

WORKING PAPER · NO. 2025-131

Monotone Ecological Inference

Hadi Elzayn, Jacob Goldin, Cameron Guage, Daniel E. Ho, and Claire Morton

SEPTEMBER 2025

Monotone Ecological Inference

Hadi Elzayn, Jacob Goldin, Cameron Guage, Daniel E. Ho, and Claire Morton

September 2025

JEL No. C1, C10, C13, C18, C21

ABSTRACT

We study monotone ecological inference, a partial identification approach to ecological inference. The approach exploits information about one or both of the following conditional associations: (1) outcome differences between groups within the same neighborhood, and (2) outcomes differences within the same group across neighborhoods with different group compositions. We show how assumptions about the sign of these conditional associations, whether individually or in relation to one another, can yield informative sharp bounds in ecological inference settings. We illustrate our proposed approach using county-level data to study differences in Covid-19 vaccination rates among Republicans and Democrats in the United States.

Hadi Elzayn
Stanford University
hselzayn@law.stanford.edu

Daniel E. Ho
Stanford University
dho@law.stanford.edu

Jacob Goldin
University of Chicago
University of Chicago Law School
and NBER
jsgoldin@gmail.com

Claire Morton
University of California, Berkeley
mortonc@berkeley.edu

Cameron Guage
Columbia University
Department of Economics
cameronguage@gmail.com

An online appendix is available at <http://www.nber.org/data-appendix/w34285>

Monotone Ecological Inference*

Hadi Elzayn Jacob Goldin Cameron Guage
Daniel E. Ho Claire Morton

September 12, 2025

Abstract

We study monotone ecological inference, a partial identification approach to ecological inference. The approach exploits information about one or both of the following conditional associations: (1) outcome differences between groups within the same neighborhood, and (2) outcomes differences within the same group across neighborhoods with different group compositions. We show how assumptions about the sign of these conditional associations, whether individually or in relation to one another, can yield informative sharp bounds in ecological inference settings. We illustrate our proposed approach using county-level data to study differences in Covid-19 vaccination rates among Republicans and Democrats in the United States.

1 Introduction

Ecological inference (EI) – the use of aggregate data to investigate individual-level associations – is a common challenge in fields such as political science, sociology, economics, epidemiology, and public health. To fix ideas, consider a researcher seeking to learn the difference in the prevalence of some outcome across two groups of individuals, but data on both the outcome and group membership are available only at some aggregated level, such as the

*For helpful comments and suggestions, we are grateful to Kosuke Imai, Gary King, Charles Manski, Shiyong Hao, and Derek Ouyang. Elzayn, Ho, and Morton: Stanford University. Goldin: University of Chicago, American Bar Foundation, and NBER. Guage: Columbia University.

individual’s neighborhood. EI is challenging because the neighborhood-level data can be consistent with divergent individual-level associations between group membership and the outcome.

There are several canonical estimation strategies that have been proposed for EI settings, but each requires strong assumptions. Ecological regression [Goodman, 1953] entails regressing neighborhood-level outcomes on neighborhood-level rates of group membership. The approach is unbiased only absent “contextual effects” – e.g., it would be biased if one group is more likely to live in neighborhoods where individuals of both groups tend to have higher values of the outcome. Alternatively, the neighborhood model [Freedman et al., 1991] uses group-weighted averages across neighborhoods to estimate group means. It yields unbiased estimates under a different, but equally strong assumption, namely that there is no association between group membership and the outcome among individuals living in the same neighborhood. A third approach to EI, proposed by King [1997], relies for identification on an assumption about the role of contextual effects within a parametric statistical model.¹

A limitation to all of these approaches is that they require strong assumptions for point identification. A potentially appealing alternative is therefore to focus on partial identification methods, such as the so-called method of bounds [Duncan and Davis, 1953], which refers to the sharp bounds obtained from imposing only that the outcome is bounded. In practice, however, the method of bounds interval is frequently wide. Depending on the application, such bounds may not shed much light on the parameter of interest.

In this paper, we propose a middle-ground approach to identification in EI settings that strikes a balance between obtaining informative results while relying on potentially more credible identifying assumptions. The identifying assumptions we consider concern two quantities: (1) the between-group within-neighborhood association – i.e., outcome differences between groups within the same neighborhood, and (2) the within-group between-neighborhood association – i.e., outcome differences within the same group across neighborhoods with different group compositions. We show how as-

¹Specifically, King models the individual-level variables of interest as independent draws from a truncated bivariate normal distribution, where truncation ensures that the estimated proportions fall in the unit interval. In doing so, King’s approach requires a similar independence assumption as ecological regression; the value of the approach is to combine the identifying power of that assumption with the constraints implied by worst-case bounds.

assumptions about the *sign* of these conditional associations – whether individually or in relation to one another – can aid identification. In particular, we derive sharp bounds for the group-level outcome means and for the difference in outcome means by group in EI settings in which the researcher can sign one or both of the conditional associations. Because our identifying assumptions relate to the signs of the conditional associations, we refer to our proposed method as *monotone ecological inference*, in the spirit of Manski and Pepper [2000].

A virtue of monotone ecological inference is that it relies on weaker identifying assumptions than canonical EI estimators, in that it does not require taking a stance on the exact magnitude of either the between- or the within-group association. In contrast, the neighborhood model estimator is unbiased only if the between-group association is exactly zero, and ecological regression is unbiased only if the within-group association is exactly zero.² These distinctions are important in practice because in many settings of interest, the researcher will not be able to entirely rule out within-neighborhood variation in group outcomes (so that the between-group association may be non-zero) nor be able to rule out the presence of all contextual effects (so that the within-group association may be non-zero). At the same time, the researcher may be able to use theory or auxiliary data to form a belief about the likely direction of any such associations if they are present, enabling monotone ecological inference.³

A related possibility is that the researcher may have reason to believe that the two conditional associations run in the same direction as one another, even without necessarily knowing the direction. For example, the researcher may expect that the contextual effects in a particular setting would amplify any group differences in the outcome that would otherwise exist, such as through social norms or local control over policy-making, so that the within-group (between-neighborhood) association would tend to have the same sign as the between-group (within-neighborhood) association. Under this condition, which we refer to as *contextual reinforcement*, we show that the data

²King’s approach also imposes that the within-group association be zero a priori, but allows for a non-zero conditional association after incorporating feasibility constraints based on the neighborhood-level data [for discussion, see Lewis, 2001, Jiang et al., 2020].

³As we illustrate below, an auxiliary sample of individual-level data may be useful to assess the sign of one or both of the conditional associations, even if the sample is not sufficiently large or representative enough to answer the primary research question on its own.

identifies the sign of the difference in group means, and we provide sharp bounds for the group means as well as for their difference.

The main reason that monotone ecological inference is appealing as a research design is that its identifying assumptions can often be reasoned about on the basis of expert institutional knowledge and/or auxiliary data. We illustrate this type of reasoning in our empirical application, where we study differences in Covid-19 vaccination rates by political party using county-level data. This topic has been the focus of substantial interest in recent years, and much of the prior evidence is ecological in nature [Albrecht, 2022, Ye, 2023].

In applying monotone ecological inference to this research question, we start by noting that the setting is one in which contextual reinforcement is likely to hold. In particular, because many public health services and policies like vaccination mandates are determined and applied on a sub-national level, we expect that any within-neighborhood association between political party and vaccination status would be amplified through neighborhood-level mediators. For example, suppose that Republicans are more vaccine hesitant for reasons that are not fully mediated through neighborhood (e.g., differences in partisan media exposure and attitude formation), so that the between-group association is negative. We would then also expect that Republicans living in “red” counties (where most of the residents are Republicans) would be more likely to face local policies and social norms inhibiting vaccinations, as compared to Republicans living in “blue” counties (where few of the residents are Republicans); hence the within-group (across-county) association would also be negative.⁴ The plausibility of such “feedback” from individual- to neighborhood-level associations supports the assumption of contextual reinforcement in this setting.

To augment the theoretical support for contextual reinforcement, we obtain an auxiliary data set consisting of individual-level vaccination records for approximately one million individuals from a registry of electronic health records. We link these records to publicly available data on voter registration and then anonymize them. For this sample, we find that both the between- and within-group associations are negative: Republicans are less likely to be vaccinated than Democrats living in the same county, and individuals of both parties are more likely to be vaccinated in counties where individuals of their

⁴Conversely, if Democrats tended to be more opposed to vaccination, we would expect the opposite pattern to obtain for both identifying assumptions.

same party have higher vaccination rates. Although the exact magnitude of these associations is unlikely to be the same in the overall population, we interpret the sign of the estimated associations for this subpopulation to support the contextual reinforcement assumption.

Analyzing the county-level data using monotone ecological inference, we find evidence that Democrats are vaccinated at higher rates than Republicans. We estimate the vaccination rate among Democrats to be between 0.61 and 0.81, compared to between 0.33 and 0.56 among Republicans. With respect to the difference in vaccination rates between Democrats and Republicans, the width of the contextual reinforcement bounds is approximately 67% narrower than the method of bounds interval. And unlike the method of bounds, the contextual reinforcement bounds exclude the possibility that there is no difference in vaccination rates by political party.

Our results contribute to a large literature that studies identification in EI settings; [for surveys see, e.g., King, 1997, King et al., 2004, Cho and Manski, 2008]. Most of this literature focuses on point-identification, with notable exceptions that include Duncan and Davis [1953], Horowitz and Manski [1995], Cross and Manski [2002], Greiner and Quinn [2009], Manski [2018], and Jiang et al. [2020]. Among these, Manski [2018] in particular shares a key feature of our approach, which is to consider sign restrictions on the joint distribution of individual-level variables to aid in identification. However, the monotonicity assumptions we consider differ substantially from those considered by Manski [2018] and aid in identifying a different parameter, leading us to a very different set of results and insights.⁵ More recently, Li et al. [2023] also consider the role of “bounded variation assumptions” for identifying personalized risk assessments from published medical studies, a setting that shares some features of ecological inference but also differs from it in important respects.

We contribute to this literature by proposing a novel partial identification strategy that can yield informative bounds under appealing identifying assumptions. A distinct contribution is to provide novel expressions for the relationship among canonical EI estimators as well as for their respective estimands, clarifying the interpretation of results when they are applied to estimate population group differences.

⁵Specifically, our primary focus is on identifying group characteristics at the population-level whereas Manski [2018] studies identification of group characteristics for particular neighborhoods. We study the latter problem in Section 5, where we relate our approach to Manski’s and introduce novel monotonicity conditions that can aid identification.

Outside of the EI setting, our results relate to a literature that infers disparities in individual-level data based on probabilistic estimates of group membership [Chen et al., 2019, Kallus et al., 2022, McCartan et al., 2024]. Closest to our approach is Elzayn et al. [2025], which applies a partial identification strategy for estimating income tax audit disparities by race using individual-level data and probabilistically inferred racial characteristics based on the sign of conditional covariance terms that are the individual-level analogs to the conditional associations we study.⁶ Our approach is also similar in spirit to prior work that studies the identification of treatment effects when the researcher substitutes identifying assumptions based on inequalities for identifying assumptions based on equalities [e.g., Manski and Pepper, 2000, Molinari, 2010, Manski and Pepper, 2018, Rambachan and Roth, 2023].

We proceed as follows. Section 2 describes our empirical setup and defines several canonical EI estimators. Section 3 derives sharp bounds for the population difference in group means using Monotone EI. Section 4 extends the same approach to identification of the levels of group means. Section 5 develops Monotone EI tools for identification of group means at the neighborhood-level. Section 6 applies our results to study partisan gaps in Covid-19 vaccinations. Section 7 concludes. An open-source R software package, MonotoneEI, is available to implement our proposed approach. The software and accompanying documentation are available at <https://github.com/reglab/MonotoneEI>.

2 Empirical Framework

A population of individuals is characterized by a triple (X, Y, N) . We interpret X to define the individual’s group membership; Y to be the outcome of interest; and N as the individual’s neighborhood. For expositional simplicity, we initially focus on the case in which $X \in \{0, 1\}$; the analysis extends naturally to settings in which there are more than two discrete groups, as we discuss below. We are primarily interested in the mean values of Y by group,

$$Y^x = \mathbb{E}[Y|X = x],$$

⁶Elzayn et al. [2025] does not relate its approach to ecological inference nor derive the sharp bounds we provide here.

as well as the difference in group means,

$$D = Y^1 - Y^0.$$

Following the presentation of our main results, we also consider the identification of the levels and difference of the group means for specific neighborhoods, $\mathbb{E}[Y|X, N = n]$ and $\mathbb{E}[Y|X = 1, N = n] - \mathbb{E}[Y|X = 0, N = n]$.

We do not observe individual-level values of X or Y , but rather observe data that has been aggregated across individuals with the same value of N . Denote the mean values of X and Y for the individuals in a particular neighborhood n by

$$X_n = \mathbb{E}[X|N = n]$$

and

$$Y_n = \mathbb{E}[Y|N = n].$$

The distribution of individuals across neighborhoods is given by

$$p_n = \Pr(N = n).$$

We assume the researcher can directly observe (X_n, Y_n, p_n) for each n , deferring issues of sampling uncertainty and statistical inference to our empirical application. We also assume that there are a finite number of neighborhoods, $N \in \mathcal{N}$ with $|\mathcal{N}| < \infty$. To avoid degenerate cases, we assume that there exists some $n \in \mathcal{N}$ for which $p_n > 0$ and $X_n \in (0, 1)$. In addition, we assume that for each group x and neighborhood n , the conditional expectation of Y with respect to x and n exists and is bounded, $\underline{Y} \leq Y_n^x \leq \bar{Y}$ for some known pair $(\underline{Y}, \bar{Y}) \in \mathbb{R}^2$, where $Y_n^x = \mathbb{E}[Y|N = n, X = x]$. This assumption will be naturally satisfied when Y itself has finite support.

