Advance Access publication January 30, 2013

*Political Analysis* (2013) 21:233–251
doi:10.1093/pan/mps043

# Causal Inference without Ignorability: Identification with Nonrandom Assignment and Missing Treatment Data

**Walter R. Mebane Jr.**
*Department of Political Science and Department of Statistics, University of Michigan, 5700 Haven Hall, Ann Arbor, MI 48109*
*e-mail: wmebane@umich.edu (corresponding author)*

**Paul Poast**
*Department of Political Science, Rutgers University, Hickman Hall, New Brunswick, NJ 08901*
*e-mail: paul.poast@rutgers.edu*

Edited by Jonathan Katz

How a treatment causes a particular outcome is a focus of inquiry in political science. When treatment data are either nonrandomly assigned or missing, the analyst will often invoke ignorability assumptions: that is, both the treatment and missingness are assumed to be as if randomly assigned, perhaps conditional on a set of observed covariates. But what if these assumptions are wrong? What if the analyst does not know why—or even if—a particular subject received a treatment? Building on Manski, Molinari offers an approach for calculating nonparametric identification bounds for the average treatment effect of a binary treatment under general missingness or nonrandom assignment. To make these bounds substantively more informative, Molinari's technique permits adding monotonicity assumptions (e.g., assuming that treatment effects are weakly positive). Given the potential importance of these assumptions, we develop a new Bayesian method for performing sensitivity analysis regarding them. This sensitivity analysis allows analysts to interpret the assumptions' consequences quantitatively and visually. We apply this method to two problems in political science, highlighting the method's utility for applied research.

## 1 Introduction

How does job loss affect vote choice? Can democracy impede war onset? Do single member districts impact political party structure? Whether using experimental data with random assignment or observational data with unknown assignment mechanisms, political scientists seek to infer the causal effect of treatments, $z$, on outcomes, $y$. Unfortunately, in many, if not most, studies involving observational data, the treatment is not randomly assigned or, even worse, is missing for some observations.

Nonrandom treatment assignment is the defining feature of observational data (Cochran and Rubin 1973; Rubin 2006). Without random assignment, variables that affect $y$ other than $z$ may be distributed differently across the treated and control groups in ways that cannot be statistically determined. In other words, in observational data the treatment is generally confounded with other covariates, and only some of the potential confounders may be observable. Consequently, "it is virtually impossible in many practical circumstances to be convinced that the estimates of the effects of treatments are in fact unbiased" (Cochran and Rubin 1973, 30). Nevertheless, analysts must sometimes rely on observational data, particularly when random treatment assignment would be impossible (e.g., the impact of a coup d'état on democratization, or studies based on historical data).

*Authors' note*: Thanks to Peter Aronow for many important contributions. Comments from Don Green, Arthur Spirling, Allison Sovey, Rory Truex, and the participants of the 2011 MPSA Annual National Conference and the 2011 Annual Meeting of the Society for Political Methodology are greatly appreciated. Supplementary materials for this article are available on the *Political Analysis* Web site. Replication materials are in study hdl:1902.1/19368 at IQSS Dataverse Network.

© The Author 2013. Published by Oxford University Press on behalf of the Society for Political Methodology.
All rights reserved. For Permissions, please email: journals.permissions@oup.com



https://doi.org/10.1093/pan/mps043 Published online by Cambridge University Press

Missing treatment data are also a pervasive problem, particularly in surveys. For example, Sharp and Joslyn (2001), in attempting to identify the impact of viewing sexually explicit material on support for regulation of pornography, note that respondents chose not to report the true amount of pornography they viewed in the past year. Similarly, Molinari (2010) highlights a study of the effect of drinking on birth outcomes, where a question asking about drinking behavior during pregnancy had a nonresponse rate of 6–14% by female respondents. Given a societal stigma associated with drinking while pregnant, nonresponses in the survey are probably related to beliefs and behaviors that affect birth outcomes.

When faced with these two limitations—nonrandom assignment and missing treatment data—analysts frequently turn to modeling assumptions. Most typically, the analyst makes an ignorability assumption using observable covariates. This assumption implies that, conditional on a covariate profile, treatment assignment can be treated as if randomly assigned and missingness is at random or ignorable. Since ignorability is a strong and, in many cases, implausible assumption, analysts may attempt to model the treatment assignment and missingness mechanisms, as with inverse probability weighting methods (Glynn and Quinn 2010).

We advocate an alternative and, in some respects, complementary approach: to embrace uncertainty. In particular, Molinari (2010) offers an approach for identifying and calculating the interval in which the average treatment effect (ATE) of a binary treatment must be contained even when treatment data are nonrandom and missing.[1] This approach requires no assumptions of ignorability and, beginning with the same basic setup as Manski's derivation of bounds for the ATE, is completely nonparametric (Manski 1990, 1995, 1997).

Like other approaches that compute identification bounds for treatment effects, the bounds, though sharp, are not always substantively informative.[2] Consequently, Molinari proposes adding two monotonicity assumptions when appropriate. One assumption holds that when the treatment has a nonzero effect, the effect is only and always in one direction: the effect is assumed to be either weakly positive (positive monotonicity) or weakly negative (negative monotonicity). For instance, in the above pregnancy example, one might assume that drinking only negatively affects birth weight. The other assumption is that the average potential outcomes of those treated will be either weakly greater (positive monotonicity) or weakly lesser (negative monotonicity) than those of the untreated. Continuing with the pregnancy example, the assumption might be that women who drink more were likely to have lower birthweight children for other reasons.

Including monotonicity assumptions can shrink the bounds, but such assumptions are purely theoretical; the motivation for these assumptions must be developed *a priori*. Hence the analyst may wish to determine the extent to which these assumptions drive the results. We develop a technique for determining the sensitivity of the bounds to the monotonicity assumptions by considering that there may be subpopulations in which there are different kinds of effects.

With motivations somewhat different from ours, others have critically studied sensitivity analyses (Rosenbaum 2002; Quinn 2009; Imai, Keele, and Yamamoto 2010; Imai et al. 2011) and bounds (Yamamoto 2012) in the assessment of causal effects. For instance, Quinn (2009) presents a Bayesian analysis of the effect of confounders on what are essentially Manski bounds for complete data without monotonicity assumptions, whereas Imai et al. (2011) develop a sensitivity analysis for key assumptions underlying mediation analysis. For the most part we do not attempt to measure the consequences of confounders nor to discriminate direct and indirect effects. Others have applied these techniques (Blattman 2009; Poast 2012). For example, Poast (2012) uses Rosenbaum sensitivity analysis to identify the extent to which unobserved confounders drive his finding that trade linkage offers increase the probability of states ending an alliance negotiation in agreement. Although confounders are always a challenge to inference with observational data, bounds can show the range of effect estimates that may be found when the data are analyzed with a model.

[1] Many of the results in Molinari (2010) hold if there are multiple treatments, and the results we use can be adapted if notions of monotonicity can be defined (c.f. Manski 1997). For instance, perhaps pairwise comparisons of treatments are of interest and there are reasons to think monotonicity holds between pairs.

[2] "Sharp" describes the largest interval that is consistent with all available information.

https://doi.org/10.1093/pan/mps043 Published online by Cambridge University Press

This article has four goals. First, given the potential for both nonignorable treatment assignment and missing treatment data in many settings, we introduce to political scientists Molinari's technique for computing the identification bounds on the ATE. We consider the case of a binary outcome.[3] This technique offers a relatively simple and nonparametric way to check the robustness of effects estimated with any given statistical model. Second, we develop a technique for determining the sensitivity of the bounds to the monotonicity assumptions. Our innovation is to treat the population of interest as being composed of two separate subpopulations—one in which a given monotonicity assumption is met and another in which the assumption is not met—even though it is unknown to which subpopulation any particular observation belongs. This innovation allows the analyst to compute regions of bounds that explicitly incorporate uncertainty and determine the extent to which these assumptions drive the results. Third, we consider complications associated with having data from a random sample and not the full population. Finally, to show how substantive assumptions are involved when using Molinari's technique and our sensitivity analysis, we apply both to two problems in political science.

## 2 Manski and Molinari Identification Bounds

When faced with nonrandom treatment assignment and missing treatment data, obtaining precise estimates may require rather strong assumptions, perhaps stronger than what one is actually willing to defend. Manski (1990, 1995) offers an approach for defining identification bounds for a treatment effect: the interval between such bounds corresponds to the set of values that are logically compatible with the data and any explicitly stated assumptions. These assumptions are often extremely weak and, at a minimum, require only the laws of probability. Using bounds in this way lays bare the influence that assumptions have on the ability to estimate an effect.

### 2.1 *Identification Bounds without Missing Treatment Data*

Manski (1990, 1995) computes identification bounds for the treatment effect by assuming a binary treatment represented by 0 or 1.[4] Manski lets $y_1$ be the outcome associated with a unit receiving treatment 1 and lets $y_0$ be the outcome associated with a unit receiving treatment 0. Thus, $\mathrm{P}[y_1|x]$ denotes the distribution of outcomes that would be realized if all units with characteristics $x$ received treatment 1 and $\mathrm{P}[y_0|x]$ denotes the distribution of outcomes that would be realized if all units with characteristics $x$ received treatment 0. If each member of the population receives one of the two treatments, let the binary variable $z$ indicate which of the two treatments a unit received, where $z = 1$ or $z = 0$ and $\mathrm{P}[z|x]$ is the distribution of the treatment variable among all units with characteristics $x$. For each unit, one observes the value of $x$, the value of $z$, and the outcome resulting from the treatment, $y$, which takes the value $y_1$ if $z = 1$ or $y_0$ if $z = 0$. For parsimony in presentation, we suppress notation indicating that we are conditioning on the covariates represented by $x$. Henceforth we write $\mathrm{P}[y_1]$ in place of $\mathrm{P}[y_1|x]$, $\mathrm{P}[y_0]$ in place of $\mathrm{P}[y_0|x]$, $\mathrm{P}[z]$ in place of $\mathrm{P}[z|x]$, and so forth.[5]

The quantity of interest is the difference $\mathrm{P}[y_1 = 1] - \mathrm{P}[y_0 = 1]$, which is the ATE. By the law of total probability, $\mathrm{P}[y_1 = 1]$ can be reexpressed as $\mathrm{P}[y_1 = 1] = \mathrm{P}[y_1 = 1|z = 1]\mathrm{P}[z = 1] + \mathrm{P}[y_1 = 1|z = 0]\mathrm{P}[z = 0]$ and $\mathrm{P}[y_0 = 1]$ can be reexpressed as $\mathrm{P}[y_0 = 1] = \mathrm{P}[y_0 = 1|z = 1]\mathrm{P}[z = 1] + \mathrm{P}[y_0 = 1|z = 0]\mathrm{P}[z = 0]$. We have no information about

[3] Molinari (2010) considers outcomes in a bounded set with known bounds. Using such outcomes adds technical complexity but no new insights relative to our purpose to develop methods for assessing monotonicity assumptions. To get a sense of the technical differences involved in using outcomes in a bounded set instead of binary outcomes, compare our equation (2) to the equation in Proposition 4 in Molinari (2010, 86).

