Political Analysis (2026), **0**, 1–20
doi:10.1017/pan.2026.10048

ARTICLE

Monotone Ecological Inference

Hadi Elzayn[1], Jacob Goldin[2], Cameron Guage[3], Daniel Ho[1] and Claire Morton[4]

[1]Law School, Stanford Law School, United States; [2]Law, The University of Chicago, United States; [3]Columbia University, United States; [4]Statistics, University of California Berkeley, United States

Corresponding author: Daniel Ho; E-mail: dho@law.stanford.edu

(Received 12 September 2025; revised 9 March 2026; accepted 11 March 2026)

Abstract

We study the identification of individual-level associations when only aggregate data are available. We characterize the biases of, and relationships among, canonical ecological inference (EI) estimators. We use these results to develop a partial identification approach: monotone EI. The approach exploits information about one or both of the following conditional associations: (1) outcome differences between groups within the same neighborhood and (2) outcome differences within the same group between 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. We illustrate our results using county-level data to study differences in COVID-19 vaccination rates among Republicans and Democrats in the United States.

Keywords: ecological inference; ecological fallacy; partial identification; ecological regression; neighborhood model; contextual effects; method of bounds; aggregate data; data combination

Edited by: Daniel J. Hopkins and Brandon M. Stewart

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. EI is challenging because absent additional assumptions, very little about the direction or magnitude of individual-level associations between the outcome and group membership can be inferred from neighborhood-level data.

There are several canonical estimation strategies that have been proposed for EI settings, but each requires strong assumptions. Ecological regression (ER; Goodman 1953) entails regressing neighborhood-level outcomes on neighborhood-level rates of group membership. The approach is unbiased only absent "contextual effects"—for example, 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 (NM; 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: that outcomes do not differ by group among individuals in the same neighborhood. Other approaches to EI, including prominent contributions by King (1997) and Wakefield (2004), rely on identification assumptions about the role of contextual effects within a parametric statistical model.[1]

[1]Specifically, 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 ER; the value of the approach is to combine the identifying power of that assumption with the constraints implied by worst-case bounds. Wakefield considers a range of Bayesian hierarchical modeling approaches.

© The Author(s), 2026. Published by Cambridge University Press on behalf of The Society for Political Methodology.
This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

https://doi.org/10.1017/pan.2026.10048 Published online by Cambridge University Press

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

In this article, we propose a middle-ground approach to identification in EI settings that strikes a balance between obtaining informative results and relying on potentially more credible identifying assumptions. The identifying assumptions we consider concern two quantities: (1) the between-group within-neighborhood association—that is, outcome differences between groups within the same neighborhood—and (2) the within-group between-neighborhood association—that is, outcome 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 aid identification. In particular, we derive sharp bounds for the outcome mean by group and for the difference in group outcome means 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 EI*, in the spirit of Manski and Pepper (2000).

A virtue of monotone EI 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, we show that the NM estimator is unbiased only if the between-group association is exactly zero, and ER 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, theory or auxiliary data may support an assumption about the likely direction of any such associations if they are present, enabling monotone EI.[3]

A related possibility is that a 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, one might 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 identify 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 EI 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). We argue for the applicability of the contextual reinforcement assumption to this setting and show that imposing it reduces the range of feasible values for the partisan vaccination gap by approximately 67%.

[2]King'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 Jiang *et al.* 2020 and Lewis 2001).

[3]As 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.

https://doi.org/10.1017/pan.2026.10048 Published online by Cambridge University Press

Our results contribute to a large literature that studies identification in EI settings; for surveys, see, for example, Cho and Manski (2008), King (1997), and King, Tanner, and Rosen (2004). 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). More recently, Li, Litvin, and Manski (2023) consider the role of "bounded variation assumptions" for identifying personalized risk assessments from published medical studies, a setting that shares some features of EI 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.

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, Mao, and Zhou 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.[4] 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, 2018; Molinari 2010; Rambachan and Roth 2023).

We proceed as follows. Section 2 describes our empirical setup. Section 3 analyzes and relates several canonical EI estimators. Section 4 derives sharp bounds on differences among group means. Section 5 extends the same approach to identifying levels of the group means. Section 6 considers neighborhood-level identification. Section 7 illustrates our approach with an EI analysis of partisan gaps in COVID-19 vaccinations. Section 8 concludes. The Supplementary Material contains proofs and additional results. 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 setting and notation

Each individual is characterized by three variables: the outcome of interest (Y), the group to which they belong (X) and the neighborhood in which they live (N). To keep things simple, we focus on the case in which group membership and the outcome of interest are each binary: $X \in \{0,1\}$ and $Y \in \{0,1\}$. The analysis extends naturally to settings in which there are more than two groups or in which the outcome is continuous, as we discuss in the Supplementary Material.

We are primarily interested in the mean values of Y by group

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

as well as the difference in group means

$$D = \mathbb{E}[Y|X = 1] - \mathbb{E}[Y|X = 0].$$

Following the presentation of our main results, we also consider the identification of these statistics for specific neighborhoods (see Section 6).

The core feature of the EI setting is that we do not observe individual-level values of X or Y. Rather, we observe one (aggregated) value of X and Y per neighborhood. We denote these aggregated values by

[4] Elzayn *et al.* (2025) does not relate its approach to EI nor show the sharpness of the bounds we provide here.

https://doi.org/10.1017/pan.2026.10048 Published online by Cambridge University Press

$X_N = \mathbb{E}[X|N]$ and $Y_N = \mathbb{E}[Y|N]$, respectively. We denote the neighborhood-level means of X and Y for a specific 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.

In our empirical application studying COVID-19 vaccinations, Y indicates whether an individual is vaccinated, X indicates whether an individual votes Republican and N indicates the county in which the individual resides. The share of vaccinated individuals in county n is denoted by Y_n, and X_n denotes the share of county residents who vote Republican. Our goal will be to use these county-level data to estimate the partisan vaccination gap, D, which we define as the difference in the mean vaccination rate of Republican relative to Democratic and third-party voters, $\mathbb{E}[Y \mid X = 1] - \mathbb{E}[Y \mid X = 0]$.

Another common application of EI methods is to study group differences in voting behavior. For example, a researcher might seek to use precinct-level voting data to measure election turnout by race. In that setting, Y would indicate whether an individual votes, X would indicate membership in a particular racial group, and N would denote the individual's precinct. The researcher would then infer differences in turnout rates by race (D) using the available precinct-level data on racial composition (X_N) and turnout (Y_N).

We next introduce a statistic summarizing the coarsening of the information on group membership (X) due to aggregation. By the law of total variance, we can decompose the total variation in X into the portions that are and are not explainable by neighborhood:[5]

$$\mathrm{Var}(X) = \mathrm{Var}(X_N) + \mathbb{E}\left[\mathrm{Var}(X) \mid N\right]. \tag{1}$$

We then define the *group aggregation ratio*

$$\gamma = \frac{\mathrm{Var}(X_N)}{\mathrm{Var}(X)}$$

to measure the share of the total variance in group membership that is attributable to neighborhood.[6] When neighborhoods do not differ much in group prevalence, $\mathrm{Var}(X_N)$—and therefore γ—will be small. In the extreme, when $\gamma = 0$, group composition does not vary at all across neighborhoods. Conversely, larger values of γ correspond to more information on group membership being preserved, with $\gamma = 1$ corresponding to the extreme scenario in which each neighborhood is entirely homogeneous by group (i.e., all individuals within the neighborhood are of the same group).[7] Because the ER estimator, and the bounds that we derive based on it, are not defined when $\gamma = 0$, we assume throughout that $\gamma \in (0, 1]$.