Figure 1 visualizes the joint relationship between X , Y , and N . The overall population association between X and Y can be decomposed into two conditional associations. We formally define the *between-group association* as

$$\delta_B := \mathbb{E}[\text{Cov}(Y, X | N)]$$

and the *within-group association* as:

$$\delta_W := \mathbb{E}[\text{Cov}(Y, X_N | X)]$$

where $X_N = \mathbb{E}[X|N]$. Intuitively, the between-group association, δ_B , refers to differences in the outcome between groups within the same neighborhood.

In turn, the within-group association, δ_W , refers to differences in the outcome across neighborhoods with different group prevalence, among individuals within the same group.⁷

It is possible to relate the overall difference in group means, D , to these conditional associations, δ_B and δ_W . To do so, define $\gamma = \frac{\text{Var}(X_N)}{\text{Var}(X)}$. Intuitively, γ reflects the loss of information due to the data being aggregated, so that $\gamma < 1$ in EI settings. We then have the following result.

Proposition 1 (Decomposition of Overall Association into Conditional Associations).

$$D = \frac{\delta_B + \delta_W}{(1 - \gamma) \text{Var}(X)}$$

We defer the proof of Proposition 1, and all subsequent results, to the Appendix. Although we are not aware of prior work that has linked the difference in group means to δ_B and δ_W in this way, the γ term that appears in Proposition 1 is familiar within the EI literature as a mediator of the discrepancy between individual and ecological estimation approaches.⁸ Because $\gamma < 1$, an immediate implication of Proposition 1 is that the sign of D is the same as that of $\delta_B + \delta_W$.

Canonical EI Estimators

We next describe three canonical approaches for estimating the group-level means in EI settings. We focus on these three methods (as opposed to other popular methods like King [1997]) because they will appear as inputs into the monotone EI bounds we derive.

The most common method for conducting EI is ecological regression (ER) [Goodman, 1953]. The ecological regression estimator for the difference in group means, $\hat{D}_{ER}$, is defined as the estimated coefficient for X_N in the weighted least squares regression of Y_N on X_N , with weights based on the

⁷These interpretations follows from the fact that with binary X , $\text{Cov}(X, Y|N = n) = \mathbb{E}[XY|N = n] - \mathbb{E}[Y|N = n]\mathbb{E}[X|N = n] = \Pr[X = 1|N = n](Y_n^1 - \mathbb{E}[Y|N]) = \text{Var}[X|N = n](Y_n^1 - Y_n^0)$. Hence, $\delta_B = \sum_{n \in \mathcal{N}} (Y_n^1 - Y_n^0) \text{Var}[X|n] p_n$. Similarly, $\delta_W = \mathbb{E}[X] \text{Cov}(Y_N, X_N|X = 1) + (1 - \mathbb{E}[X]) \text{Cov}(Y_N, X_N|X = 0)$.

⁸For example, γ corresponds to $\eta_{X,A}^2$ in Equation 1 of Robinson [1950].

share of the population in each neighborhood:

$$\hat{D}_{ER} = \frac{\sum_{n \in \mathcal{N}} p_n \tilde{Y}_n \tilde{X}_n}{\sum_{n \in \mathcal{N}} p_n (\tilde{X}_n)^2}$$

where $\tilde{X}_n = X_n - \frac{\sum_n p_n X_n}{\sum_n p_n}$ denotes the demeaned value of X_n , and similarly for $\tilde{Y}_n = Y_n - \frac{\sum_n p_n Y_n}{\sum_n p_n}$. We will focus on the ecological regression estimator's asymptotic limit under an *iid* sampling process:

$$D_{ER} = \frac{\text{Cov}(Y_N, X_N)}{\text{Var}(X_N)}$$

where $X_N = \mathbb{E}[X|N]$ and $Y_N = \mathbb{E}[Y|N]$.

In turn, the ecological regression estimates for the group-level means are given by

$$\hat{Y}_{ER}^0 = \sum_{n \in \mathcal{N}} p_n Y_n - \left(\frac{\sum_{n \in \mathcal{N}} p_n \tilde{Y}_n \tilde{X}_n}{\sum_{n \in \mathcal{N}} p_n (\tilde{X}_n)^2} \right) \left(\sum_{n \in \mathcal{N}} p_n X_n \right)$$

and

$$\hat{Y}_{ER}^1 = \sum_{n \in \mathcal{N}} p_n Y_n + \left(\frac{\sum_{n \in \mathcal{N}} p_n \tilde{Y}_n \tilde{X}_n}{\sum_{n \in \mathcal{N}} p_n (\tilde{X}_n)^2} \right) \left(\sum_{n \in \mathcal{N}} p_n (1 - X_n) \right)$$

which respectively converge to

$$Y_{ER}^0 = \mathbb{E}[Y_N] - \left(\frac{\text{Cov}(Y_N, X_N)}{\text{Var}(X_N)} \right) \mathbb{E}[X_N]$$

and

$$Y_{ER}^1 = \mathbb{E}[Y_N] + \left(\frac{\text{Cov}(Y_N, X_N)}{\text{Var}(X_N)} \right) (1 - \mathbb{E}[X_N]).^9$$

An alternative method for estimating group-level differences in ecological inference settings is the so-called Neighborhood Model (NM) [Freedman

⁹One can equivalently define the ecological regression estimates for the group means as the estimated coefficients of the weighted regression of Y_N on X_N and $1 - X_N$.

et al., 1991].¹⁰ The neighborhood model estimators for the group-level means are weighted averages of the neighborhood-level outcomes, with weights given by the group's prevalence in the neighborhood:

$$\hat{Y}_{NM}^1 = \frac{\sum_{n \in \mathcal{N}} p_n X_n Y_n}{\sum_{n \in \mathcal{N}} p_n X_n}$$

and

$$\hat{Y}_{NM}^0 = \frac{\sum_{n \in \mathcal{N}} p_n (1 - X_n) Y_n}{\sum_{n \in \mathcal{N}} p_n (1 - X_n)}$$

The neighborhood model estimator for the difference in group means, $\hat{D}_{NM}$, is correspondingly defined as

$$\hat{D}_{NM} = \frac{\sum_{n \in \mathcal{N}} p_n X_n Y_n}{\sum_{n \in \mathcal{N}} p_n X_n} - \frac{\sum_{n \in \mathcal{N}} p_n (1 - X_n) Y_n}{\sum_{n \in \mathcal{N}} p_n (1 - X_n)}$$

As above, we will focus on the neighborhood model estimators' asymptotic limits:

$$Y_{NM}^1 = \frac{\mathbb{E}[X_N Y_N]}{\mathbb{E}[X_N]}$$

$$Y_{NM}^0 = \frac{\mathbb{E}[(1 - X_N) Y_N]}{\mathbb{E}[1 - X_N]}$$

and

$$D_{NM} = \frac{\mathbb{E}[X_N Y_N]}{\mathbb{E}[X_N]} - \frac{\mathbb{E}[(1 - X_N) Y_N]}{\mathbb{E}[1 - X_N]}$$

It will be useful in our subsequent results to observe that ecological regression and the neighborhood model estimates follow a close mechanical relationship.¹¹

¹⁰Freedman et al. [1991] proposes two model variants: the non-linear neighborhood model and the linear neighborhood. Our focus is on the former.

¹¹Surprisingly, the relationship described in Lemma 1 has not received attention in the EI literature. The only prior reference to it that we could identify appears in an unpublished monograph Ansolabehere and Rivers [1995].

Lemma 1 (Relationship Between Neighborhood Model and Ecological Regression).

$$D_{NM} = \gamma D_{ER}$$

where $\gamma = \frac{\text{Var}(X_N)}{\text{Var}(X)} < 1$, as above. Whereas ecological regression and the neighborhood model yield point estimates for the group-level means and their difference, an alternative approach is to calculate the most extreme values of these quantities that are consistent with the observed data. In EI settings, consistency with the data requires that, for each $n \in \mathcal{N}$, (i) $Y_n^0 \in [\underline{Y}, \bar{Y}]$; (ii) $Y_n^1 \in [\underline{Y}, \bar{Y}]$; and (iii) $X_n Y_n^1 + (1 - X_n) Y_n^0 = Y_n$. The Method of Bounds (MOB) interval is defined by the minimum and maximum group-level means that satisfy these constraints.¹²

Proposition 2 (Method of Bounds for Group Means). *Suppose that for each x and n , the conditional expectation of Y with respect to x and n exists, $\underline{Y} \leq \mathbb{E}[Y|X = x, N = n] \leq \bar{Y}$ for some pair $(\underline{Y}, \bar{Y}) \in \mathbb{R}^2$. Define the following parameters:*

$$\begin{aligned} Y_{MOB}^{1+} &:= \frac{\mathbb{E}[\min\{Y_N - \underline{Y}(1 - X_N), \bar{Y}X_N\}]}{\mathbb{E}[X]} \\ Y_{MOB}^{0-} &:= \frac{\mathbb{E}[\max\{Y_N - \bar{Y}X_N, \underline{Y}(1 - X_N)\}]}{1 - \mathbb{E}[X]} \\ Y_{MOB}^{1-} &:= \frac{\mathbb{E}[\max\{Y_N - \bar{Y}(1 - X_N), \underline{Y}X_N\}]}{\mathbb{E}[X]} \\ Y_{MOB}^{0+} &:= \frac{\mathbb{E}[\min\{Y_N - \underline{Y}X_N, \bar{Y}(1 - X_N)\}]}{1 - \mathbb{E}[X]} \end{aligned}$$

It follows that

$$Y^1 \in [Y_{MOB}^{1-}, Y_{MOB}^{1+}] \tag{1}$$

$$\text{and } Y^0 \in [Y_{MOB}^{0-}, Y_{MOB}^{0+}] \tag{2}$$

The bounds in (1) and (2) are sharp.

¹²In Proposition 2 and throughout, we refer to an interval $[a, b]$ as a sharp bound for Y^1 if, for every $y^1 \in [a, b]$, there exists some joint distribution of (X, Y) that is consistent with the observed marginal distribution of (p_n, X_n, Y_n) as well as any additional imposed assumptions, and that implies $Y^1 = y^1$, and similarly for bounds for Y^0 and D .

The insight underlying Proposition 2 is originally due to Duncan and Davis [1953]; it has been extended and formalized by Horowitz and Manski [1995], Cross and Manski [2002], Cho and Manski [2008]. Those authors showed how one could derive sharp bounds on group means for outcomes with bounded support and for neighborhoods in which the group’s prevalence is known. Proposition 2 provides a novel closed-form expression for the aggregation of those neighborhood-specific bounds to the population-level.

The same approach yields bounds for the difference in group means:

Proposition 3 (Method of Bounds for Difference in Group Means). *Suppose that for each x and n , the conditional expectation of Y with respect to x and n exists, $\underline{Y} \leq \mathbb{E}[Y|X = x, N = n] \leq \bar{Y}$ for some pair $(\underline{Y}, \bar{Y}) \in \mathbb{R}^2$. Define the following parameters:*

$$D_{MOB}^+ = \frac{\mathbb{E}[\min\{Y_N - \underline{Y}(1 - X_N), \bar{Y}X_N\}] - \mathbb{E}[X]\mathbb{E}[Y]}{\text{Var}[X]}$$

$$\text{and } D_{MOB}^- = \frac{\mathbb{E}[Y](1 - \mathbb{E}[X]) - \mathbb{E}[\min\{Y_N - \underline{Y}X_N, \bar{Y}(1 - X_N)\}]}{\text{Var}(X)}$$

It follows that

$$D \in [D_{MOB}^-, D_{MOB}^+] \quad (3)$$

The bounds in (3) are sharp.

As a corollary to Propositions 2 and 3, observe that the estimands associated with the neighborhood model will always be feasible.

Corollary 1. *Absent additional information, the neighborhood model estimates, Y_{NM}^0 , Y_{NM}^1 , and D_{NM} , are feasible given the observed neighborhood-level data.*

In the next section, we consider the asymptotic bias of these estimators and show that they define sharp bounds under various assumptions about the sign of the within- and between-group associations.

3 Identification of the Difference in Group Means

In this section we study how assumptions about the sign of the within- and between-group associations provide identifying power for learning about the

difference in group means. We first establish that the bias of the ecological regression estimate for the difference in group means depends on the within-group association.¹³

Proposition 4 (Bias of Ecological Regression for Difference in Group Means).

$$D_{ER} - D = \frac{\delta_W}{\text{Var}(X_N)}$$

We also have the following direct corollary:

Corollary 2. *The ecological regression estimator, $\hat{D}_{ER}$, is unbiased if and only if $\delta_W = 0$.*

When the within-group association is non-zero, ecological regression is biased because differences in the prevalence of groups across neighborhoods are conflated by other neighborhood-level contextual effects. This phenomenon was recognized as early as Goodman [1953]; it is sometimes referred to as aggregation bias in the EI literature.

Because our focus is on identification of group outcomes at the population level (as opposed to group differences for a particular neighborhood), the condition provided in Corollary 1 differs from the condition typically described in the EI literature under which ecological regression is unbiased, which is that

$$\text{Cov}(Y_N^1, X_N) = \text{Cov}(Y_N^0, X_N) = 0 \quad (4)$$


where $Y_N^x = \mathbb{E}[Y|N, X = x]$ for $x = 0, 1$.¹⁴ Although condition (4) implies $\delta_W = 0$, the converse is not true. This distinction is significant because condition (4) can be empirically tested with the residuals from the ecological regression, as proposed, for instance, by Loewen and Grofman [1989] and Gelman et al. [2001]. However, δ_W may be zero (so that ecological regression identifies our parameter of interest), even when the ecological regression

¹³Throughout, we use “bias” to refer to the asymptotic bias of the specified estimator, i.e., the difference between what the estimator identifies and the target estimand. Propositions 4 and 6 are the EI analogues to Proposition 1.1 and 1.2 in Elzayn et al. [2025], which apply to individual-level data in which group membership is probabilistically observed.