[4] Supplementary appendix A formally describes the Manski framework, including a discussion of the assumption of "individualistic treatment response," an assumption analogous to Rubin's (1978) stable unit treatment value assumption.

[5] All of the calculations in this article may be made conditional on a covariate profile without any meaningful complications to the mathematics. The variable $x$ would represent the covariates capturing the characteristics of all members of the population under study. For all probabilities, P[.], we would write P[.|$x$] in its place.

https://doi.org/10.1093/pan/mps043 Published online by Cambridge University Press

$\mathrm{P}[y_1 = 1|z = 0]$ and $\mathrm{P}[y_0 = 1|z = 1]$, other than that each lies between 0 and 1. Thus, by Manski (1995, 39–41),

$$\mathrm{P}[y_1 = 1|z = 1]\mathrm{P}[z = 1] \le \mathrm{P}[y_1 = 1] \le \mathrm{P}[y_1 = 1|z = 1]\mathrm{P}[z = 1] + \mathrm{P}[z = 0],$$
$$\mathrm{P}[y_0 = 1|z = 0]\mathrm{P}[z = 0] \le \mathrm{P}[y_0 = 1] \le \mathrm{P}[z = 1] + \mathrm{P}[y_0 = 1|z = 0]\mathrm{P}[z = 0].$$

Without further information, sharp bounds for $\mathrm{P}[y_1 = 1] - \mathrm{P}[y_0 = 1]$ are defined by

$$\begin{aligned}\mathrm{P}[y_1 = 1|z = 1]\mathrm{P}[z = 1] - \mathrm{P}[y_0 = 1|z = 0]\mathrm{P}[z = 0] - \mathrm{P}[z = 1] \le \mathrm{P}[y_1 = 1] - \mathrm{P}[y_0 = 1] \\ \le \mathrm{P}[y_1 = 1|z = 1]\mathrm{P}[z = 1] + \mathrm{P}[z = 0] - \mathrm{P}[y_0 = 1|z = 0]\mathrm{P}[z = 0].\end{aligned} \tag{1}$$

### 2.2 *Identification Bounds with Missing Treatment Data*

Molinari (2010) extends Manski bounds to develop a procedure for identifying treatment effects when treatment data are missing.[6] The notation is the same as above, except Molinari introduces a new variable, $d$, which is a binary variable equal to 1 if the treatment received by a unit is observed, 0 otherwise. The analyst knows two distributions: $\mathrm{P}[d = 1]$, the probability that the treatment is observed, and $\mathrm{P}[z|y,d = 1]$, the conditional distribution of realized treatments given realized outcomes and the observability of the realized treatment.

Given this setup, what can be learned about the ATE, $\mathrm{P}[y_1 = 1] - \mathrm{P}[y_0 = 1]$? Using the law of total probability, Molinari shows[7] that the sharp bounds for the treatment effect in the absence of knowledge of $\mathrm{P}[z = 1|d = 0]$ are $\mathrm{LB}_{\mathrm{TE}} \le \mathrm{P}[y_1 = 1] - \mathrm{P}[y_0 = 1] \le \mathrm{UB}_{\mathrm{TE}}$, where

$$\begin{aligned}\mathrm{LB}_{\mathrm{TE}} = -\mathrm{P}[d = 0] - \mathrm{P}[d = 1](\mathrm{P}[z = 1|d = 1] \\ + \mathrm{P}[y = 1|d = 1, z = 0]\mathrm{P}[z = 0|d = 1] - \mathrm{P}[y = 1|d = 1, z = 1]\mathrm{P}[z = 1|d = 1]),\end{aligned}$$

$$\begin{aligned}\mathrm{UB}_{\mathrm{TE}} = \mathrm{P}[d = 0] + \mathrm{P}[d = 1](\mathrm{P}[z = 0|d = 1] \\ + \mathrm{P}[y = 1|d = 1, z = 1]\mathrm{P}[z = 1|d = 1] - \mathrm{P}[y = 1|d = 1, z = 0]\mathrm{P}[z = 0|d = 1]).\end{aligned}$$

These bounds are always at least 100% points wide (e.g., $[-0.25, 0.75]$), and always include 0. We can narrow the identification interval with further assumptions but can never rule out an ATE of 0. Although perhaps unsatisfactory for some analysts, this possibility represents a fundamental characteristic of the data and the assumptions that we are imposing. It is for this reason that Manski (1995, 8) calls for analysts to have "a greater tolerance for ambiguity."[8]

### 2.3 *Monotonicity Assumptions*

To identify the ATE more precisely, Molinari, following Manski (1997) and Manski and Pepper (2000), suggests assuming monotonicity in treatment response and treatment selection, referred to, respectively, as the *Weak Monotone Treatment Response* (MTR) assumption and the *Weak Monotone Treatment Selection* (MTS) assumption. MTR implies that the outcomes associated with treatments for each unit satisfy a particular ordering. MTS implies that units that have particular kinds of outcomes select (or are selected) to have particular kinds of treatments.

[6]Molinari's method for treatment effect identification with missing treatment data assumes Individualistic Treatment Response (ITR): the unit's treatment response is a function only of the value of its own treatment and not of the treatment realized for any other unit, which is a close analog to the stable-unit-treatment-value-assumption (SUTVA) in the Rubin (1978, 1991) Causal Model. Molinari assumes that each member $j$ of a population has "a specific response function $y_j(\cdot) : T \to Y$ mapping treatments $t \in T$ into outcomes" (Molinari 2010, 83) and so focuses on individualistic effects. Molinari's method and the sensitivity tests we introduce may still be useful in the absence of ITR if the data can be arranged so that an assessment of individualistic treatment effects is meaningful. In this case, all of $\mathrm{P}[y(t^J)]$ will not be identifiable.

[7]For a fuller discussion, see supplementary appendix B.

[8]To help intuition, supplementary appendix C works through a simple example of calculating Manski and Molinari bounds.

https://doi.org/10.1093/pan/mps043 Published online by Cambridge University Press

The MTR assumption refers to the treatment effect for each unit. With MTR, one assumes that the treatment has a monotone effect on the outcome. This means that the lower (upper) bound of the treatment effect is 0 when one assumes positive (negative) MTR. Formally, positive MTR implies that $y^j(z^j = 1) \geq y^j(z^j = 0)$ for each unit $j$. MTR refers to the different possible outcome values that exist for each unit.

The MTS assumption relies not just on unit level characteristics but describes a property of sets of units. With MTS, one assumes that if the population of observed units is divided into two groups according to the treatment, then the average potential outcome of the group that received the treatment is greater than the average potential outcome of the group that did not receive the treatment (Morgan and Winship 2007, 177). Formally, following Molinari (2007, 10), positive MTS is the assumption that $P[y_1 = 1|z = 1, d = i] \geq P[y_1 = 1|z = 0, d = i]$ and $P[y_0 = 1|z = 1, d = i] \geq P[y_0 = 1|z = 0, d = i]$, $i \in \{0, 1\}$. MTS refers to the outcomes for different units: units that receive the treatment valued $t = 1$ tend to have higher outcomes than those who receive treatment $t = 0$ even if, counterfactually, both types of units receive the same kind of treatment. A classic application of this assumption is found in Manski and Pepper (2000, 1003), to identify the causal effect of education on wages. MTS implies we assume that individuals with higher education would have, on average, received higher wages than individuals with lower education even if both had, in a counterfactual world, the same levels of education. Even though the treatment variable distinguishes those with higher outcomes from those with lower outcomes, the treatment variable may not be the reason for the outcome difference.

These monotonicity assumptions can be positive or negative. For the sake of simplicity, we will show the results only for weakly positive monotonicity. Following Molinari (2007), if the analyst first assumes only MTR, then the bounds become

$$\begin{aligned} 0 \leq \mathrm{P}[y_1 = 1] - \mathrm{P}[y_0 = 1] \leq \mathrm{P}[d = 0] + \mathrm{P}[d = 1](\mathrm{P}[z = 0|d = 1] \\ + \mathrm{P}[y = 1|d = 1, z = 1]\mathrm{P}[z = 1|d = 1] - \mathrm{P}[y = 1|d = 1, z = 0]\mathrm{P}[z = 0|d = 1]). \end{aligned} \tag{2}$$

Under the combination of the weak MTR and weak MTS assumptions, the bounds become

$$0 \leq \mathrm{P}[y_1 = 1] - \mathrm{P}[y_0 = 1] \leq \mathrm{P}[d = 1](\mathrm{P}[y = 1|d = 1, z = 1] - \mathrm{P}[y = 1|d = 1, z = 0]) + \mathrm{P}[d = 0].$$

These assumptions together imply a narrower identification interval.

The data sometimes rule out both MTR and MTS holding (Manski and Pepper 2000). If

$$\mathrm{P}[d = 1](\mathrm{P}[y = 1|d = 1, z = 1] - \mathrm{P}[y = 1|d = 1, z = 0]) < 0, \tag{3}$$

then it is logically impossible for the outcomes for units under treatment to be as large as those for units under control (for observed units, on average). Since MTR implies that the treatment potential outcomes are at least as large as the control potential outcomes, invoking both MTR and MTS in such a case implies a logical contradiction.

## 3 Extending Molinari Bounds: Sensitivity Analysis

A limitation of Molinari bounds is that while MTR and MTS are critical for obtaining informative bounds, analysts may not know how sensitive their results are to violations of the assumptions. We generalize Molinari bounds by incorporating uncertainty over the MTR and MTS assumptions. It is worth noting that all sensitivity analyses developed in this section may also be applied to Manski bounds through application to a population with no missing treatment data.

