Ecological Inference with Instrumental Variables

Jörg L. Spenkuch
Northwestern University

First Draft: February 2017
This Version: July 2017

Abstract

This note observes that instrumental variables techniques solve the ecological inference problem under essentially the same set of conditions under which they allow researchers to draw conclusions about causality. It, therefore, provides a theoretical justification for generalizing from causally identified estimates to the behavior of individuals, even if only aggregate data are available.

*I have benefitted from helpful conversations with Pablo Montagnes and Philipp Tillmann. Correspondence can be addressed to the author at MEDS Department, Kellogg School of Management, 2211 Campus Drive, Evanston, IL 60208, or by email: j-spenkuch@kellogg.northwestern.edu.

1. Introduction

Much empirical work in the social sciences aims to estimate causal effects, i.e., the (probabilistic) change in some outcome y that is attributable to a change in x, holding all else equal. Even with access to modern microdata, producing valid *ceteris paribus* comparisons often presents a challenge. To complicate matters even further, many interesting questions can only be addressed with aggregate data. If the behavioral relationship between x and y is theoretically defined at the individual level, then researchers using such data to draw inferences about causal effects must also worry about the ecological fallacy.

The ecological fallacy arises because ecological correlations, i.e., correlations in aggregated data, may systematically deviate from their individual-level counterparts. Robinson (1950), for instance, studies the relationship between literacy and nativity in the 1930 Census. The true correlation between both variables on the individual level is 0.1, meaning that native-born Americans were more likely to be literate than the foreign-born. On the state level, however, the correlation between the share of native-borns and the share of adults who are literate is -0.5. Based on such contradictory results, Robinson (1950, p. 357) concludes that aggregate correlations are "meaningless" and should not be used to draw inferences about individuals. Yet, in many applications—especially historical ones—microdata are simply not available. Scholars studying these settings cannot help but rely on aggregate information to learn about individual behavior.

In light of this predicament, a large subsequent literature proposes ways to avoid the ecological fallacy, including imposing homogeneity assumptions (Goodman 1953, 1959), calculating deterministic bounds (Duncan and Davis 1953), as well as procedures that combine deterministic and probabilistic information (King 1997; King et al. 2004). All of these methods work well in some situations but are known to yield either uninformative or, worse, misleading estimates in others (see, e.g., Freedman et al. 1991, 1998; Gelman et al. 2001). In addition, existing approaches were designed for *descriptive* rather than causal inferences.

That is, they recover raw individual-level differences, *without* holding all else equal.[1] Even extensions of King's method that allow for covariate adjustment cannot account for unobserved confounders (King et al. 1999, 2004). If there are any (unmeasured) variables that correlate with both x and y, then estimates from conventional approaches to the ecological inference problem do not have a causal interpretation. They do not correspond to valid *ceteris paribus* comparisons.

At first blush, it may seem that recovering causal, individual-level estimates from aggregate data is a near-hopeless endeavor. This note shows that such a conclusion is premature. In fact, as demonstrated below, standard instrumental variables techniques solve the ecological inference problem under almost the same assumptions that would be required to claim causality in microdata. As a consequence, drawing inferences about individual-level effects from aggregate data is often *not* more difficult than obtaining causal estimates in the first place. In many applications, addressing the problem of causal identification is sufficient to extrapolate from aggregate data to the behavior of individuals.

2. Preliminaries: From Individual-Level Behavior to Ecological Correlations

To clarify why aggregate correlations can generally not be interpreted as individual-level effects, and to convey the intuition behind ecological inference with instrumental variables (EI-IV), it is useful to start with a statistical description of individual behavior. To this end, consider a data-generating process of the following form:

$$(1) \qquad y_i = \beta x_i + \mathbf{W}_i{}'\theta + \mathbf{Q}_a{}'\gamma + \epsilon_i,$$

where i indexes individuals, and y_i denotes the outcome of interest. $\mathbf{W}_i$ is a vector of all individual-level covariates that affect behavior, while $\mathbf{Q}_a$ contains contextual factors, which