¹⁴A stronger condition that is sometimes discussed, which implies (4), is the constancy model, under which each group’s outcomes are the same in each neighborhood, $Y_n^x = Y^x$ for all n and for $x = 0, 1$.

residuals indicate that the conditional expectation function of Y_N given X_N is highly non-linear (indicating that Condition (4) is violated). Figure 1 provides an illustration.

Figure 1: Ecological Regression with Contextual Effects



Notes: The figure illustrates a setting in which $\text{Cov}(Y_N, X_N|X) \neq 0$ for each X , so that $\mathbb{E}[Y_N|X_N]$ is nonlinear. However, $\mathbb{E}[\text{Cov}(Y_N, X_N|X)] = \delta_W = 0$, so ecological regression is still unbiased. While the ecological regression line (red dashed) does not travel through $Y(0)$ or $Y(1)$ (green), the slope of the ecological regression line is equal to the difference between $Y(1)$ and $Y(0)$.

Our next result characterizes the identifying power of assumptions that sign the within-group (across neighborhood) association.

Proposition 5 (Bounds Based on the Sign of the Within-Group Association).

(i) If $\delta_W \geq 0$, then

$$D \in [D_{MOB}^-, \text{Min} \{D_{MOB}^+, D_{ER}\}]$$

(ii) If $\delta_W \leq 0$, then

$$D \in [\text{Max} \{D_{MOB}^-, D_{ER}\}, D_{MOB}^+]$$

(iii) The bounds in (i) and (ii) are sharp.

We now proceed analogously to derive bounds based on the sign of the between-group (within-neighborhood) association. First, Proposition 6 characterizes the bias of the neighborhood model estimator:

Proposition 6 (Bias of the Neighborhood Model Estimator for the Difference in Group Means).

$$D_{NM} - D = \frac{-\delta_B}{\text{Var}(X)}$$

The following Corollary directly follows:

Corollary 3. *The neighborhood model estimator, $\hat{D}_{NM}$, is unbiased if and only $\delta_B = 0$.*

Freedman et al. [1991] noted that the neighborhood model point-identifies $\mathbb{E}[Y | X, N]$ under a related but stronger condition, that there is no systematic difference in the outcome between groups within any neighborhood, i.e., $\mathbb{E}[Y|0, N] = \mathbb{E}[Y|1, N]$ for all N .

Together with Proposition 1, Corollaries 2 and 3 establish that except in degenerate cases, the assumptions justifying the ecological regression and neighborhood model estimators are mutually exclusive:

Corollary 4. *If $D \neq 0$, then either: (i) the ecological regression estimator, $\hat{D}_{ER}$, is biased; (ii) the neighborhood model estimator, $\hat{D}_{NM}$, is biased; or (iii) both ecological regression and the neighborhood model estimator are biased.*

The following result uses Proposition 6 to derive sharp bounds for D based on the sign of the between-group association.

Proposition 7 (Bounds Based on the Sign of the Between-Group Association).

(i) *If $\delta_B \geq 0$, then*

$$D \in [D_{NM}, D_{MOB}^+]$$

(ii) *If $\delta_B \leq 0$, then*

$$D \in [D_{MOB}^-, D_{NM}]$$

(iii) *The bounds in (i) and (ii) are sharp.*

It is interesting to compare the bounds in Proposition 7 based on δ_B to the bounds in Proposition 5 based on δ_W . Both bounds make use of information about the sign of a conditional association to narrow the method of bounds interval in one direction. However, whereas the bounds based on δ_B are strictly within the interior of the method of bounds interval, the same is not true of δ_W . The explanation for the difference is that the estimates from ecological regression are not guaranteed to be feasible, whereas the estimates from the neighborhood model are necessarily within the method of bounds interval.

The bounds in Proposition 7 are conceptually related to the “bounded variation” approach considered in Manski [2018], which can entail imposing an assumption on the sign of $\mathbb{E}[Y|X = 1, N] - \mathbb{E}[Y|X = 0, N]$ to derive bounds for $\mathbb{E}[Y|X, N]$.¹⁵ Because our focus is on $\mathbb{E}[Y|X]$ rather than $\mathbb{E}[Y|X, N]$, Proposition 7 requires only that this sign restriction hold *on average* across values of N . See Section 5 for additional discussion.

Whereas the bounds in Propositions 5 and 7 are based on assumptions about the respective signs of δ_W and δ_B , in some settings it may be more credible to assume that δ_W and δ_B share the same sign as one another, without taking a stance on what that shared sign is. For example, the researcher may have theoretical or institutional reasons to expect that contextual effects reinforce the individual-level (within-neighborhood) difference, such as through mutually reinforcing social norms. We refer to this condition as contextual reinforcement, and formally define it as follows:

Definition 1. *Contextual Reinforcement* is satisfied if $\delta_W \cdot \delta_B \geq 0$.

Proposition 8. (*Identification of Group Differences with Contextual Reinforcement*)

If $\delta_W \cdot \delta_B \geq 0$, then:

(i) If $D_{ER} \geq D_{NM}$, then

$$0 \leq D_{NM} \leq D \leq \text{Min}\{D_{MOB}^+, D_{ER}\}$$

(ii) If $D_{ER} \leq D_{NM}$, then

$$\text{Max}\{D_{MOB}^-, D_{ER}\} \leq D \leq D_{NM} \leq 0$$

(iii) The bounds in (i) and (ii) are sharp.

¹⁵More generally, Manski [2018] considers identification of $p(Y, X, N)$ given assumptions that constrain the possible values of $p(Y|X = 1, N) - p(Y|X = 0, N)$. When Y is binary, this term is equal to $\text{Cov}(Y, X|N)$, the analogue of δ_B for a specific N .

The following Corollary follows directly:

Corollary 5. *If contextual reinforcement holds, then*

$$\text{Sign}(D) = \text{Sign}(D_{ER} - D_{NM})$$

The following Theorem summarizes our results so far and provides corresponding results for cases in which both conditional associations have known signs but contextual reinforcement does not hold.

Theorem 1 (Identification of D with Monotone EI). *Table 1 provides sharp bounds for D under the specified information. If it is known that contextual reinforcement holds, the bounds on D in one of the shaded table cells will apply.*

Proof. The results for each cell in the table have been proven by the previous propositions except for those in which δ_B and δ_W have opposite signs. For the cell in which $\delta_B \geq 0$ and $\delta_W \leq 0$, the claim is that:

$$D \in [\max\{D_{NM}, D_{ER}, D_{MOB}^-\}, D_{MOB}^+].$$

Proposition 4 and Proposition 6 together show that $D \geq D_{ER}$ and $D \geq D_{NM}$; Proposition 3 gives that $D \geq D_{MOB}^-$ – which, combined with the previous inequalities, gives the lower bound – and that $D \leq D_{MOB}^+$; the bound is thus proved. To see that the bound is sharp, note that first of all it is a subset of $[D_{MOB}^-, D_{MOB}^+]$, so by Lemma 2 every $\tilde{D}$ within the range is feasible (with respect to the constraints of $[\underline{Y}, \bar{Y}]$ and Y_n) with some distribution $\tilde{Y}$. But note that $\delta_B \geq 0 \iff D \geq D_{NM}$ and $\delta_W \geq 0 \iff D \geq D_{ER}$; both these must hold for all $\tilde{D}$ in the interval, so all such $\tilde{D}$ are feasible. The bound is thus sharp. A similar approach gives the sharp bound for the $\delta_B \leq 0$ and $\delta_W \geq 0$ cell. $\square$

4 Identification of Group Means

This section provides results relating to the use of monotone ecological inference for identifying the group means, Y^0 and Y^1 . The ecological regression estimates for the *levels* of the group means are biased under the same condition, and in a related manner, as the ecological regression estimate for the *difference* in group means described in Proposition 4. A similar parallel exists

for the neighborhood model estimates of the group means compared to the neighborhood level estimate for the difference in group means. The nature of these relationships are formalized in the following two propositions.

Proposition 9 (Bias of Ecological Regression for Group Means).

$$Y_{ER}^1 - Y^1 = \frac{\delta_W (1 - \mathbb{E}[X])}{\text{Var}(X_N)}$$

$$Y_{ER}^0 - Y^0 = -\frac{\delta_W \mathbb{E}[X]}{\text{Var}(X_N)}$$

Proposition 10 (Bias of Neighborhood Model Estimator for Group Means).

$$Y_{NM}^1 - Y^1 = -\frac{\delta_B}{\mathbb{E}[X]}$$

$$Y_{NM}^0 - Y^0 = \frac{\delta_B}{1 - \mathbb{E}[X]}$$

Applying these results in the manner employed in the prior subsection yields the following identification results on Y^1 and Y^0 .

Theorem 2 (Identification of Y^1 with Monotone EI). *Table 2 provides sharp bounds for Y^1 under the specified information. If it is known that contextual reinforcement holds, the bounds on Y^1 in one of the shaded table cells will apply.*

Theorem 3 (Identification of Y^0 with Monotone EI). *Table 3 provides sharp bounds for Y^0 under the specified information. If it is known that contextual reinforcement holds, the bounds on Y^0 in one of the shaded table cells will apply.*

Finally, we have so far assumed that group membership is binary. Suppose instead that there are $G + 1$ discrete groups, indexed by $g \in \{0, 1, \dots, G\}$. To identify $\mathbb{E}[Y|G = g]$, we can define $X^g \in \{0, 1\}$ to indicate whether or not an individual belongs to group g ,

$$X^g = 1 \iff G = g$$

and then apply the bounds in Table 2 with respect to X^g instead of X . Note that unlike the binary case, the relevant conditional associations will differ

for each group. For example, the within-group association that determines the bias of the ecological regression estimator for $\mathbb{E}[Y|G = g]$ is given by

$$\delta_W^g = \mathbb{E}[\text{Cov}(Y, X_N^g | X^g)],$$

where $X_N^g = \mathbb{E}[X^g | N]$. Similarly, the between-group association that determines the bias of the neighborhood model estimator for $\mathbb{E}[Y|G = g]$ is given by

$$\delta_B^g = \mathbb{E}[\text{Cov}(Y, X^g | N)].$$

Table 1: Identification of the Difference in Group Means

		$\delta_B = ?$	$\delta_B \geq 0$	$\delta_B \leq 0$	$\delta_B = 0$
Between					
Within					
$\delta_W = ?$		$D \in [D_{MOB}^-, D_{MOB}^+]$	$D \in [D_{NM}^-, D_{MOB}^+]$	$D \in [D_{MOB}^-, D_{NM}]$	$D = D_{NM}$
$\delta_W \geq 0$		$D \in [D_{MOB}^-, \text{Min}\{D_{MOB}^+, D_{ER}\}]$	$0 \leq D_{NM} \leq D \leq \text{Min}\{D_{MOB}^+, D_{ER}\}$	$D \in [D_{MOB}^-, \text{min}\{D_{NM}, D_{ER}, D_{MOB}^+\}]$	$D = D_{NM}$
$\delta_W \leq 0$		$D \in [\text{Max}\{D_{MOB}^-, D_{ER}\}, D_{MOB}^+]$	$D \in [\text{Max}\{D_{MOB}^-, D_{ER}, D_{NM}\}, D_{MOB}^+]$	$\text{Max}\{D_{MOB}^-, D_{ER}\} \leq D \leq D_{NM} \leq 0$	$D = D_{NM}$
$\delta_W = 0$		$D = D_{ER}$	$D = D_{ER}$	$D = D_{ER}$	$D = 0$

Table 2: Identification of the Group 1 Mean

		$\delta_B = ?$	$\delta_B \geq 0$	$\delta_B \leq 0$	$\delta_B = 0$
Between					
Within					
$\delta_W = ?$		$Y^1 \in [Y_{MOB}^{1-}, Y_{MOB}^{1+}]$	$Y^1 \in [Y_{NM}^{1-}, Y_{MOB}^{1+}]$	$Y^1 \in [Y_{MOB}^{1-}, Y_{MOB}^{1+}]$	$Y^1 = Y^{NM}$
$\delta_W \geq 0$		$Y^1 \in [Y_{MOB}^{1-}, \text{Min}\{Y_{ER}^{1+}, Y_{MOB}^{1+}\}]$	$Y_{NM}^{1-} \leq Y^1 \leq \text{Min}\{Y_{ER}^{1+}, Y_{MOB}^{1+}\}$	$Y^1 \in [Y_{MOB}^{1-}, \text{Min}\{Y_{NM}^{1-}, Y_{ER}^{1+}, Y_{MOB}^{1+}\}]$	$Y^1 = Y^{NM}$
$\delta_W \leq 0$		$Y^1 \in [\text{Max}\{Y_{MOB}^{1-}, Y_{ER}^{1-}\}, Y_{MOB}^{1+}]$	$Y^1 \in [\text{Max}\{Y_{ER}^{1-}, Y_{NM}^{1-}, Y_{MOB}^{1-}\}, Y_{MOB}^{1+}]$	$\text{Max}\{Y_{ER}^{1-}, Y_{MOB}^{1-}\} \leq Y^1 \leq Y_{NM}^{1+}$	$Y^1 = Y^{NM}$
$\delta_W = 0$		$Y^1 = Y^{ER}$	$Y^1 = Y^{ER}$	$Y^1 = Y^{ER}$	$Y^1 = Y^0$