What if the analyst is not confident about the MTR or MTS assumptions, or both, for some proportion of the data? Our innovation is to consider that the population of interest (not just the missing treatment data) is composed of two separate subpopulations—one in which a given monotonicity assumption is met and another in which the assumption is not met—but it is not known which observations are in which subpopulation. This is one way to assess how robust inferences are to the two assumptions (cf. Arnborg 2006). One computes bounds that explicitly incorporate the analyst's uncertainty.

https://doi.org/10.1093/pan/mps043 Published online by Cambridge University Press

### 3.1 *MTR Assumption*

We begin with the case where the analyst is unsure of applying the MTR assumption on the population of data, but applies no MTS assumption. A vector of five parameters is observable:

$$\begin{aligned}\xi = [&\mathrm{P}[d=0], \mathrm{P}[y=1|z=1,d=1]\mathrm{P}[z=1|d=1]\mathrm{P}[d=1],\\ &\mathrm{P}[y=0|z=1,d=1]\mathrm{P}[z=1|d=1]\mathrm{P}[d=1], \mathrm{P}[y=1|z=0,d=1]\mathrm{P}[z=0|d=1]\mathrm{P}[d=1],\\ &\mathrm{P}[y=0|z=0,d=1]\mathrm{P}[z=0|d=1]\mathrm{P}[d=1]]'.\end{aligned}$$

$\xi$ is a realization of a random variable $\Xi$ that is distributed with some probability over the standard 4-simplex $\Delta^4$.[9] $\xi$ measures the configuration of the observable data. We add the parameter $t^j_{\mathrm{MTR}} \in \{0,1\}$. If unit $j$ follows the MTR assumption, $t^j_{\mathrm{MTR}} = 1$, if not, $t^j_{\mathrm{MTR}} = 0$. Define a 5-element vector corresponding to $\xi$,

$$\begin{aligned}\upsilon = [&\mathrm{P}[t_{\mathrm{MTR}}=0|d=0], \mathrm{P}[t_{\mathrm{MTR}}=0|y=1,z=1,d=1], \mathrm{P}[t_{\mathrm{MTR}}=0|y=0,z=1,d=1],\\ &\mathrm{P}[t_{\mathrm{MTR}}=0|y=1,z=0,d=1], \mathrm{P}[t_{\mathrm{MTR}}=0|y=0,z=0,d=1]]',\end{aligned}$$

which is a realization of a random variable $\Upsilon$ that is distributed over the unit 5-cube $\Lambda$.[10] $\upsilon$ represents the conditional distribution of $t_{\mathrm{MTR}}$ given the configuration of the observable data. The unconditional probability that a unit does not follow the MTR assumption is the inner product

$$\omega = \xi'\upsilon. \tag{4}$$

Given $\upsilon$, it is straightforward to compute identification bounds. The upper bound is unchanged by the MTR assumption, so the upper bound associated with any $\upsilon$ is $\mathrm{UB}_\upsilon = \mathrm{UB}_{\mathrm{TE}}$. The MTR assumption changes the lower bound, with its value depending on the components of $\upsilon$. The lower bound is a convex combination of the bounds for the data that follow the MTR assumption and for the data that do not follow the MTR assumption. The former bound is 0; therefore, given $\omega$ and $\upsilon$, the lower bound is

$$\begin{aligned}\mathrm{LB}_\upsilon = \omega\{&-\mathrm{P}[d=0|t_{\mathrm{MTR}}=0] - \mathrm{P}[d=1|t_{\mathrm{MTR}}=0](\mathrm{P}[z=1|d=1,t_{\mathrm{MTR}}=0]\\ &+ \mathrm{P}[y=1|d=1,z=0,t_{\mathrm{MTR}}=0]\mathrm{P}[z=0|d=1,t_{\mathrm{MTR}}=0]\\ &- \mathrm{P}[y=1|d=1,z=1,t_{\mathrm{MTR}}=0]\mathrm{P}[z=1|d=1,t_{\mathrm{MTR}}=0])\}\\ &= -(\upsilon_1\xi_1 + \upsilon_3\xi_3 + \upsilon_4\xi_4),\end{aligned}$$

where $\upsilon_1$ denotes the first element of $\upsilon$.[11]

But what if we only had firm knowledge of $\omega$, the proportion of the data that do not follow the MTR assumption, but not any particular $\upsilon$? Prior to the data and $\xi$ being observed, knowledge of $\omega$ does not provide any constraints on $\upsilon$. However, once $\xi$ is observed, knowledge of $\omega$ does imply constraints on $\upsilon$, through equation (4). Equation (4) can imply a likelihood function for any given $\omega$: using $\mathrm{I}(\cdot)$ to denote the indicator function,

$$L_\omega(\upsilon|\Xi=\xi) = \mathrm{I}(\omega = \xi'\upsilon).$$

Given $\omega$ and $\xi$, we may rule out many possibilities for $\upsilon$.

By expressing our prior uncertainty over the components in $\upsilon$ as a prior probability distribution, we may construct expected values and posterior intervals for the lower bound. One option, which we adopt, is to assume a prior where $\upsilon$ is distributed uniformly over $\Lambda$, such that all values in its support are equally likely.[12] The posterior probability of $\upsilon$ given $\xi$ is

[9] $\Delta^4$ contains all sets of real numbers $(u_0,u_1,u_2,u_3,u_4)$ such that $\sum_{i=0}^4 u_i = 1$ and, for all $i$, $u_i \geq 0$.
[10] $\Lambda$ contains all sets of real numbers $(u_1,u_2,u_3,u_4,u_5)$ such that, for all $i = 1,\ldots,5$, $0 \leq u_i \leq 1$.
[11] Derivations for the sensitivity analyses for the MTR assumptions are in supplementary appendix D.
[12] One could use other distributions if there were appropriate prior information, but in general we believe no such information would exist.

https://doi.org/10.1093/pan/mps043 Published online by Cambridge University Press

$$f_{\Lambda|\Xi=\xi}(\upsilon) = \frac{L_\omega(\upsilon|\Xi=\xi)}{\int_\Lambda L_\omega(\upsilon|\Xi=\xi)\,d\upsilon}.$$

With knowledge of the posterior distribution of bounds, one may construct posterior intervals and expected values for the bounds. We discuss computation of these quantities in Section 3.3.

### 3.2 *MTS Assumption*

The logic behind assessing the MTS assumption directly follows from that of the MTR assumption. We now assume that the entire population follows the MTR assumption ($\omega = 0$), but that some proportion $\zeta$ of the population does not follow the MTS assumption. We add a parameter $t^j_{\text{MTS}} \in \{0, 1\}$ such that $t^j_{\text{MTS}} = 1$ if unit $j$ follows the MTS assumption and $t^j_{\text{MTS}} = 0$ if not. Let $\zeta = \text{P}[t_{\text{MTS}} = 0|\omega = 0]$. Define a vector of five elements,

$$\begin{aligned}\theta =&[\text{P}[t_{\text{MTS}} = 0|d = 0, \omega = 0], \text{P}[t_{\text{MTS}} = 0|y = 1, z = 1, d = 1, \omega = 0],\\ &\text{P}[t_{\text{MTS}} = 0|y = 0, z = 1, d = 1, \omega = 0], \text{P}[t_{\text{MTS}} = 0|y = 1, z = 0, d = 1, \omega = 0],\\ &\text{P}[t_{\text{MTS}} = 0|y = 0, z = 0, d = 1, \omega = 0]]',\end{aligned}$$

which is a realization of a random variable that is also distributed over $\Lambda$. Now, $\zeta = \xi'\theta$. The method for articulating bounds is fundamentally unchanged from that described in Section 3.1, and we will therefore avoid much of the exposition from Section 3.1.

For all $\theta$, since the MTR assumption holds, the lower bound is $\text{LB}_\theta = \text{LB}_{\text{MTR}} = 0$. That is, by assumption, the lower bound is 0. The upper bound is a convex combination of the upper bounds for the data that do not follow the MTS assumption and for the data that follow the MTS assumption. Assuming $0 < \text{P}[z = 1|d = 1, t_{\text{MTS}} = 1] < 1$, this gives[13]

$$\begin{aligned}\text{UB}_\theta =&\zeta\{\text{P}[d = 0|\, t_{\text{MTS}} = 0, \omega = 0] + (\text{P}[z = 0|d = 1, t_{\text{MTS}} = 0, \omega = 0]\\ &+ \text{P}[y = 1|d = 1, z = 1, t_{\text{MTS}} = 0, \omega = 0]\text{P}[z = 1|d = 1, t_{\text{MTS}} = 0, \omega = 0]\\ &- \text{P}[y = 1|d = 1, z = 0, t_{\text{MTS}} = 0, \omega = 0]\text{P}[z = 0|d = 1, t_{\text{MTS}} = 0, \omega = 0])\\ &\times \text{P}[d = 1|t_{\text{MTS}} = 0, \omega = 0]\}\\ &+ (1 - \zeta)\{(\text{P}[y = 1|d = 1, z = 1, t_{\text{MTS}} = 1, \omega = 0]\\ &- \text{P}[y = 1|d = 1, z = 0, t_{\text{MTS}} = 1, \omega = 0])\text{P}[d = 1|\, t_{\text{MTS}} = 1, \omega = 0]\\ &+ \text{P}[d = 0|t_{\text{MTS}} = 1, \omega = 0]\}\\ =&\,\kappa_1 + \kappa_2,\end{aligned}$$

where

$$\begin{aligned}\kappa_1 &= \theta_1\xi_1 + \theta_5\xi_5 + \theta_2\xi_2\\ \kappa_2 &= (1 - \theta_1)\xi_1 + \kappa_3((1 - \theta_2)\xi_2 + (1 - \theta_3)\xi_3 + (1 - \theta_4)\xi_4 + (1 - \theta_5)\xi_5)\end{aligned}$$

with

$$\kappa_3 = \frac{(1 - \theta_2)\xi_2}{(1 - \theta_2)\xi_2 + (1 - \theta_3)\xi_3} - \frac{(1 - \theta_4)\xi_4}{(1 - \theta_4)\xi_4 + (1 - \theta_5)\xi_5}.$$

Much like the MTR case, a likelihood function identifies plausible values of $\theta$:

$$L_\zeta(\theta|\Xi = \xi) = \text{I}(\zeta = \xi'\theta)\text{I}(\kappa_3 \geq 0).$$

Recalling (3), the condition $\kappa_3 \geq 0$ reflects the fact that the observed treatment outcomes must be, on average, larger for the observed treated outcomes for the MTS-respecting subpopulation if MTR and MTS jointly hold. If this condition is violated, the posterior probability on $\theta$ is 0.

[13] Derivations are in supplementary appendix E.

https://doi.org/10.1093/pan/mps043 Published online by Cambridge University Press

### 3.3 *Computation for Sensitivity Analysis*

We now discuss computation of salient quantities from the posterior distribution. To estimate the posterior distributions of bounds, we perform a procedure analogous to rejection sampling. To perform a sensitivity analysis for the MTR assumption, we sample a random value of $\upsilon$ from the uniform prior distribution on $\Lambda$, then compute $\omega = \xi'\upsilon$ and $\mathrm{LB}_\upsilon$. For the MTS assumption, we sample a random draw of $\theta$ from the prior distribution, reject the draw if $\kappa_3 < 0$, and, if not rejected, we record $\zeta = \xi'\theta$ and $\mathrm{UB}_\theta$. We repeat these procedures for a large number of replications.

We use the fact that $\mathrm{LB}_\upsilon$ and $\mathrm{UB}_\theta$ are continuous and smooth functions of arguments $(\omega, \upsilon)$ and $(\zeta, \theta)$, respectively, to compute expected values and quantiles of the posterior distributions. The expectations and quantiles are also continuous and smooth functions of the arguments, so we may estimate the expected values and quantiles for each $\omega$ or $\zeta$ using nonparametric curve fitting. To estimate an expectation for the lower bound at $\omega$ or the upper bound at $\zeta$, we use smoothing splines in a generalized additive model (Hastie 1992).[14] To estimate quantiles, we use neural network quantile regression (Taylor 2000). The predicted values from these fits can then be presented graphically or examined numerically.

## 4 Inference under Random Sampling

If we have sample survey data and not population data, then the probabilities $\mathrm{P}(\cdot)$ or $\xi$ must be estimated instead of simply observed. Bounds on the ATE may be estimated using the sample analog of the population bounds, but construction of confidence intervals (CIs) is somewhat more involved. We describe the method, based on a bootstrap, that is developed by Molinari (2010) and gives an identification interval that asymptotically has probability $1-\alpha$ of including the population identification interval. One might describe this as a $100(1-\alpha)\%$ CI.

