in turn will allow us to calculate the variance associated with the intersection random effect.

We derive the MAIHDA Model 2 regression coefficients implied by the MAIHDA Model 3 DGP by simulating a very large data set according to MAIHDA Model 3, then fitting MAIHDA Model 2 to these data. We then used these estimates as the implied values for Model 2. We obtained $\beta_0 = \beta = -.3576237$. We note that we also obtained $\beta_z = 0.2598076$ which is, as expected, identical to the MAIHDA Model 3 β_z.

Thus, the variance of the fixed-part of the model is given by:

$$Var(\beta_0 + \beta_1 x_{1j} + \beta_2 x_{2j} + \beta_z x_{zj}) = Var(\beta + \beta x_{1j} + \beta x_{2j} + \beta_z x_{zj})$$
$$= Var(\beta x_{1j}) + Var(\beta x_{2j}) + Var(\beta_z x_{zj})$$
$$= \beta^2 Var(x_{1j}) + \beta^2 Var(x_{2j}) + \beta_z^2 Var(\mathrm{z_j})$$
$$= 0.25\beta^2 + 0.25\beta^2 + 1\beta_z^2$$
$$= 0.5\beta^2 + \beta_z^2$$

Substituting in the implied parameter values gives

$$0.5\beta^2 + \beta_z^2 = 0.5(-0.3576237)^2 + (0.2598076)^2$$

$$= .13144734$$

It follows that the MAIHDA Model 2 intersection variance is:

$$\sigma_u^2 = 0.15 - .13144734 = .01855266$$

This then allows us to calculate the MAIHDA Model 2 VPC, reliability coefficient, and PCV versions Model 1.

STATISTICAL PROPERTIES OF THE DISTRIBUTION OF PREDICTED INTERSECTION MEANS

Recall from the main text that our goal in this section of the results is the big picture—how the approaches vary in capturing the distribution of intersection means.

Variance

The expression for the variance of MAIHDA Model 3 means takes the same form as that for Model 2. Specifically, it equals the variance of the true means minus one minus the conditional reliability for that model, multiplied by the intersection variance for that model. For ease of comparison, we reproduce below the expressions for the variance of the Simple Means, MAIHDA Model 1 means, and MAIHDA Model 2 means, as presented in the main text.

Simple: $$\text{Var}(\bar{y}_{.j}) = \sigma_\mu^2 + \frac{\sigma_e^2}{n}$$

Model 1: $$\text{Var}(\tilde{\mu}_{1j}) = R_1\sigma_\mu^2$$

Model 2: $$\text{Var}(\tilde{\mu}_{2j}) = \sigma_\mu^2 - (1 - R_2)\sigma_{u2}^2$$

Model 3: $$\text{Var}(\tilde{\mu}_{3j}) = \sigma_\mu^2 - (1 - R_3)\sigma_{u3}^2$$

Substituting the true and implied parameter values from the overview table yields the variance associated with each approach.

Approach	Value
True	0.1500
Simple means	0.2350
MAIHDA Model 1 means	0.0957
MAIHDA Model 2 means	0.1348
MAIHDA Model 3 means	0.1373

The variance of the true intersection means is 0.1500. The variance of the Simple Means is higher, at 0.2350, indicating an upward bias. In contrast, the variance of the MAIHDA Model 1 means is lower, at 0.0957, reflecting a downward bias. The variance of the MAIHDA Model 2 means is 0.1348; although larger than that of Model 1, it still slightly underestimates the true variance. The variance of the MAIHDA Model 3 means is slightly higher again, at 0.1370, representing only a very modest reduction in the downward bias.

Thus, as expected, in the presence of a fixed two-way interaction in the DGP, MAIHDA Model 3—which accounts for this interaction—produces the most accurate predicted means in terms of matching the variance of the true means. However, the improvement over Model 2 is marginal. Our central findings remain: MAIHDA Model 2 means are substantially more accurate than those from Model 1, which in turn outperform the Simple Means.

Correlation

The expression for the correlation of the MAIHDA Model 3 means with the true means takes the same form as that for Model 2. Specifically, it is equal to the square root of the sum of the PCV and the conditional reliability, multiplied by one minus the conditional PCV. For ease of comparison, we also repeat the corresponding expressions for the Simple Means, MAIHDA Model 1 means, and MAIHDA Model 2 means.

Simple: $\mathrm{Corr}(\bar{y}_{.j}, \mu_j) = \sqrt{R_1}$

Model 1: $\mathrm{Corr}(\tilde{\mu}_{1j}, \mu_j) = \sqrt{R_1}$

Model 2: $\mathrm{Corr}(\tilde{\mu}_{2j}, \mu_j) = \sqrt{\mathrm{PCV}_2 + R_2(1 - \mathrm{PCV}_2)}$

Model 3: $\mathrm{Corr}(\tilde{\mu}_{3j}, \mu_j) = \sqrt{\mathrm{PCV}_3 + R_3(1 - \mathrm{PCV}_3)}$

Substituting the true and implied parameter values from the overview table yields the correlation associated with each approach.

Approach	Values
True	1
Simple means	0.7989
MAIHDA Model 1 means	0.7989
MAIHDA Model 2 means	0.9479
MAIHDA Model 3 means	0.9566

The correlation between the Simple Means and the true means, as well as between the MAIHDA Model 1 means and the true means, is 0.7989. These values are identical in this case because the data are balanced. In unbalanced settings, however, the correlation between the MAIHDA Model 1 means and the true means will generally be higher than that between the Simple Means and the true means, as shown in the main text. The correlation for the MAIHDA Model 2 means is notably higher, at 0.9479. The correlation for the MAIHDA Model 3 means is slightly higher still, at 0.9566—though this represents only a marginal improvement over Model 2.

Thus, as expected, in the presence of a fixed two-way interaction in the DGP, MAIHDA Model 3—which accounts for this interaction—produces the most accurate predicted means in terms of correlation with the true means. However, the improvement over Model 2 is minimal. Our central finding remains: MAIHDA Model 2 means are much more strongly correlated with the true means than those from Model 1, which in turn outperform the Simple Means.

STATISTICAL PROPERTIES OF PREDICTED INTERSECTION MEAN FOR A GIVEN INTERSECTION

Recall from the main text that our goal here is to examine the statistical properties of the predicted mean *for a given intersection* across repeated random samples of individuals from that intersection.

Bias

The expression for the bias of the MAIHDA Model 3 mean for a given intersection takes the same form as that for Models 1 and 2. Specifically, it is the product of two terms: the first depends on the model's reliability, and the second is the deviation of the model-implied mean (i.e., the mean implied by the fixed effects) from the true mean. For ease of comparison, we also present the bias expressions for the Simple Means, MAIHDA Model 1 Means, and MAIHDA Model 2 Means as given in the main text.

Simple: $\text{bias}(\bar{y}_{.j}|\mu_j) = 0$

Model 1: $\text{bias}(\tilde{\mu}_{1j}|\mu_j) = -(1 - R_{1j})(\mu_j - \beta_{1,0})$

Model 2: $\text{bias}(\tilde{\mu}_{2j}|\mu_j) = -(1 - R_{2j})\{\mu_j - (\beta_{2,0} + \beta_{2,1}x_{1j} + \cdots + \beta_{2,p}x_{pj})\}$

Model 3: $\text{bias}(\tilde{\mu}_{2j}|\mu_j) = -(1 - R_{3j})\{\mu_j - (\beta_{2,0} + \beta_{2,1}x_{1j} + \cdots + \beta_{2,p}x_{pj} + \beta_{2,p+1}x_1x_2)\}$

The deviation of the model-implied mean from the true mean in Models 1, 2, and 3 is specified through a value of σ_u, chosen so that the model-implied mean is 1 standard deviation below the true mean in terms of the distribution of deviations across all intersections. This implies that approximately 32% of intersections exhibit deviations more extreme than this (so more than 1SD away from the true mean in either direction), while 68% have less extreme deviations ($2(1 - \Phi(1)) = 0.32$). Importantly, the actual magnitude of this deviation decreases as we move from Model 1 to Model 2 to Model 3, reflecting improved accuracy of the model-implied means with the inclusion of more fixed covariates.

Substituting in the true and implied parameter values from the overview table gives the bias associated with each approach.

Approach	Values
Simple means	0
MAIHDA Model 1 means	-0.1401
MAIHDA Model 2 means	-0.1118
MAIHDA Model 3 means	-0.1041

The bias associated with the model-implied mean being 1 standard deviation below the true mean decreases from a downward bias of 0.1401 points in Model 1, to 0.118 points in Model 2, and further to 0.1041 points in Model 3, reflecting the effect of shrinkage towards the model-implied mean.

Thus, as expected, when the DGP includes a fixed two-way interaction, MAIHDA Model 3—which explicitly accounts for this interaction—produces the least biased predicted means among the three MAIHDA models. However, the improvement over Model 2 is minimal. Our central finding remains: MAIHDA Model 2 means are substantially less biased than Model 1 means, while only the simple means are truly unbiased.