[1] For instance, conventional methods for ecological inference are sometimes used to determine whether, in a particular election, blacks were more or less likely to turn out to vote than whites. These methods do not statistically adjust for the well-known socioeconomic disparities between both groups. A *ceteris paribus* comparison, however, controls for all differences, i.e., it compares turnout among otherwise identical blacks and whites. Depending on the research question, unadjusted individual-level differences may or may not be the object of interest.

are the same for everyone in aggregate unit a. For example, depending on the application, $\mathbf{W}_i$ may include age, gender, income, etc., whereas $\mathbf{Q}_a$ might consist of controls for socioeconomic inequality, segregation, or, say, the quality of schools. Since equation (1) includes all variables that have a systematic impact on y_i, x_i and ϵ_i are conditionally uncorrelated and β denotes the causal effect of x.[2]

Clearly, if individual-level data were available and if $\mathbf{W}_i$ and $\mathbf{Q}_a$ were fully observed by the econometrician, then estimating β would be straightforward. Often, however, microdata are not available, and researchers may or may not have access to all of the behaviorally relevant variables. In the typical application, the available data take the form averages, i.e.,

$$(2) \qquad \begin{array}{ccccccccccc} & \frac{1}{N_c}\sum_{i=1}^{N_c} y_i & = & \beta\frac{1}{N_c}\sum_{i=1}^{N_c} x_i & + & \frac{1}{N_c}\sum_{i=1}^{N_c}\mathbf{W}_i'\theta & + & \frac{1}{N_c}\sum_{i=1}^{N_c}\mathbf{Q}_a'\gamma & + & \frac{1}{N_c}\sum_{i=1}^{N_c}\epsilon_i & , \\ \Rightarrow & \overline{y}_c & = & \beta\overline{x}_c & + & \overline{\mathbf{W}}_c'\theta & + & \overline{\mathbf{Q}}_c'\gamma & + & \overline{\epsilon}_c & \end{array}$$

where c refers to the unit of aggregation, say, counties, N_c is the number of individuals within each unit, and upper bars denote unit-level means.[3]

Simply regressing $\overline{y}_c$ on $\overline{x}_c$ yields a point estimate with probability limit

$$(3) \qquad \text{plim}\,\widehat{\beta}_{OLS} = \beta + \frac{\text{Cov}\,(\overline{x}_c, \overline{\eta}_c)}{\text{Var}\,(\overline{x}_c)},$$

where $\overline{\eta}_c \equiv \overline{\mathbf{W}}_c'\theta + \overline{\mathbf{Q}}_c'\gamma + \overline{\epsilon}_c$. As a result, an ecological regression of $\overline{y}_c$ on $\overline{x}_c$ recovers β, the true individual-level parameter, *if and only if* $\text{Cov}\,(\overline{x}_c, \overline{\eta}_c) = 0$ (see also Langbein and Lichtman 1978, ch. 2; King 1997, ch. 3). Put differently, ecological regressions produce biased results whenever the regressor of interest is correlated with the error term.

This basic observation holds regardless of whether all, some, or none of the variables in $\overline{\mathbf{W}}_c$ and $\overline{\mathbf{Q}}_c$ are statistically controlled for. Suppose for instance, that the econometrician fully

[2]The linear specification in (1) might be perceived as very restrictive. This is not necessarily so. The data-generating process can be generalized to include interaction effects and nonlinearities without affecting any of the results below. In fact, by the Weierstrass Approximation Theorem, it is theoretically possible to uniformly approximate any continuous functional relationship to any desired degree of accuracy by including higher-order polynomials of the relevant variables. Thus, linear models need not be restrictive at all.