Many of our subsequent results relate to two conditional associations summarizing the relationship between X, Y and N. We define the *between-group association* as

$$\delta_B = \mathbb{E}\left[\mathrm{Cov}\left(Y, X \mid N\right)\right].$$

Intuitively, the between-group association measures differences in the outcome between individuals of different groups who live in the same neighborhood as one another. It reflects the portion of the association between X and Y that is not explained by neighborhood. In the context of our empirical application, for example, δ_B captures average within-county differences in vaccination uptake between

[5] We define all population moments, including those over neighborhood-level variables, with respect to the distribution of those variables among individuals. This implies, for example, that $\mathbb{E}[X_N] = \mathbb{E}[X]$.

[6] Equivalently, one can interpret γ as the R^2 from the regression of X on neighborhood fixed effects.

[7] Related measures of the information loss due to aggregation have been considered within the EI literature. For example, γ corresponds to $\eta^2_{X,A}$ in Equation (1) of Robinson (1950).

https://doi.org/10.1017/pan.2026.10048 Published online by Cambridge University Press

Republican and Democratic voters. Similarly, in a voting turnout application, δ_B would reflect average racial differences in turnout among individuals in the same precinct.

The second conditional association on which we focus is the *within-group association*

$$\delta_W = \mathbb{E}\left[\mathrm{Cov}\left(Y, X_N \mid X\right)\right].$$

The within-group association captures within-group differences in the outcome across neighborhoods with different group compositions. That is, it measures the portion of the association between X and Y that is attributable to neighborhoods, holding individual group membership constant. To illustrate, in our empirical application, δ_W captures differences in vaccination uptake among individuals who are registered for the same political party who live in neighborhoods with differing concentrations of Republican voters. It would tend to be negative if Republican voters living in mostly-Republican neighborhoods were vaccinated at lower rates than Republican voters living in mostly-Democratic neighborhoods, and similarly, if Democratic voters living in mostly-Republican neighborhoods were vaccinated at lower rates than Democratic voters living in mostly-Democratic neighborhoods. In the context of a voter turnout study, δ_W would reflect differences in turnout among individuals of the same racial group living in precincts with differing concentrations of minority residents. For example, it would be non-zero if both minority and non-minority residents tended to be less likely to turn out for an election when their precinct has a high fraction of minority residents.

In the next section, we study canonical approaches to EI estimation through the lens of the group aggregation ratio and these conditional associations.

3. Canonical EI estimators

This section focuses on three common approaches for studying group-level differences in EI settings: ER, the NM and the MOB.[8] We present these familiar estimators informally; the Supplementary Material contains additional detail. For ease of exposition, we initially focus on identifying the difference in group means (i.e., D), deferring consideration of the levels of the group means (i.e., Y^0 and Y^1) to Section 5.

3.1. Ecological regression

The most common approach to EI is ER (Goodman 1953). The ER estimate for the difference in group means is the coefficient for X_N from the weighted neighborhood-level regression of Y_N on X_N, with weights based on neighborhood population (P_n). In our empirical application, for example, ER involves regressing county-level vaccination rates on county-level Republican vote shares; the ER estimate for the partisan vaccination gap is the estimated slope from this regression.

We focus on the ER estimator's asymptotic limit under an *iid* sampling process[9]:

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

The following proposition relates the bias of ER to the within-group association.[10]

[8]We focus on these approaches (as opposed to other popular methods like King (1997)) because they will appear as inputs into the monotone EI bounds we derive.

[9]Recall that all moments are defined with respect to the population of individuals, which is why the regression weights P_N do not appear in the D_{ER} formula. Computing the sample analog to this quantity would involve weighting each observed neighborhood-level quantity by its population share, P_n, as formalized in the Supplementary Material.

[10]Throughout, we use "bias" to refer to the asymptotic bias of the specified estimator, that is, the difference between what the estimator identifies and the target estimand. Propositions 1 and 2 are the EI analogs to Propositions 1.1 and 1.2 in Elzayn *et al.* (2025), which apply to individual-level data in which group membership is replaced by probabilistic proxies.

https://doi.org/10.1017/pan.2026.10048 Published online by Cambridge University Press

Proposition 1 (Bias of ER for difference in group means).

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

where $\delta_W = \mathbb{E}\left[\mathrm{Cov}\left(Y, X_N \mid X\right)\right]$.

The proof of Proposition 1, and all subsequent results, is provided in the Supplementary Material.

When δ_W is non-zero, ER is biased because differences in the prevalence of groups across neighborhoods are conflated by other neighborhood-level contextual effects (Ansolabehere and Rivers 1995; Goodman 1953). Proposition 1 provides a simple expression for how such contextual effects shape this bias.

Proposition 1 implies the following corollary.

Corollary 1. *The ER estimand, D_{ER}, is unbiased if and only if $\delta_W = 0$.*

A related condition that is often discussed in the EI literature 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 each $x \in \{0,1\}$. Although this condition implies $\delta_W = 0$, the converse is not true.

3.2. The neighborhood model

An alternative EI method for estimating group differences is the so-called NM (Freedman *et al.* 1991).[11] To estimate D, the NM first constructs estimates of Y^0 and Y^1. Each group-specific estimate is formed from the weighted average of the outcome across neighborhoods, with weights equal to the specified group's prevalence in each neighborhood. In our empirical application, for example, the NM estimate for the Republican vaccination rate is the average of county-level vaccination rates, weighted by the prevalence of Republicans in each county. The NM estimate for D is then formed as the difference in the NM estimates for Y^0 and Y^1.

As with ER, we will focus on the NM estimator's asymptotic limit:

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

The following proposition shows that the bias of the NM is determined by the between-group association.

Proposition 2 (Bias of the NM estimator for the difference in group means).

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

where $\delta_B = \mathbb{E}\left[\mathrm{Cov}(X, Y|N)\right]$.

Intuitively, δ_B measures how much groups differ within the same neighborhood. If $\delta_B > 0$, group 1 tends to have higher outcomes than group 0 within each neighborhood. In that case, the NM, which effectively averages neighborhood outcomes weighting by group presence, tends to understate the true difference in group means because it ignores that group 1 individuals are disproportionately in higher-outcome positions within each neighborhood. Conversely, when $\delta_B < 0$, the NM overstates the difference in group means because it ignores that group 1 individuals tend to have lower outcomes than group 0 individuals within the neighborhoods being averaged.

Proposition 2 also implies the following corollary.

[11] (Freedman *et al.* 1991) propose two model variants: the non-linear NM and the linear neighborhood model. Our focus is on the former.

https://doi.org/10.1017/pan.2026.10048 Published online by Cambridge University Press

Corollary 2. *The NM estimand, D_{NM}, is unbiased if and only if $\delta_B = 0$.*

Freedman *et al.* (1991) noted that the NM point-identifies $\mathbb{E}[Y \mid X, N = n]$ for a specific neighborhood n under the related condition that there is no systematic difference in the outcome between groups within that neighborhood, that is, $\mathbb{E}[Y|0, N = n] = \mathbb{E}[Y|1, N = n]$. Because our focus is on group differences over the entire population (rather than for a specific neighborhood), Corollary 2 requires only that there be no within-neighborhood differences on average.