Variance

The expression for the variance of the MAIHDA Model 3 mean for a given intersection takes the same form as that for Models 1 and 2. Specifically, it equals the squared conditional reliability coefficient multiplied by the variance of the simple means. For ease of comparison, we also present the expressions for the variance of the Simple Means, MAIHDA Model 1 Means, and MAIHDA Model 2 Means as given in the main text.

Simple: $\mathrm{Var}\left(\bar{y}_{.j} \middle| \mu_j\right) = \frac{\sigma_e^2}{n_j}$

Model 1: $\mathrm{Var}\left(\tilde{\mu}_{1j} \middle| \mu_j\right) = R_{1j}^2 \frac{\sigma_e^2}{n_j}$

Model 2: $\mathrm{Var}\left(\tilde{\mu}_{2j} \middle| \mu_j\right) = R_{2j}^2 \frac{\sigma_e^2}{n_j}$

Model 3: $\mathrm{Var}\left(\tilde{\mu}_{3j} \middle| \mu_j\right) = R_{3j}^2 \frac{\sigma_e^2}{n_j}$

Substituting in the true and implied parameter values from the overview table gives the bias associated with each approach.

Approach	Values
Simple means	0.0850
MAIHDA Model 1 means	0.0346

MAIHDA Model 2 means	0.0027
MAIHDA Model 3 means	0.0019

The variance of the Simple Means across repeated samples is 0.0850. The variance of the MAIHDA Model 1 means is substantially lower, at 0.0346. This variance decreases further for the MAIHDA Model 2 means, reaching 0.0027. The variance for the MAIHDA Model 3 means is slightly lower still, at 0.0019.

As expected, when the DGP includes a fixed two-way interaction, MAIHDA Model 3—which explicitly accounts for this interaction—yields the most consistent predictions of the true means by minimizing the variance of predicted means across repeated samples. However, the improvement over Model 2 is marginal. Our central findings hold: MAIHDA Model 2 means are substantially better than those from Model 1, which in turn outperform the simple means.

Mean Squared Error

The expression for the mean squared error (MSE) of the MAIHDA Model 3 mean for a given intersection takes the same form as for Models 1 and 2. Specifically, it equals the sum of the variance and the squared bias. For ease of comparison, we also present the expressions for the MSE of the Simple Means, MAIHDA Model 1 Means, and MAIHDA Model 2 Means as given in the main text.

Simple: $\mathrm{MSE}\left(\bar{y}_{.j}\middle|\mu_j\right) = \frac{\sigma_e^2}{n_j}$

Model 1: $\mathrm{MSE}\left(\tilde{\mu}_{1j}\middle|\mu_j\right) = R_{1j}^2\frac{\sigma_e^2}{n_j} + \left(1 - R_{1j}\right)^2\left(\mu_j - \beta_{1,0}\right)^2$

Model 2: $\mathrm{MSE}\left(\tilde{\mu}_{2j}\middle|\mu_j\right) = R_{2j}^2\frac{\sigma_e^2}{n_j} + \left(1 - R_{2j}\right)^2\left\{\mu_j - \left(\beta_{2,0} + \beta_{2,1}x_{1j} + \cdots + \beta_{2,p}x_{pj}\right)\right\}^2$ (1)

Model 3: $\text{MSE}(\tilde{\mu}_{3j}|\mu_j) = R_{3j}^2 \frac{\sigma_e^2}{n_j} + (1 - R_{3j})^2 \{\mu_j - (\beta_{2,0} + \beta_{2,1}x_{1j} + \cdots + \beta_{2,p}x_{pj} + \beta_{2,p+1}x_1x_2)\}^2$

We continue to consider the case where the model-implied mean is 1 standard deviation lower than the true mean in each model, based on the distribution of deviations for that model.

Approach	Values
Simple means	0.0850
MAIHDA Model 1 means	0.0543
MAIHDA Model 2 means	0.0152
MAIHDA Model 3 means	0.0128

The MSE associated with the model-implied mean being 1 standard deviation below the true mean is 0.0850 for the Simple Mean, substantially lower at 0.0543 for Model 1, further reduced to 0.0152 for Model 2, and then marginally decreased to 0.0128 for Model 3.

Thus, as expected, when the DGP includes a fixed two-way interaction, MAIHDA Model 3—which explicitly accounts for this interaction—produces the most accurate predicted means across repeated samples. However, the improvement over Model 2 is minimal. Our central finding remains: MAIHDA Model 2 means achieve substantially higher predictive accuracy than Model 1 means, which themselves outperform the simple means.

CONCLUSION

In this section, we used a specific example to show that even when the true data-generating process (DGP) is more complex than MAIHDA Model 2, our key finding still holds: MAIHDA

Model 2 means are more predictively accurate measures of the true means than Model 1 means, which themselves outperform the Simple Means. We focused on one particular deviation from MAIHDA Model 2—specifically, the addition of a single, intentionally large two-way interaction as a fixed effect rather than representing it implicitly through intersection random effects. Although this example supports our argument, it is only one among many possible departures from the MAIHDA Model 2 DGP that merit further investigation. The key point is that MAIHDA Model 2 is expected to better approximate most real-world DGPs than either Model 1 or the Simple Means. Our main conclusion thus remains robust under this assumption.

REFERENCES

McCulloch, Charles E., Shayle R. Searle, and John M. Neuhaus. 2008. Generalized, Linear, and Mixed Models. 2nd ed. Hoboken, NJ: John Wiley & Sons.

Rabe-Hesketh, Sophia, and Anders Skrondal. 2021. Multilevel and Longitudinal Modeling Using Stata, Fourth Edition: Volume I: Continuous Responses. College Station, TX: Stata Press.

Raudenbush, Stephen W., and Anthony S. Bryk. 2002. *Hierarchical Linear Models: Applications and Data Analysis Methods*. 2nd ed. Thousand Oaks, CA: Sage.

Skrondal, Anders, and Sophia Rabe-Hesketh. 2004. Generalized Latent Variable Modeling: Multilevel, Longitudinal, and Structural Equation Models. Boca Raton, FL: Chapman & Hall/CRC.

Snijders, Tom A. B., and Roel J. Bosker. 2012. *Multilevel Analysis: An Introduction to Basic and Advanced Multilevel Modeling*. 2nd ed. London: SAGE Publications.

TABLES

Table S1

Summary table for the three prediction methods: simple means, MAIHDA Model 1, and MAIHDA Model 2

Description	Simple Mean	Model 1	Model 2
MODELS AND STATISTICS			
Model equation for individual outcome		$y_{ij} = \mu_j + e_{ij}$	$y_{ij} = \mu_j + e_{ij}$
Model equation for the intersection mean (model-implied mean)		$\mu_j = \beta_{0,1} + u_{1j}$	$\mu_j = \beta_{0,2} + \beta_{1,2}x_{1j} + \cdots + \beta_{1,2}x_{1j} + u_{2j}$
Variance Partition		$\text{VPC}_1 = \frac{\sigma_{u1}^2}{\sigma_{u1}^2 + \sigma_{e1}^2}$	

coefficient (VPC)			
Proportion Change in Variance (PCV)			$\text{PCV} = \frac{\sigma_{u1}^2 - \sigma_{u2}^2}{\sigma_{u1}^2}.$
Reliability		$R_{1j} = \frac{\sigma_{u1}^2}{\sigma_{u1}^2 + \frac{\sigma_e^2}{n_j}}$	$R_{2j} = \frac{\sigma_{u2}^2}{\sigma_{u2}^2 + \frac{\sigma_e^2}{n_j}}$

PREDICTED MEANS

Predicted mean	$\bar{y}_{.j} = \frac{1}{n_j}\sum_{i=1}^{n_j} y_{ij}$	$\tilde{\mu}_{1j} = R_{1j}\bar{y}_{.j} + (1 - R_{1j})\beta_{1,0}$	$\tilde{\mu}_{2j} = R_{2j}\bar{y}_{.j} + (1 - R_{2j})(\beta_{2,0} + \beta_{2,1}x_{1j} + \cdots + \beta_{2,p}x_{pj})$
Degree to which simple mean is shrunk		$\tilde{\mu}_{1j} - \bar{y}_{.j} = (R_{1j} - 1)(\bar{y}_{.j} - \beta_{1,0})$	$\tilde{\mu}_{2j} - \bar{y}_{.j} = (R_{2j} - 1)\{\bar{y}_{.j} - (\beta_{2,0} + \beta_{2,1}x_{1j} + \cdots + \beta_{2,p}x_{pj})\}$