In the case of our sensitivity analysis, sample inference needs to be combined with the Bayesian inference that underpins our assessment of the two monotonicity assumptions. We combine Molinari's bootstrap method with a method for estimating the sampling uncertainty in the spline and quantile curve estimates. At each of a large number of values of $\omega, \zeta \in [0, 1]$, we compute a CI for the identification interval for the ATE. The same theory that motivates Molinari's CIs shows that the closure of the union of all such CIs gives a confidence region that asymptotically has probability $1-\alpha$ of including the corresponding region for the population.

### 4.1 *Estimating CIs*

Citing results from Beresteanu and Molinari (2008) for set-valued random variables, Molinari (2010, 91) emphasizes that the focus of identification is on the identification interval $[\vartheta_\mathrm{L}, \vartheta_\mathrm{U}]$ between lower bound $\vartheta_\mathrm{L}$ and upper bound $\vartheta_\mathrm{U}$, not on the bounds per se. To estimate a CI for this identification interval using sample data, the theory of Beresteanu and Molinari (2008) allows the sample inference problem to focus on the bounds. Consistent estimates for the bounds are based on the sample expected values of the probabilities. Added to these initial estimates are quantities determined by a bootstrap method. If $\vartheta_\mathrm{L}$, $\vartheta_\mathrm{U}$, $\hat{\vartheta}_\mathrm{L}$, and $\hat{\vartheta}_\mathrm{U}$ are, respectively, population and estimated lower and upper bounds for the interval $[\vartheta_\mathrm{L}, \vartheta_\mathrm{U}]$, and if for a sequence of sample sizes $n$

$$\sqrt{n}\begin{pmatrix} \hat{\vartheta}_\mathrm{L} - \vartheta_\mathrm{L} \\ \hat{\vartheta}_\mathrm{U} - \vartheta_\mathrm{U} \end{pmatrix} \xrightarrow{d} \begin{pmatrix} w_1 \\ w_2 \end{pmatrix} \sim N(\mathbf{0}, \Pi),$$

then $\hat{c}_{\alpha n}$ is an estimator for $c_\alpha = \inf\{t : \Pr\{\max\{(w_1)_+, (w_2)_-\} > t\} \geq \alpha\}$ that may be obtained using a bootstrap, such that a CI that asymptotically covers $[\vartheta_\mathrm{L}, \vartheta_\mathrm{U}]$ with probability $1-\alpha$ is $[\hat{\vartheta}_\mathrm{L} - \hat{c}_{\alpha n}/\sqrt{n}, \hat{\vartheta}_\mathrm{U} + \hat{c}_{\alpha n}/\sqrt{n}]$ (Molinari 2010, 91).[15]

[14] Wahba (1990, 16–8) discusses smoothing splines as Bayes estimates.
[15] $(w_1)_+$ denotes $\max\{0, w_1\}$ and $(w_2)_-$ denotes $\max\{0, -w_2\}$.

https://doi.org/10.1093/pan/mps043 Published online by Cambridge University Press

### 4.2 *Confidence Regions for Sensitivity Analysis*

Our method for sample data combines our Bayesian method for weakening the monotonicity assumptions with the asymptotic sample-based method derived by Beresteanu and Molinari (2008) (a bootstrap). The combination can be considered analogous to a fully Bayesian method for a set of special cases. First, we implicitly invoke an assumption that units are exchangeable, given $\xi$, in the way we use the same exchangeable prior uniform distributions for $\upsilon$ and $\theta$ with all observations. Moreover, we assume observations are exchangeable in the way we make our sample inferences. Because we have exchangeable priors, treat the samples as simple random samples, and use only symmetric statistics (like proportions and functions of proportions), the posterior distributions are independent of the sampling design (Scott and Smith 1973, 59). If we had informative sampling designs, or perhaps if we knew more about how the sample data we are using were produced, we might gain from using different sample inference methods (Sugden and Smith 1984). As it is, the independence assumptions and the fact that all possible values of the variables are observed can support a Bayesian motivation for a bootstrap (Rubin 1981).

The complication that stands in the way of a fully Bayesian treatment is the fact that, following Beresteanu and Molinari (2008), we are dealing with set-valued random variables. Although Molinari (2010) considers CIs for identification intervals, we consider confidence regions for identification regions (defined immediately below). Bayesian bootstrap methods such as those that Newton and Raftery (1994) and Efron (2011) discuss have so far considered only scalar- and vector-valued random variables. We say more about this complication after defining the identification regions and confidence regions we use.

We associate $\mathrm{LB}_\upsilon$ and $\mathrm{UB}_\theta$ with four identification regions to connect with Molinari's bootstrap method: $[E_\omega(\mathrm{LB}_\upsilon), \mathrm{UB_{TE}}]$, where $E_\omega(\mathrm{LB}_\upsilon)$ denotes the Bayesian expectation of $\mathrm{LB}_\upsilon$ evaluated at each $\omega \in [0,1]$; $[0, E_\zeta(\mathrm{UB}_\theta)]$, where $E_\zeta(\mathrm{UB}_\theta)$ denotes the Bayesian expectation of $\mathrm{UB}_\theta$ evaluated at each $\zeta \in [0,1]$; $[Q_{\omega,.025}(\mathrm{LB}_\upsilon), Q_{\omega,.975}(\mathrm{LB}_\upsilon)]$, where $Q_{\omega,q}(\mathrm{LB}_\upsilon)$ denotes the Bayesian $q$-quantile of $\mathrm{LB}_\upsilon$ evaluated at each $\omega \in [0,1]$; and $[Q_{\zeta,.025}(\mathrm{UB}_\theta), Q_{\zeta,.975}(\mathrm{UB}_\theta)]$, where $Q_{\zeta,q}(\mathrm{UB}_\theta)$ denotes the Bayesian $q$-quantile of $\mathrm{UB}_\theta$ evaluated at each $\zeta \in [0,1]$.[16] So for MTR, for instance, we are interested in using the sample data to estimate the region between the spline curve that represents $E_\omega(\mathrm{LB}_\upsilon)$ along $\omega \in [0,1]$ and the line $[0,1] \times \mathrm{UB_{TE}}$ associated with the upper bound point.

Using arguments from Wang and Wahba (1994) and Beresteanu and Molinari (2008), we define a bootstrap based on elements $\hat{c}_{i\alpha n}$, $i \in [0,1]$, that gives a confidence region that asymptotically with probability $1-\alpha$ covers the region that would arise if we had population data. Let the smoothing spline model be written $b_i = f(s_i) + \epsilon_i$ for bound values $b_i$, function $f$,[17] design points $s_i$, and disturbance $\epsilon_i$, so that the smoothing spline in the original sample is $\hat{f}_{\hat{\lambda}}$ for estimate $\hat{\lambda}$ of smoothing parameter $\lambda$ (Wang and Wahba 1994). We can implement Molinari's bootstrap corrections in our extension of her model by using $B$ bootstrap samples to compute replicates $\check{\xi}^k$ of the original sample estimate $\hat{\xi}$, $k = 1, \ldots, B$, and then to estimate smoothing splines using $\check{\xi}^k$ from each of the $k$ bootstrap replicates (with the draws of $\upsilon$ or $\theta$ from the prior and the design points held constant) to give bootstrap estimates $\check{f}^k_{\check{\lambda}^k}$. Values $\hat{c}_{\alpha n}$ can then be estimated at each design point: call each of the estimated values $\hat{c}_{i\alpha n} \equiv \hat{c}_{s_i \alpha n}$, where each $i$ is identified with its design point, $i \equiv s_i \in [0,1]$; if $\hat{\vartheta}_{\mathrm{L}i}$ and $\hat{\vartheta}_{\mathrm{U}i}$ denote the upper and lower bound values estimated using the original sample at each design point, $\check{\vartheta}^k_{\mathrm{L}i}$ and $\check{\vartheta}^k_{\mathrm{U}i}$ denote estimates computed using the $k$th bootstrap resample and $\check{w}^k_{1i} = \sqrt{n}(\check{\vartheta}^k_{\mathrm{L}i} - \hat{\vartheta}_{\mathrm{L}i})$ and $\check{w}^k_{2i} = \sqrt{n}(\check{\vartheta}^k_{\mathrm{U}i} - \hat{\vartheta}_{\mathrm{U}i})$, then $\hat{c}_{i\alpha n}$ is the $1-\alpha$ order statistic of $\max\{(\check{w}^k_{1i})_+, (\check{w}^k_{2i})_-\}$.[18] Across all $i \in [0,1]$ the $\hat{c}_{i\alpha n}$ values estimated at each design point give bootstrap CI estimates, $[\hat{\vartheta}_{\mathrm{L}i} - \hat{c}_{i\alpha n}/\sqrt{n}, \hat{\vartheta}_{\mathrm{U}i} + \hat{c}_{i\alpha n}/\sqrt{n}]$, that have the average coverage probability property that Wang and Wahba (1994) demonstrate for several kinds of bootstrap intervals for smoothing splines.

[16] Bounds are appropriately reversed if the assumptions concern negative monotonicity. See Section 5.2.

[17] $f \in \mathbf{W}_m$ where $\mathbf{W}_m = \{f : f, f', \cdots, f^{(m-1)}$ are absolutely continuous, $\int_0^1 (f^{(m)})^2 dt < \infty\}$ (Wang and Wahba 1994).

[18] See Beresteanu and Molinari (2008, 773, Algorithm 2.1). Our function $\max\{(\check{w}^k_{1i})_+, (\check{w}^k_{2i})_-\}$ corresponds to Beresteanu and Molinari's $\sqrt{n}H(\cdot)$.

https://doi.org/10.1093/pan/mps043 Published online by Cambridge University Press

The "confidence collection" result of Beresteanu and Molinari (2008, 775) then implies that the confidence region that is the closure of the union of these CIs has asymptotically the desired coverage probability. Denote the interval estimated at each design point after effecting the closure operation by $[\tilde{\vartheta}_{\mathrm{L}i}, \tilde{\vartheta}_{\mathrm{U}i}]$. Not pointwise but on average across all $i \in [0,1]$, $[\tilde{\vartheta}_{\mathrm{L}i}, \tilde{\vartheta}_{\mathrm{U}i}]$ asymptotically covers with probability $1-\alpha$ the corresponding population interval $[\vartheta_{\mathrm{L}i}, \vartheta_{\mathrm{U}i}]$.

For instance, for MTR, we find a $100(1-\alpha)\%$ confidence region $[\tilde{E}_\omega(\mathrm{LB}_\upsilon), \widetilde{UB}_{\mathrm{TE}}]$ that covers the region that would arise with population data between the spline curve over $\omega \in [0,1]$ that represents $E_\omega(\mathrm{LB}_\upsilon)$ and the line $[0,1] \times \mathrm{UB}_{\mathrm{TE}}$. The quantiles may be bootstrapped similarly (Molinari 2005; supplementary appendix B). In this case, we find confidence regions that asymptotically with probability $1-\alpha$ cover the set of 95% posterior intervals for $\omega,\zeta \in [0,1]$. To ensure that the confidence region for expectations is compatible with the region for the quantiles, the bootstrap method produces $\hat{c}_{i\alpha n}$ values that give conservative confidence regions that simultaneously cover both regions: asymptotically with probability at least $1-\alpha$, the confidence regions cover the posterior expectation and quantile regions that would arise if we had population data.[19]

Viewed as a conventional estimation problem, Wang and Wahba (1994) suggest that the bootstrap method we use produces CIs comparable to and as good as the Bayesian "confidence intervals" derived for smoothing splines by Wahba (1983). From that perspective, the posterior expectations we compute are likely similar to the ones we would derive if we used a nominally Bayesian "confidence interval." Of course we do not have a conventional estimation problem, because Molinari's bootstrap method is connected to the set-valued inference problem considered by Beresteanu and Molinari (2008), and the spline estimates concern only edges of the identification region that is of actual interest to us.