We are now in a position to highlight some connections between the ER and NM approaches. First, note that the two estimands follow a close mechanical relationship, mediated by the group aggregation ratio.[12]

Lemma 1 (Relationship between NM and ER).

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

where $\gamma = \frac{\mathrm{Var}(X_N)}{\mathrm{Var}(X)}$.

Using Lemma 1, we can next derive the following identity.

Proposition 3 (Relationship of conditional associations and group aggregation ratio).

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

Because δ_W and δ_B control the respective biases of the ER and NM estimators, Proposition 3 shows that the performance of these estimators is fundamentally linked and related to the amount of information lost due to aggregation (via γ). In particular, when group prevalence does not vary much by neighborhood (so that most of the variation in X is within rather than between neighborhoods), will be small, and the sum will be large relative to the actual difference in group means, D. Finally, in conjunction with Corollaries 1 and 2, Proposition 3 establishes that except in special cases (i.e., $\gamma = 1$ or $D = 0$), the assumptions justifying the ER and NM estimators are mutually exclusive.

3.3. Method of bounds

Whereas ER and the NM yield point estimates for D, an alternative approach is to calculate the most extreme values of D that are consistent with the observed data. In EI settings with binary Y, consistency with the data requires that, for each neighborhood n, (i) $Y_n^0 \in [0,1]$; (ii) $Y_n^1 \in [0,1]$; and (iii) $X_n Y_n^1 + (1 - X_n) Y_n^0 = Y_n$, where $Y_n^x = \mathbb{E}[Y|N = n, X = x]$. The MOB interval for the group means is defined by the minimum and maximum values that satisfy these constraints. The interval for the difference in group means is then derived by comparing the minimum and maximum values that each group mean can take on.

Proposition 4 (MOB for difference in group means). *Define the following parameters:*

$$D_{MOB}^{+} = \frac{\mathbb{E}[\min\{Y_N, X_N\}] - \mathbb{E}[X_N]\mathbb{E}[Y_N]}{\mathrm{Var}(X)}$$

$$\textit{and } D_{MOB}^{-} = \frac{\mathbb{E}[Y_N](1 - \mathbb{E}[X_N]) - \mathbb{E}[\min\{Y_N, 1 - X_N\}]}{\mathrm{Var}(X)}.$$

It follows that

$$D \in \left[D_{MOB}^{-}, D_{MOB}^{+}\right].$$

These bounds are sharp.

[12] Surprisingly, the relationship 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).

https://doi.org/10.1017/pan.2026.10048 Published online by Cambridge University Press

The insight underlying Proposition 4 is originally due to Duncan and Davis (1953); it has been extended and formalized by Horowitz and Manski (1995), Cross and Manski (2002) and Cho and Manski (2008), who focus on neighborhood-level bounds. Proposition 4 provides a closed-form expression for the aggregation of those neighborhood-specific bounds to the population-level.

We next relate the MOB interval to the ER and NM point estimates. Beginning with the latter, note that the aggregate data can never rule out the possibility that outcomes are homogeneous within neighborhoods, that is, that $\delta_B = 0$. Hence, the NM estimate is guaranteed to fall within the MOB interval.

Corollary 3. $D_{NM} \in \left[D^-_{MOB}, D^+_{MOB}\right]$.

In contrast, no such guarantee is available for δ_W; as a result, the ER estimate may sometimes be infeasible. Comparing the MOB endpoints to the expression for D_{ER} in Proposition 1 sheds light on when this occurs.

Corollary 4.

$$D_{ER} > D^+_{MOB} \iff \delta_W > \gamma \operatorname{Var}(X)\left(D^+_{MOB} - D\right), \text{ and}$$
$$D_{ER} < D^-_{MOB} \iff \delta_W < -\gamma \operatorname{Var}(X)\left(D - D^+_{MOB}\right).$$

Corollary 4 shows that ER is more likely to yield an infeasible estimate when γ is small. Intuitively, when neighborhoods barely differ in their group composition, ER infers the slope based on small variations in X_N. Any residual correlation between X_N and Y_N (reflected by a non-zero δ_W) is thus amplified, leading to extreme values of D_{ER}.

Because the bounds in Proposition 4 are sharp, they highlight the limits of what can be learned from the data in EI settings without imposing additional assumptions. In particular, the following lemma shows that the aggregate data are more informative about D as neighborhoods are more (internally) homogeneous with respect to either group membership or outcome.

Lemma 2 (Width of MOB interval).

$$D^+_{MOB} - D^-_{MOB} = \frac{1}{\operatorname{Var}(X)} \mathbb{E}\left[\min\{X_N, 1 - X_N, Y_N, 1 - Y_N\}\right].$$

To illustrate, suppose that individuals are perfectly sorted across neighborhoods according to their group, that is, $X_n \in \{0,1\}$ for all n. In this case, Lemma 2 shows that the MOB interval collapses to a single point (i.e., D). Moreover, because this setting (i.e., homogeneity of groups within neighborhoods) corresponds to $\gamma = 1$, Lemma 1 implies that $D_{NM} = D_{ER}$. Finally, because the MOB interval always contains D_{NM} (Corollary 3), it must be the case that all three of the canonical EI approaches converge on the true value of D. We can therefore conclude that $\gamma = 1$ implies $D_{ER} = D_{NM} = D^-_{MOB} = D^+_{MOB} = D$.

A different limiting case occurs when individuals are sorted across neighborhoods according to outcomes rather than groups, that is, $Y_n \in \{0,1\}$ for all n. Here too, Lemma 2 implies that the MOB interval collapses to a single point, which again must coincide with both D and D_{NM}. However, unless neighborhoods also happen to be perfectly homogeneous with respect to group, we will not have $\gamma = 1$, so from Lemma 1, we know that $D_{ER} \neq D_{NM} = D$. Hence, whereas selection of individuals into neighborhoods according to group membership guarantees unbiasedness for all three of the canonical approaches, selection into neighborhoods according to outcomes implies that only the NM and the MOB will be unbiased.

Beyond these limiting cases, the informativeness of the MOB bounds for D can be related to the degree of information about X and Y lost through aggregation, as illustrated by the following proposition.

https://doi.org/10.1017/pan.2026.10048 Published online by Cambridge University Press

Proposition 5 (Maximum width of the MOB interval). *The MOB width is bounded as follows:*

$$0 \leq D^{+}_{MOB} - D^{-}_{MOB} \leq 2\min\left\{1-\gamma, \frac{\mathrm{Var}(Y)}{\mathrm{Var}(X)}\left(1-\frac{\mathrm{Var}(Y_N)}{\mathrm{Var}(Y)}\right)\right\}.$$

Proposition 5 provides an upper bound on the width of the MOB interval; the actual interval may be more informative in a particular application. Interestingly, the proposition shows that D can be recovered with relative precision when *either* γ or $\frac{\mathrm{Var}(Y_N)}{\mathrm{Var}(Y)}$ is close to 1. This is because extreme values of X_n and Y_n can independently constrain the range of feasible group means for a given neighborhood. In contrast, when aggregation eliminates most of the information about both X and Y, then the bounds on D may be relatively wide.

4. Signing the conditional associations

In this section, we study how assumptions about the sign of the within- and between-group associations can aid with identification.