Degree to which model-implied mean is shrunk		$\tilde{\mu}_{1j} - \beta_{1,0} = R_{1j}\left(\bar{y}_{.j} - \beta_{1,0}\right)$	$\tilde{\mu}_{1j} - \beta_{2,0} + \beta_{2,1}x_{1j} + \cdots + \beta_{2,p}x_{pj}$ $= R_{2j}\{\bar{y}_{.j}$ $- \left(\beta_{2,0} + \beta_{2,1}x_{1j} + \cdots + \beta_{2,p}x_{pj}\right)\}$
STATISTICAL PROPERTIES OF THE DISTRIBUTION OF PREDICTED INTERSECTION MEANS			
Variance	$\sigma_\mu^2 + \dfrac{\sigma_e^2}{n}$	$R_1\sigma_\mu^2$	$\sigma_\mu^2 - (1 - R_2)\sigma_{u2}^2$
Correlation with true means	$\sqrt{R_1}$	$\sqrt{R_1}$	$\sqrt{\text{PCV} + R_2(1 - \text{PCV})}$
STATISTICAL PROPERTIES OF PREDICTED INTERSECTION MEAN FOR A GIVEN INTERSECTION			
Bias	0	$-\left(1 - R_{1j}\right)\left(\mu_j - \beta_{1,0}\right)$	$-\left(1 - R_{2j}\right)\{\mu_j - \left(\beta_{2,0} + \beta_{2,1}x_{1j} + \cdots + \beta_{2,p}x_{pj}\right)\}$
Variance	$\dfrac{\sigma_e^2}{n_j}$	$R_{1j}^2\dfrac{\sigma_e^2}{n_j}$	$R_{2j}^2\dfrac{\sigma_e^2}{n_j}$

Mean Squared Error (MSE)	$\frac{\sigma_e^2}{n_j}$	$R_{1j}^2 \frac{\sigma_e^2}{n_j} + \left(1 - R_{1j}\right)^2 \left(\mu_j - \beta_{1,0}\right)^2$	$R_{2j}^2 \frac{\sigma_e^2}{n_j} + \left(1 - R_{2j}\right)^2 \{\mu_j - \left(\beta_{2,0} + \beta_{2,1} x_{1j} + \cdots + \beta_{2,p} x_{pj}\right)\}^2$

Note.

FIGURES

Figure S1.

Variance of the simple means (17), MAIHDA Model 1 means (18), and MAIHDA Model 2 means (19) across the distribution of intersections, plotted against intersection size. The plot assumes a Model 2 PCV of 0.90, and equal intersection sizes. The plot is repeated for three Model 1 VPC values of 0.01 (small intersectional inequalities), 0.15 (medium) and 0.50 (large).

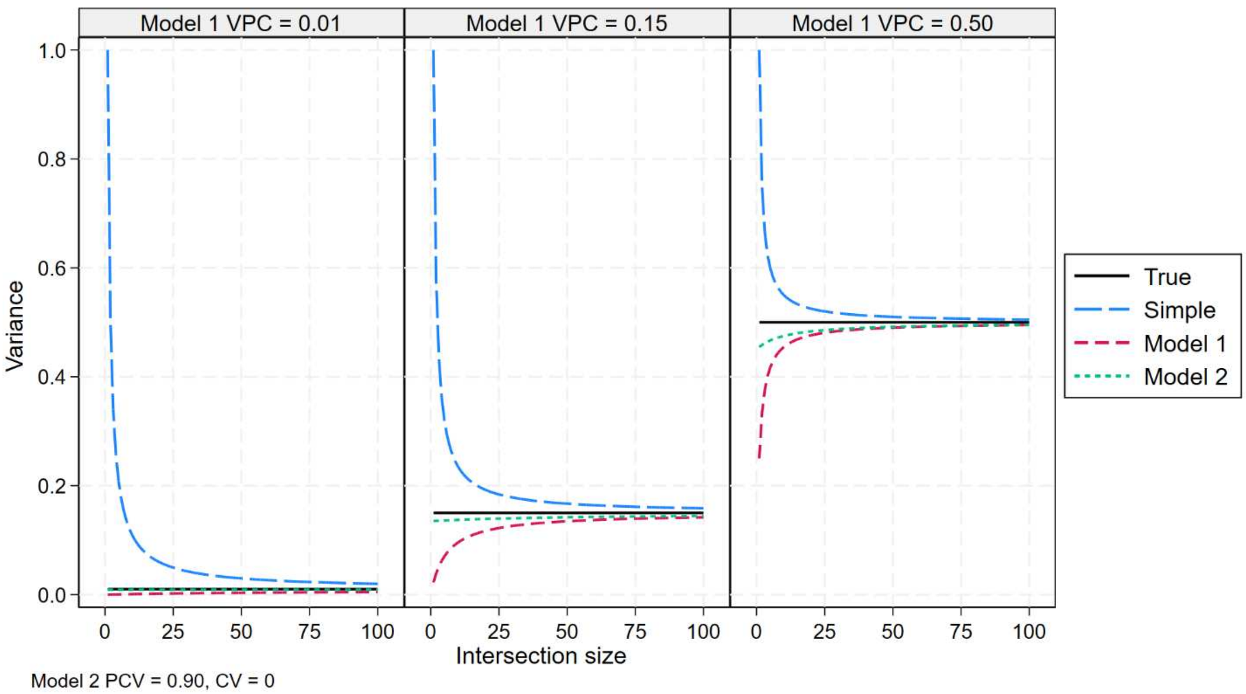

Figure S2.

Variance of the simple means (17), MAIHDA Model 1 means (18), and MAIHDA Model 2 means (19) across the distribution of intersections, plotted against intersection size. The plot assumes a Model 1 VPC of 0.15, and equal intersection sizes. The plot is repeated for four Model 2 PCV values of 0.00 (no additive patterning in the intersectional inequalities), 0.50 (low), 0.90 (medium), and 0.99 (high).

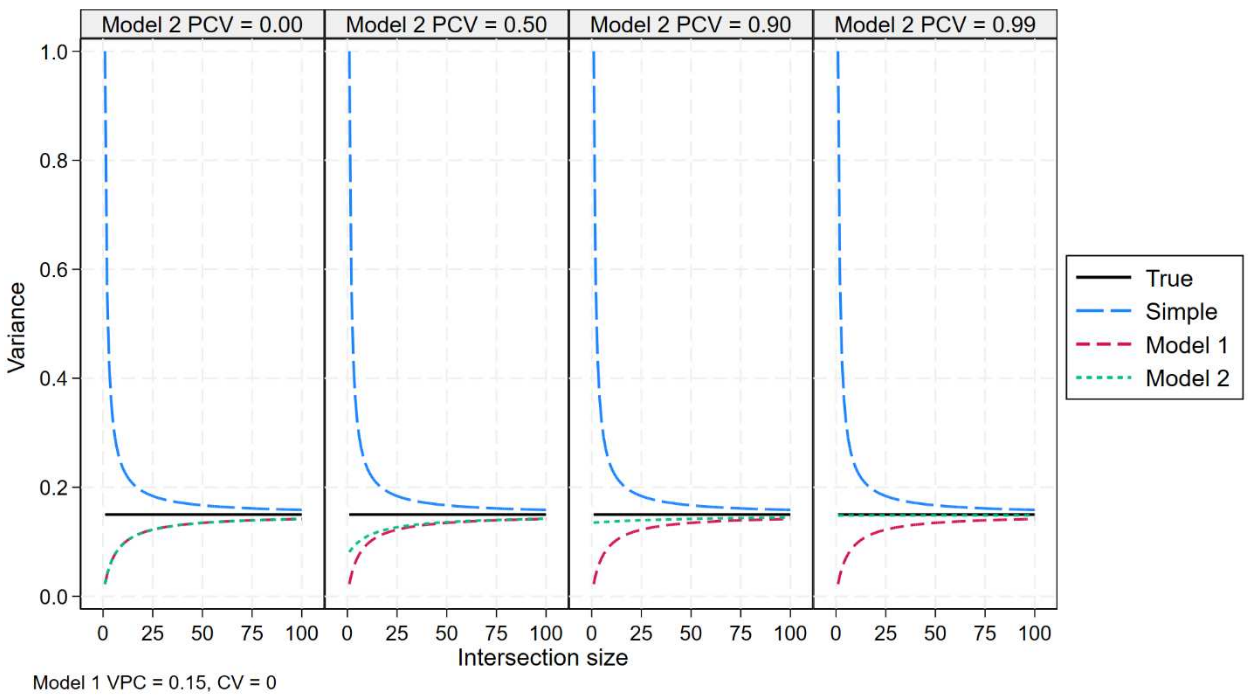

Figure S3.

Variance of the simple means (17), MAIHDA Model 1 means (18), and MAIHDA Model 2 means (19) across the distribution of intersections, plotted against intersection size. The plot assumes a Model 1 VPC of 0.15, and a Model 2 PCV of 0.90. The plot is repeated for three coefficients of variation of intersection sizes of 0 (equal intersection sizes), 1 (medium variation in intersection sizes), and 2 (high variation in intersection sizes).