Table 3: Identification of the Group 0 Mean

		$\delta_B = ?$	$\delta_B \geq 0$	$\delta_B \leq 0$	$\delta_B = 0$
Between		$Y^0 \in [Y_{MOB}^{0-}, Y_{MOB}^{0+}]$ $Y^0 \in [\text{Max}\{Y_{ER}^{0-}, Y_{MOB}^{0-}\}, Y_{MOB}^{0+}]$ $Y^0 \in [Y_{MOB}^{0-}, \text{Min}\{Y_{ER}^{0+}, Y_{MOB}^{0+}\}]$ $Y^0 = Y_{ER}^{0-}$	$Y^0 \in [Y_{MOB}^{0-}, Y_{NM}^{0+}]$ $\text{Max}\{Y_{ER}^{0-}, Y_{MOB}^{0-}\} \leq Y^0 \leq Y_{NM}^{0+}$ $Y^0 \in [Y_{MOB}^{0-}, \text{Min}\{Y_{NM}^{0+}, Y_{ER}^{0+}, Y_{MOB}^{0+}\}]$ $Y^0 = Y_{ER}^{0-}$	$Y^0 \in [Y_{NM}^{0-}, Y_{MOB}^{0+}]$ $Y^0 \in [\text{max}\{Y_{MOB}^{0-}, Y_{ER}^{0-}, Y_{NM}^{0-}\}, Y_{MOB}^{0+}]$ $Y_{NM}^{0-} \leq Y^0 \leq \text{Min}\{Y_{ER}^{0+}, Y_{MOB}^{0+}\}$ $Y^0 = Y_{ER}^{0-}$	$Y^0 = Y_{NM}^{0-}$ $Y^0 = Y_{NM}^{0-}$ $Y^0 = Y_{NM}^{0-}$ $Y^0 = Y^1$
Within					
$\delta_W = ?$					
$\delta_W \geq 0$					
$\delta_W \leq 0$					
$\delta_W = 0$					

5 Identification of Neighborhood-Specific Group Means

Our focus so far has been on identification of the overall group means, $Y^x = \mathbb{E}[Y|X = x]$, as well as the difference in overall group means, $D = Y^1 - Y^0$. In some settings, a researcher may seek to learn a group mean within a specific neighborhood, $Y_n^x = \mathbb{E}[Y | N = n, X = x]$, or the difference in group means within that neighborhood, $D_n = Y_n^1 - Y_n^0$.¹⁶ In this section, we consider the identifying power of monotone ecological inference for studying neighborhood-specific group means.

As a starting point, note that the method of bounds provides sharp bounds on Y_n^x and D_n for each n and x under the same assumptions imposed by Propositions 2 and 3 [Duncan and Davis, 1953, Cho and Manski, 2008]. We will refer to the endpoints of these intervals for a specific neighborhood n using the following notation:

$$Y_n^x \in [Y_{MOB,n}^{x-}, Y_{MOB,n}^{x+}]$$

and

$$D_n \in [D_{MOB,n}^-, D_{MOB,n}^+]$$

where $Y_{MOB,n}^{1+} = \min\left\{\frac{Y_n}{X_n}, \bar{Y}\right\}$, $Y_{MOB,n}^{0-} = \frac{Y_n - \min\{Y_n, X_n \bar{Y}\}}{1 - X_n}$, $Y_{MOB,n}^{0+} = \min\left\{\frac{Y_n}{1 - X_n}, \bar{Y}\right\}$, $Y_{MOB,n}^{1-} = \frac{Y_n - \min\{Y_n, \bar{Y}(1 - X_n)\}}{X_n}$, $D_{MOB,n}^+ = Y_{MOB,n}^{1+} - Y_{MOB,n}^{0-}$, and $D_{MOB,n}^- = Y_{MOB,n}^{1-} - Y_{MOB,n}^{0+}$.

Monotone ecological inference entails sharpening these bounds at the neighborhood-level by imposing sign restrictions on various aspects of the unobserved individual-level relationship between Y , X , and N . Consider first the neighborhood-level analog to the between-group association,

$$\delta_{B,n} := \text{Cov}(Y, X | N = n).$$

Whereas δ_B describes the *average* between-group variation within neighborhoods, $\delta_{B,n}$ depends on the between-group variation for a specific neighborhood.

¹⁶Using the terminology of Cross and Manski [2002], the overall group mean corresponds to the “short regression” of Y on X whereas the neighborhood-specific group mean corresponds to the “long regression” of Y on X and N .

Manski [2018] previously considered the identifying power of the sign of $\delta_{B,n}$ for studying Y_n^x . Because we will draw on his result below, we reproduce it using our notation in the following proposition:

Proposition 11 (Bounds Based on the Sign of the Neighborhood-Specific Between-Group Association [Manski, 2018]).

(i) If $\delta_{B,n} \geq 0$ for some $n \in \mathcal{N}$ then:

$$\begin{aligned} D_n &\in [0, D_{MOB,n}^+], \\ Y_n^1 &\in [Y_n, Y_{MOB,n}^{1+}], \\ \text{and } Y_n^0 &\in [Y_{MOB,n}^{0-}, Y_n]. \end{aligned}$$

(ii) If $\delta_{B,n} \leq 0$ for all $n \in \mathcal{N}$ then:

$$\begin{aligned} D_n &\in [D_{MOB,n}^-, 0], \\ Y_n^1 &\in [Y_{MOB,n}^{1-}, Y_n], \\ \text{and } Y_n^0 &\in [Y_n, Y_{MOB,n}^{0+}]. \end{aligned}$$

(iii) The bounds in (i) and (ii) are sharp.

We next consider the identifying power of assumptions on the neighborhood-specific analog to the within-neighborhood association, which has not to our knowledge been considered in prior research. Let $\mu_x(x_n)$ denote the conditional expectation of Y for a member of group x in a neighborhood with group-prevalence x_n

$$\mu_x(x_n) = \mathbb{E}[Y \mid X = x, X_n = x_n]$$

We restrict our focus to neighborhoods with group-prevalence values for which the first derivative of $\mu_x(\cdot)$ exists for $x = 0$ and $x = 1$, and we denote those derivatives by $\mu'_x(\cdot)$

$$\mu'_x(x_n) = \lim_{u \rightarrow 0} \frac{1}{u} (\mathbb{E}[Y \mid X = x, X_n = x_n + u] - \mathbb{E}[Y \mid X = x, X_n = x_n])$$

For neighborhood n , define the local within-neighborhood association $\delta_{W,n}$ as

$$\delta_{W,n} = X_n \mu'_1(X_n) + (1 - X_n) \mu'_0(X_n)$$

where $X_n = \mathbb{E}[X \mid N = n]$. Conceptually, $\delta_{W,n}$ captures the within-group association between the mean of the outcome, Y , and the group-prevalence

of the neighborhood, X_N , among neighborhoods with similar levels of group-prevalence. It differs from δ_W in that δ_W depends on a summary measure of the relationship between Y and X_N across all neighborhoods, whereas $\delta_{W,n}$ reflects the “local” relationship between Y and X_N among a set of neighborhoods with similar levels of group prevalence.

Although μ_0 and μ_1 are unobserved, we do observe the overall regression function, $\mu(x_n) := \mathbb{E}[Y \mid X_n = x_n]$, which is the mixture of the group-specific regression functions

$$\mu(x_n) = x_n \mu_1(x_n) + (1 - x_n) \mu_0(x_n). \quad (5)$$

We also observe the derivative of the overall regression function, $\mu'(x_n)$, which, by construction, is guaranteed to exist.

Differentiating (5) yields

$$\begin{aligned} \mu'(x_n) &= \mu_1(x_n) - \mu_0(x_n) + x_n \mu'_1(x_n) + (1 - x_n) \mu'_0(x_n) \\ &= D_n + \delta_{W,n} \end{aligned}$$

Thus, given knowledge of $\delta_{W,n}$, the derivative of the conditional expectation function, $\mu'(x_n)$, provides information about the group means in neighborhoods with group-prevalence x_n .

In addition to assuming knowledge of the sign of $\delta_{W,n}$, it will be convenient to assume that each neighborhood with the same prevalence X_n has the same average outcome - i.e. $\mathbb{E}[Y \mid X = x, N = n] = \mathbb{E}[Y \mid X = x, X_N = X_n]$.

Assumption 1. *For every neighborhood n ,*

$$\mathbb{E}[Y \mid X = x, N = n] = \mathbb{E}[Y \mid X = x, X_N = X_n]$$

In words, Assumption 1 requires that any two neighborhoods that have the same group-prevalence will also share the same mean outcomes by group (at least in expectation). Under this assumption, knowledge of the sign of $\delta_{W,n}$ facilitates identification as follows:

Proposition 12.

(i) *Suppose that $\delta_{W,n} \geq 0$. Then:*

$$\begin{aligned} D_n &\in [D_{MOB,n}^-, \min \{ \mu'(X_n), D_{MOB,n}^+ \}], \\ Y_n^1 &\in [Y_{MOB,n}^{1-}, \min \{ Y_{MOB,n}^{1+}, Y_n + (1 - X_n) \cdot \mu'(X_n) \}], \\ \text{and } Y_n^0 &\in [\max \{ Y_{MOB,n}^{0-}, Y_n - X_n \cdot \mu'(X_n) \} , Y_{MOB,n}^{0+}]. \end{aligned}$$

(ii) Suppose that $\delta_{W,n} \leq 0$. Then:

$$\begin{aligned} D_n &\in [\max \{ \mu'(X_n), D_{MOB,n}^- \}, D_{MOB,n}^+], \\ Y_n^1 &\in [\max \{ Y_{MOB,n}^{1-}, Y_n + (1 - X_n) \cdot \mu'(X_n) \}, Y_{MOB,n}^{1+}], \\ \text{and } Y_n^0 &\in [Y_{MOB,n}^{0-}, \min \{ Y_{MOB,n}^{0+}, Y_n - X_n \cdot \mu'(X_n) \}]. \end{aligned}$$

(iii) The bounds in (i) and (ii) are sharp in the absence of further information.

When Assumption 1 does not hold, the bounds in Proposition 12 identify the group means or differences for the set of neighborhoods with the same group prevalence as the target neighborhood n . We formalize this claim and provide a proof in Proposition 16 of the Online Appendix.

Finally, consider the possibility that the researcher has information not about the individual signs of $\delta_{W,n}$ and $\delta_{B,n}$ but rather that these two (local) associations share the same sign as one another. This local analog to contextual reinforcement can substantially aid in identification:

Proposition 13. Suppose that $\delta_{B,n} \cdot \delta_{W,n} \geq 0$. Then either:

(i) $\mu'(X_n) \geq 0$, and

$$\begin{aligned} D_n &\in [0, \min \{ \mu'(X_n), D_{MOB,n}^+ \}], \\ Y_n^1 &\in [Y_n, \min \{ Y_n + (1 - X_n) \mu'(X_n), Y_{MOB,n}^{1+} \}], \\ \text{and } Y_n^0 &\in [\max \{ Y_{MOB,n}^{0-}, Y_n - X_n \mu'(X_n) \}, Y_n]; \end{aligned}$$

or

(ii) $\mu'(X_n) \leq 0$, and

$$\begin{aligned} D_n &\in [\max \{ \mu'(X_n), D_{MOB,n}^- \}, 0], \\ Y_n^1 &\in [\max \{ Y_{MOB,n}^{1-}, Y_n + (1 - X_n) \mu'(X_n) \}, Y_n], \\ \text{and } Y_n^0 &\in [Y_n, \min \{ Y_n - X_n \mu'(X_n), Y_{MOB,n}^{0+} \}]. \end{aligned}$$

(iii) The bounds in (i) and (ii) are sharp.

6 Empirical Application

To illustrate the benefits of monotone ecological inference in a concrete setting, we investigate the question of partisan polarization in COVID-19 vaccine uptake. This question has been of acute interest to policymakers, academics, and the media, both for COVID-19 response specifically and broader questions of trust in public health authorities [Milligan, 2021, Collins, 2024].¹⁷ Unfortunately, most research in the U.S. context has had to rely on ecological inference, as joint individual data on vaccination status and partisan membership is not available for a large nationally representative set of individuals [e.g., Albrecht, 2022, Ye, 2023].¹⁸

Our primary data set consists of county-level data from 3,115 counties on COVID-19 vaccination uptake and partisanship. We measure vaccination uptake as the share of county residents who had received one or more COVID-19 vaccination as of December 31, 2021, obtained from the Centers for Disease Control and Prevention [2024]. We measure partisanship as the fraction of voters in the county who cast their ballot for the Republican candidate in the 2020 presidential election, obtained from the MIT Election Data and Science Lab [2020].¹⁹

Using our previous notation, Y indicates whether an individual is vaccinated, X indicates whether an individual is Republican, and N indicates the county in which the individual resides. The share of vaccinated individuals in a county is between $\underline{Y} = 0$ and $\bar{Y} = 1$. Our goal is to use the county-level data to estimate the partisan vaccination gap, which we define as the difference in the mean vaccination rate of Republicans relative to Democrats and third party voters, $D = \mathbb{E}[Y|X = 1] - \mathbb{E}[Y|X = 0]$.

Figure 2 plots the binned county-level data. The figure shows a clear downward trend: counties with more Republicans tend to have lower vacci-


¹⁷Jones and McDermott [2022] provide a helpful conceptual discussion and review of the literature.

¹⁸An alternative methodological approach is to collect individual-level survey data on vaccine uptake and political affiliation; such research designs avoid the need for ecological inference but often face limitations based on sample size, representativeness, and response bias.