We begin with the within-group association. When $\delta_W \geq 0$, we know from Proposition 1 that ER will overstate the true difference in group means, that is, $D_{ER} > D$. We can tighten our preexisting bounds for D (i.e., the MOB interval) by intersecting it with this new one-sided bound on D to obtain

$$\delta_W \geq 0 \implies D \in \left[D^{-}_{MOB}, \mathrm{Min}\left\{D^{+}_{MOB}, D_{ER}\right\}\right]. \tag{2}$$

The direction of the ER bias is reversed when $\delta_W \leq 0$:

$$\delta_W \leq 0 \implies D \in \left[\mathrm{Max}\left\{D^{-}_{MOB}, D_{ER}\right\}, D^{+}_{MOB}\right]. \tag{3}$$

We can proceed analogously to derive bounds based on the sign of the between-group association using Proposition 2:

$$\delta_B \geq 0 \implies D \in \left[D_{NM}, D^{+}_{MOB}\right] \tag{4}$$

and

$$\delta_B \leq 0 \implies D \in \left[D^{-}_{MOB}, D_{NM}\right]. \tag{5}$$

Because the NM estimate is guaranteed to be within the MOB interval (Corollary 3), the bounds in (4) and (5) have a simpler form than those appearing in (2) and (3).

Just as knowledge of the sign of one conditional association can aid in identification, knowledge of the signs of both δ_W and δ_B can further narrow the feasible range of values that D can take on. For example, if δ_B and δ_W are both known to be positive, one can intersect the bounds in (2) and (4) to obtain:

$$\delta_W \geq 0 \text{ and } \delta_B \geq 0 \implies D \in \left[D_{NM}, \mathrm{Min}\left\{D^{+}_{MOB}, D_{ER}\right\}\right]. \tag{6}$$

Table 1 applies the logic reflected in Equations (2)–(6) to describe what can be concluded about D based on what is assumed about δ_W and δ_B. When nothing is assumed (top-left cell), the identification region is equal to the MOB interval. At the other extreme, when either δ_W or δ_B is assumed to be zero (top row or rightmost column), the bounds reduce to either D_{NM} or D_{ER}.

Theorem 1 (Identification of *D* with monotone EI). *Suppose* $\gamma \in (0,1)$. *Table 1 provides sharp bounds for D based on the specified assumptions.*

Whereas the bounds in Equations (2)–(6) 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. We refer to this condition, that is,

https://doi.org/10.1017/pan.2026.10048 Published online by Cambridge University Press

Table 1. Identification of the difference in group means

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

that $\delta_B \cdot \delta_W \geq 0$, as *contextual reinforcement* and illustrate the type of reasoning that might support it in Section 7.

Proposition 6. *(Identification of group differences with contextual reinforcement)*
If contextual reinforcement holds (i.e., if $\delta_W \cdot \delta_B \geq 0$), then either:
(i)

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

(ii)

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

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

The following corollary follows directly.

Corollary 5. *If contextual reinforcement holds, then*

$$Sign(D) = Sign\left(D_{ER}\right) = Sign\left(D_{NM}\right).$$

Thus, when contextual reinforcement holds, a researcher can look to the sign of D_{NM} or D_{ER} to learn which bounds apply, as well as to identify the sign of D.

Finally, like the MOB interval, the width of the contextual reinforcement bounds is shaped by γ; this is because γ controls the divergence between D_{NM} and D_{ER} (see Lemma 1). Hence, when γ is close to 1, $D_{ER} - D_{NM}$ will be small, and the contextual reinforcement bounds are guaranteed to be relatively informative about D.

5. Identification of group mean levels

This section maps our results for identifying the difference in group means, D, to the task of identifying the group means themselves, Y^0 and Y^1. Each of the canonical EI approaches we have considered for estimating D can also be applied to estimate the levels of the group means.

With ER, this amounts to incorporating the estimated intercept term. Doing so yields the following estimands:

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

and

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

https://doi.org/10.1017/pan.2026.10048 Published online by Cambridge University Press

Table 2. Identification of the group 1 mean

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

For the NM, the asymptotic limits of the group mean estimates are given by

$$Y^1_{NM} = \frac{\mathbb{E}[X_N\, Y_N]}{\mathbb{E}[X_N]}$$

$$Y^0_{NM} = \frac{\mathbb{E}\left[(1-X_N)\, Y_N\right]}{\mathbb{E}\left[1-X_N\right]}.$$

Finally, for the MOB, intervals corresponding to the group mean levels are given by

$$Y^1 \in \left[Y^{1-}_{MOB}, Y^{1+}_{MOB}\right]$$

and

$$Y^0 \in \left[Y^{0-}_{MOB}, Y^{0+}_{MOB}\right],$$

where

$$Y^{1-}_{MOB} := \frac{\mathbb{E}[\max\{Y_N - (1-X_N), 0\}]}{\mathbb{E}[X]}\,,\; Y^{1+}_{MOB} := \frac{\mathbb{E}\left[\min\{Y_N, X_N\}\right]}{\mathbb{E}[X]}$$

and

$$Y^{0-}_{MOB} := \frac{\mathbb{E}[\max\{Y_N - X_N, 0\}]}{1-\mathbb{E}[X]}\,,\; Y^{0+}_{MOB} := \frac{\mathbb{E}[\min\{Y_N, 1-X_N\}]}{1-\mathbb{E}[X]}.$$

We can relate the biases of the ER and NM estimates of Y^0 and Y^1 to δ_W and δ_B in a similar manner as we did for the corresponding estimators for D, as formalized in the following proposition.

Proposition 7 (Bias of ER and NM for group means).

$$Y^1_{ER} - Y^1 = \frac{\delta_W\,(1-\mathbb{E}[X])}{\text{Var}(X_N)}\,,\; Y^0_{ER} - Y^0 = -\frac{\delta_W\,\mathbb{E}[X]}{\text{Var}(X_N)}$$

$$Y^1_{NM} - Y^1 = -\frac{\delta_B}{\mathbb{E}[X]}\,,\; Y^0_{NM} - Y^0 = \frac{\delta_B}{1-\mathbb{E}[X]}.$$

We can now use the results in Proposition 7 to obtain bounds for the group means just as we used Propositions 1 and 2 to bound D in the prior section.

Theorem 2 (Identification of Y with monotone EI). *Suppose $\gamma \in (0,1)$. Tables 2 and 3 provide sharp bounds for Y^1 (Table 2) and Y^0 (Table 3) under the specified assumptions.*

Finally, if contextual reinforcement is assumed, one can first use Proposition 6 to learn the (common) sign of δ_B and δ_W, and then use that information to determine the applicable bounds on Y^0 and Y^1 from Tables 2 and 3.

https://doi.org/10.1017/pan.2026.10048 Published online by Cambridge University Press

Table 3. Identification of the group 0 mean

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

6. Neighborhood-specific group means and differences

Our focus so far has been on group means for the overall population. In this section, we consider group means for individuals in a specific neighborhood, $Y^x_n = \mathbb{E}[Y \mid N = n, X = x]$, or the difference in group means for individuals within that neighborhood, $D_n = Y^1_n - Y^0_n$.

To start, recall that the MOB provides sharp bounds on Y^x_n and D_n for each n and x (Cho and Manski 2008; Duncan and Davis 1953)[13]:

$$Y^x_n \in \left[Y^{x-}_{MOB,n}, Y^{x+}_{MOB,n}\right]$$

and

$$D_n \in \left[D^{-}_{MOB,n}, D^{+}_{MOB,n}\right].$$

As above, we can narrow the MOB interval 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, previously studied by Manski (2018). Define

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

Whereas δ_B describes the *average* between-group variation within all neighborhoods, $\delta_{B,n}$ refers to the between-group variation for a specific neighborhood n. Because X is binary, $\delta_{B,n} = D_n \cdot \text{Var}(X|N = n)$. It follows that

$$\delta_{B,n} \geq 0 \iff \text{(i) } D_n \geq 0, \text{ (ii) } Y^1_n \geq Y_n \text{ and (iii) } Y^0_n \leq Y_n.$$

As in Section 4, intersecting these one-sided bounds with the applicable MOB interval yields sharp bounds (Manski 2018).