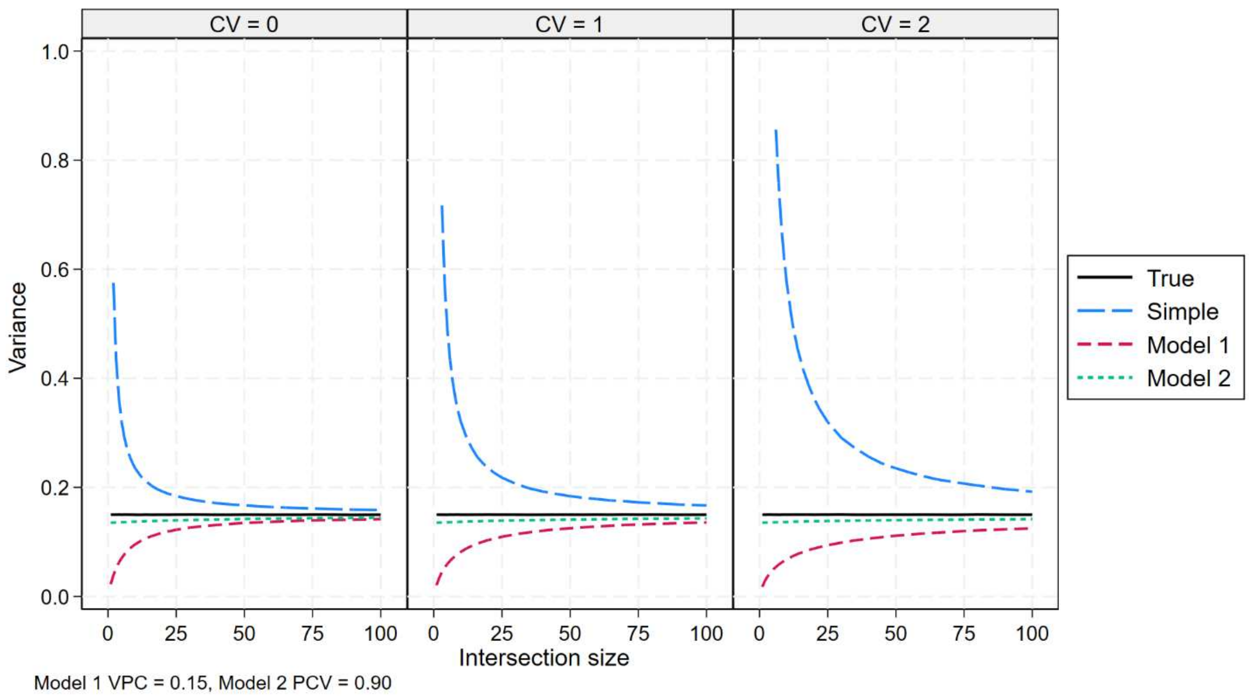

Figure S4.

Correlation between the simple means and the true means (20), MAIHDA Model 1 means and the true means (21), and MAIHDA Model 2 means and the true means (22), across the distribution of intersections, plotted against intersection size. The plot assumes a Model 2 PCV of 0.90, and equal intersection sizes. The plot is repeated for three Model 1 VPC values of 0.01 (small intersectional inequalities), 0.15 (medium) and 0.50 (large).

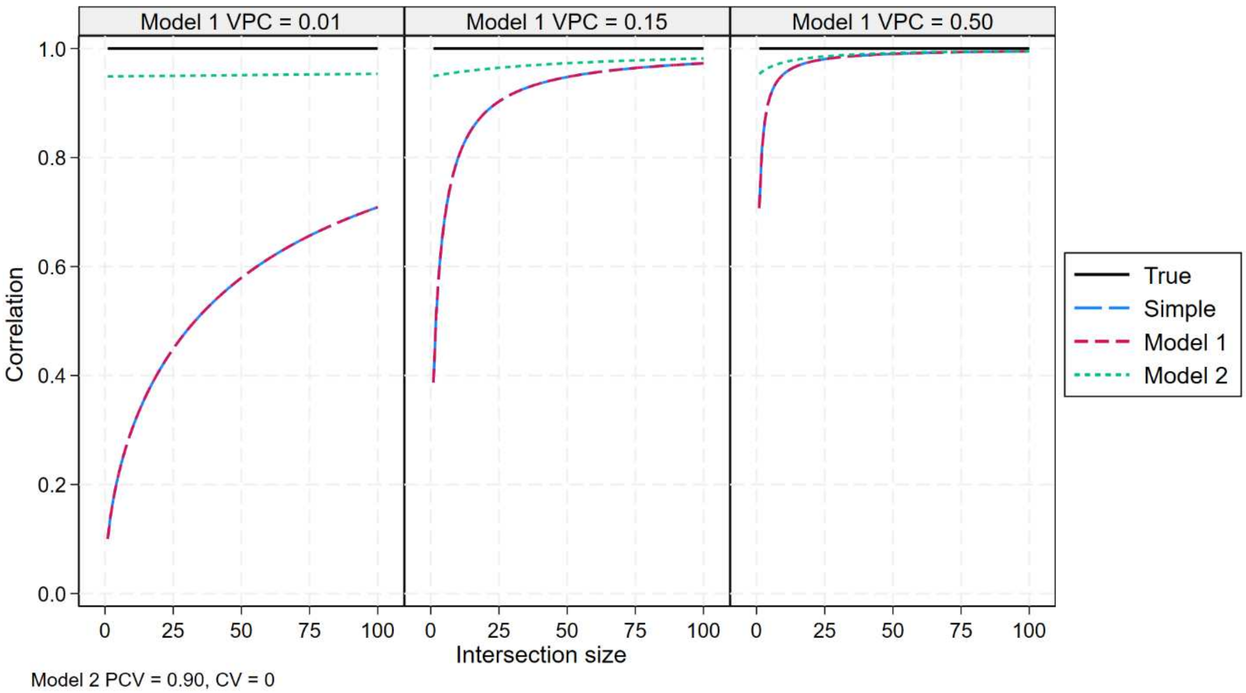

Figure S5.

Correlation between the simple means and the true means (20), MAIHDA Model 1 means and the true means (21), and MAIHDA Model 2 means and the true means (22), across the distribution of intersections, plotted against intersection size. The plot assumes a Model 1 VPC of 0.15, and equal intersection sizes. The plot is repeated for three Model 2 PCV values of 0.00 (no additive patterning in the intersectional inequalities), 0.50 (low), 0.90 (medium), and 0.99 (high).

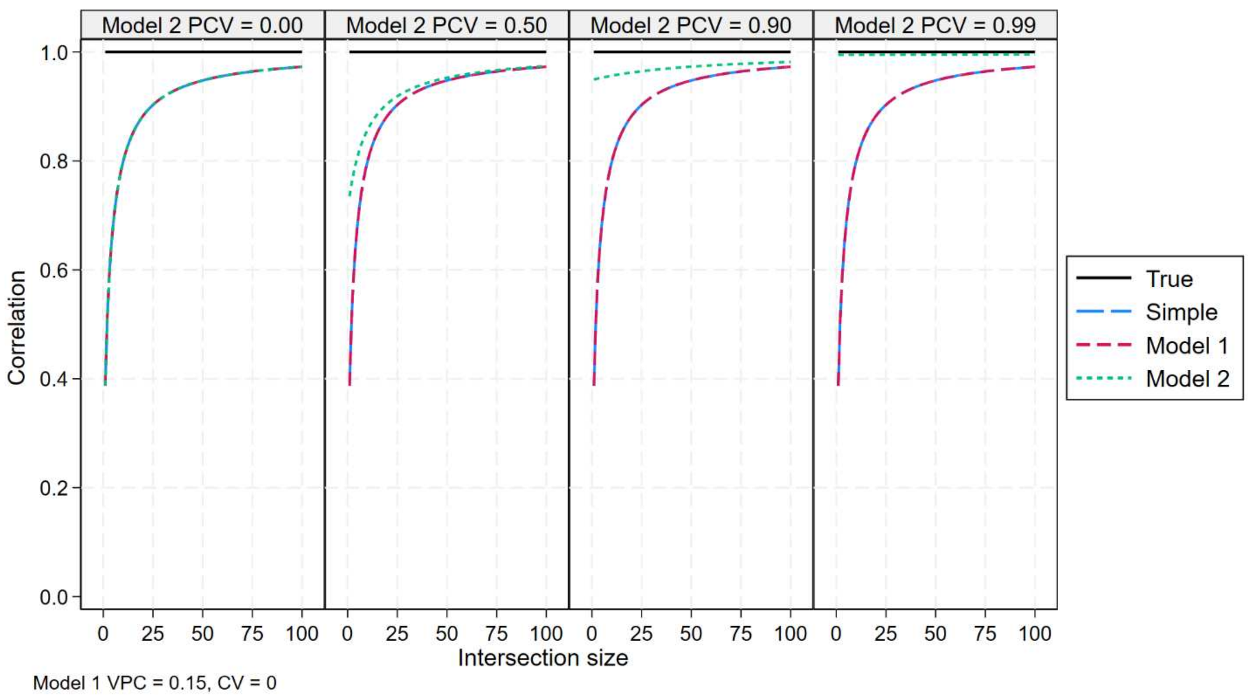

Figure S6.