## 5 Applications

To demonstrate the applicability of Molinari bounds, we consider two examples. We use these examples because (1) they are emblematic of data used in political science; (2) the treatment is not randomly assigned; and (3) some values for the treatment variable are missing. The first example draws from observational data on government torture and civil war, using population data. The second example utilizes random samples collected as part of the popular World Values Survey (WVS).

### 5.1 *Effect of Torture on Civil War*

A growing segment of civil-war scholars focus on state repression as a causal factor (Lichbach 1987; Tarrow 1989; Petersen 2002; Sambanis and Zinn 2006; Davenport 2007a, 2007b). State repression includes acts such as direct rule of ethnic minorities, extrajudicial killings, and torture. According to Sambanis and Zinn (2006), government repression can lead to civil war via a number of channels (Sambanis and Zinn 2006, 8). For one, state repression lowers the opportunity cost of violent actions against the state (if one is going to be repressed anyhow, there is less to lose by attacking the state). Alternatively, state repression bolsters the resolve of society's more radical elements. In either case, repression can motivate the taking up of arms against the state.

To test these claims in a systematic way, scholars draw upon extensive data sets that measure civil war onset and state repression. The Cingranelli–Richards (CIRI) Human Rights Dataset (Cingranelli and Richards 2010) includes as an indicator of state repression the variable *Torture*, which records whether torture is practiced frequently, occasionally, or has not occurred in a given year within a given country from 1981 to 2006. The UCDP/PRIO Armed Conflict Dataset

[19] To ensure that the confidence region for posterior expectations is compatible with the region for the posterior quantiles, the bootstraps for the two regions are combined. Let $(\check{w}^k_{1iE}, \check{w}^k_{2iE})$ denote bootstrap differences for expectations and $(\check{w}^k_{1iQ}, \check{w}^k_{2iQ})$ differences for the corresponding quantiles. We use $\hat{c}_{i\alpha n}$ values given by the $1-\alpha$ order statistics of $\max\{(\check{w}^k_{1iE})_+, (\check{w}^k_{2iE})_-, (\check{w}^k_{1iQ})_+, (\check{w}^k_{2iQ})_-\}$ to obtain confidence regions that have asymptotically probability at least $1-\alpha$ of covering the population regions.

https://doi.org/10.1093/pan/mps043 Published online by Cambridge University Press

**Table 1** Joint distribution of torture and civil war

| | Civil war in year $t$ | | |
|---|---|---|---|
| Torture in year $t-1$ | No (0) | Yes (1) | Total |
| No (0) | 1840 | 150 | 1998 |
| Yes (1) | 948 | 399 | 1347 |
| Missing | 474 | 67 | 541 |
| Total | 3262 | 624 | 3886 |

**Table 2** Bounds for ATE of torture on subsequent civil war

| Effect of torture on civil war | | |
|---|---|---|
| Method | ATE bounds | |
| Manski (excluding missing data) | −0.33 | 0.67 |
| Molinari | −0.42 | 0.72 |
| Molinari (MTR) | 0 | 0.72 |
| Molinari (MTR and MTS) | 0 | 0.33 |

(Gleditsch et al. 2002) measures instances of civil war onset from 1946 to 2004. We examine the effect that torture in one year has on the onset of civil war in the following year during the years these data sets overlap: 1981–2004.

We address two of the major impediments standing in the way of testing the causal relationship between torture and civil war onset. First, the existence of torture within countries is not randomly assigned. Political leaders systematically assess the relative costs and benefits associated with repression, the availability of other options, and the likelihood of success (Davenport 2007b, 490). Second, data on the incidence of torture are sometimes missing, and such missing data are likely not missing at random. Table 1, which presents the joint distribution of the *Torture* and *Civil War* variables, shows that 541 of the 3886 values for the *Torture* variable are missing. In contrast, there are no missing values for the *Civil War* dependent variable. We let $d = 1$ (meaning we observe the treatment) when *Torture* has a value of either 0 or 1, and $d = 0$ for those observations for which *Torture* is coded as missing.

We begin our analysis by computing the Manski bounds based on ignoring the missing torture data. Because Table 1 shows that 541 observations have missing treatment data, these observations are excluded. The Manski bounds, [−0.33,0.67], are displayed in Table 2. Because these bounds ignore the missing observations, it describes identification bounds for the causal effect in a substantively meaningless subpopulation: country-years with nonmissing torture data.

Because the population of interest includes both missing and nonmissing data, we compute Molinari bounds. The Molinari bounds for the ATE, reported in Table 2, are [−0.42,0.72]. Because they incorporate greater uncertainty, the Molinari bounds are wider than the Manski bounds.

We may achieve narrower bounds through additional assumptions. Because theory suggests that torture will have either a positive or no impact on the probability of civil war, it is plausible to apply a weakly positive MTR assumption. Alone, this assumption merely raises the lower bound on the treatment effect from −0.42 to 0. A further assumption that may be plausible is that countries in which torture occurs are more likely to experience civil war (MTS) regardless of whether torture actually occurs. Given the MTR assumption, this amounts to the weakly positive MTS assumption. One argument that can support the MTS assumption is based on the likelihood that governments apply torture when there are individuals or groups that hold grievances against the government. If the presence of such groups importantly contributes to civil war onset (Cederman, Weidmann, and

https://doi.org/10.1093/pan/mps043 Published online by Cambridge University Press

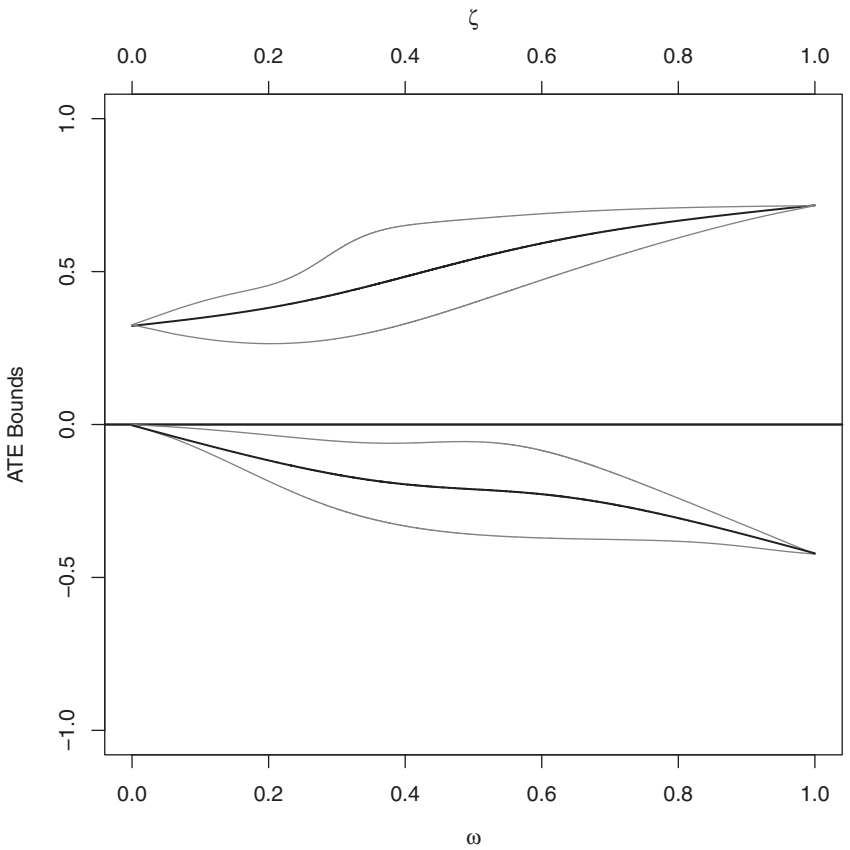



**Fig. 1** Sensitivity analysis for effect of torture on civil war. Solid black lines represent expected values of bounds assuming uniform priors. Gray lines represent 95% posterior intervals for bounds assuming uniform priors. Twenty-five thousand draws from the prior used to simulate posterior distribution.

Gleditsch 2011), then the MTS assumption is plausible. Table 2 shows that making both monotonicity assumptions implies that torture increases the probability of civil war by no more than 0.33.

Because we may have uncertainty about the MTR and MTS assumptions, it is useful to assess how sensitive the apparent causal effect is to deviations from the MTR and MTS assumptions. Figure 1 displays results from the sensitivity analysis. The vertical axis of Fig. 1 shows the values of the upper and lower bounds on the ATE: the value of the upper bound is shown above 0; the value of the lower bound is shown below 0. The top horizontal axis reports the values of $\zeta$, whereas the bottom horizontal axis reports the values of $\omega$. In Fig. 1, the solid black lines represent expected values of bounds assuming uniform priors for $\upsilon$ and $\theta$ over $\Lambda$, and the solid gray lines represent the equal-tailed 95% posterior intervals.

The graph shows how the identification bounds for the ATE are sensitive to the monotonicity assumptions.[20] The part of the graph below 0 shows the lower bound assuming that $\omega$ proportion of all the data do not follow the MTR assumption (without making any MTS assumption). As $\omega$ rises, the posterior expected value first gradually declines, and the posterior 95% intervals associated with $\mathrm{LB}_\upsilon$ gradually expand then gradually narrow as $\omega$ approaches 1. The lower bound of the posterior 95% interval ($Q_{\omega,0.025}(\mathrm{LB}_\upsilon)$) is only slightly greater than $\mathrm{LB}_{\mathrm{TE}}$ when $\omega = 0.5$. The upper bound of the 95% posterior interval ($Q_{\omega,0.975}(\mathrm{LB}_\upsilon)$) does not vary much for $\omega \lessapprox 0.6$, but declines steadily as $\omega$ increases beyond that. The part of the graph above 0 shows the upper bounds when MTR is assumed for all the data ($\omega = 0$) and $\zeta$ proportion of the data do not

[20]Smoothing spline and quantile estimates are based on 25,000 draws from the prior distributions of $\upsilon$ and $\theta$.

https://doi.org/10.1093/pan/mps043 Published online by Cambridge University Press

satisfy the MTS assumption. The expected value of the upper bound gradually rises with $\zeta$. For example, when $\zeta = 0.5$, the expected value is 0.55. As $\zeta$ increases, the 95% posterior intervals associated with $\text{UB}_\theta$ gradually expand to admit a wide range of ATE values, and as the probability of the MTS assumption being true approaches 0, the bounds gradually shrink to admit a narrower range of values. The upper bound of the 95% interval ($Q_{\zeta,0.975}(\text{UB}_\theta)$) is near $\text{UB}_{\text{TE}}$ for $\zeta \gtrapprox 0.35$.