¹⁹We focus on this measure of partisan ideology rather than Republican party voter registration to be consistent with the prior EI research on the topic. A second reason is that a meaningful share of voters are registered as independent, and the ideologies associated with voters in this group is likely to vary widely across different counties around the country.

nation rates. The pattern is consistent with the possibility that Republicans are vaccinated at lower rates than Democrats. However, this interpretation is potentially subject to the ecological fallacy; counties with more Republicans may have lower vaccination rates for reasons unrelated to partisan composition [Ye, 2023]. Indeed, the method of bounds interval for the partisan vaccination gap ranges from -77.4 to 52.8 percentage points. The width of the MOB interval implies that the county-level data does not provide much information about the magnitude or even direction of differences in vaccination rates between Republicans and Democrats, at least without further assumptions. Point identification approaches also diverge sharply: the neighborhood model and ecological regression imply respective partisan vaccination gaps of -5.5 and -47.9 percentage points.

Figure 2: County Vaccination Rate by County Republican Vote Share



Notes: The figure reports county-level COVID-19 vaccination rates by the share of voters in the county who voted for the Republican candidate in the 2020 presidential election. Counties are grouped into 100 equal-population bins. The neighborhood model estimates for the mean vaccination rates among Republicans and non-Republicans are respectively denoted by the lower and upper red dotted lines. The ecological regression line is in black.

Sharpening identification through monotone ecological inference involves

making assumptions about two conditional associations. First, the between-group association, $\mathbb{E}[\text{Cov}(Y, X|N)]$, refers to differences in vaccination uptake between Republicans and Democrats living in the same county. Prior research provides some basis for expecting the between-group association to be negative; there are well-documented partisan differences in information sources that are not fully mediated through neighborhood [Iyengar and Hahn, 2009, Peterson et al., 2021], and the prominent Republican politicians featured on more conservative media outlets were more likely to espouse anti-vaccination beliefs and/or downplay the health risks associated with the COVID-19 virus [Hornsey et al., 2020, Gollwitzer et al., 2020, Albrecht, 2022]. Second, the within-group association, $\mathbb{E}[\text{Cov}(Y, X_N|X)]$, refers to differences in vaccination uptake among individuals of the same political party who live in neighborhoods with differing concentrations of Republican voters.²⁰ Like the between-group association, there is some reason to believe that the within-group association is negative. For example, Republican counties tend to be lower income and more rural (Figure S2), which are factors associated with lower access to public health services like vaccinations [Sun and Monnat, 2022, Hernandez et al., 2022, Parolin and Lee, 2022].

More generally, there are reasons to expect that the two conditional associations share the same sign as one another, whether that sign is positive or negative. One mechanism through which such contextual reinforcement may operate is network effects, such as social norms or peer effects. In particular, people’s health behaviors are known to be influenced by the people around them [Sato and Takasaki, 2019, Klaesson et al., 2023], and in Republican counties, a larger share of the people with whom one interacts are likely to be Republican. Thus if the Republicans in a neighborhood tend to be more skeptical of COVID-19 vaccinations (i.e., the between-group association is negative), that is likely to reduce the vaccination rate among both Democrats and Republicans living in that neighborhood. In the words of one author, “In many communities, wearing a mask or getting a [COVID-19] vaccine became a political statement, with many Republicans arguing that these actions violated their individual freedoms and were unnecessary anyway” [Albrecht, 2022]. Along similar lines, for many people, vaccine uptake may depend in part on local policies, such as whether vaccines are man-

²⁰That is, the within-group association refers to vaccination uptake gaps between Republicans living in counties with more Republicans compared to Republicans living in counties with fewer Republicans, and between Democrats living in counties with more Republicans compared to Democrats living in counties with fewer Republicans.

dated for public sector employees [Howard-Williams et al., 2022]. Thus, if Republicans exhibit more vaccine hesitancy, we would expect that counties in which more Republicans live would be more likely to elect leaders that do not adopt pro-vaccine policies, leading to lower vaccine rates for county residents, whether Democrat or Republican.²¹

The foregoing discussion provides a theoretical basis for the contextual reinforcement assumption in this setting. We empirically validate the assumption by drawing on an auxiliary dataset that contains individual-level data on vaccination status, political party registration, and neighborhood. We construct this dataset by matching a national dataset of voter registration records [L2, 2024], which contain individual-level data on political party, to a large dataset of electronic health care records [Balraj et al., 2023], which contain individual-level data on COVID-19 vaccination status. Appendix B provides a further description of the underlying data sources and of our matching procedure. The final matched dataset contains approximately 1.3 million registered Republicans and Democrats in 2,576 counties and 49 states, plus the District of Columbia.²²


Figure 3 uses the auxiliary data to plot vaccination rates by (binned) county-level Republican vote share, separately for Republicans and Democrats. The figure provides visual support for contextual reinforcement: the Republican bins tend to lie below the Democratic bins with similar partisan makeup (so that the between-group association is negative) and both the Republican bins and Democratic bins exhibit a downward sloping trend (so that the within-group association is also negative). Formal statistical tests regarding the sign of these quantities yield the same conclusion (see Table S2). Based on these results, we adopt the contextual reinforcement assumption to interpret the county-level analyses.²³

²¹Analogously, Patterson [2022] finds that states with Republican governors were less likely to quickly adopt stay-at-home orders during the pandemic, and that individuals in those states exhibited less social distancing as a result.

²²While this is a large linked dataset, most of the literature has not been able to secure such individual-level data due to data restrictions, hence relying primarily on aggregate data [see Albrecht, 2022, Ye, 2023].

²³A limitation of this analysis for assessing contextual reinforcement is that the set of individual included in the auxiliary dataset may not be representative of the overall population due to the nature of selection into our electronic health records data or non-random match rates with the voter registration data. We obtain similar results when we replicate the analysis using alternative matching criteria to construct the auxiliary sample (Table S3) and when we re-weight the auxiliary data based on individuals' observable

Figure 3: Vaccination Rate by Vote Share and Political Party Membership



Notes: The figure uses the matched auxiliary data set to report Covid-19 vaccination rates by county-level Republican vote share for individuals registered as Republicans (red) and individuals registered as Democrats (blue). Individuals are grouped into ten equal-sized bins based on the Republican vote share for the county in which they live. The average vertical difference between the blue and red points reflects the estimated sign of δ_B . The average slope of the linear best fit lines reflects the estimated sign of δ_W .

When contextual reinforcement holds, Corollary 5 establishes that we can identify the sign of the difference in group means based on the sign of the difference between the ecological regression and neighborhood model estimators. As shown in Figure 4, we find that this difference is positive ($p < .01$), implying that Democrats are vaccinated at higher rates than

characteristics to more closely match the national population (Table S4). A different limitation is the potential for discordance between the partisanship measure contained in the auxiliary data (party registration) and the one employed in our main county-level analysis (presidential vote share).


Republicans. In turn, the contextual reinforcement bounds from Proposition 8 imply that the partisan vaccination gap is between -47.9 and -5.5 percentage points (95% CI: -51.8 to -5.1).

Figure 4 summarizes our results under various identifying assumptions.²⁴ If δ_W or δ_B was assumed to be zero, the partisan vaccination gap would be point-identified at -47.9 or -5.5 percentage points, respectively. Conversely, without any assumptions beyond the observed data, the method of bounds interval for the partisan vaccination gap ranges from -77.4 to 52.8 percentage points. Turning to monotone EI, as discussed above there are plausible reasons to expect $\delta_B \leq 0$ as well as $\delta_W \leq 0$. Imposing these sign restrictions individually respectively tightens the method of bounds interval by 45% and 22%. Finally, under contextual reinforcement, our preferred identifying assumption, we can conclude that the partisan vaccination gap is between -47.9 and -5.5 percentage points, an interval that is 67% smaller than the one obtained from the method of bounds.

We also demonstrate how our method may be used to more precisely bound county-specific vaccine disparities. For purposes of this exercise, we assume that the sign assumptions discussed above hold not only for δ_B and δ_W but locally for each county as well. To calculate the implied bounds, we estimate $\mu'(X_n)$ using population-weighted local linear approximation with an Epanechnikov kernel, with bandwidth chosen through 10-fold cross validation (Figure S3). Based on this estimated derivative, we use Proposition 13(ii) to calculate county-level bounds on the Democrat and Republican vaccination rates. For Contra Costa County, which leans left-of-center politically, we estimate that the Republican vaccination rate is between 0.50 and 0.77 and that the Democrat vaccination rate is between 0.77 and 0.87. For Galveston county, which leans right-of-center politically, the respective bounds for Republicans and Democrats are from 0.36 to 0.57 and from 0.57 to 0.89. Each of these intervals is substantially narrower than the corresponding interval derived from the method of bounds (see Figure 5).

²⁴Figure S1 provides corresponding results for mean vaccination rates among Democrats and Republicans.

Figure 4: Identification of Partisan Vaccination Gap Using Monotone EI




Notes: The partisan vaccination gap is defined as the proportion of vaccinated Democrats subtracted from the proportion of vaccinated Republicans. A negative gap indicates that a higher proportion of Democrats than Republicans are vaccinated. Red bars are 95% confidence intervals following [Imbens and Manski, 2004]; the confidence intervals are based on standard errors from a county-level bootstrap with 1000 bootstrap replicates.

7 Conclusion

We study the classic statistical challenge of ecological inference (EI). Our results clarify the nature of the biases associated with canonical EI methods for point-identifying group outcomes and differences at the population-level. We use those results to derive novel partial identification results based on assumptions about the sign of the conditional associations between the outcome of interest and group membership or neighborhood. Although our approach requires additional structure relative to assumption-free tools like the method of bounds, the payoff to that additional structure can be substantially tighter bounds for the parameter of interest. In our empirical application, we illustrate how one can reason about the sign of the relevant conditional associations based on institutional knowledge and/or auxiliary individual-level data. Under plausible assumptions, the county-level data we rely on allows us to conclude that the Covid-19 vaccination rate among Re-

Figure 5: County-Specific Bounds



Notes: This figure shows the Method of Bounds for the partisan vaccination gap for each county (red). The county-level bounds, based on the assumptions in Proposition 13, are shown in gray, and the estimated derivative of the conditional expectation function of the vaccination rate by county partisan makeup is shown in black. The implied bounds for Contra Costa County, California, and Galveston County, Texas, are highlighted in green.

publicans is between 5.5 and 47.9 (95% CI: 5.1 and 51.8) percentage points lower than the corresponding rate among Democrats.

References

- Leo A Goodman. Ecological regressions and behavior of individuals. *American sociological review*, 1953.
- David A Freedman, Stephen P Klein, Jerome Sacks, Charles A Smyth, and Charles G Everett. Ecological regression and voting rights. *Evaluation review*, 15(6):673–711, 1991.

- Gary King. *A solution to the ecological inference problem: Reconstructing individual behavior from aggregate data*. Princeton University Press, 1997.
- Otis Dudley Duncan and Beverly Jeanne Davis. An alternative to ecological correlation. *American Sociological Review*, 18:665, 1953.
- Charles F. Manski and John V. Pepper. Monotone instrumental variables: With an application to the returns to schools. *Econometrica*, 68(4):997–1010, 2000.
- Jeffrey B Lewis. Understanding king’s ecological inference model a method-of-moments approach. *Historical Methods: A Journal of Quantitative and Interdisciplinary History*, 34(4):170–188, 2001.
- Wenxin Jiang, Gary King, Allen Schmalz, and Martin A Tanner. Ecological regression with partial identification. *Political Analysis*, 28(1):65–86, 2020.
- Don Albrecht. Vaccination, politics and covid-19 impacts. *BMC Public Health*, 22(1):96, 2022.
- Xinyuan Ye. Exploring the relationship between political partisanship and covid-19 vaccination rate. *Journal of Public Health*, 45(1):91–98, 2023.
- Gary King, Martin A Tanner, and Ori Rosen. *Ecological inference: New methodological strategies*. Cambridge University Press, 2004.
- Wendy Tam Cho and Charles F Manski. 24 cross-level/ecological inference. 2008.
- Joel L Horowitz and Charles F Manski. Identification and robustness with contaminated and corrupted data. *Econometrica: Journal of the Econometric Society*, pages 281–302, 1995.
- Philip J Cross and Charles F Manski. Regressions, short and long. *Econometrica*, 70(1):357–368, 2002.
- James Greiner and Kevin M Quinn. $R \times c$ ecological inference: bounds, correlations, flexibility and transparency of assumptions. *Journal of the Royal Statistical Society Series A: Statistics in Society*, 172(1):67–81, 2009.
- Charles F Manski. Credible ecological inference for medical decisions with personalized risk assessment. *Quantitative Economics*, 9(2):541–569, 2018.

- Sheyu Li, Valentyn Litvin, and Charles F Manski. Partial identification of personalized treatment response with trial-reported analyses of binary subgroups. *Epidemiology*, 34(3):319–324, 2023.
- Jiahao Chen, Nathan Kallus, Xiaojie Mao, Geoffrey Svacha, and Madeleine Udell. Fairness under unawareness: Assessing disparity when protected class is unobserved. In *Proceedings of the conference on fairness, accountability, and transparency*, pages 339–348, 2019.
- Nathan Kallus, Xiaojie Mao, and Angela Zhou. Assessing algorithmic fairness with unobserved protected class using data combination. *Management Science*, 68(3):1959–1981, 2022.
- Cory McCartan, Robin Fisher, Jacob Goldin, Daniel E Ho, and Kosuke Imai. Estimating racial disparities when race is not observed. Technical report, National Bureau of Economic Research, 2024.
- Hadi Elzayn, Evelyn Smith, Thomas Hertz, Cameron Guage, Arun Ramesh, Robin Fisher, Daniel E. Ho, and Jacob Goldin. Measuring and mitigating racial disparities in tax audits. *The Quarterly Journal of Economics*, 2025.
- Francesca Molinari. Missing treatments. *Journal of Business & Economic Statistics*, 28(1):82–95, 2010.
- Charles F Manski and John V Pepper. How do right-to-carry laws affect crime rates? coping with ambiguity using bounded-variation assumptions. *Review of Economics and Statistics*, 100(2):232–244, 2018.
- Ashesh Rambachan and Jonathan Roth. A more credible approach to parallel trends. *Review of Economic Studies*, 90(5):2555–2591, 2023.
- WS Robinson. Ecological correlations and the behavior of individuals. *American Sociological Review*, 15(3), 1950.
- Stephen Ansolabehere and Douglas Rivers. Bias in ecological regression. Technical report, Stanford University and Massachusetts Institute of Technology, 1995. URL <https://web.stanford.edu/class/polisci353/2004spring/reading/erbias.pdf>.
- James W Loewen and Bernard Grofman. Recent developments in methods used in vote dilution litigation. *Urb. Law.*, 21:589, 1989.