Correlation between the simple means and the true means (20), MAIHDA Model 1 means and the true means (21), and MAIHDA Model 2 means and the true means (22), across the distribution of intersections, plotted against intersection size. The plot assumes a Model 1 VPC of 0.15, and a Model 2 PCV of 0.90. The plot is repeated for three coefficients of variation of intersection sizes of 0 (equal intersection sizes), 1 (medium variation in intersection sizes), and 2 (high variation in intersection sizes).

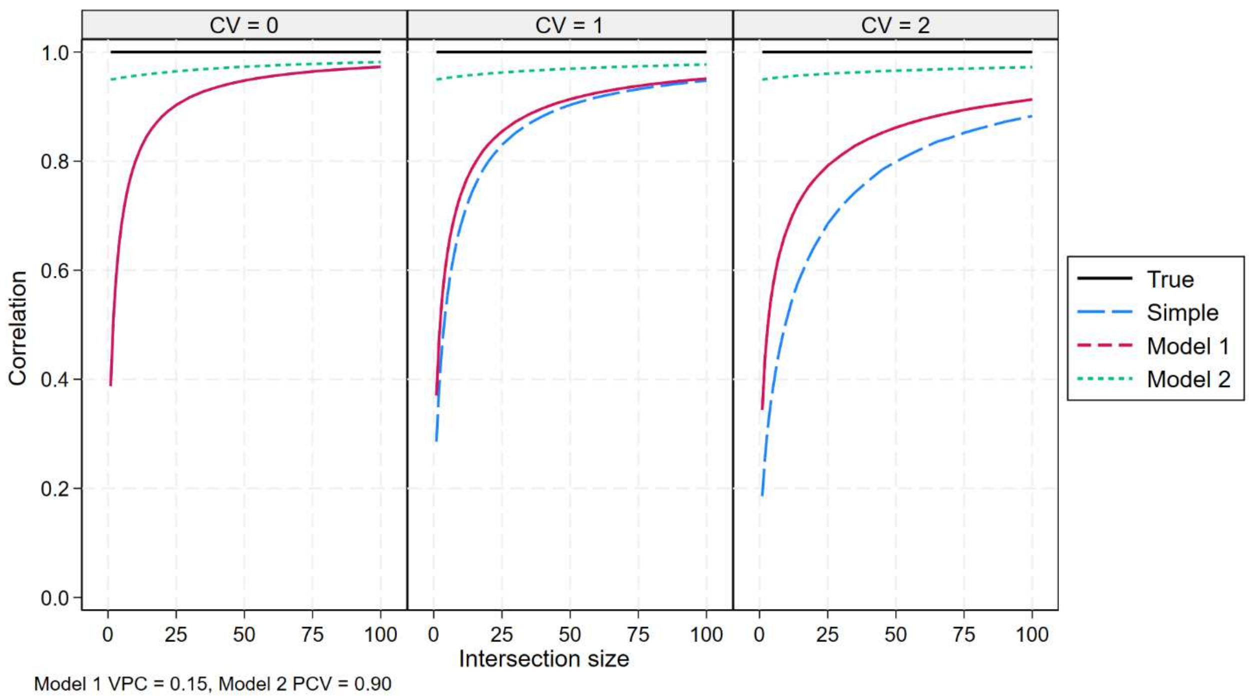

Figure S7. Bias of the simple mean (23), MAIHDA Model 1 mean (24), and MAIHDA Model 2 mean (25) across repeated samples of individuals for a given intersection, plotted against the difference between the true mean and the model-implied mean for that intersection. The plot assumes an intersection size of 10 individuals, and a Model 1 VPC of 0.15. The plot is repeated three times for intersection sizes of 5 (very small), 10 (small), and 50 (medium) individuals. The Model 1 model-implied mean corresponds to the overall mean, while the Model 2 model-implied mean reflects the additive effects of the social identities defining the intersection. The bias of the MAIHDA Model 2 mean is shown over a narrower range due to smaller differences between the true and model-implied means in this model.

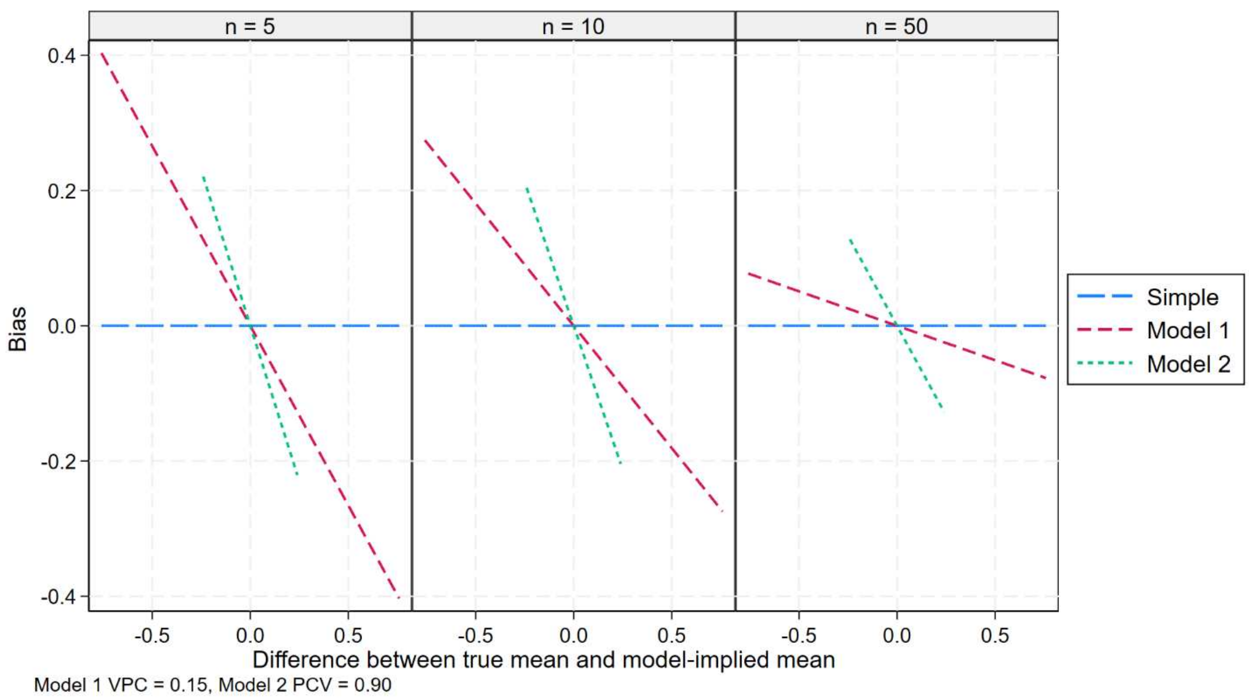

Figure S8. Bias of the simple mean (23), MAIHDA Model 1 mean (24), and MAIHDA Model 2 mean (25) across repeated samples of individuals for a given intersection,, plotted against the difference between the true mean and the model-implied mean for that intersection. The plot assumes an intersection size of 10 individuals, and a Model 2 PCV of 0.90. The plot is repeated for three Model 1 VPC values of 0.01 (small intersectional inequalities), 0.15 (medium) and 0.50 (large). The Model 1 model-implied mean corresponds to the overall mean, while the Model 2 model-implied mean reflects the additive effects of the social identities defining the intersection. The bias of the MAIHDA Model 2 mean is shown over a narrower range due to smaller differences between the true and model-implied means in this model.