[3]Note that a and c may but need not be at the same level of aggregation. If $c = a$, then $\overline{\mathbf{Q}}_c = \mathbf{Q}_a$.

observed $\overline{\mathbf{W}}_c$ and $\overline{\mathbf{Q}}_c$, and included both vectors in the regression. In such a scenario, the probability limit of the OLS estimator would become plim $\widehat{\beta}_{OLS} = \beta + \text{Cov}\left(\overline{x}^*_c, \overline{\epsilon}_c\right) / \text{Var}\left(\overline{x}^*_c\right)$, with $\overline{x}^*_c$ denoting the residual from projecting $\overline{x}_c$ on the set of controls (see Frisch and Waugh 1933).[4] Even if x is uncorrelated with the error term in individual-level data, it need not be the case that plim $\widehat{\beta}_{OLS} = \beta$. Since $\text{Cov}\left(\overline{x}^*_c, \overline{\epsilon}_c\right) = \frac{1}{N^2}\sum_i \sum_j \text{Cov}\left(x^*_i, \epsilon_j\right)$, bias may arise whenever $\text{Cov}\left(x^*_i, \epsilon_j\right) \neq 0$ for some subset of individuals $i \neq j$ within the same unit of aggregation.

Note, the bias is *not* due to misspecification of the regression model. If equation (2) could be estimated on individual data, OLS would recover the true β, i.e., the causal effect of x_i on y_i. Rather, the ecological fallacy arises because the behavior of individuals correlates with the observed characteristics of *others* within the same unit.

Of course, in many applications, data limitations preclude scholars from measuring all relevant controls, so that some components of $\overline{\mathbf{W}}_c$ and $\overline{\mathbf{Q}}_c$ are unobserved. If x is correlated with any of these unobservables, then to obtain unbiased individual-level parameter estimates, the econometrician needs to deal with issues of causal identification *and* ecological inference.

3. EI-IV

Conveniently, standard instrumental variables techniques solve both problems. To see this, suppose the econometrician has access to an instrument, z_c, that is *(i)* relevant, i.e., it predicts x conditional on covariates, and *(ii)* excludable, i.e., it affects individual outcomes only through observables. Given such an instrument, two-stage least squares recovers the individual-level effect of x on y.

PROPOSITION 1: *Applied to (2), the two-stage least squares estimator of β is asymptotically consistent and equal to the true, individual-level parameter whenever the instrument, z_c, satisfies (i)* $\text{Cov}\left(z^*_c, \overline{x}_c\right) \neq 0$ *and (ii)* $\text{Cov}\left(z^*_c, \varepsilon_i\right) = 0$*, where ε_i is the individual-level error term and * denotes the residual from projecting z_c on the space spanned by the controls.*

[4] Formally, $\overline{x}^*_c = \overline{x}_c - \overline{\mathbf{W}}'_c\widehat{\omega} - \overline{\mathbf{Q}}'_c\widehat{\chi}$, where $\widehat{\omega}$ and $\widehat{\chi}$ denote the ordinary least squares coefficients from regressing $\overline{x}_c$ on $[\overline{\mathbf{W}}_c, \overline{\mathbf{Q}}_c]$.

Proof: The probability limit of the two-stage least squares estimator equals $\text{plim}\, \widehat{\beta}_{2SLS} = \text{Cov}\,(z_c^*, \overline{y}_c) \,/\, \text{Cov}\,(z_c^*, \overline{x}_c,) \;=\; \beta + \text{Cov}\,(z_c^*, \overline{\varepsilon}_c) \,/\, \text{Cov}\,(z_c^*, \overline{x}_c,)$. It thus suffices to show that $\text{Cov}\,(z_c^*, \overline{\varepsilon}_c) = 0$. By the usual properties of covariances, $\text{Cov}\,(z_c^*, \overline{\varepsilon}_c) = \frac{1}{N}\sum_{i=1}^{N} \text{Cov}\,(z_c^*, \varepsilon_i)$. Given that individual indices are exchangeable within each c, we immediately see that $\text{Cov}\,(z_c^*, \varepsilon_i) = 0$ implies $\text{Cov}\,(z_c^*, \overline{\varepsilon}_c) = 0$, as desired. *Q.E.D.*

Loosely speaking, the proposition states that if the econometrician observes a variable that would satisfy the exclusion restriction in individual-level data, then the same variable is also a valid instrument in the aggregated, i.e., averaged, data. The mathematics behind this result is strikingly simple. If an instrument is uncorrelated with all individual error terms, then it must also be uncorrelated with their mean. As a result, the two-stage least squares estimator solves both the causal as well as the ecological inference problem.