- Andrew Gelman, David K Park, Stephen Ansolabehere, Phillip N Price, and Lorraine C Minnite. Models, assumptions and model checking in ecological regressions. *Journal of the Royal Statistical Society Series A: Statistics in Society*, 164(1):101–118, 2001.
- Susan Milligan. A deadly political divide. *US News and World Reports*, 2021.
- Francis Collins. Why didn’t facing a common enemy bring us together? *The Atlantic*, 2024.
- David R Jones and Monika L McDermott. Partisanship and the politics of covid vaccine hesitancy. *Polity*, 54(3):408–434, 2022.
- Centers for Disease Control and Prevention. COVID-19 Vaccinations in the United States, County, 2024.
- MIT Election Data and Science Lab. County Presidential Election Returns 2000-2020, 2020.
- Shanto Iyengar and Kyu S Hahn. Red media, blue media: Evidence of ideological selectivity in media use. *Journal of communication*, 59(1):19–39, 2009.
- Erik Peterson, Sharad Goel, and Shanto Iyengar. Partisan selective exposure in online news consumption: Evidence from the 2016 presidential campaign. *Political science research and methods*, 9(2):242–258, 2021.
- Matthew J Hornsey, Matthew Finlayson, Gabrielle Chatwood, and Christopher T Begeny. Donald trump and vaccination: The effect of political identity, conspiracist ideation and presidential tweets on vaccine hesitancy. *Journal of Experimental Social Psychology*, 88:103947, 2020.
- Anton Gollwitzer, Cameron Martel, William J Brady, Philip Pärnamets, Isaac G Freedman, Eric D Knowles, and Jay J Van Bavel. Partisan differences in physical distancing are linked to health outcomes during the covid-19 pandemic. *Nature human behaviour*, 4(11):1186–1197, 2020.
- Yue Sun and Shannon M Monnat. Rural-urban and within-rural differences in covid-19 vaccination rates. *The Journal of Rural Health*, 38(4):916–922, 2022.

- Inmaculada Hernandez, Sean Dickson, Shangbin Tang, Nico Gabriel, Lucas A Berenbrok, and Jingchuan Guo. Disparities in distribution of covid-19 vaccines across us counties: A geographic information system–based cross-sectional study. *Plos Medicine*, 19(7):e1004069, 2022.
- Zachary Parolin and Emma K Lee. The role of poverty and racial discrimination in exacerbating the health consequences of covid-19. *The Lancet Regional Health–Americas*, 7, 2022.
- Ryoko Sato and Yoshito Takasaki. Peer effects on vaccination behavior: experimental evidence from rural nigeria. *Economic Development and Cultural Change*, 68(1):93–129, 2019.
- Johan Klaesson, José Lobo, and Charlotta Mellander. Social interactions and covid-19 vaccine hesitancy: Evidence from a full population study in sweden. *PLoS One*, 18(11):e0289309, 2023.
- Mara Howard-Williams, Rieza H Soelaeman, Leah S Fischer, Russell McCord, Robin Davison, and Christopher Dunphy. Association between state-issued covid-19 vaccine mandates and vaccine administration rates in 12 us states and the district of columbia. In *JAMA Health Forum*, volume 3, pages e223810–e223810. American Medical Association, 2022.
- Shawn Patterson. The politics of pandemics: The effect of stay-at-home orders on covid-19 mitigation. *State Politics & Policy Quarterly*, 22(1): 1–23, 2022.
- L2. Data mapping, 2024.
- Deepa Balraj, Ayin Vala, Shiyong Hao, Melanie Philofsky, Anna Tsvetkova, Elena Trach, Shravani Priya Narra, Oleg Zhuk, Mary Shamkhorskaya, Jim Singer, et al. American family cohort, a data resource description. *arXiv preprint arXiv:2309.13175*, 2023.
- Guido W Imbens and Charles F Manski. Confidence intervals for partially identified parameters. *Econometrica*, 72(6):1845–1857, 2004.
- Robert Phillips. The prime registry helps thousands of primary care clinicians liberate ehr data and prepare for mips. 2017.

Shiying Hao, David H Rehkopf, Esther Velasquez, Ayin Vala, Andrew W Bazemore, and Robert L Phillips Jr. Covid-19 vaccine strategy left small primary care practices on the sidelines. *Health Affairs*, 42(8):1147–1151, 2023.

Leslie L Roos Jr, Andre Wajda, and J Patrick Nicol. The art and science of record linkage: methods that work with few identifiers. *Computers in Biology and Medicine*, 16(1):45–57, 1986.

Appendix: Additional Proofs of Propositions

Proof of Lemma 1. We can write

$$D_{NM} = \frac{\mathbb{E}[X_N Y_N]}{\mathbb{E}[X_N]} - \frac{\mathbb{E}[(1 - X_N) Y_N]}{\mathbb{E}[1 - X_N]} = \frac{\mathbb{E}[X_N Y_N] - \mathbb{E}[Y_N] \mathbb{E}[X_N]}{\mathbb{E}[X_N](1 - \mathbb{E}[X_N])} = \frac{\text{Cov}(X_N, Y_N)}{\mathbb{E}[X_N](1 - \mathbb{E}[X_N])}.$$

Multiplying and dividing by $\text{Var}(X_N)$, recognizing that $\mathbb{E}[X_N](1 - \mathbb{E}[X_N]) = \mathbb{E}[X](1 - \mathbb{E}[X]) = \text{Var}[X]$, and noting that $D_{ER} = \frac{\text{Cov}(X_N, Y_N)}{\text{Var}(X_N)}$ yields the claim. $\square$

Lemma 2 (Feasibility of Intermediates). *Suppose that we have two feasible marginal distributions $(Y^1, Y^0) = (A^1, A^0)$ and $(Y^1, Y^0) = (B^1, B^0)$, where feasible means that $Y_n^j \in [\underline{Y}, \overline{Y}]$ for j in $\{0, 1\}$ and $Y_n = X_n Y_n^1 + (1 - X_n) Y_n^0$. Then:*

(i) *for any $\alpha \in [0, 1]$, the pair of marginals $(Y^1, Y^0) = (\alpha A^1 + (1 - \alpha) B^1, \alpha A^0 + (1 - \alpha) B^0)$ is feasible in the same sense and has disparity $D^\alpha := \alpha D_A + (1 - \alpha) D_B$ where $D_A = A^1 - A^0$ and $D_B = B^1 - B^0$.*

(ii) *If $D_A \leq D_B$ are two feasible values for disparity in the sense above with $(Y^1, Y^0) = (A^1, A^0)$ and $(Y^1, Y^0) = (B^1, B^0)$ respectively, then for any $D \in [D_A, D_B]$, D is feasible as well, and a (possibly non-unique) feasible distribution which gives rise to D is the α -mixture over (A^1, A^0) and (B^1, B^0) where $\alpha = \frac{D_B - D}{D_B - D_A}$.*

Proof. (i) To show feasibility, we need that for each $n \in \mathcal{N}$:

$$\begin{aligned} X_n Y_n^1 + (1 - X_n) Y_n^0 &= Y_n \\ Y_n^0, Y_n^1 &\in [\underline{Y}, \overline{Y}] \end{aligned}$$

So consider the α -mixture over the distributions A and B , i.e.:

$$\begin{aligned} Y^1 &= \alpha A^1 + (1 - \alpha) B^1, \\ Y^0 &= \alpha A^0 + (1 - \alpha) B^0. \end{aligned}$$

We designate the neighborhood-level distributions by Y_n^{1A} and Y_n^{0A} for (A^1, A^0) and Y_n^{1B} and Y_n^{0B} for (B^1, B^0) . Define

$$\begin{aligned} Y_n^{1\alpha} &= \alpha Y_n^{1A} + (1 - \alpha) Y_n^{1B} \\ Y_n^{0\alpha} &= \alpha Y_n^{0A} + (1 - \alpha) Y_n^{0B}. \end{aligned}$$

Since $Y_n^{1\alpha}$ is in between Y_n^{1A} and Y_n^{1B} , we must have that $Y_n^{1\alpha}, Y_n^{0\alpha} \in [\underline{Y}, \bar{Y}]$, and similarly for $Y_n^{0\alpha}$. We also have that

$$\begin{aligned} X_n Y_n^{1\alpha} + (1 - X_n) Y_n^{0\alpha} &= X_n (\alpha Y_n^{1A} + (1 - \alpha) Y_n^{1B}) + (1 - X_n) (\alpha Y_n^{0A} + (1 - \alpha) Y_n^{0B}) \\ &= \alpha (X_n Y_n^{1A} + (1 - X_n) Y_n^{0A}) + (1 - \alpha) (X_n Y_n^{1B} + (1 - \alpha) Y_n^{0B}) \\ &= \alpha Y_n + (1 - \alpha) Y_n \\ &= Y_n. \end{aligned}$$

Thus, the choice of $Y_n^{1\alpha}, Y_n^{0\alpha}$ is feasible for each n , and $(Y^{1\alpha}, Y^{0\alpha})$ is feasible overall.

For disparity, note that:

$$\begin{aligned} Y^{1\alpha} - Y^{0\alpha} &= \alpha A^1 + (1 - \alpha) A^0 - (\alpha B^1 + (1 - \alpha) B^0) \\ &= \alpha D^A + (1 - \alpha) D^B. \end{aligned}$$

(ii) First note that $\alpha = \frac{D_B - D}{D_B - D_A} \in [0, 1]$. (i) then implies that the α -mixture over (A^1, A^0) and (B^1, B^0) is feasible, and has disparity:

$$D^\alpha = \alpha D^A + (1 - \alpha) D^B = \frac{D_B - D}{D_B - D_A} \cdot D_A + \frac{D - D_A}{D_B - D_A} \cdot D_B = D.$$

Thus, D is feasible, and the α -mixture gives rise to it, as desired. $\square$

Proof of Proposition 1. Applying $D_{NM} = \gamma D_{ER}$ to Propositions 4 and 6 gives:

$$D - \frac{\delta_B}{\text{Var}(X)} = \gamma \left(D + \frac{\delta_W}{\text{Var}(X_N)} \right) \implies D(1 - \gamma) = \frac{\delta_B}{\text{Var}(X)} + \frac{\gamma \delta_W}{\text{Var}(X_N)} = \frac{\delta_B + \delta_W}{\text{Var}(X)}.$$

Dividing both sides by $(1 - \gamma)$ yields the claim. $\square$

Proof Sketch of Proposition 2. First consider the constraints that the observed data impose on the neighborhood-level means. In particular, we know that:

$$Y_n = X_n Y_n^1 + (1 - X_n) Y_n^0 \implies Y_n^0 = \frac{Y_n - X_n \cdot Y_n^1}{1 - X_n}; Y_n^1 = \frac{Y_n - (1 - X_n) \cdot Y_n^0}{X_n}.$$

Thus, increasing Y_n^1 necessarily requires decreasing Y_n^0 , and vice versa. Additionally, we know that $Y_n^1, Y_n^0 \in [\underline{Y}, \bar{Y}]$. We can therefore rewrite the constraints as limits on Y^1 :

$$Y_n^1 \geq \frac{Y_n - \bar{Y}(1 - X_n)}{X_n}; Y_n^1 \leq \frac{Y_n - \underline{Y}(1 - X_n)}{X_n}; Y_n^1 \leq \bar{Y}, Y_n^1 \geq \underline{Y},$$

where the first two inequalities correspond to setting Y_n^0 to the highest and lowest possible values, respectively. Consolidating these constraints yields tight bounds for Y_n^1 :

$$Y_{MOB,n}^{1-} := \max \left\{ \frac{Y_n - \bar{Y}(1 - X_n)}{X_n}, \underline{Y} \right\} \leq Y_n^1 \leq \min \left\{ \frac{Y_n - \underline{Y}(1 - X_n)}{X_n}, \bar{Y} \right\} =: Y_{MOB,n}^{1+}.$$

A similar calculation yields the corresponding bounds for Y_n^0 :

$$Y_{MOB,n}^{0-} := \max \left\{ \frac{Y_n - \bar{Y} X_n}{1 - X_n}, \underline{Y} \right\} \leq Y_n^0 \leq \min \left\{ \frac{Y_n - \underline{Y} X_n}{1 - X_n}, \bar{Y} \right\} =: Y_{MOB,n}^{0+}.$$

Turning from the neighborhood-level to the population, note that applying the law of iterated expectations followed by Bayes rule allows us to write:

$$Y^1 = \sum_{n \in \mathcal{N}} \Pr[N = n | X = 1] Y_n^1 = \sum_{n \in \mathcal{N}} \frac{\Pr[X = 1 | N = n] \Pr[N = n]}{\Pr[X = 1]} Y_n^1 = \sum_{n \in \mathcal{N}} \frac{X_n p_n}{\mathbb{E}[X]} Y_n^1.$$

Next, note that there are no additional constraints that limit the distribution of Y across neighborhoods. The maximum possible value of Y^1 is therefore obtained when Y_n^1 is as high as possible for all n and Y_n^0 is as low as possible, and vice versa for the minimum possible value of Y^1 . Substituting