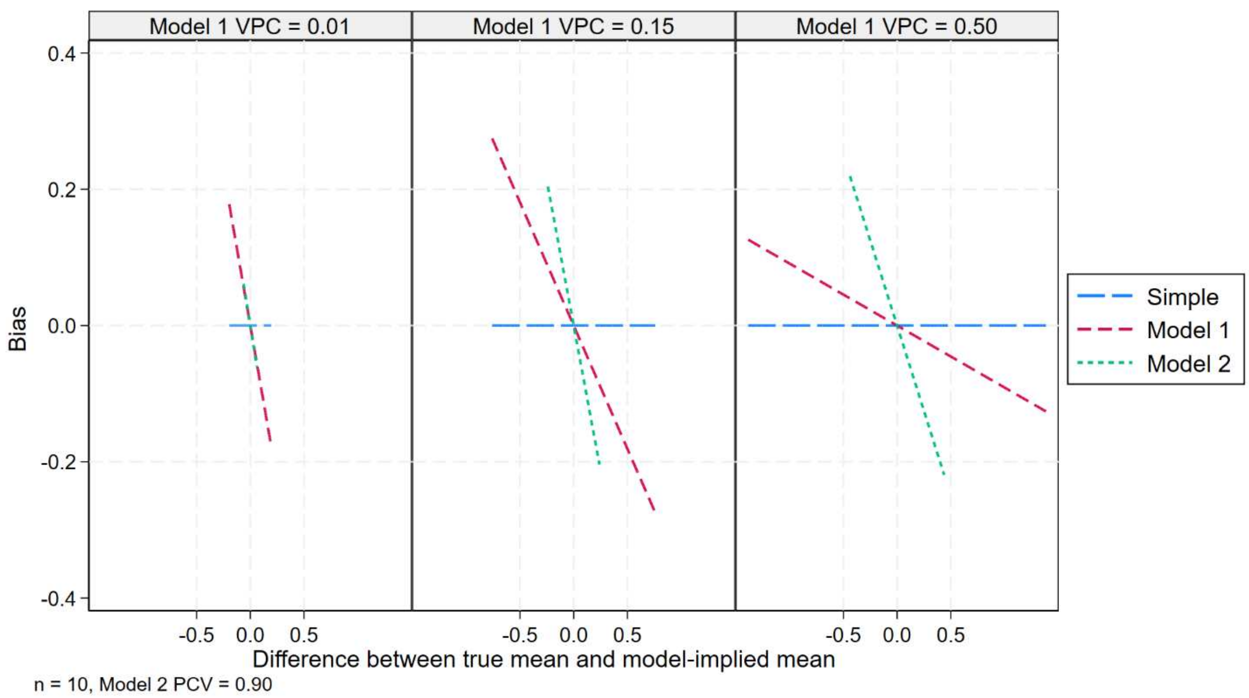

Figure S9. Bias of the simple mean (23), MAIHDA Model 1 mean (24), and MAIHDA Model 2 mean (25) across repeated samples of individuals for a given intersection,, plotted against the difference between the true mean and the model-implied mean for that intersection. The plot assumes an intersection size of 10 individuals, and a Model 1 VPC of 0.15. The plot is repeated for three Model 2 PCV values of 0.00 (no additive patterning in the intersectional inequalities), 0.50 (low), 0.50 (low), 0.90 (medium), and 0.99 (high). The Model 1 model-implied mean corresponds to the overall mean, while the Model 2 model-implied mean reflects the additive effects of the social identities defining the intersection. The bias of the MAIHDA Model 2 mean is shown over a narrower range due to smaller differences between the true and model-implied means in this model.

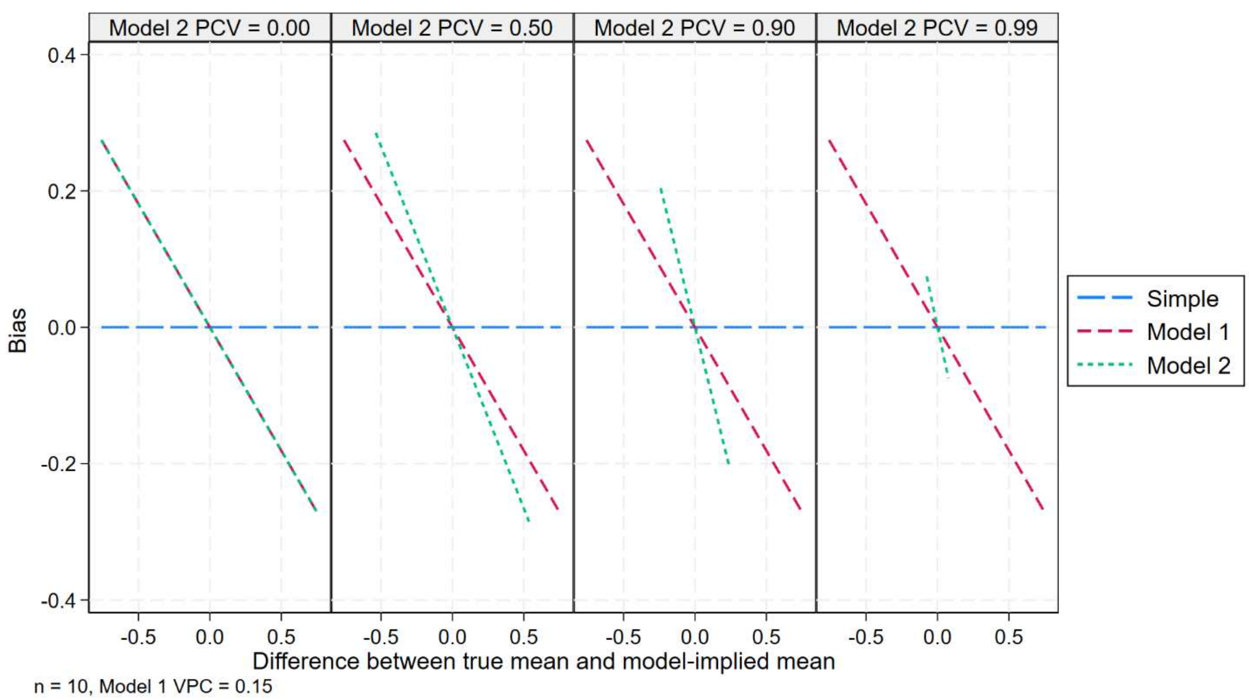

Figure S10. Variance of the simple mean (26), MAIHDA Model 1 mean (27), and MAIHDA Model 2 mean (28) across repeated samples of individuals for a given intersection, plotted against the difference between the true mean and the model-implied mean for that intersection. The plot assumes a Model 1 VPC of 0.15, and a Model 2 PCV of 0.90. The plot is repeated three times for intersection sizes of 5 (very small), 10 (small), and 50 (medium) individuals. The Model 1 model-implied mean corresponds to the overall mean, while the Model 2 model-implied mean reflects the additive effects of the social identities defining the intersection. The bias of the MAIHDA Model 2 mean is shown over a narrower range due to smaller differences between the true and model-implied means in this model.

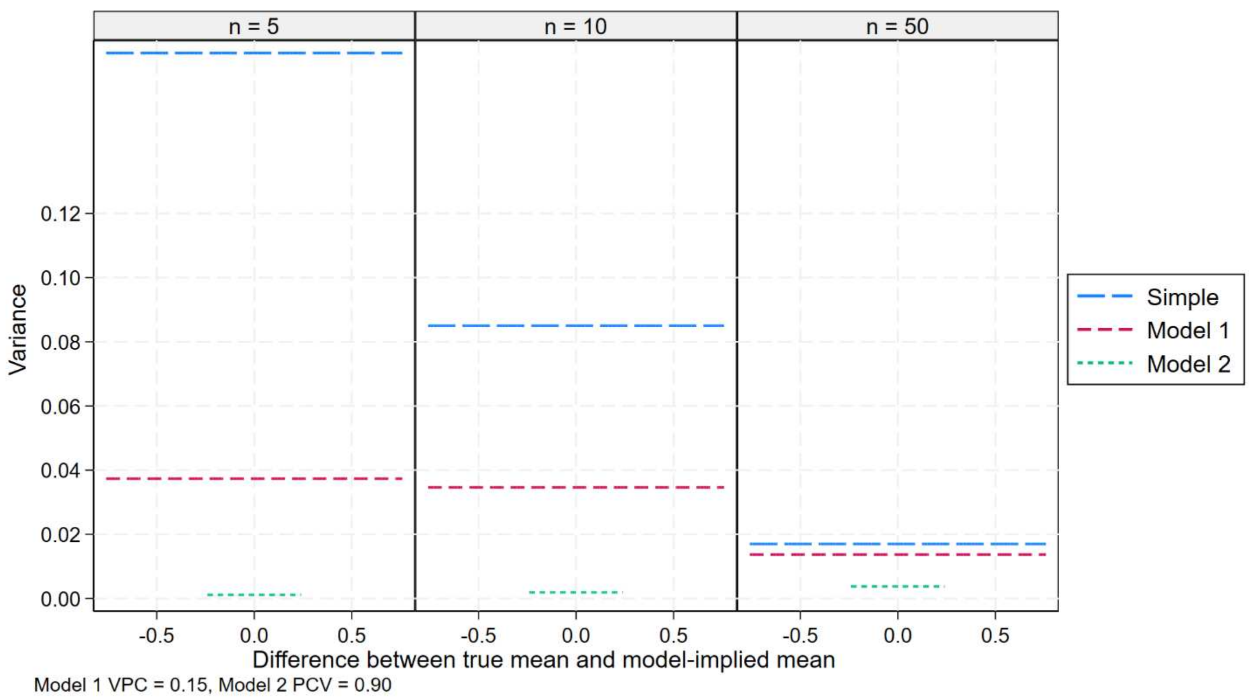

Figure S11. Variance of the simple mean (26), MAIHDA Model 1 mean (27), and MAIHDA Model 2 mean (28) across repeated samples of individuals for a given intersection, plotted against the difference between the true mean and the model-implied mean for that intersection. The plot assumes an intersection size of 10 individuals, and a Model 2 PCV of 0.90. The plot is repeated for three Model 1 VPC values of 0.01 (small intersectional inequalities), 0.15 (medium) and 0.50 (large). The Model 1 model-implied mean corresponds to the overall mean, while the Model 2 model-implied mean reflects the additive effects of the social identities defining the intersection. The bias of the MAIHDA Model 2 mean is shown over a narrower range due to smaller differences between the true and model-implied means in this model.