Importantly, EI-IV relies on essentially the same set of assumptions that are required for IV analyses in microdata. Hence, in a broad class of applications, drawing causal individual-level inferences from aggregated data is *not* substantially more difficult than obtaining causal estimates in the first place. IV estimates from aggregate data generalize to the individual level whenever the ecological regression model is directly implied by a statistical description of individual behavior—as in equations (1) and (2).

Naturally, the validity of these inferences is contingent on the excludability of z_c. The assumption that the instrument and the error term are uncorrelated always requires justification—even if the individual-level data were available!

4. An Illustrative Application with Sensitivity Analysis

Religion and the Nazi Vote An enormous Nazi voting literature identifies religion as a key correlate of NSDAP support. However, existing scholarship is almost exclusively concerned with the ecological inference problem rather than establishing causality (see, e.g., Brown 1982; Falter 1991; King et al. 2008; O'Loughlin 2002). In fact, Falter (1991) explicitly acknowledges that the assumptions required for his estimates to have a valid *ceteris paribus* interpretation are "in many cases unrealistic" (p. 443).

Since the average Catholic in Weimar Germany differed greatly from the average Protestant, it is unclear whether the well-known difference in Nazi support is due to religion itself or to other, socioeconomic disparities, for which conventional methods of ecological inference do not account. To better understand the effect of religion, Spenkuch and Tillmann (2017) use EI-IV in conjunction with the following econometric specification:

$$v_c = \mu_d + \beta \overline{Catholic}_c + \overline{\mathbf{W}}_c' \theta + \xi_c, \tag{4}$$

where, v_c denotes NSDAP vote shares in county c in the November election of 1932, $\overline{Catholic}_c$ measures the share of Catholics, $\overline{\mathbf{W}}_c$ is a vector of county-level controls, and μ_d marks an electoral district fixed effect. To ensure that equation (4) corresponds to the aggregation of an individual-level model, v_c is measured with respect to all adults in a particular county rather than the total number of votes.

Due to the scarcity of data from the end of the Weimar Republic, there may be many unobservables that are correlated with both NSDAP votes and constituents' religion, implying that $\text{Cov}\left(\overline{Catholic}_c^*, \xi_c\right) \neq 0$. Thus, the simple OLS estimates in Table 1 can be given neither a causal nor an individual-level interpretation.

[Table 1 about here.]

In their search for an instrument, Spenkuch and Tillmann (2017) turn to a stipulation in a sixteenth-century peace treaty. According to the principle *cuius regio, eius religio* ("whose realm, his religion"), the Peace of Augsburg granted more than a thousand local lords the right to determine their territories' official faith and, therefore, the religion of all their subjects. As the evidence in the middle two columns of Table 1 demonstrates, a territory's official faith at the end of the sixteenth century correlates strongly with the religion of Germans living in the same area more than 300 years later. Thus, princes' choices in the aftermath of the Peace satisfy the criterion of instrument relevance, i.e., $\text{Cov}\left(Z_c^*, \overline{Catholic}_c\right) \neq 0$. Importantly, the historical record indicates that rulers' decisions were often based on their personal convictions and other idiosyncratic factors, suggesting that, conditional on socio-economic controls, excludability might also hold, i.e., $\text{Cov}\left(Z_c^*, \xi_c\right) = 0$.

If one believes that, conditional on the set controls, princes' decisions affected Nazi vote shares in 1932 only through the constituents' religion, then the two-stage least squares estimates in the rightmost columns of Table 1 can be interpreted as causal individual-level effects. According to these estimates, Catholics were more than 25 percentage points less likely to vote for the Nazis than their Protestant counterparts—an even *greater* difference than in the raw data.

Sensitivity Analysis Depending on the application, scholars may or may not have much faith in the exclusion restriction. Fortunately, it is easy to assess the sensitivity of the estimated *individual-level* effect with respect to potential violations of this key assumption.