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 studied 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:

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

It will be convenient to assume that if two neighborhoods have the same group prevalence, they also share the same group means[14]:

$$Y^x_n = \mu_x(X_n) \text{ for all } x, n.$$

We also restrict our focus to neighborhoods with group prevalence values for which the first derivative of $\mu_x(\cdot)$ exists for each x, and we denote those derivatives by $\mu'_x(\cdot)$.[15]

[13] These intervals are defined by $Y^{1+}_{MOB,n} = \min\left\{\frac{Y_n}{X_n}, 1\right\}$, $Y^{0-}_{MOB,n} = \frac{Y_n - \min\{Y_n, X_n\}}{1 - X_n}$, $Y^{0+}_{MOB,n} = \min\left\{\frac{Y_n}{1-X_n}, 1\right\}$, $Y^{1-}_{MOB,n} = \frac{Y_n - \min\{Y_n, 1-X_n\}}{X_n}$, $D^{+}_{MOB,n} = Y^{1+}_{MOB,n} - Y^{0-}_{MOB,n}$, and $D^{-}_{MOB,n} = Y^{1-}_{MOB,n} - Y^{0+}_{MOB,n}$.

[14] We consider the case where this assumption does not hold in the Supplementary Material.

[15] Formally, we define $\mu'_x(x_n) = \lim_{u \to 0} \frac{1}{u}\left(\mathbb{E}[Y \mid X = x, X_n = x_n + u] - \mathbb{E}[Y \mid X = x, X_n = x_n]\right)$.

https://doi.org/10.1017/pan.2026.10048 Published online by Cambridge University Press

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

$$\delta_{W,n} = X_n \mu_1'(X_n) + (1 - X_n)\, \mu_0'(X_n).$$

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 individuals in neighborhoods with a particular level 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 at a particular group prevalence level.

Although μ_0 and μ_1 are unobserved, we do observe the overall conditional expectation function of Y_N given X_N, which is the mixture of the group-specific conditional expectations functions

$$\mu(x_n) = x_n \mu_1(x_n) + (1 - x_n)\, \mu_0(x_n). \tag{7}$$

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

Differentiating (7) yields

$$\begin{aligned} \mu'(x_n) &= \mu_1(x_n) - \mu_0(x_n) + x_n \mu_1'(x_n) + (1 - x_n)\, \mu_0'(x_n) \\ &= D_n + \delta_{W,n}. \end{aligned} \tag{8}$$

Equation (8) links D_n to the derivative of the conditional expectation function, which is observable, and to $\delta_{W,n}$, which is not. An assumption about the sign of $\delta_{W,n}$ can therefore be combined with observation of $\mu'(X_n)$ to bound D_n, and by extension, Y_n^0 and Y_n^1, as formalized in the following proposition.

Proposition 8. (i) *If* $\delta_{W,n} \geq 0$*:*

$$\begin{aligned} D_n &\in \left[D^-_{MOB,n}, \min\left\{\mu'(X_n), D^+_{MOB,n}\right\}\right], \\ Y_n^1 &\in \left[Y^{1-}_{MOB,n}, \min\left\{Y^{1+}_{MOB,n}, Y_n + (1 - X_n) \cdot \mu'(X_n)\right\}\right], \\ \textit{and } Y_n^0 &\in \left[\max\left\{Y^{0-}_{MOB,n}, Y_n - X_n \cdot \mu'(X_n)\right\}, Y^{0+}_{MOB,n}\right]. \end{aligned}$$

(ii) *If* $\delta_{W,n} \leq 0$*:*

$$\begin{aligned} D_n &\in \left[\max\left\{\mu'(X_n), D^-_{MOB,n}\right\}, D^+_{MOB,n}\right], \\ Y_n^1 &\in \left[\max\left\{Y^{1-}_{MOB,n}, Y_n + (1 - X_n) \cdot \mu'(X_n)\right\}, Y^{1+}_{MOB,n}\right], \\ \textit{and } Y_n^0 &\in \left[Y^{0-}_{MOB,n}, \min\left\{Y^{0+}_{MOB,n}, Y_n - X_n \cdot \mu'(X_n)\right\}\right]. \end{aligned}$$

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

Intuitively, the bounds in Proposition 8 are the neighborhood-specific analog to the population-level bounds based on the sign of δ_W that were derived in Sections 4 and 5. The main difference is that the neighborhood-specific version relies on the derivative of the conditional expectation function $\mu'(x_n)$, which serves as a local version of the ER slope.

These bounds can be further narrowed by combining assumptions on the signs of the (local) conditional associations, as in the earlier sections.

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) conditional associations share the same sign as one another. This is the local analog to contextual reinforcement and yields an analogous identification result.

https://doi.org/10.1017/pan.2026.10048 Published online by Cambridge University Press

Proposition 9. *Suppose that* $\delta_{B,n} \cdot \delta_{W,n} \geq 0$. *Then:*

(i) $\mu'(X_n) \geq 0$ *implies*

$$D_n \in \left[0, \min\left\{\mu'(X_n), D^{+}_{MOB,n}\right\}\right],$$
$$Y^1_n \in \left[Y_n, \min\left\{Y_n + (1-X_n)\mu'(X_n), Y^{1+}_{MOB}\right\}\right],$$
$$\text{and } Y^0_n \in \left[\max\left\{Y^{0-}_{MOB,n}, Y_n - X_n\mu'(X_n)\right\}, Y_n\right].$$

(ii) $\mu'(X_n) \leq 0$ *implies*

$$D_n \in \left[\max\left\{\mu'(X_n), D^{-}_{MOB,n}\right\}, 0\right],$$
$$Y^1_n \in \left[\max\left\{Y^{1-}_{MOB,n}, Y_n + (1-X_n)\mu'(X_n)\right\}, Y_n\right],$$
$$\text{and } Y^0_n \in \left[Y_n, \min\left\{Y_n - X_n\mu'(X_n), Y^{0+}_{MOB,n}\right\}\right].$$

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

(iv) $Sign(D_n) = Sign\left(\mu'(X_n)\right)$.

We will illustrate these bounds in our empirical application, described in the following section.

7. Empirical application

To illustrate an application of monotone EI, we investigate partisan polarization in COVID-19 vaccine uptake. This question has been of interest to policymakers, academics and the media (Collins 2024; Jones and McDermott 2022; Milligan 2021). Most research in the U.S. context has had to rely on EI, as joint individual data on vaccination status and partisan membership are not available for a large nationally representative set of individuals (e.g., Albrecht 2022; Ye 2023).

Our primary dataset 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 vaccinations 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).[16] 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 Republican voters relative to Democratic and third-party voters.

Figure 1 plots the binned county-level data. The figure shows a clear downward trend: counties with higher Republican vote share tend to have lower vaccination rates. The pattern is consistent with the possibility that Republican voters are vaccinated at lower rates than Democratic and third-party voters. However, this interpretation is potentially subject to the ecological fallacy; counties with higher Republican voter share may have lower vaccination rates for reasons unrelated to partisan composition (Ye 2023). Indeed, the MOB interval for the partisan vaccination gap ranges from −77.4 to 52.8 percentage points. Without further assumptions, the county-level data do not provide much information about the magnitude or even direction of partisan differences in vaccinations. Point identification approaches also diverge sharply: the NM and ER imply respective partisan vaccination gaps of −5.5 and −47.9 percentage points.