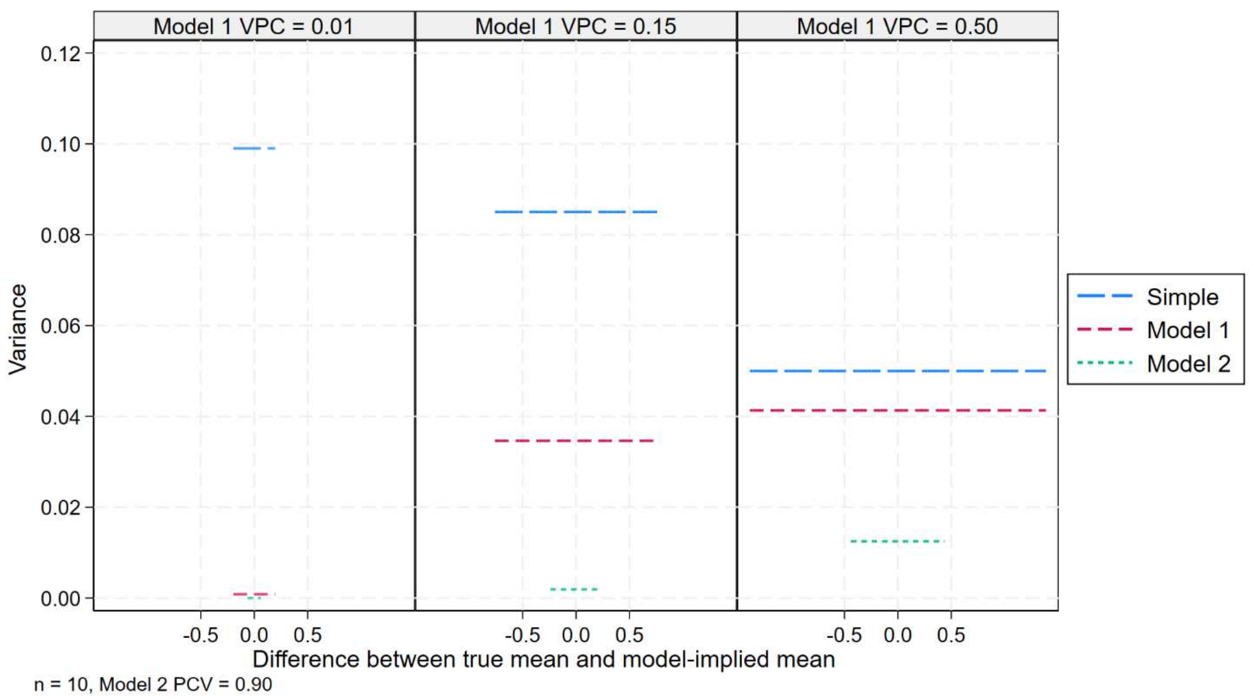

Figure S12. Variance of the simple mean (26), MAIHDA Model 1 mean (27), and MAIHDA Model 2 mean (28) across repeated samples of individuals for a given intersection, plotted against the difference between the true mean and the model-implied mean for that intersection. The plot assumes an intersection size of 10 individuals, and a Model 1 VPC of 0.15. The plot is repeated for three Model 2 PCV values of 0.00 (no additive patterning in the intersectional inequalities), 0.50 (low), 0.90 (medium), and 0.99 (high). The Model 1 model-implied mean corresponds to the overall mean, while the Model 2 model-implied mean reflects the additive effects of the social identities defining the intersection. The bias of the MAIHDA Model 2 mean is shown over a narrower range due to smaller differences between the true and model-implied means in this model.

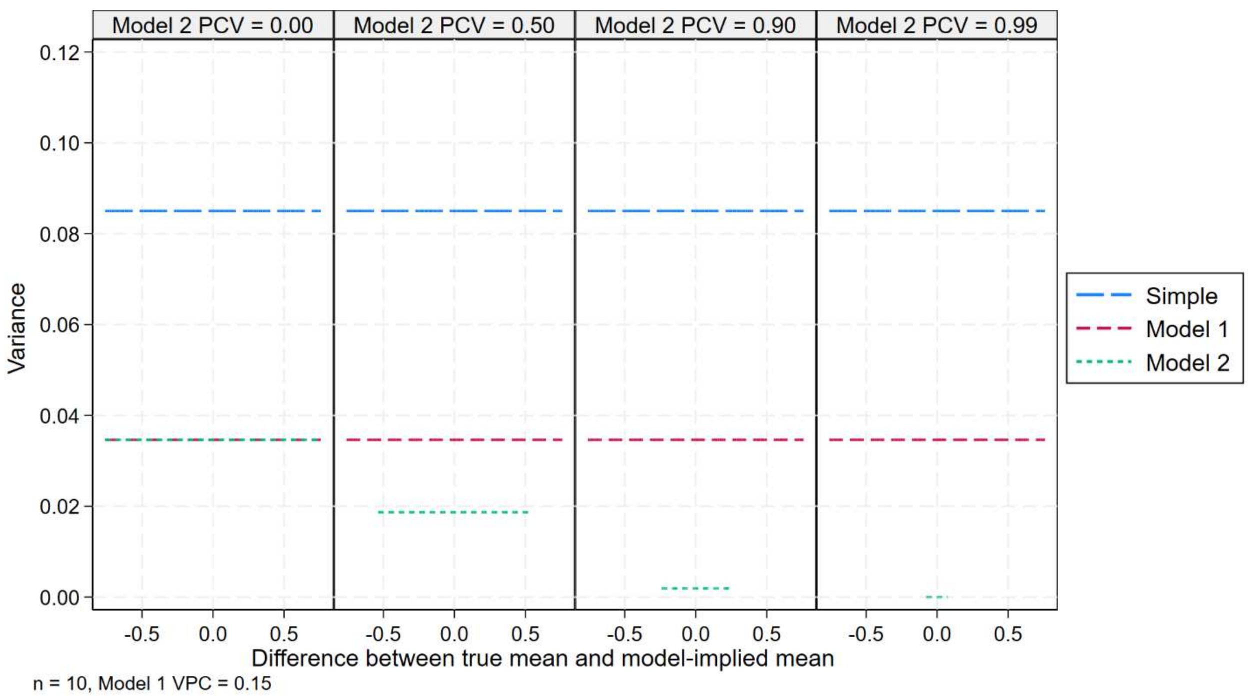

Figure S13.

Mean squared error of the simple mean (29), MAIHDA Model 1 mean (30), and MAIHDA Model 2 mean (31) across repeated samples of individuals for a given intersection, plotted against the difference between the true mean and the model-implied mean for that intersection. The plot assumes a Model 1 VPC of 0.15, and a Model 2 PCV of 0.90. The plot is repeated three times for intersection sizes of 5 (very small), 10 (small), and 50 (medium) individuals. The Model 1 model-implied mean corresponds to the overall mean, while the Model 2 model-implied mean reflects the additive effects of the social identities defining the intersection. The bias of the MAIHDA Model 2 mean is shown over a narrower range due to smaller differences between the true and model-implied means in this model.