Building on Conley et al. (2012), consider the augmented model:

$$(5) \qquad v_c = \mu_d + \beta \overline{Catholic}_c + \overline{\mathbf{W}}_c' \theta + \gamma_0 Historically\,Catholic_c + \gamma_1 Historically\,Mixed_c + \xi_c.$$

Here, $\gamma = [\gamma_0, \gamma_1]$ parameterizes the extent to which excludability fails. If the exclusion restriction does, in fact, hold, then $\gamma_0 = \gamma_1 = 0$ (i.e., princes' decisions had no independent impact on Nazi votes). Since $\overline{Catholic}_c$ is endogenous, β and γ cannot be separately identified. Conley et al. (2012), however, show how to estimate β and conduct inference conditional on specifying the *support* of γ.[5]

Without prior information on the independent impact of the instrument, one obtains identical point estimates as in the EI-IV setup. The confidence intervals, however, widen. Figure 1 displays confidence intervals imposing only the assumption that each element of γ falls within $[-\delta, \delta]$—so that δ denotes the maximal theoretical violation of the exclusion restriction.

[Figure 1 about here.]

[5]A user-friendly STATA routine that implements the necessary calculations can be downloaded from the Boston College Statistical Software Components Archive by typing `ssc install plausexog` at the STATA prompt.

Reassuringly, as long as one is willing to rule out a direct impact of rulers' choices greater or equal to about nine percentage points, one would still reject the null hypothesis of no individual-level effect of religion on Nazi support. To put this number into perspective, nine percentage points corresponds to more than one-third of all NSDAP supporters in the November elections of 1932. In light of this finding, Spenkuch and Tillmann (2017) conclude that religion almost certainly exerted a genuine effect. Furthermore, appealing to the result above, they extrapolate to the individual level without imposing assumptions that are stronger than those required to claim causality in a microanalysis. Relying on EI-IV rather than conventional methods for ecological inference, Spenkuch and Tillmann (2017) address the ecological fallacy all the while estimating causal effects.

5. Discussion

This note observes that standard instrumental variables techniques have the potential to solve the ecological inference problem. While the mathematics behind this result is strikingly simple, the implications are nonetheless powerful. EI-IV allows researchers to draw conclusions about individual-level effects from aggregate data under essentially the same set of assumptions that are required for causal identification in microdata. Put differently, the derivations above show that whenever an ecological regression specification is directly implied by a statistical description of individual behavior, then the resulting estimates can be generalized to the individual level if and only if they are causal.

EI-IV can easily be implemented with standard statistical software. It is also straightforward to assess the robustness of the individual-level inferences with respect to the (untestable) exclusion restriction. Furthermore, a number of research designs, such as randomized experiments or RDDs, can be recast in an instrumental variables framework (see, e.g., Hahn et al. 2001). The insights in this note are, therefore, *not* limited to the two-stage least squares estimator. Given the emphasis on causal identification in modern research, EI-IV offers an attractive methodological approach to a wide range of questions in the social sciences and beyond.

References

Brown, C. (1982). "The Nazi Vote: A National Ecological Study." *American Political Science Review*, 76(2): 285–302.

Conley, T., C. Hansen, and P. Rossi (2012). "Plausibly Exogenous." *Review of Economics and Statistics*, 94(1): 260–272.

Duncan, O., and B. Davis (1953). "An Alternative to Ecological Correlation." *American Sociological Review*, 18(6): 665–666.

Falter, J. (1991). *Hitlers Wähler*. Munich: C.H. Beck.

Freedman, D., S. Klein, J. Sacks, C. Smyth, and C. Everett (1991). "Ecological Regression and Voting Rights." *Evaluation Review*, 15(6): 673–711.

——— , ——— , M. Ostland, and M. Roberts (1998). "On 'Solutions' to the Ecological Inference Problem." Technical Report No. 515, Statistics Departmen, UC Berkeley.

Frisch, R., and F. Waugh (1933). "Partial Time Regressions as Compared with Individual Trends." *Econometrica*, 1(4): 387–401.

Gelman, A., D. Park, S. Ansolabehere, P. Rice, and L. Minnite (2001). "Models, Assumptions and Model Checking in Ecological Regressions." *Journal of the Royal Statistical Society: A*, 164: 101–118.