To sharpen identification through monotone EI, consider first the sign of the between-group association, $\mathbb{E}[\text{Cov}(Y, X|N)]$, which here refers to differences in vaccination uptake between Republican and Democratic voters 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

[16] We 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 are likely to vary widely across different counties around the country.

https://doi.org/10.1017/pan.2026.10048 Published online by Cambridge University Press

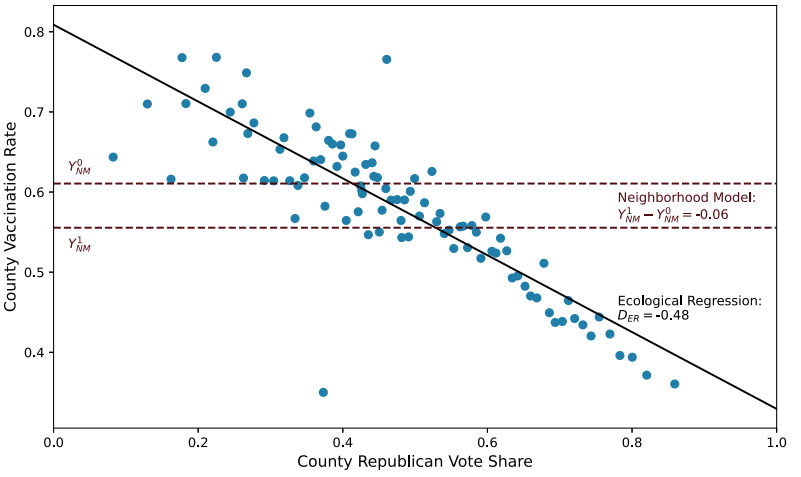


Figure 1. County vaccination rate by county Republican vote share. *Note:* The figure reports county-level COVID-19 vaccination rates by the share of voters in the county who voted for the Republican candidate in the 2020 presidential election. Counties are grouped into 100 equal-population bins. The neighborhood model estimates for the mean vaccination rates among Republican and non-Republican voters are, respectively, denoted by the lower and upper red dotted lines. The ecological regression line is in black.

information sources that are not fully mediated through neighborhood (Iyengar and Hahn 2009; Peterson, Goel, and Iyengar 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 as vaccinations were made available (Albrecht 2022; Gollwitzer *et al.* 2020; Hornsey *et al.* 2020).

The within-group association, $\mathbb{E}[\mathrm{Cov}(Y, X_N|X)]$, refers to differences in vaccination uptake among individuals who vote for the same political party who live in neighborhoods with differing concentrations of Republican voters. Like the between-group association, there is some reason to expect the within-group association to be negative. For example, Republican counties tend to be lower income and more rural (Figure S2 in the Supplementary Material), which are factors associated with lower access to public health services like vaccinations (Hernandez *et al.* 2022; Parolin and Lee 2022; Sun and Monnat 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 (Klaesson, Lobo, and Mellander 2023; Sato and Takasaki 2019), and in mostly 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 mandated 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 Democratic or Republican.[17]

[17] Analogously, 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.

https://doi.org/10.1017/pan.2026.10048 Published online by Cambridge University Press

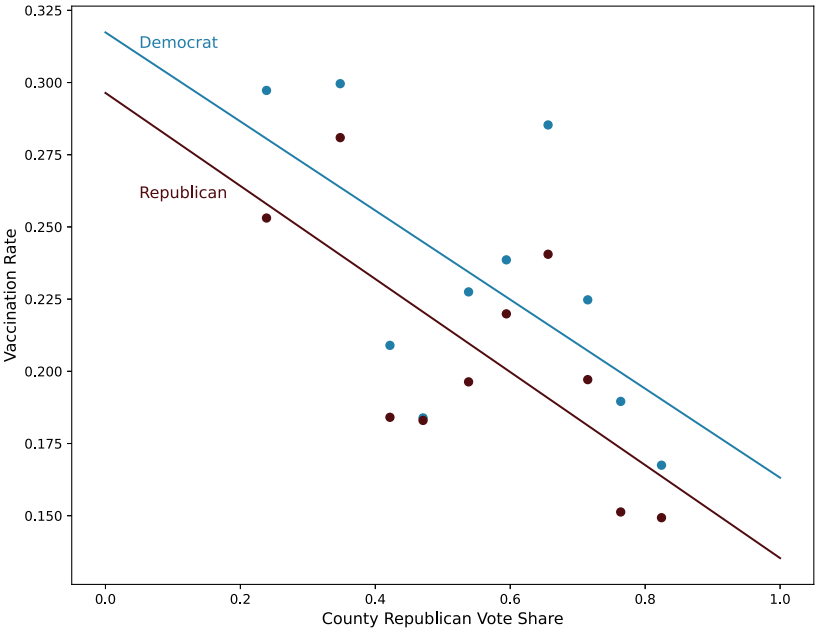


Figure 2. Vaccination rate by vote share and political party membership. *Note:* The figure uses the matched auxiliary dataset to report COVID-19 vaccination rates by county-level Republican vote share for individuals registered as Republicans (red) and individuals registered as Democrats (blue). Individuals 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.

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.[18] We construct this dataset by matching a national dataset of voter registration records (L2 2024), which contains individual-level data on political party, to a large dataset of electronic health care records (Balraj *et al.* 2023), which contains individual-level data on COVID-19 vaccination status. Appendix E of the Supplementary Material 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.[19]

Figure 2 uses the auxiliary data to plot vaccination rates by (binned) county-level Republican vote share, separately for registered 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 S3 in the Supplementary Material). Based on these results, we adopt the contextual reinforcement assumption to interpret the county-level analyses.[20]

[18]For this validation analysis, we measure partisanship using party registration rather than voting behavior, as we obviously cannot observe the latter at the individual level due to secret ballots.

[19]While 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).

[20]A limitation of this analysis for assessing contextual reinforcement is that the set of individuals 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 S4 in the Supplementary Material) and when we re-weight the auxiliary data to match known county-level political party composition (Table S5 in the Supplementary Material).

https://doi.org/10.1017/pan.2026.10048 Published online by Cambridge University Press

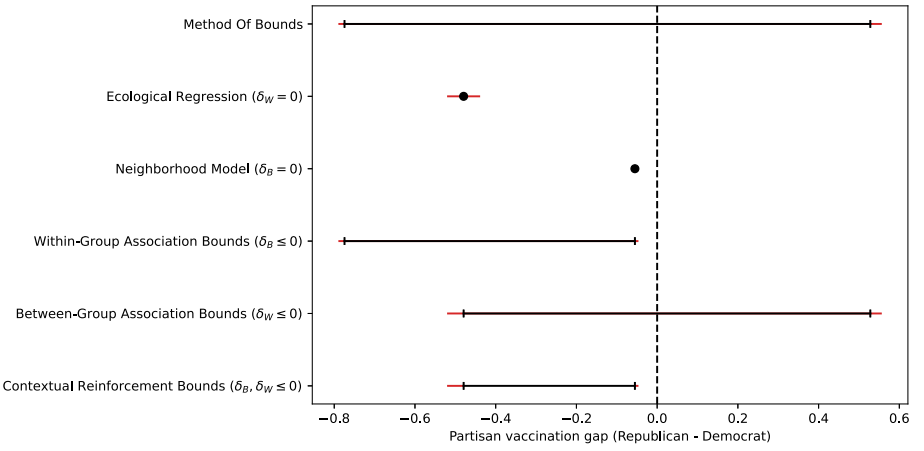


Figure 3. Identification of partisan vaccination gap using monotone EI. *Note:* The partisan vaccination gap is defined as the proportion of vaccinated Democratic voters (and third-party supporters) subtracted from the proportion of vaccinated Republican voters. A negative gap indicates that a lower share of Republican voters 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 1,000 bootstrap replicates.