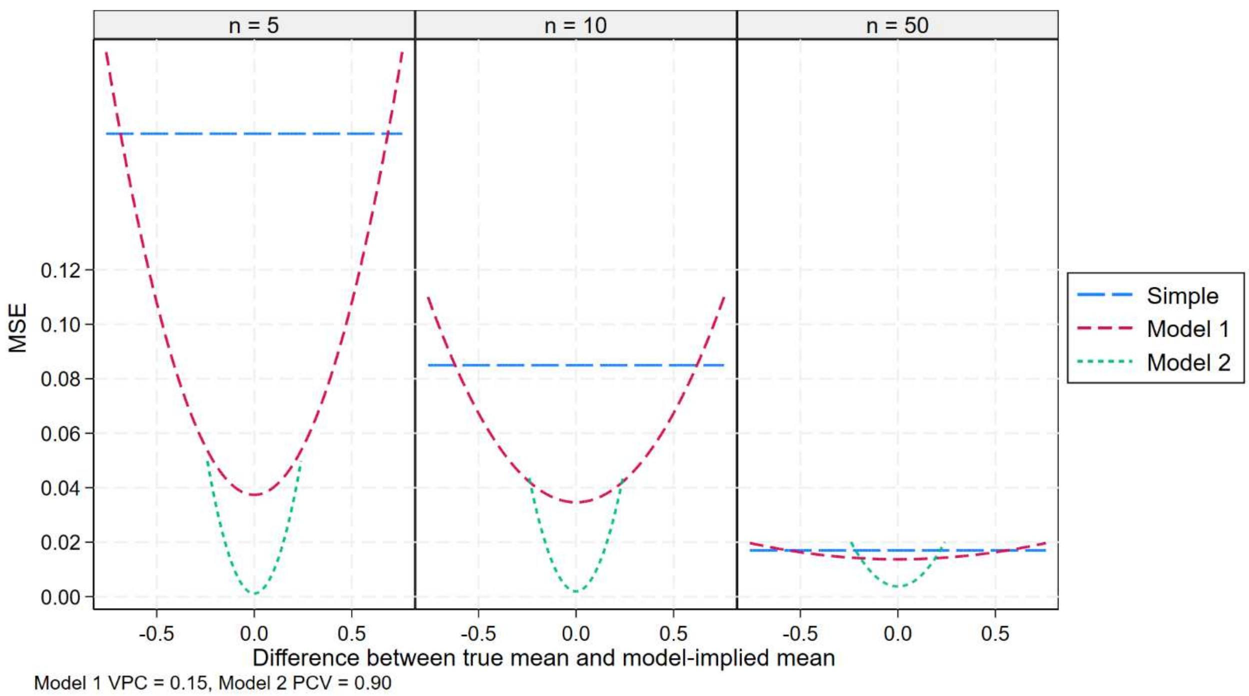

Figure S14. Mean squared error of the simple mean (29), MAIHDA Model 1 mean (30), and MAIHDA Model 2 mean (31) across repeated samples of individuals for a given intersection, plotted against the difference between the true mean and the model-implied mean for that intersection. The plot assumes an intersection size of 10 individuals, and a Model 2 PCV of 0.90. The plot is repeated for three Model 1 VPC values of 0.01 (small intersectional inequalities), 0.15 (medium) and 0.50 (large). The Model 1 model-implied mean corresponds to the overall mean, while the Model 2 model-implied mean reflects the additive effects of the social identities defining the intersection. The bias of the MAIHDA Model 2 mean is shown over a narrower range due to smaller differences between the true and model-implied means in this model.

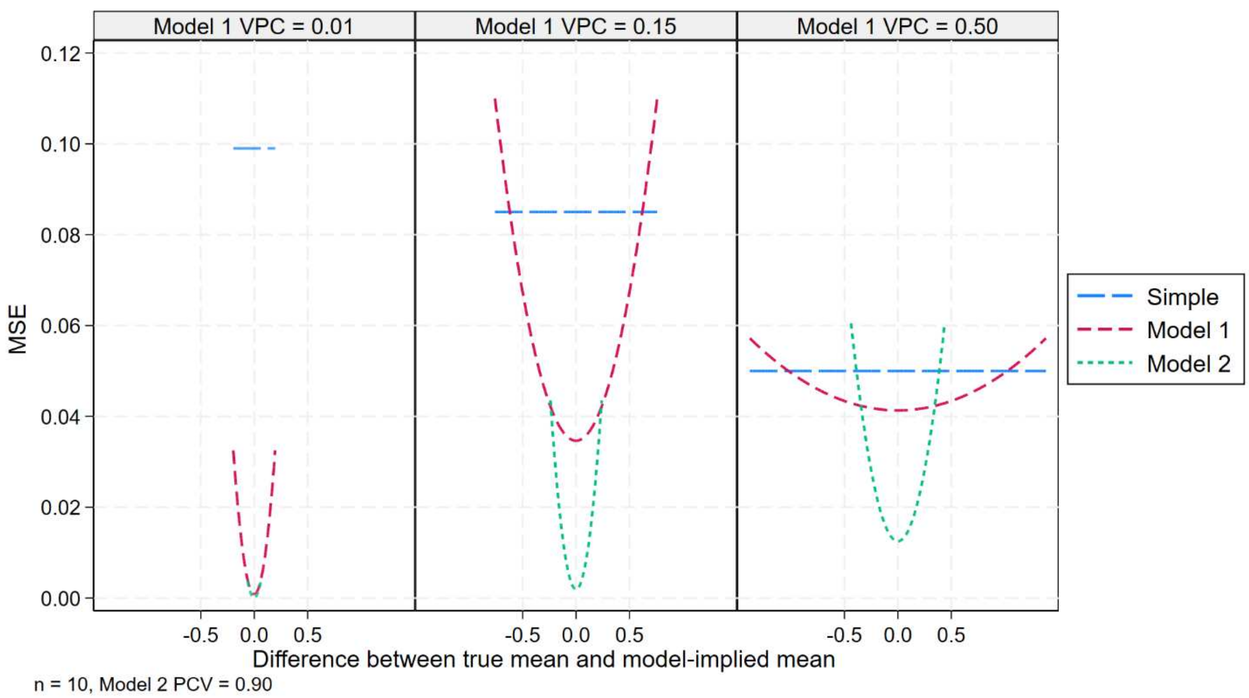

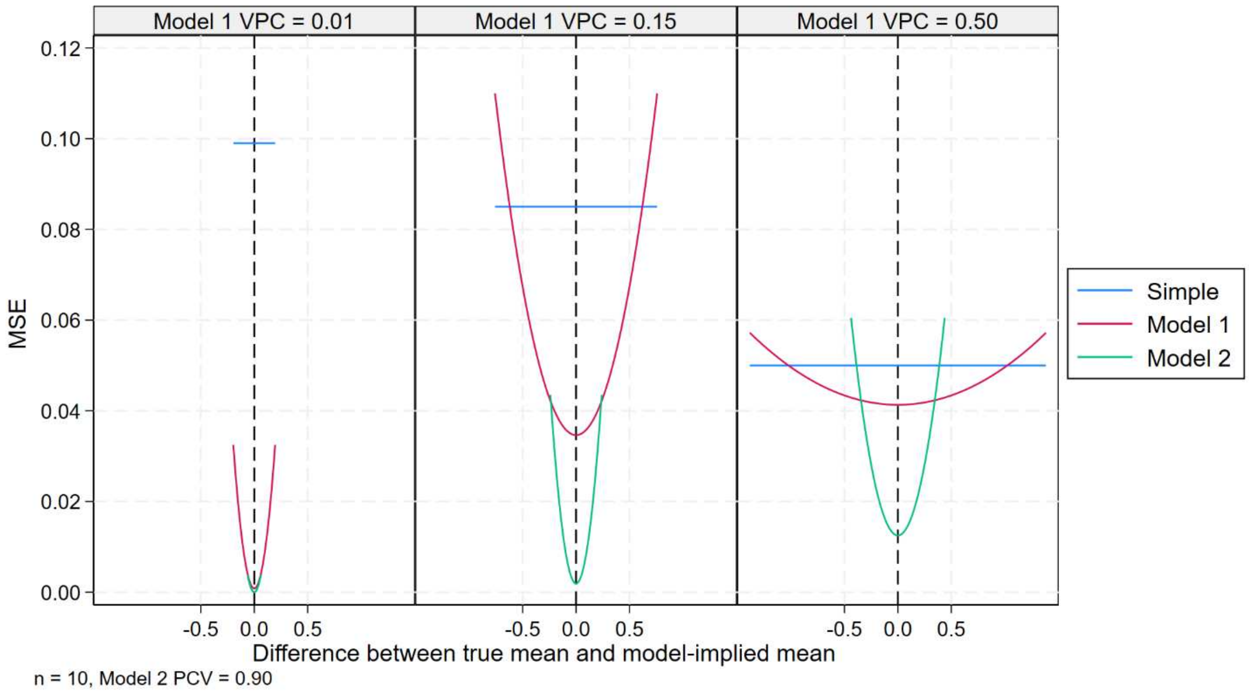
Model 1 VPC = 0.01
Model 1 VPC = 0.15
Model 1 VPC = 0.50
0.12
0.10
0.08
0.06
0.04
0.02
0.00
MSE
-0.5
0.0
0.5
Difference between true mean and model-implied mean
Simple
Model 1
Model 2
n = 10, Model 2 PCV = 0.90

Figure S15. Mean squared error of the simple mean (29), MAIHDA Model 1 mean (30), and MAIHDA Model 2 mean (31) across repeated samples of individuals for a given intersection, plotted against the difference between the true mean and the model-implied mean for that intersection. The plot assumes an intersection size of 10 individuals, and a Model 1 VPC of 0.15. The plot is repeated for three Model 2 PCV values of 0.00 (no additive patterning in the intersectional inequalities), 0.50 (low), 0.90 (medium), and 0.99 (high). The Model 1 model-implied mean corresponds to the overall mean, while the Model 2 model-implied mean reflects the additive effects of the social identities defining the intersection. The bias of the MAIHDA Model 2 mean is shown over a narrower range due to smaller differences between the true and model-implied means in this model.