Considering both sets of bounds, inference about the ATE may not depend all that critically on the monotonicity assumptions. Fairly weak assumptions about MTR suffice to give similar results: the lower bound of the ATE decreases gradually as the proportion of the data believed to satisfy MTR decreases; the lower bound does not vary much for $0.25 < \omega < 0.60$. The upper bound ATE results are slightly more sensitive to the MTS assumption: the lower bound of the 95% interval for $\zeta \gtrapprox 0.3$ is greater than the upper bound of the 95% interval for $\zeta \approx 0$.

### 5.2 *Effect of Familial Immigrants on Immigration Policy Attitudes*

Several recent studies consider the policy implications of attitudes toward immigrant workers (Hainmueller and Hiscox 2007, 2010). Focusing primarily on how skill level affects a respondent's views toward immigration policy, these studies also consider how having friends who are immigrants affect the respondent's views (Hainmueller and Hiscox 2007, 430). In a probit regression model estimated using survey data, Hainmueller and Hiscox (2007, 432) find that individuals who have immigrant friends tend to hold more favorable views toward immigration.

To examine the relationship between having immigrant friends or family and views on immigration policy, we use data from the 2005 wave of the WVS, a comprehensive survey of political and sociocultural attitudes covering more than 50 countries (World Values Survey Association 2009). We consider separately WVS data from Germany ($n = 1659$) and from Italy ($n = 817$). To measure the outcome variable—attitudes toward immigrants—we use question 45 of the survey, which asks whether respondents agree, disagree, neither, or could provide no response to the question "When jobs are scarce, employers should give priority to [British] people over immigrants," where the nationality of citizens of the referent country is used in place of the word "[British]." We restrict our population to respondents who either agreed or disagreed with question 45.

The treatment variable is whether the respondent has an immigrant parent, measured using question 215 ("Is your father an immigrant of this country or not?") and question 216 ("Is your mother an immigrant of this country or not?"). To each of these questions the respondent could provide answers of "not an immigrant," "immigrant," or "don't know/could not provide an answer." A respondent who answered yes to either was recorded as having an immigrant parent, and a respondent who answered no to both was recorded as not having an immigrant parent. A respondent who did not answer yes to either, but had at least one missing value, was recorded as having missing treatment data.

Analysts may have two concerns regarding their ability to draw causal inferences in this application. First, the immigration status of a respondent's parents is not randomly assigned. Second, some individuals might be unwilling to divulge such information (perhaps, e.g., out of fear of deportation). As highlighted by a recent study on how vague legal status affects the well-being of children in Canada, children living with an undocumented parent viewed teachers, nurses, police officers, and other authority figures with distrust. The children feared that such figures could have their parents detained or deported, and, if that occurred, "they will never be reunited with their parents" (Bernhard et al. 2007, 103). Such motivations might lead to nonresponses to the parental immigration status questions. This concern appears supported by the WVS data. Table 3 shows that in each country 13 of the individuals who responded to question 45 have missing treatment data, even though these individuals provided other demographic information. To identify the effect of having immigrant parents on an individual's immigration policy views, one must account for both the nonrandom assignment of the treatment and the missing treatment data.

As with the previous example, we begin by computing the Manski bounds which exclude the observations with missing treatment data. Estimates of the Manski bounds and $100(1-\alpha)\%$ CIs ($\alpha = 0.05$, $B = 500$ bootstrap replications) are reported in Table 4. ATE bounds are similar for

https://doi.org/10.1093/pan/mps043 Published online by Cambridge University Press

**Table 3** Joint distribution of immigrant parent and support for anti-immigrant policies

| | Germany | | | Italy | | |
|---|---|---|---|---|---|---|
| | Support for anti-immigrant policies | | | Support for anti-immigrant policies | | |
| Immigrant parent | No (0) | Yes (1) | Total | No (0) | Yes (1) | Total |
| No (0) | 450 | 1053 | 1503 | 184 | 608 | 792 |
| Yes (1) | 78 | 65 | 143 | 3 | 9 | 12 |
| Missing | 5 | 8 | 13 | 1 | 12 | 13 |
| Total | 533 | 1126 | 1659 | 188 | 629 | 817 |

**Table 4** Bounds for ATE of immigrant parent on support for anti-immigrant policies. $B = 500$ bootstrap replications used to estimate $100(1 - \alpha)\%$ CIs with $\alpha = 0.05$

| | Germany | | | | Italy | | | |
|---|---|---|---|---|---|---|---|---|
| Method | ATE Bounds | | 95% CI | | ATE Bounds | | 95% CI | |
| Effect of immigrant parent on support for anti-immigrant policies | | | | | | | | |
| Manski (excluding missing data) | −0.69 | 0.31 | −0.72 | 0.35 | −0.76 | 0.24 | −0.81 | 0.29 |
| Molinari | −0.69 | 0.32 | −0.73 | 0.36 | −0.76 | 0.25 | −0.81 | 0.30 |
| Molinari (MTR) | −0.69 | 0 | −0.73 | 0 | −0.76 | 0 | −0.81 | 0 |
| Molinari (MTR and MTS) | −0.25 | 0 | −0.38 | 0 | −0.03 | 0 | −0.46 | 0 |

both Germany and Italy, but the CI is wider for Italy ($[-0.81, 0.29]$) than for Germany ($[-0.72, 0.35]$), a reflection of the difference in sample sizes. To account for the missing treatment data, we compute the Molinari bounds (Table 4). The Molinari bounds are very slightly wider, due to the fact that Molinari bounds incorporate greater uncertainty, and again they are very similar between the two countries. The CI is again wider for Italy ($[-0.81, 0.30]$) than for Germany ($[-0.73, 0.36]$).

We may again achieve narrower bounds through additional assumptions. Because we expect that having an immigrant parent can only make the respondent *less* likely to support government policies that discriminate against immigrants, we apply a weakly negative MTR assumption. Alone, this assumption merely reduces the upper bound on the ATE from 0.36 or 0.30 to 0. To add the corresponding negative MTS assumption, one must assume that, on average, individuals with at least one immigrant parent would have been less likely to support government policies that favor nationals even if they did not have that immigrant parent. This counterfactual may be more difficult to comprehend, particularly if one thinks of an individual's parentage as an attribute of the individual (Holland 1986), but it is at least conceivable to conjecture that, even in the absence of an immigrant parent, a counterfactual respondent who was otherwise in the same social position as the respondent would be more likely to have connections to immigrants and this exposure would reduce support for policies that discriminate against immigrants.

The intervals for the Molinari bounds shrink considerably once we invoke the MTR and MTS assumptions. Making both of these weakly negative monotonicity assumptions and using the CI estimates, Table 4 shows that having immigrant parents decreases the probability of supporting government policies that favor nationals by no more than 0.38 in Germany and no more than 0.46 in Italy. Accounting for sampling error is particularly important. The ATE lower bound estimate ignoring sampling error is −0.25 in Germany and −0.03 in Italy, dramatically different from the CI lower bounds of −0.38 and −0.46. Clearly the MTS assumption makes a big difference, but not taking sampling error into account would lead to substantially overestimating the difference.

https://doi.org/10.1093/pan/mps043 Published online by Cambridge University Press

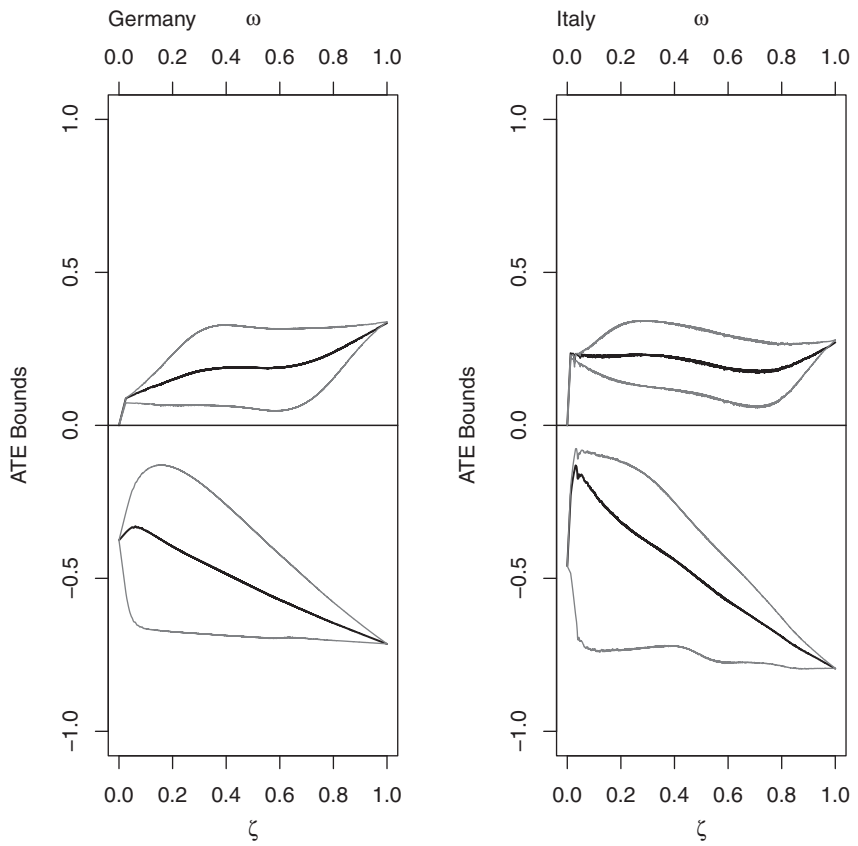



**Fig. 2** Sensitivity analysis for effect of immigrant parent on anti-immigrant policies, German and Italian World Values Survey samples. Solid black lines represent boundaries of $100(1-\alpha)\%$ confidence regions for expected values of bounds assuming uniform priors. Gray lines represent boundaries of $100(1-\alpha)\%$ confidence regions for 95% posterior intervals assuming uniform priors. Twenty-five thousand draws from the prior used to simulate posterior distribution. Five hundred bootstrap replications used to estimate boundaries for $\alpha = 0.05$.

The results from using sensitivity analysis to reflect uncertainty about the MTR and MTS assumptions appear in Fig. 2. Figure 2 presents the sets of bounds associated with the conservative $100(1-\alpha)\%$ confidence regions defined for the sensitivity analysis, with $\alpha = 0.05$. Adapting previously defined notation to refer to the negative monotonicity case, the four regions we consider when estimating confidence regions are $[\text{LB}_{\text{TE}}, E_{\omega}(\text{UB}_{\upsilon})]$ and $[Q_{\omega,0.025}(\text{UB}_{\upsilon}), Q_{\omega,0.975}(\text{UB}_{\upsilon})]$, pertaining to MTR, and $[E_{\zeta}(\text{LB}_{\theta}), 0]$ and $[Q_{\zeta,0.025}(\text{LB}_{\theta}), Q_{\zeta,0.975}(\text{LB}_{\theta})]$, pertaining to MTR-MTS. As in Fig. 1, the vertical axis of Fig. 2 shows the values of the upper and lower bounds on the ATE: the value of the upper bound is shown above 0; the value of the lower bound is shown below 0. Because negative and not positive monotonicity assumptions are of interest, the upper part of Fig. 2 now reports deviations from MTR, whereas the lower part reports deviations from MTS. Therefore, unlike Fig. 1, the top horizontal axis now reports the values of $\omega$ whereas the bottom horizontal axis reports the values of $\zeta$. As with Fig. 1, the solid black lines represent expected values of bounds assuming uniform priors for $\upsilon$ and $\theta$ over $\Lambda$, and the solid gray lines represent the equal-tailed 95% posterior intervals.