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 NM and ER estimators. As shown in Figure 3, we find that this difference is positive ($p < 0.01$), implying that Democrats are vaccinated at higher rates than Republicans. In turn, the contextual reinforcement bounds from Proposition 6 imply that the partisan vaccination gap is between −47.9 and −5.5 percentage points (95% CI: −51.8 to −5.1).

Figure 3 summarizes our results under various identifying assumptions.[21] Imposing $\delta_B \leq 0$ tightens the MOB interval by 45%, whereas imposing $\delta_W \leq 0$ tightens the MOB interval by 22%. Imposing contextual reinforcement, our preferred assumption, implies 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 MOB.

Finally, monotone EI may also be used to more precisely bound county-specific vaccine disparities. For purposes of this exercise, we assume that contextual reinforcement holds locally for each county: that is, $\delta_{B,n} \cdot \delta_{W,n} \geq 0$ for all n. To calculate the implied bounds from Proposition 9, we estimate $\mu'(X_n)$ using population-weighted local linear approximation (Figure S3 in the Supplementary Material). For Contra Costa County, which leans left-of-center politically, we estimate that the vaccination rate for Republican voters is between 0.50 and 0.77 and that the vaccination rate for Democratic voters 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 MOB (see Figure 4).

8. Conclusion

We study the classic statistical challenge of EI. Our results clarify the biases associated with, and relationship among, several canonical EI methods for identifying group means and differences at the population level. We use those results to derive a partial identification approach based on assumptions

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).

[21] Figure S1 in the Supplementary Material provides corresponding results for mean vaccination rates separately among Democrats and Republicans.

https://doi.org/10.1017/pan.2026.10048 Published online by Cambridge University Press

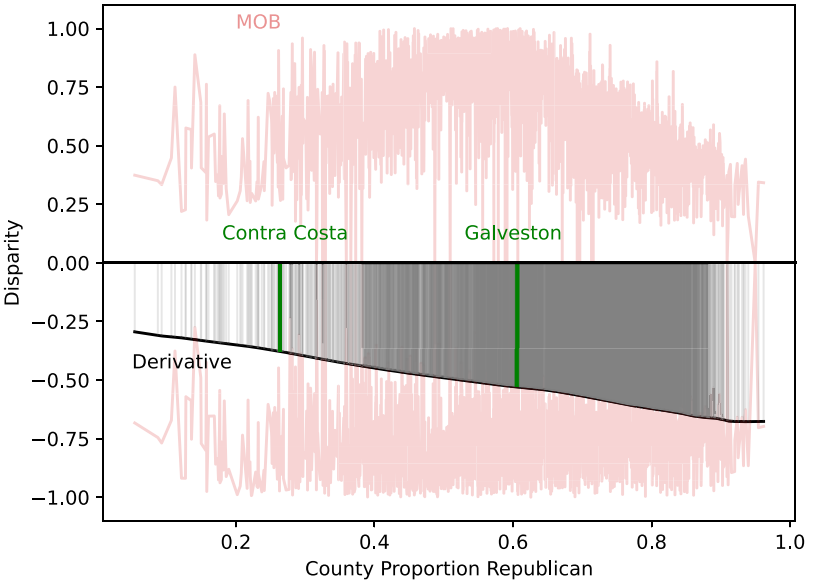


Figure 4. County-specific bounds. *Note:* This figure shows the method of bounds for the partisan vaccination gap for each county (red). The county-level bounds, based on the assumptions in Proposition 9, are shown in gray, and the estimated derivative of the conditional expectation function of the vaccination 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.

about the sign of either or both of the conditional associations between the outcome of interest and group membership or neighborhood group composition. Although our approach requires additional structure relative to assumption-free tools like the MOB, the payoff to that additional structure can be substantially tighter bounds for the parameter of interest.

For researchers seeking to apply our approach, Tables 1–3 can serve as a reference for the applicable sharp bounds that apply based on what, if anything, one is willing to assume about the signs of δ_B and δ_W. Once the relevant table cell is determined, calculating the bounds is straightforward because they are derived from ER, the NM and the MOB endpoints. Those statistics are straightforward to calculate individually or can be jointly calculated using the MonotoneEI R package.

Supplementary material. The supplementary material for this article can be found at https://doi.org/10.1017/pan.2026.10048.

Data availability statement. Replication code for this article has been published on Code Ocean at https://codeocean.com/capsule/4440816/tree/v1.

Elzayn, H., Goldin, J., Guage, C., Ho, D. E., and Morton, C. M. 2026. Monotone Ecological Inference. Code Ocean, Version 1.0. https://doi.org/10.24433/CO.6554779.v1.

Acknowledgements. For helpful comments and suggestions, we are grateful to Kosuke Imai, Gary King, Charles Manski, Shiying Hao and Derek Ouyang.

Funding statement. This work was supported by the Stanford Institute for Human-Centered Artificial Intelligence.

Competing interests. The authors have no competing interests to disclose.

References

Albrecht, D. 2022. "Vaccination, Politics and COVID-19 Impacts." *BMC Public Health* 22 (1): 96.

Ansolabehere, S., and D. Rivers. 1995. "Bias in Ecological Regression." Technical report, Stanford University and Massachusetts Institute of Technology. https://web.stanford.edu/class/polisci353/2004spring/reading/erbias.pdf.

https://doi.org/10.1017/pan.2026.10048 Published online by Cambridge University Press