$Y_{MOB,n}^{1+}$ into the above expression yields:

$$\begin{aligned}
Y^1 &\leq \sum_{n \in \mathcal{N}} \frac{X_n p_n}{\mathbb{E}[X]} Y_{MOB}^{1+} \\
&= \sum_{n \in \mathcal{N}} \frac{X_n p_n}{\mathbb{E}[X]} \min \left\{ \frac{Y_n - \underline{Y}(1 - X_n)}{X_n}, \bar{Y} \right\} \\
&= \sum_{n \in \mathcal{N}} p_n \frac{\min \{Y_n - \underline{Y}(1 - X_n), \bar{Y} X_n\}}{\mathbb{E}[X]} \\
&= \frac{\mathbb{E}[\min \{Y_N - \underline{Y}(1 - X_N), \bar{Y} X_N\}]}{E[X]} := Y_{MOB}^{1+}
\end{aligned}$$

The claimed value for Y_{MOB}^{1-} and the method of bounds interval for Y^0 follow similarly. Sharpness follows by construction: the endpoints were selected to be (jointly) consistent with feasible neighborhood-level values of the outcome and are thus themselves feasible. Finally, feasibility of the endpoints of the interval implies the feasibility of each interior point; see Lemma 2. $\square$

Proof Sketch of Proposition 3. The largest differences in group means is obtained when Y^1 takes on its largest feasible value and Y^0 takes on its smallest feasible value, and vice-versa with respect to the smallest difference in group means:

$$D \in [Y_{MOB}^{1+} - Y_{MOB}^{0-}, Y_{MOB}^{1-} - Y_{MOB}^{0+}]$$

The claimed bounds follow from substituting in the values of Y_{MOB}^{1+} , Y_{MOB}^{1-} , Y_{MOB}^{0-} , and Y_{MOB}^{0+} from Proposition 2 and simplifying. Sharpness follows from the sharpness of the bounds for the outcome means in Proposition 2 the pairs $(Y_{MOB}^{1+}, Y_{MOB}^{0-})$ and $(Y_{MOB}^{1-}, Y_{MOB}^{0+})$ are both feasible (since Y_{MOB}^{1+} is the maximum value that leaves Y_{MOB}^{0-} still feasible, and so on), hence the endpoints of the claimed bounds for D , $Y_{MOB}^{1+} - Y_{MOB}^{0-}$ and $Y_{MOB}^{1-} - Y_{MOB}^{0+}$ are feasible values D can take on. Lemma 2 then implies that all points within the bounds are also feasible; the bounds are thus sharp. $\square$

Proof of Corollary 1. We will provide a proof for D_{NM} (the proofs for Y_{NM}^1 and Y_{NM}^0 are analogous). We will show that D_{NM} lies within the interior of the method of bounds interval for D provided in Proposition 3. Feasibility then follows from the sharpness of that interval. To show that D_{NM}

lies within the method of bounds interval, recall that we can express $D_{NM} = \frac{\mathbb{E}[X_N Y_N] - \mathbb{E}[X_N] \mathbb{E}[Y_N]}{\text{Var}(X)}$, so that $D_{MOB}^+ - D_{NM} \geq 0 \iff \mathbb{E}[\min \{Y_N - \underline{Y}(1 - X_N), \bar{Y}X_N\} - X_N Y_N] \geq 0$. A sufficient condition under which this inequality holds for the expectation is that it holds for each n ,

$$\begin{aligned} 0 &\leq \min \{Y_n - \underline{Y}(1 - X_n), \bar{Y}X_n\} - X_n Y_n \\ &= \min \{Y_n - \underline{Y}(1 - X_n) - X_n Y_n, \bar{Y}X_n - X_n Y_n\} \\ &= \min \{(Y_n - \underline{Y})(1 - X_n), (\bar{Y} - Y_n)X_n\}. \end{aligned}$$

The inequality follows from the fact that $Y_n \in [\underline{Y}, \bar{Y}]$. The proof that $D_{NM} \geq D_{MOB}^-$ is analogous. $\square$

Proof of Proposition 4. To begin, note that:

$$\begin{aligned} D_{ER} &= \frac{\text{Cov}(Y, X_N)}{\text{Var}[X_N]} \\ &= \frac{\mathbb{E}[\text{Cov}(Y, X_N|X)] + \text{Cov}(\mathbb{E}[Y|X], \mathbb{E}[X_N|X])}{\text{Var}[X_N]} \\ &= \frac{\delta_W}{\text{Var}[X_N]} + \frac{\text{Cov}(\mathbb{E}[Y|X], \mathbb{E}[X_N|X])}{\text{Var}[X_N]}, \end{aligned}$$

where the first equality follows from definition of D^{ER} , the second from the law of total covariance, and the third from the definition of δ_W . It thus suffices to show that $\text{Cov}(\mathbb{E}[Y|X], \mathbb{E}[X_N|X]) = D \cdot \text{Var}[X_N]$.

To do so, first note that:

$$\mathbb{E}[X_N|X = 1] = \sum_{n \in \mathcal{N}} \Pr[N = n|X = 1] \cdot X_n = \sum_{n \in \mathcal{N}} \frac{X_n \cdot p_n}{\Pr[X = 1]} X_n = \frac{\mathbb{E}[X_N^2]}{\mathbb{E}[X]}$$

where the first equality follows by the definition of conditional expectation and the second by Bayes' Rule. Similarly,

$$\mathbb{E}[X_N|X = 0] = \sum_{n \in \mathcal{N}} \Pr[N = n|X = 0] \cdot X_n = \sum_{n \in \mathcal{N}} \frac{(1 - X_n)p_n}{\Pr[X = 0]} \cdot X_n = \frac{\mathbb{E}[X] - \mathbb{E}[X_N^2]}{1 - \mathbb{E}[X]}$$

Turning to the main issue:

$$\begin{aligned} \text{Cov}(\mathbb{E}[Y|X], \mathbb{E}[X_N|X]) &= \mathbb{E}[\mathbb{E}[Y|X] \cdot \mathbb{E}[X_N|X]] - \mathbb{E}[\mathbb{E}[Y|X]] \cdot \mathbb{E}[\mathbb{E}[X_N|X]] \\ &= \mathbb{E}[\mathbb{E}[Y|X] \cdot \mathbb{E}[X_N|X]] - \mathbb{E}[Y] \cdot \mathbb{E}[X] \end{aligned}$$

where the second equality follows by the law of iterated expectations. Now note that:

$$\begin{aligned}
\mathbb{E}[\mathbb{E}[Y|X] \cdot \mathbb{E}[X_N|X]] &= \Pr[X = 1] \cdot (\mathbb{E}[Y|X = 1] \cdot \mathbb{E}[X_N|X = 1]) \\
&\quad + (1 - \Pr[X = 1]) \cdot (\mathbb{E}[Y|X = 0] \cdot \mathbb{E}[X_N|X = 0]) \\
&= \mathbb{E}[X] \cdot \left(\mathbb{E}[Y|X = 1] \frac{\mathbb{E}[X_N^2]}{\mathbb{E}[X]} \right) + (1 - \mathbb{E}[X]) \cdot \left(\mathbb{E}[Y|X = 0] \cdot \frac{\mathbb{E}[X] - \mathbb{E}[X_N^2]}{1 - \mathbb{E}[X]} \right) \\
&= \mathbb{E}[Y|X = 1]\mathbb{E}[X_N^2] + \mathbb{E}[Y|X = 0] \cdot \mathbb{E}[X] - \mathbb{E}[X_N^2] \cdot \mathbb{E}[Y|X = 0]
\end{aligned}$$

Collecting terms, this equals:

$$\mathbb{E}[X_N^2] (\mathbb{E}[Y|X = 1] - \mathbb{E}[Y|X = 0]) + \mathbb{E}[Y|X = 0] \cdot \mathbb{E}[X] = \mathbb{E}[X_N^2] \cdot D + \mathbb{E}[X] \cdot \mathbb{E}[Y|X = 0]$$

Substituting this result back into $\text{Cov}(\mathbb{E}[Y|X], \mathbb{E}[X_N|X])$, we have:

$$\begin{aligned}
\text{Cov}(\mathbb{E}[Y|X], \mathbb{E}[X_N|X]) &= \mathbb{E}[X_N^2] \cdot D + \mathbb{E}[X] \cdot \mathbb{E}[Y|X = 0] - \mathbb{E}[X] \cdot \mathbb{E}[Y] \\
&= \mathbb{E}[X_N^2] \cdot D + \mathbb{E}[X] \cdot (\mathbb{E}[Y|X = 0] - \mathbb{E}[Y])
\end{aligned}$$

But note that:

$$\begin{aligned}
\mathbb{E}[Y|X = 0] - \mathbb{E}[Y] &= \mathbb{E}[Y|X = 0] - \mathbb{E}[X] \cdot \mathbb{E}[Y|X = 1] - (1 - \mathbb{E}[X]) \cdot \mathbb{E}[Y|X = 0] \\
&= \mathbb{E}[X] \cdot \mathbb{E}[Y|X = 0] - \mathbb{E}[X] \cdot \mathbb{E}[Y|X = 1] \\
&= -\mathbb{E}[X] \cdot D
\end{aligned}$$

Substituting this back into the previous equation, we have:

$$\begin{aligned}
\text{Cov}(\mathbb{E}[Y|X], \mathbb{E}[X_N|X]) &= \mathbb{E}[X_N^2] \cdot D + \mathbb{E}[X] \cdot (-\mathbb{E}[X] \cdot D) \\
&= D \cdot (\mathbb{E}[X_N^2] - \mathbb{E}[X]^2)
\end{aligned}$$

Finally, since $\mathbb{E}[X_N] = \mathbb{E}[X]$, the latter term is $\mathbb{E}[X_N^2] - \mathbb{E}[X_N]^2 = \text{Var}[X_N]$. We therefore have that:

$$\text{Cov}(\mathbb{E}[Y|X], \mathbb{E}[X_N|X]) = D \cdot \text{Var}(X_N)$$

proving the result. $\square$

Proof of Proposition 5. (i) If $\delta_W \geq 0$, then by Proposition 4 we have that

$$D = D_{ER} - \frac{\delta_W}{\text{Var}(X_N)} \leq D_{ER}.$$

By Proposition 3, we have that $D \in [D_{MOB}^-, D_{MOB}^+]$ as well. Since $D \leq D_{MOB}^+$ and $D \leq D_{ER}$, $D \leq \min\{D_{ER}, D_{MOB}^+\}$. Combining these yields the claimed bound.

(ii) Proof is similar to (i): if $\delta_W \leq 0$, then Proposition 4 shows that $D \geq D_{ER}$, and combining this with Proposition 3 yields the claimed bound.

(iii) We prove sharpness of (i) and (ii) follows similarly. First note that under (i), it must be that $D_{MOB}^- \leq D_{ER}$, because $\delta_W \geq 0 \implies D \leq D_{ER}$ by Proposition 4 and $D_{MOB}^- \leq D$ by Proposition 3. Thus, $[D_{MOB}^-, \min\{D_{MOB}^+, D_{ER}\}] \subseteq [D_{MOB}^-, D_{MOB}^+]$. Proposition 3 implies that for every $D' \in [D_{MOB}^-, D_{MOB}^+]$, there exists a distribution $\tilde{Y}$ resulting in $D = \tilde{D}$ consistent with the constraints on Y_n , i.e. $\tilde{Y}_n^x \in [\underline{Y}, \bar{Y}]$ for all $x \in \{0, 1\}$ and all n , and $X_n \cdot \tilde{Y}_n^1 + (1 - X_n)\tilde{Y}_n^0 = Y_n$ for all n . Since $\tilde{D} \leq D_{ER}$, such a $\tilde{Y}$ must also have $\delta_W \geq 0$ (again by Proposition 4). But this means that for all $D \in [D_{MOB}^-, \min\{D_{MOB}^+, D_{ER}\}]$, there exists a distribution $\tilde{Y}$ consistent with all constraints. The bounds are thus sharp. $\square$

Proof of Proposition 6. Recall that:

$$D_{NM} = \frac{\text{Cov}(X_N, Y_N)}{\text{Var}(X)}$$

Apply the Law of Iterated Expectations to rewrite $\text{Cov}(X_N, Y_N)$ as $\text{Cov}(X, Y) - \mathbb{E}[\text{Cov}(X, Y|N)]$. Note that the latter term is δ_B . Substitute this into D_{NM} , which gives

$$D_{NM} = \frac{\text{Cov}(X, Y) - \delta_B}{\text{Var}(X)}.$$

Recognizing $\frac{\text{Cov}(X, Y)}{\text{Var}(X)} = D$ and re-arranging yields the claim. $\square$

Proof of Proposition 7. (i) Note first that Corollary 1 shows that $D_{NM} \in [D_{MOB}^-, D_{MOB}^+]$. Now, by Proposition 3, we always have that $D \in [D_{MOB}^-, D_{MOB}^+]$; Proposition 6 shows that $\delta_B \geq 0 \implies D \geq D_{NM}$. Thus, we must have $D \geq D_{NM} \geq D_{MOB}^-$ and $D \leq D_{MOB}^+$; combining these yields the claimed bound.

(ii) follows similarly: Proposition 6 shows that $\delta_B \leq 0 \implies D \leq D_{NM}$, and combining with Proposition 3 and Corollary 1 yields the claimed bound.