Goodman, L. (1953). "Ecological Regressions and Behavior of Individuals." *American Sociological Review*, 18(6): 663–664.

——— (1959) "Some Alternatives to Ecological Correlation." *American Journal of Sociology*, 64(6): 610–625.

Hahn, J., P. Todd, and W. van der Klaauw (2001). "Identication and Estimation of Treatment Effects with a Regression-Discontinuity Design." *Econometrica*, 69(1): 201–209.

King, G. (1997). *A Solution to the Ecological Inference Problem*. Princeton, NJ: Princeton University Press.

——— , O. Rosen, M. Tanner (1999). "Binomial-Beta Hierarchical Models for Ecological Inference." *Sociological Methods and Research*, 28(1): 61–90.

——— , ——— , ——— (eds.) (2004). *Ecological Inference: New Methodological Strategies*. Cambridge, UK: Cambridge University Press.

——— , ——— , ——— , and A. Wagner (2008). "Ordinary Economic Voting Behavior in the Extraordinary Election of Adolf Hitler." *Journal of Economic History*, 68(4): 951–996.

Langbein, L., and A. Lichtman (1978) *Ecological Inference*. Beverly Hills, CA: Sage.

O'Loughlin, J. (2002). "The Electoral Geography of Weimar Germany." *Political Analysis*, 10(3): 217–243.

Robinson, P. (1950). "Ecological Correlation and the Behavior of Individuals." *American Sociological Review*, 15(3): 351–357.

Spenkuch, J., and P. Tillmann (2017) "Elite Influence? Religion and the Electoral Success of the Nazis." *American Journal of Political Science*, forthcoming.

ONLINE APPENDIX

Appendix A: Aggregate vs. Individual-Level Instruments

To avoid potential confusion, it may be worth elaborating on a subtle point. Proposition 1 defines the instrument on the level at which the econometrician actually observes the data. To see why this is important, note that averaging a variable that would be excludable if applied to individual-level data may but need not yield an instrument that is also valid in the aggregate. Specifically, let $\overline{z}_c \equiv \frac{1}{N_c} \sum_{i=1}^{N_c} z_i$. If z_i does not vary across individuals within the same c, then $\overline{z}_c = z_i$ and $\text{Cov}\left(z_i, \varepsilon_i\right) = 0$ implies that $\text{Cov}\left(\overline{z}_c, \overline{\varepsilon}_c\right) = 0$, as required for $\text{plim}\, \widehat{\beta}_{2SLS} = \beta$. Thus, whenever the instrument is logically constant within a unit of aggregation, EI-IV relies on *exactly* the same set of assumptions that researchers impose anyway in order to obtain causal estimates. Extrapolating from the results to individual behavior is thus without any further loss.

If, however, $z_i \neq z_j$ for some $i, j \in c$, then $\text{Cov}\left(z_i, \varepsilon_i\right) = 0$ does not automatically imply $\text{Cov}\left(\overline{z}_c, \overline{\varepsilon}_c\right) = 0$.[1] A sufficient assumption for $\overline{z}_c$ to be excludable is $\text{Cov}\left(z_i, \varepsilon_j\right) = 0$ for all i and j within the same unit of aggregation. Or, in words, the instrument must affect *all* individuals within the same unit only through covariates that are already included in the regression. In effect, scholars need to argue that, defined on the individual level, the instrument would be excludable in the usual sense *and* exhibit no unobserved spillover effects on others.[2] With such an instrument, EI-IV recovers causal individual-level estimates from aggregate data.

[1] Again, this is due to the possibility that $\text{Cov}\left(z_i, \varepsilon_j\right) = 0$ for some $i \neq j$ within the same c.

[2] That is, $\text{Cov}\left(z_i, \varepsilon_i\right) = 0$ for all i, and $\text{Cov}\left(z_i, \varepsilon_j\right) = 0$ for all $i \neq j$ within the same unit of aggregation.