On the left in Fig. 2—for Germany—the estimated bounds on the ATE show a gradual response to changes in beliefs about the MTR assumption. The upper bound first jumps rapidly up from 0. This reflects the effect of sampling: the upper bound at $\omega = 0$ is exactly 0, but for every value of $\omega > 0$ the confidence region's boundary increases by $\hat{c}_{ian}$ as a consequence of sampling uncertainty. After the estimated expected value of the upper bound, $\tilde{E}_{\omega}(\text{UB}_{\upsilon})$, initially jumps to a value of

https://doi.org/10.1093/pan/mps043 Published online by Cambridge University Press

$\tilde{E}_{0.026}(\mathrm{UB}_\upsilon) \approx 0.088$ at $\omega = 0.026$, both $\tilde{E}_\omega(\mathrm{UB}_\upsilon)$ and the estimated upper boundary of the 95% posterior interval, $\tilde{Q}_{\omega,0.975}(\mathrm{UB}_\upsilon)$, gradually increase as $\omega$ increases, whereas the estimated lower boundary of the 95% posterior interval, $\tilde{Q}_{\omega,0.025}(\mathrm{UB}_\upsilon)$, slightly decreases from its post-jump value of $\tilde{Q}_{0.026,0.025}(\mathrm{UB}_\upsilon) \approx 0.074$. $\tilde{E}_\omega(\mathrm{UB}_\upsilon)$ increases until $\omega \approx 0.42$, when $\tilde{E}_{0.42}(\mathrm{UB}_\upsilon) \approx 0.19$, then remains relatively constant until $\omega \approx 0.57$ at which point $\tilde{E}_\omega(\mathrm{UB}_\upsilon)$ resumes increasing up to $\tilde{E}_1(\mathrm{UB}_\upsilon) = 0.36$. $\tilde{Q}_{\omega,0.975}(\mathrm{UB}_\upsilon)$ increases up to $\omega \approx 0.40$, where $\tilde{Q}_{0.40,0.975}(\mathrm{UB}_\upsilon) \approx 0.33$, after which it remains relatively level up to $\omega = 1$. $\tilde{Q}_{\omega,0.025}(\mathrm{UB}_\upsilon)$ continues decreasing down to $\omega \approx 0.59$, when $\tilde{Q}_{0.59,0.025}(\mathrm{UB}_\upsilon) \approx 0.047$, then steadily increases up to $\omega = 1$. Necessarily, $\tilde{E}_1(\mathrm{UB}_\upsilon) = \tilde{Q}_{1,0.025}(\mathrm{UB}_\upsilon) = \tilde{Q}_{1,0.975}(\mathrm{UB}_\upsilon)$. The posterior expected value and quantiles are nearly constant for $0.42 \lessapprox \omega \lessapprox 0.59$, change gradually outside that range, and the initial jump up right above 0 is not all that large.

In contrast, the graph on the right in Fig. 2, for Italy, shows that ATE bounds respond suddenly to changes in beliefs about the MTR assumption. The jump in the upper bound just above $\omega = 0$ is large: $\tilde{E}_\omega(\mathrm{UB}_\upsilon)$ jumps to $\tilde{E}_{0.0125}(\mathrm{UB}_\upsilon) \approx 0.236$, and $\tilde{E}_\omega(\mathrm{UB}_\upsilon)$ then remains relatively constant until $\omega \approx 0.284$, when it starts to decline; but at the point where $\tilde{E}_\omega(\mathrm{UB}_\upsilon)$ reaches its minimum, at $\omega \approx 0.84$, we find $\tilde{E}_{0.84}(\mathrm{UB}_\upsilon) \approx 0.200$, not much smaller than the post-jump value $\tilde{E}_{0.0125}(\mathrm{UB}_\upsilon)$. Moreover, $\tilde{E}_1(\mathrm{UB}_\upsilon) = 0.30$, again not much bigger than $\tilde{E}_{0.0125}(\mathrm{UB}_\upsilon)$. The 95% posterior intervals expand around $\tilde{E}_\omega(\mathrm{UB}_\upsilon)$ gradually as $\omega$ increases from 0 and converge back gradually to $\tilde{E}_\omega(\mathrm{UB}_\upsilon)$ as $\omega$ approaches 1. The distance between $\tilde{Q}_{\omega,0.025}(\mathrm{UB}_\upsilon)$ and $\tilde{Q}_{\omega,0.975}(\mathrm{UB}_\upsilon)$ remains fairly constant for $0.2 \lessapprox \omega \lessapprox 0.7$, with the maximum of $\tilde{Q}_{\omega,0.975}(\mathrm{UB}_\upsilon)$ being $\tilde{Q}_{0.293,0.975}(\mathrm{UB}_\upsilon) \approx 0.34$ and the minimum of $\tilde{Q}_{\omega,0.025}(\mathrm{UB}_\upsilon)$ being $\tilde{Q}_{0.708,0.025}(\mathrm{UB}_\upsilon) \approx 0.058$. So, the jump from the upper bound of 0 at $\omega = 0$ to the value of the estimated upper bound at $\omega = 0.0125$ is large relative to the magnitude the upper bound has throughout the range of $\omega > 0$. Slightly weakening the MTR assumption would produce almost as large an effect on the upper bound as giving it up entirely. The MTR assumption would be more influential—and the consequent analysis less robust—for Italy than for Germany.

Given MTR, the implication of invoking MTS seems sudden in both countries, albeit sudden in a different sense: uncertainty about the location of the lower bound increases substantially and immediately as $\zeta$ increases above 0. For Germany, referring to the left part of Fig. 2 below 0, the estimated expected value, $\tilde{E}_\zeta(\mathrm{LB}_\theta)$, first rises from $\tilde{E}_0(\mathrm{LB}_\theta) = -0.38$ through $\tilde{E}_{0.026}(\mathrm{LB}_\theta) \approx -0.356$ to $\tilde{E}_{0.061}(\mathrm{LB}_\theta) \approx -0.330$, then declines steadily down to $\tilde{E}_1(\mathrm{LB}_\theta) = -0.73$. But at $\zeta = 0.061$, the estimated upper boundary of the posterior 95% interval, $\tilde{Q}_{\zeta,0.975}(\mathrm{LB}_\theta)$, is $\tilde{Q}_{0.061,0.975}(\mathrm{LB}_\theta) \approx -0.188$—which is near the maximum value $\tilde{Q}_{0.159,0.975}(\mathrm{LB}_\theta) \approx -0.130$—and the estimated lower boundary of the posterior 95% interval, $\tilde{Q}_{\zeta,0.025}(\mathrm{LB}_\theta)$, is $\tilde{Q}_{0.061,0.025}(\mathrm{LB}_\theta) \approx -0.644$, not that far from the minimum of $\tilde{Q}_{1,0.025}(\mathrm{LB}_\theta) = -0.73$.

Uncertainty about the lower bound increases rapidly and substantially in both countries as $\zeta$ grows even slightly greater than 0, but in Italy even $\tilde{E}_\zeta(\mathrm{LB}_\theta)$ changes substantially for slight departures from MTS. For Italy, referring to the right part of Fig. 2, the estimated expected value, $\tilde{E}_\zeta(\mathrm{LB}_\theta)$, first rises from $\tilde{E}_0(\mathrm{LB}_\theta) = -0.46$ through $\tilde{E}_{0.0125}(\mathrm{LB}_\theta) \approx -0.272$ to $\tilde{E}_{0.031}(\mathrm{LB}_\theta) \approx -0.132$, then declines steadily down to $\tilde{E}_1(\mathrm{LB}_\theta) = -0.81$. But at $\zeta = 0.031$ $\tilde{Q}_{\zeta,0.975}(\mathrm{LB}_\theta)$ is $\tilde{Q}_{0.031,0.975}(\mathrm{LB}_\theta) \approx -0.077$—near the maximum value of $\tilde{Q}_{\zeta,0.975}(\mathrm{LB}_\theta)$—and $\tilde{Q}_{\zeta,0.025}(\mathrm{LB}_\theta)$ is $\tilde{Q}_{0.031,0.025}(\mathrm{LB}_\theta) \approx -0.619$, well on the way toward the minimum of $\tilde{Q}_{1,0.025}(\mathrm{LB}_\theta) = -0.81$.[21]

Even though CI estimates for the ATE bounds do not differ all that much between the two countries, assessing the effect of weakening the MTR and MTS assumptions shows the countries to be very different. For Germany, slight departures from a homogeneous MTR assumption would not affect inferences about the ATE all that much, whereas for Italy such departures would have a sudden and large effect. Given MTR, in both countries slight departures from a homogeneous MTS assumption would immediately widen the scope for modeling assumptions to affect ATE estimates.

[21] Sensitivity analysis figures are based on 25,000 draws from the prior distributions. Using instead 100,000 prior draws produces a slightly smoother-looking graph for Italy that covers the same regions except the stairstep pattern in $\tilde{Q}_{\zeta,0.025}(\mathrm{LB}_\theta)$ for $0.05 \lessapprox \zeta \lessapprox 0.55$ is smoothed out. See supplementary appendix F.

https://doi.org/10.1093/pan/mps043 Published online by Cambridge University Press

In light of the posterior 95% intervals, assuming MTS prevails throughout almost all of the population would not be much more restrictive than assuming MTS does not hold at all.

## 6 Discussion

Causal inference is the product of both assumptions and data. We introduced a variation of Manski identification bounds for the ATE developed by Molinari (2010) that offers a relatively simple and nonparametric way to check the robustness of an estimated effect when treatment data are both nonrandomly assigned and missing. Because monotonicity assumptions may be critical for making Molinari bounds informative, we also developed a technique for determining the sensitivity of the bounds to these assumptions. This provides a systematic approach for determining the extent to which the monotonicity assumptions drive the results.

The implications of several notions of monotonicity have been studied in situations when treatments or outcomes are not binary (Manski 1997), but even with merely binary treatments and outcomes combinations of assumptions other than MTR or MTR-MTS may be considered. For instance, Manski and Pepper (2000, 1001) demonstrate bounds when MTS alone is assumed. One can imagine other patterns in which the monotonicity assumptions might be weakened besides the ones considered here. For example, within the context of the torture and civil war application, this might amount to assuming that the government applies torture when there are individuals or groups that hold grievances against the government, that these groups importantly contribute to civil war onset, but that torture can either effectively reduce the probability of civil war (perhaps because it provides usable intelligence to the government) or increase the probability of civil war. The substantive scientific basis for such alternative combinations would need to be clarified in order to motivate these investigations.