(iii) For sharpness, we focus on (i); (ii) follows similarly. Note that $D_{NM} \leq D_{MOB}^- \implies [D_{NM}, D_{MOB}^+] \subseteq [D_{MOB}^-, D_{MOB}^+]$; the sharpness result of Proposition 3 thus implies that for all $\tilde{D} \in [D_{NM}, D_{MOB}^+]$, there exists a distribution $\tilde{Y}$ such that $\tilde{Y}_n^x \in [\underline{Y}, \bar{Y}]$ for $x \in \{0, 1\}$ and all n , and $X_n \tilde{Y}_n^1 + (1 - X_n) \tilde{Y}_n^0 = Y_n$ for all n resulting in $D = \tilde{D}$. But by Proposition 6, $\tilde{Y}$ results in $\delta_B \geq 0$ as required. $\tilde{Y}$ is thus feasible; since this is true for all $\tilde{D}$ in the bounds, the bounds are thus sharp. $\square$

Proof of Proposition 8. First note that since $D_{NM} = D_{ER} \cdot \gamma$ (Lemma 1), the fact that $0 \leq \gamma < 1$ implies that $\text{Sign}(D_{NM}) = \text{Sign}(D_{ER})$ and $|D_{ER}| \geq |D_{NM}|$. So it must always be true that either $0 \leq D_{NM} \leq D_{ER}$ or $D_{ER} \leq D_{NM} \leq 0$. At the same time, applying Propositions 4 and 6 and the fact that $\delta_W \cdot \delta_B \geq 0$ shows that D must be in between D_{NM} and D_{ER} .

There are thus two cases. In the first case, $0 \leq D_{NM} \leq D_{ER}$ and $D \in [D_{NM}, D_{ER}]$; since $D \leq D_{MOB}^+$ by Proposition 3, we also have that $D \leq \min\{D_{ER}, D_{MOB}^+\}$. Combining these bounds yields (i). In the second case, $D_{ER} \leq D_{NM} \leq 0$ and $D \in [D_{ER}, D_{NM}]$, and Proposition 3 shows that $D \geq D_{MOB}^-$. Combining these inequalities yields (ii).

We now turn to (iii). We prove sharpness for (i), and (ii) follows similarly. First note that by Corollary 1, D_{NM} is feasible in terms of the constraints on Y_n , and thus in particular $D_{NM} \geq D_{MOB}^-$; at the other end, trivially $\min\{D_{MOB}^+, D_{ER}\} \leq D_{MOB}^+$. The range $[D_{NM}, \min\{D_{MOB}^+, D_{ER}\}]$ is thus a subset of $[D_{MOB}^-, D_{MOB}^+]$, and the sharpness guaranteed by Proposition 3 thus implies that for all $\tilde{D}$ in the range, there is a feasible distribution $\tilde{Y}$ giving rise to $D = \tilde{D}$ with $\tilde{Y}_n^1, \tilde{Y}_n^0 \in [\underline{Y}, \bar{Y}]$ and $X_n \tilde{Y}_n^1 + (1 - X_n) \tilde{Y}_n^0 = Y_n$. Thus the range is consistent with Y_n constraints. At the same time, Propositions 4 and 6 together imply that for any $\tilde{D} \in [D_{NM}, D_{ER}]$, any $\tilde{Y}$ that gives rise to $D = \tilde{D}$ will also have $\text{Sign}(\delta_W) = \text{Sign}(\delta_B)$ and thus $\delta_W \cdot \delta_B \geq 0$. But $[D_{NM}, \min\{D_{MOB}^+, D_{ER}\}]$ is also a subset of $[D_{NM}, D_{ER}]$. Thus, $\delta_W \cdot \delta_B \geq 0$ for all $\tilde{D}$ in the range. The bounds given are thus sharp. $\square$

Proof of Proposition 9. Consider the claim for the bias of Y_{ER}^1 . It is enough

to show that $\delta_W = \frac{\text{Var}(X_N)}{1 - \mathbb{E}[X]} \cdot (Y_{ER}^1 - Y^1)$. We have:

$$Y_{ER}^1 = \mathbb{E}[Y_N] + \left(\frac{\text{Cov}(Y_N, X_N)}{\text{Var}(X_N)} \right) (1 - \mathbb{E}[X_N]).$$

Multiplying both sides by $\gamma \mathbb{E}[X]$, simplifying, and using $Y^1 = \mathbb{E}[XY]/\mathbb{E}[X]$ shows that:

$$\text{Cov}(X_N, Y_N) - \gamma \text{Cov}(X, Y) = \frac{\text{Var}(X_N)}{1 - \mathbb{E}[X]} \cdot (Y_{ER}^1 - Y^1).$$

Applying the definition of γ , the fact that $D = \text{Cov}(X, Y)/\text{Var}(X)$, and the bias formula from Proposition 5 shows that $\text{Cov}(X_N, Y_N) - \gamma \text{Cov}(X, Y) = \delta_W$; and chaining these equalities together proves the claim. The claim for $Y_{ER}^0 - Y^0$ follows similarly. $\square$

Proof of Proposition 10. Similar to Proposition 9, we rely on algebraic calculation: if we substitute in $\mathbb{E}[X_N Y_N] = \text{Cov}(X_N, Y_N) + \mathbb{E}[X] \mathbb{E}[Y]$ into the equation $Y_{NM}^1 - Y^1 = \frac{\mathbb{E}[X_N Y_N] - \mathbb{E}[X] \mathbb{E}[Y]}{\mathbb{E}[X]}$ and apply the Law of Total Covariance, we can show that:

$$Y_{NM}^1 - Y^1 = \frac{\text{Cov}(X_N, Y_N) - \text{Cov}(X, Y)}{\mathbb{E}[X]} = \frac{-\delta_B}{\mathbb{E}[X]},$$

which proves the claim. A similar approach shows the claim for $Y_{NM}^0 - Y^0$. $\square$

Proof of Proposition 11. (i) Note first of all that $D_n \leq D_{MOB,n}^+$ by construction; note further that $D_{MOB,n}^+ \geq 0$, so $[0, D_{MOB,n}^+]$ is non-empty. Now, note that:

$$\begin{aligned} \delta_{B,n} &:= \text{Cov}(Y, X|N = n) = \mathbb{E}[Y \cdot X|N = n] - \mathbb{E}[Y|N = n] \cdot \mathbb{E}[X|N = n] \\ &= \mathbb{E}[Y|X = 1, N = n] \cdot X_n - \mathbb{E}[Y|N = n] X_n \\ &= X_n \cdot (\mathbb{E}[Y|X = 1, N = n] - \mathbb{E}[Y|N = n]) \end{aligned}$$

Since Y_n is a convex combination of Y_n^1 and Y_n^0 , $\delta_{B,n} \geq 0 \iff Y_n^1 \geq Y_n^0 \iff D_n \geq 0$. Thus, $\delta_{B,n} \geq 0 \implies D_n \in [0, D_{MOB,n}^+]$. For Y_n^1 , note that $Y_n \leq Y_{MOB,n}^{1+}$ by construction, and $Y_n^1 \geq Y_n^0 \implies Y_n^1 \geq Y_n$ (again since Y_n is a convex combination of Y_n^1 and Y_n^0); thus, $Y_n^1 \in [Y_n, Y_{MOB,n}^{1+}]$. For Y_n^0 again $Y_n^0 \geq Y_{MOB,n}^{0-}$ by construction, and similarly $Y_n^0 \leq Y_n^1 \implies Y_n^0 \leq Y_n$.

(ii) follows similarly, with inequalities reversed as implied by $\delta_{B,n} \leq 0$.

(iii) We focus on (i), and (ii) follows similarly. First note that $Y_n^1 = Y_n^0 = Y_n \in [\underline{Y}, \bar{Y}]$ results in $\mathbb{E}[Y|N = n] = Y_n$ and $\delta_{B,n} = D_n = 0$. At the same time, the choice $Y_n^1 = Y_{MOB,n}^{1+} \in [\underline{Y}, \bar{Y}]$ and $Y_n^0 = Y_{MOB,n}^{0-} \in [\underline{Y}, \bar{Y}]$ results in $\mathbb{E}[Y|N = n] = Y_n$ and $D = D_{MOB,n}^+$ by construction. Thus, the endpoints of each range are feasible. As for the interior of the range, note that for any $\tilde{D} \in [0, D_{MOB,n}^+]$, choosing $\alpha = (D_{MOB,n}^+ - \tilde{D})/D_{MOB,n}^+$ and letting $Y_n^1 = \alpha \cdot Y_n + (1 - \alpha) \cdot Y_{MOB,n}^{1+}$ and $Y_n^0 = \alpha \cdot Y_{MOB,n}^{0-} + (1 - \alpha) \cdot Y_n$ gives $D = \tilde{D}$ (as described in the proof of Lemma 2). A similar argument can be made for the interiors of $[Y_n, Y_{MOB,n}^{1+}]$ and $[Y_{MOB,n}^{0-}, Y_n]$. The bounds are thus sharp. $\square$

Proof of Proposition 12. (i) By construction, $D_n \in [D_{MOB,n}^-, D_{MOB,n}^+]$. Now, $D_n = \mu'(X_n) - \delta_{W,n}$, so $\delta_{W,n} \geq 0 \implies D_n \leq \mu'(X_n)$. Thus, $D_n \leq \min\{D_{MOB,n}^+, \mu'(X_n)\}$, yielding the claimed bound for D_n . For Y_n^1 , re-arrange $D_n = Y_n^1 - Y_n^0$ as $Y_n^0 = D_n - Y_n^1$, substitute this into $Y_n = X_n Y_n^1 + (1 - X_n) Y_n^0$, and re-arrange to see that $Y_n^1 = Y_n + (1 - X_n) \cdot D_n$. Combining this with the just-proved result that $\delta_{W,n} \geq 0 \implies D_n \leq \mu'(X_n)$ shows that $\delta_{W,n} \geq 0 \implies Y_n^1 \leq Y_n + (1 - X_n) \cdot \mu'(X_n)$. Since $Y_{MOB,n}^{1+} \geq Y_n^1$ by construction, we thus have that $Y_n^1 \leq \min\{Y_{MOB,n}^{1+}, Y_n + (1 - X_n) \cdot \mu'(X_n)\}$; combining this with $Y_n \geq Y_{MOB,n}^{1-}$ by construction yields the claimed bound. For $Y - n^0$, proceed similarly by substituting $Y_n^1 = D_n + Y_n^0$ into $Y_n = X_n \cdot Y_n^1 + (1 - X_n) \cdot Y_n^0$ and solve for $Y_n^0 = Y_n - X_n \cdot D_n$; the bound follows analogously.

(ii) follows by reversing the inequalities with similar substitution.

(iii) We focus on (i), and (ii) follows similarly. First note that $[D_{MOB,n}^-, \min\{\mu'(X_n), D_{MOB,n}^+\}] \subseteq [D_{MOB,n}^-, D_{MOB,n}^+]$; any $\tilde{D}_n$ in the interior of the interval then corresponds to some α -mixture over the $(Y_n^1, Y_n^0) = (Y_{MOB,n}^{1+}, Y_{MOB,n}^{0-})$ and $(Y_n^1, Y_n^0) = (Y_{MOB,n}^{1-}, Y_{MOB,n}^{0+})$, which are feasible in terms of being in $[\underline{Y}, \bar{Y}]$ and result in Y_n by construction. As shown in Lemma 2, any $\tilde{D}_n$ in the interior can be feasibly obtained by some α -mixture over the endpoints. And for the all $\tilde{D}_n$ in the interior, $\tilde{D}_n \leq \mu'(X_n)$, so $\delta_{W,n} \geq 0$. The bounds for D_n are thus sharp. For Y_n^1 , similarly note that the endpoints $Y_{MOB,n}^{1-}$ and $Y_{MOB,n}^{1+}$ (paired with their appropriate complements $Y_{MOB,n}^{0-}$ and $Y_{MOB,n}^{0+}$) correspond to $D_n = D_{MOB,n}^+$ and $D_{MOB,n}^-$, and are feasible in terms of $[\underline{Y}, \bar{Y}]$ and resulting in Y_n ; by similar logic as before, any $\tilde{Y}_n^1$ in the interior can be constructed

as some mixture over the endpoints and result in the equivalent mixture over $D_{MOB,n}^-$ and $D_{MOB,n}^+$. Since $Y_n^1 = Y_n + (1 - X_n) \cdot \mu'(X_n)$ is exactly the choice Y_n that makes $D_n = \mu'(X_n)$, then for all $Y_n \leq \{Y_{MOB,n}^{1+}, Y_n + (1 - X_n) \cdot \mu'(X_n)\}$, the resulting $D_n \leq \mu'(X_n)$ and $D_{MOB,n}^+$, so $\delta_{W,n} \geq 0$. The bound for Y_n^1 is thus sharp as well. Y_n^1 can be handled similarly. $\square$

Proof Sketch of Proposition 13. The bounds themselves follow by applying Propositions 11 and 12 simultaneously. For sharpness we again focus on (i), as (ii) is similar. Suppose $\mu'(X_n) \geq 0$. First note that $D_n = 0$ is achieved by $Y_n^1 = Y_n^0 = Y_n \in [\underline{Y}, \bar{Y}]$ and since $\min\{\mu'(X_n), D_{MOB,n}^+\} \leq D_{MOB,n}^+$, which can be achieved by $Y_{MOB,n}^{1+}, Y_{MOB,n}^{0-}$, each $\tilde{D}_n$ in the interval be achieved by some distribution that is feasible in terms of $[\underline{Y}, \bar{Y}]$ and resulting in Y_n . So it is enough to show that any such distribution will result in $\text{Sign}(\delta_{B,n}) = \text{Sign}(\delta_{W,n})$. But this holds because $\delta_{W,n} \geq 0 \iff D_n \leq \mu'(X_n)$, and $\delta_{B,n} \geq 0 \iff D_n \geq 0$, and both hold on the interior of the interval. A similar argument can be made for the bounds on Y_n^1 and Y_n^0 . $\square$