Balraj, D., A. Vala, S. Hao, M. Philofsky, A. Tsvetkova, E. Trach, S. P. Narra, O. Zhuk, M. Shamkhorskaya, J. Singer, J. Mesterhazy, S. Datta, I. Chu, and D. Rehkopf. 2023. "American Family Cohort, a Data Resource Description." Preprint, arXiv:2309.13175.
Centers for Disease Control and Prevention. 2024. "COVID-19 Vaccinations in the United States, County."
Chen, J., N. Kallus, X. Mao, G. Svacha, and M. Udell. 2019. "Fairness under Unawareness: Assessing Disparity when Protected Class Is Unobserved." In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 339–348. New York: Association for Computing Machinery.
Cho, W. T., and Manski, C. F. 2008. "24 Cross-Level/Ecological Inference." In *The Oxford Handbook of Political Methodology*, edited by J. M. Box-Steffensmeier, H. E. Brady, and D. Collier, 530–569. Oxford: Oxford University Press.
Collins, F. 2024. "Why Didn't Facing a Common Enemy Bring us Together?" The Atlantic.
Cross, P. J., and C. F. Manski. 2002. "Regressions, Short and Long." *Econometrica* 70 (1): 357–368.
Duncan, O. D., and B. J. Davis. 1953. "An Alternative to Ecological Correlation." *American Sociological Review* 18: 665.
Elzayn, H., E. Smith, T. Hertz, C. Guage, A. Ramesh, R. Fisher, D. E. Ho, and J. Goldin. 2025. "Measuring and Mitigating Racial Disparities in Tax Audits." *The Quarterly Journal of Economics* 140 (1): 113–163.
Freedman, D. A., S. P. Klein, J. Sacks, C. A. Smyth, and C. G. Everett. 1991. "Ecological Regression and Voting Rights." *Evaluation Review* 15 (6): 673–711.
Gollwitzer, A., et al. 2020. "Partisan Differences in Physical Distancing are Linked to Health Outcomes during the COVID-19 Pandemic." *Nature Human Behaviour* 4 (11): 1186–1197.
Goodman, L. A. 1953. "Ecological Regressions and the Behavior of Individuals." *American Sociological Review* 18 (6): 663–664.
Greiner, J., and K. M. Quinn. 2009. "R×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.
Hao, S., D. H. Rehkopf, E. Velasquez, A. Vala, A. W. Bazemore, and R. L. Phillips Jr. 2023. "COVID-19 Vaccine Strategy Left Small Primary Care Practices on the Sidelines." *Health Affairs* 42 (8): 1147–1151.
Hernandez, I., S. Dickson, S. Tang, N. Gabriel, L. A. Berenbrok, and J. Guo. 2022. "Disparities in Distribution of COVID-19 Vaccines across US Counties: A Geographic Information System–Based Cross-Sectional Study." *PLoS Medicine* 19 (7): e1004069.
Hornsey, M. J., M. Finlayson, G. Chatwood, and C. T. Begeny. 2020. "Donald Trump and Vaccination: The Effect of Political Identity, Conspiracist Ideation and Presidential Tweets on Vaccine Hesitancy." *Journal of Experimental Social Psychology* 88: 103947.
Horowitz, J. L., and C. F. Manski. 1995. "Identification and Robustness with Contaminated and Corrupted Data." *Econometrica: Journal of the Econometric Society* 63 (2): 281–302.
Howard-Williams, M., R. H. Soelaeman, L. S. Fischer, R. McCord, R. Davison, and C. Dunphy. 2022. "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* 3 (10): e223810.
Imbens, G. W., and C. F. Manski. 2004. "Confidence Intervals for Partially Identified Parameters." *Econometrica* 72 (6): 1845–1857.
Iyengar, S., and K. S. Hahn. 2009. "Red Media, Blue Media: Evidence of Ideological Selectivity in Media Use." *Journal of Communication* 59 (1): 19–39.
Jiang, W., G. King, A. Schmaltz, and M. A. Tanner. 2020. "Ecological Regression with Partial Identification." *Political Analysis* 28 (1): 65–86.
Jones, D. R., and M. L. McDermott. 2022. "Partisanship and the Politics of COVID Vaccine Hesitancy." *Polity* 54 (3): 408–434.
Kallus, N., X. Mao, and A. Zhou. 2022. "Assessing Algorithmic Fairness with Unobserved Protected Class Using Data Combination." *Management Science* 68 (3): 1959–1981.
King, G. 1997. *A Solution to the Ecological Inference Problem: Reconstructing Individual Behavior from Aggregate Data*. Princeton: Princeton University Press.
King, G., M. A. Tanner, and O. Rosen. 2004. *Ecological Inference: New Methodological Strategies*. Cambridge: Cambridge University Press.
Klaesson, J., J. Lobo, and C. Mellander. 2023. "Social Interactions and COVID-19 Vaccine Hesitancy: Evidence from a Full Population Study in Sweden." *PLoS One* 18 (11): e0289309.
L2. 2024. "Data Mapping." L2, Inc.
Lewis, J. B. 2001. "Understanding King's Ecological Inference Model a Method-of-Moments Approach." *Historical Methods: A Journal of Quantitative and Interdisciplinary History* 34 (4): 170–188.
Li, S., V. Litvin, and C. F. Manski. 2023. "Partial Identification of Personalized Treatment Response with Trial-Reported Analyses of Binary Subgroups." *Epidemiology* 34 (3): 319–324.
Manski, C. F. 2018. "Credible Ecological Inference for Medical Decisions with Personalized Risk Assessment." *Quantitative Economics* 9 (2): 541–569.
Manski, C. F., and J. V. Pepper. 2000. "Monotone Instrumental Variables: With an Application to the Returns to Schools." *Econometrica* 68 (4): 997–1010.
Manski, C. F., and J. V. Pepper. 2018. "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.

https://doi.org/10.1017/pan.2026.10048 Published online by Cambridge University Press

McCartan, C., R. Fisher, J. Goldin, D. E. Ho, and K. Imai. 2024. "Estimating Racial Disparities when Race Is Not Observed." Technical report, National Bureau of Economic Research.
Milligan, S. 2021. "A Deadly Political Divide." US News and World Report.
MIT Election Data and Science Lab. 2020. "County Presidential Election Returns 2000-2020." Harvard Dataverse.
Molinari, F. 2010. "Missing Treatments." *Journal of Business & Economic Statistics* 28 (1): 82–95.
Parolin, Z., and E. K. Lee. 2022. "The Role of Poverty and Racial Discrimination in Exacerbating the Health Consequences of COVID-19." *The Lancet Regional Health–Americas* 7: 100178.
Patterson, S. 2022. "The Politics of Pandemics: The Effect of Stay-at-Home Orders on COVID-19 Mitigation." *State Politics & Policy Quarterly* 22 (1): 1–23.
Peterson, E., S. Goel, and S. Iyengar. 2021. "Partisan Selective Exposure in Online News Consumption: Evidence from the 2016 Presidential Campaign." *Political Science Research and Methods* 9 (2): 242–258.
Phillips, R. 2017. "The PRIME Registry Helps Thousands of Primary Care Clinicians Liberate EHR Data and Prepare for MIPS." *Journal of the American Board of Family Medicine* 30 (4): 559.
Rambachan, A., and J. Roth. 2023. "A more Credible Approach to Parallel Trends." *Review of Economic Studies* 90 (5): 2555–2591.
Robinson, W. S. 1950. "Ecological Correlations and the Behavior of Individuals." *American Sociological Review* 15 (3): 351–357.
Roos Jr., L. L., A. Wajda, and J. Patrick Nicol. 1986. "The Art and Science of Record Linkage: Methods that Work with Few Identifiers." *Computers in Biology and Medicine* 16 (1): 45–57.
Sato, R., and Y. Takasaki. 2019. "Peer Effects on Vaccination Behavior: Experimental Evidence from Rural Nigeria." *Economic Development and Cultural Change* 68 (1): 93–129.
Sun, Y., and S. M. Monnat. 2022. "Rural-Urban and within-Rural Differences in COVID-19 Vaccination Rates." *The Journal of Rural Health* 38 (4): 916–922.
Wakefield, J. 2004. "Ecological Inference for 2 × 2 Tables (with Discussion)." *Journal of the Royal Statistical Society Series A: Statistics in Society* 167 (3): 385–445.
Ye, X. 2023. "Exploring the Relationship between Political Partisanship and COVID-19 Vaccination Rate." *Journal of Public Health* 45 (1): 91–98.

https://doi.org/10.1017/pan.2026.10048 Published online by Cambridge University Press

Cite this article: Elzayn, H., Goldin, J., Guage, C., Ho, D. and Morton, C., (2026). Monotone Ecological Inference. *Political Analysis*, 1–20. https://doi.org/10.1017/pan.2026.10048