Referring to the monotonicity assumptions we consider, Manski (1997) presents a thorough discussion of MTR. One point worth noting is that MTR seems inherently more plausible than the assumption used in common practice with, for instance, regression analysis, in which one commits to the idea that an effect is the same for every observation but then consults the data to learn the sign and magnitude of the effect. If the previous supporting science is strong enough to show an effect is the same for every observation, then it ought to be strong enough to say whether the homogeneous effect is positive or negative. MTR does not entail any commitment to homogeneity, but using it does imply that prior relevant science can pin down an effect's sign.

A basic challenge to inferences about treatment effects with observational data is the near certainty that there is confounding, meaning that other covariates besides the treatment variable of interest are associated with both the treatment and the outcome. Indeed, this basic problem arises whenever treatments are not randomly assigned. The goal of the identification methods we have considered is to show the range of all possible effects that might be found with a model that takes relevant confounders into account. The goal is not to assess the effects of such confounders. The fact that identification regions can be computed after conditioning on a covariate profile (recall footnote 5) means that, if data are sufficiently extensive, observable distinctions in identification regions can be addressed. An illustration of this occurs in our analysis of World Values Survey data, when we separate the data by country. The range of effects in Italy does appear importantly different from the effects in Germany: the MTR assumption seems less brittle in Germany.

Another concern may be that dependencies between observations are being ignored. In the civil war example, for instance, one may think there is temporal or spatial dependence in civil war, in torture practices, or in both. In the other example, beliefs about immigration policy may be contagious in communities such that some people's policy beliefs depend on the immigrant-parent status of their neighbors. To some extent these possibilities are ruled out by the ITR assumption. If that assumption is dubious, treatment effects are much more difficult to assess (Manski 2011; Aronow and Samii 2012; Bowers, Fredrickson, and Panagopoulos 2012).

Relying on Wahba (1990) and Wang and Wahba (1994), our method for randomly sampled data can be considered Bayesian in the same sense as Wahba's (1983) Bayesian "confidence intervals" with their average coverage property, but some may nonetheless object to using the bootstrap in this way. Is it possible to avoid use of the bootstrap while retaining Beresteanu and Molinari's

https://doi.org/10.1093/pan/mps043 Published online by Cambridge University Press

(2008) approach of associating identification regions with set-valued random variables? One approach may be to use credal set theory (Levi 1983; Karlsson 2010) to combine methods for weakening monotonicity assumptions with empirical updates from the sample data. A credal set is a convex set of probability functions or of likelihood functions (Karlsson 2010, 22), and so it is framing information with a topology comparable to the convex sets studied in Beresteanu and Molinari (2008)—the kinds of intervals and regions that arise when there is partial identification. Bayesian updating can dominate credal operator updating when central tendencies are in focus (Karlsson, Johansson, and Andler 2011), but with partial identification one is concerned with the entirety of identification regions—with their convex hulls in Beresteanu and Molinari's (2008) "confidence collection" result—so the "second-order" properties of credal sets (Karlsson, Johansson, and Andler 2011) are likely to be important. Such developments are likely to have a practically important effect with sample data only when the sampling design is informative or when observations are not treated as exchangeable. Developing such ideas goes beyond the scope of our work here.

## References


Arnborg, Stefan. 2006. Robust Bayesianism: Relation to evidence theory. *Journal of Advances in Information Fusion* 1:75–90.

Aronow, Peter, and Cyrus Samii. 2012. Estimating average causal effects under general interference. Paper presented at the 2012 Summer Meeting of the Society for Political Methodology, University of North Carolina, Chapel Hill, July 19–21.

Beresteanu, Arie, and Francesca Molinari. 2008. Asymptotic properties for a class of partially identified models. *Econometrica* 76:763–814.

Bernhard, Judith K., Luin Goldring, Julie Young, Carolina Berinstein, and Beth Wilson. 2007. Living with precarious legal status in Canada: Implications for the well-being of children and families. *Refuge: Canada's Periodical on Refugees* 24:101–14.

Blattman, C. 2009. From violence to voting: War and political participation in Uganda. *American Political Science Review* 103:231–47.

Bowers, Jake, Mark Fredrickson, and Costas Panagopoulos. 2012. Interference is interesting: Statistical inference for interference in social network experiments. Paper presented at the 2012 Summer Meeting of the Society for Political Methodology, University of North Carolina, Chapel Hill, July 19–21.

Cederman, Lars-Erik, Nils B. Weidmann, and Kristian Skrede Gleditsch. 2011. Horizontal inequalities and ethnonationalist civil war: A global comparison. *American Political Science Review* 105:478–95.

Cingranelli, David, and David Richards. 2010. *The Cingranelli-Richards (CIRI) human rights dataset*. http://www.humanrightsdata.org.

Cochran, William G., and Donald B. Rubin. 1973. Controlling bias in observational studies: A review. *Sankya: The Indian Journal of Statistics* 35:417–66.

Davenport, Christian. 2007a. State repression and political order. *Annual Review of Political Science* 10:1–23.

———. 2007b. State repression and the tyrannical peace. *Journal of Peace Research* 4:485–504.

Efron, Bradley. 2011. *Bayesian inference and the parametric bootstrap*. Working paper.

Gleditsch, Nils Petter, Peter Wallensteen, Mikael Eriksson, Margareta Sollenberg, and Håvard Strand. 2002. Armed conflict 1946–2001: A new dataset. *Journal of Peace Research* 39:615–37.

Glynn, Adam N., and Kevin M. Quinn. 2010. An introduction to the augmented inverse propensity weighted estimator. *Political Analysis* 18:36–56.

Hainmueller, Jens, and Michael J. Hiscox. 2007. Educated preferences: Explaining attitudes toward immigration in Europe. *International Organization* 61:399–442.

———. 2010. Attitudes toward highly skilled and low-skilled immigration: Evidence from a survey experiment. *American Political Science Review* 104:61–84.

Hastie, Trevor J. 1992. Generalized additive models. In *Statistical models in S*, eds. John M. Chambers and Trevor J. Hastie. New York: Wadsworth & Brooks/Cole.

Holland, Paul W. 1986. Statistics and causal inference. *Journal of the American Statistical Association* 81:945–60.

Imai, Kosuke, Luke Keele, Dustin Tingley, and Teppei Yamamoto. 2011. Unpacking the black box of causality: Learning about causal mechanisms from experimental and observational studies. *American Political Science Review* 105:765–89.

Imai, Kosuke, Luke Keele, and Teppei Yamamoto. 2010. Identification, inference, and sensitivity analysis for causal mediation effects. *Statistical Science* 25:51–71.

Karlsson, Alexander. 2010. Evaluating credal set theory as a belief framework in high-level information fusion for automated decision-making. Technical Report. Örebro Studies in Technology 38.

Karlsson, Alexander, Ronnie Johansson, and Sten F. Andler. 2011. Characterization and empirical evaluation of Bayesian and credal combination operators. *Journal of Advances in Information Fusion* 6:150–66.

Levi, Isaac. 1983. *The enterprise of knowledge*. Cambridge, MA: MIT Press.


https://doi.org/10.1093/pan/mps043 Published online by Cambridge University Press

Lichbach, Mark Irving. 1987. Deterrence or escalation? The puzzle of aggregate studies of repression and dissent. *Journal of Conflict Resolution* 31:266–97.
Manski, Charles. 1990. Nonparametric bounds on treatment effects. *American Economic Review Papers and Proceedings* 80:319–23.
———. 1995. *Identification problems in the social sciences*. Cambridge, MA: Harvard University Press.
———. 1997. Monotone treatment response. *Econometrica* 65:1311–34.
———. 2011. *Identification of treatment response with social interactions*. Working paper.
Manski, Charles, and John V. Pepper. 2000. Monotone instrumental variables: With an application to the returns to schooling. *Econometrica* 68:997–1010.
Molinari, Francesca. 2005. *Missing treatments*. CAE Working paper #05-11.
———. 2007. *Missing treatments*. Working paper.
———. 2010. Missing treatments. *Journal of Business and Economic Statistics* 28:82–95.
Morgan, Stephen L., and Christopher Winship. 2007. *Counterfactuals and causal inference: Methods and principles for social research (Analytical Methods for Social Research)*. New York: Cambridge University Press.
Newton, Michael A., and Adrian E. Raftery. 1994. Approximate Bayesian inference with the weighted likelihood bootstrap. *Journal of the Royal Statistical Society, Series B* 56:3–48.
Petersen, Roger D. 2002. *Understanding ethnic violence: Fear, hatred, and resentment in twentieth-century Eastern Europe*. Cambridge, UK: Cambridge University Press.
Poast, Paul. 2012. Does issue linkage work? Evidence from European alliance negotiations, 1860 to 1945. *International Organization* 66:277–310.
Quinn, Kevin M. 2009. *What Can Be Learned from a Simple Table? Bayesian inference and sensitivity analysis for causal effects from 2x2 and 2x2xK tables in the presence of unmeasured confounding*. Working paper.
Rosenbaum, Paul. 2002. Covariance adjustment in randomized experiments and observational studies (with discussion). *Statistical Science* 17:286–327.
Rubin, Donald B. 1978. Bayesian inference for causal effects: The role of randomization. *Annals of Statistics* 6:34–58.
———. 1981. The Bayesian bootstrap. *Annals of Statistics* 9:130–4.
———. 1991. Practical implications of modes of statistical inference for causal effects and the critical role of the assignment mechanism. *Biometrics* 47:1213–34.
———. 2006. *Matched sampling for causal effects*. New York: Cambridge University Press.
Sambanis, Nicholas, and Annalisa Zinn. 2006. *From protest to violence: An analysis of conflict escalation with an application to self-determination movements*. Working paper.
Scott, Alastair, and T. M. F. Smith. 1973. Survey design, symmetry and posterior distributions. *Journal of the Royal Statistical Society, Series B* 35:57–60.
Sharp, Elaine B., and Mark Joslyn. 2001. Individual and contextual effects on attributions about pornography. *Journal of Politics* 63:501–19.
Sugden, R. A., and T. M. F. Smith. 1984. Ignorable and informative designs in survey sampling inference. *Biometrika* 71:495–506.
Tarrow, Sidney. 1989. *Democracy and disorder: Social conflict, political protest and democracy in Italy*. New York: Oxford University Press.
Taylor, James W. 2000. A quantile regression neural network approach to estimating the conditional density of multiperiod returns. *Journal of Forecasting* 19:299–311.
Wahba, Grace. 1983. Bayesian "confidence intervals" for the cross-validated smoothing spline. *Journal of the Royal Statistical Society, Series B* 45:133–50.
———. 1990. *Spline models for observational data*. Philadelphia: Society for Industrial and Applied Mathematics.
Wang, Yuedong, and Grace Wahba. 1994. Bootstrap confidence intervals for smoothing splines and their comparison to Bayesian "confidence intervals." *Journal of Statistical Computation and Simulation* 51:263–79.
World Values Survey Association. 2009. *WORLD VALUES SURVEY 2005 OFFICIAL DATA FILE v.20090901*. Aggregate File Producer: ASEP/JDS, Madrid.
Yamamoto, Teppei. 2012. Understanding the past: Statistical analysis of causal attribution. *American Journal of Political Science* 56:237–56.

https://doi.org/10.1093/pan/mps043 Published online by Cambridge University Press