Table 1: Estimating the Effect of Religion on Nazi Support

	A. OLS		*B. First Stage*		*C. 2SLS*	
	NSDAP Vote Share		Percent Catholic		NSDAP Vote Share	
	(1)	(2)	(3)	(4)	(5)	(6)
Percent Catholic	-.190	-.287			-.254	-.273
	(.019)	(.025)			(.018)	(.028)
Official Hist. Religion:						
Catholic			66.666	42.117		
			(3.232)	(3.681)		
Mixed			39.270	22.005		
			(4.320)	(3.322)		
Controls	No	Yes	No	Yes	No	Yes
Electoral District FE	No	Yes	No	Yes	No	Yes
First Stage F-Statistic	--	--	--	--	212.74	71.38
Over-ID Test [p-value]	--	--	--	--	.275	.581
Number of Observations	982	982	982	982	982	982

Notes: Entries in columns (1), (2), (5), and (6) are OLS and 2SLS point estimates from regression models akin to equation (4). Columns (3) and (4) show the first stage results, i.e., the partial correlation between instruments and the endogenous variables. Standard errors account for clustering by electoral district and are reported in parentheses. The set of controls is the same as in the most inclusive specification in Tables 1 and 2 of Spenkuch and Tillmann (2017). For a precise description of the instrument as well as all other variables, see the aforementioned paper.
Source: Spenkuch and Tillmann (2017), Tables 1 and 2

Figure 1: Inference Allowing for Violations of the Exclusion Restriction

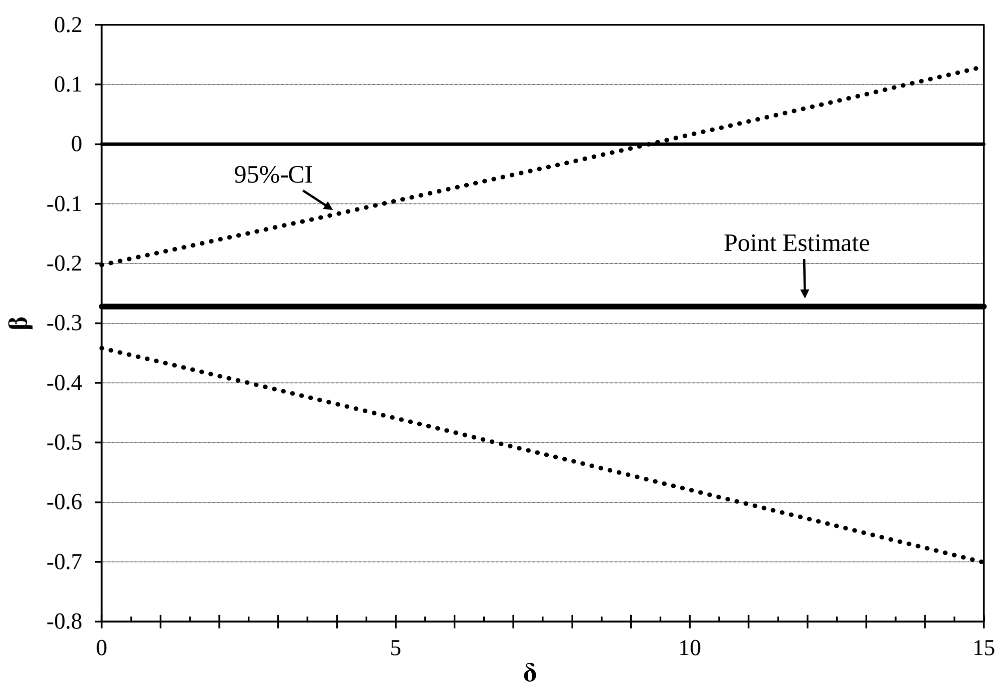

Notes: Figure depicts point estimates and 95%-confidence intervals for the effect of Catholicism on NSDAP vote shares in the November elections of 1932. The confidence intervals impose only the prior information that the support of each element of γ in equation (5) falls within $[-\delta,\delta]$. Intuitively, δ parameterizes the maximal allowable violation of the exclusion restriction. For details on the estimation procedure, see Conley et al. (2012).
Source: Spenkuch and Tillmann (2017), Appendix Figure A.4