

**DETERRENCE AND THE DEATH PENALTY:
PARTIAL IDENTIFICATION ANALYSIS USING REPEATED CROSS SECTIONS**

Charles F. Manski
Department of Economics and Institute for Policy Research
Northwestern University

and

John V. Pepper
Department of Economics
University of Virginia

forthcoming in the *Journal of Quantitative Criminology*

Abstract: Researchers have long used repeated cross sectional observations of homicide rates and sanctions to examine the deterrent effect of the adoption and implementation of death penalty statutes. The empirical literature, however, has failed to achieve consensus. A fundamental problem is that the outcomes of counterfactual policies are not observable. Hence, the data alone cannot identify the deterrent effect of capital punishment. How then should research proceed? It is tempting to impose assumptions strong enough to yield a definitive finding, but strong assumptions may be inaccurate and yield flawed conclusions. Instead, we study the identifying power of relatively weak assumptions restricting variation in treatment response across places and time. The results are findings of partial identification that bound the deterrent effect of capital punishment. By successively adding stronger identifying assumptions, we seek to make transparent how assumptions shape inference. We perform empirical analysis using state-level data in the United States in 1975 and 1977. Under the weakest restrictions, there is substantial ambiguity: we cannot rule out the possibility that having a death penalty statute substantially increases or decreases homicide. This ambiguity is reduced when we impose stronger assumptions, but inferences are sensitive to the maintained restrictions. Combining the data with some assumptions implies that the death penalty increases homicide, but other assumptions imply that the death penalty deters it.

An early version of this research was prepared for presentation at the National Research Council Workshop on Deterrence and the Death Penalty, April 2011. We have benefitted from comments received at the workshop, the opportunity to present the work in seminars at the University of Bern and the University of Virginia, and the comments of Daniel Nagin. Manski's research was supported in part by National Science Foundation grant SES-0911181.

Ehrlich (1975, p. 398): “In fact, the empirical analysis suggests that on the average the tradeoff between the execution of an offender and the lives of potential victims it might have saved was of the order of 1 for 8 for the period 1933–1967 in the United States.”

Blumstein, Cohen, and Nagin (1978, p. 62): “The current evidence on the deterrent effect of capital punishment is inadequate for drawing any substantive conclusion.”

1. Introduction

Researchers have long used data on homicide rates and sanctions to examine the deterrent effect of capital punishment. There is now a large body of work addressing this controversial question, yet the literature has failed to achieve consensus on even the most basic matters. Donohue and Wolfers (2005), who review a set of recent studies, provide a striking illustration. They find that a seemingly trivial change to the model estimated by Dezhbakhsk, Rubin and Shepard (2003) “flips the sign of the original estimates: instead of saving eighteen lives, each execution leads to eighteen lives lost.”

Numerous shortcomings of the research were documented over thirty years ago in the report of the National Research Council (NRC) Panel on Research on Deterrent and Incapacitative Effects (Blumstein, Cohen, and Nagin, 1978). These problems persist in more recent work. A fundamental difficulty is that the outcomes of counterfactual policies are unobservable. Data alone cannot reveal what the homicide rate in a state without (with) a death penalty would have been had the state (not) adopted a death penalty statute. Here, as always when analyzing treatment response, data must be combined with assumptions to enable inference on counterfactual outcomes. Hence, research

confronts the *selection problem*.

If available data alone cannot reveal the deterrent effects of capital punishment, how should research proceed? It is tempting to impose assumptions strong enough to yield a definitive finding. When this is achieved, a deterrent effect is said to be *point-identified*. Researchers often recognize that strong assumptions may have little foundation, but they apply them nonetheless. They may defend their strong assumptions as necessary to “provide answers.” However, strong assumptions may be inaccurate, yielding flawed and conflicting conclusions. One of us has cautioned against the imposition of untenable strong assumptions as follows (Manski, 2003, p. 1):

The Law of Decreasing Credibility: The credibility of inference decreases with the strength of the assumptions maintained.

We have seen this repeatedly in the empirical literature on the death penalty.

With this in mind, we study inference under various relatively weak assumptions that may possess greater credibility. The assumptions we study do not point-identify deterrent effects, but they do *partially* identify them, yielding bounds rather than point estimates. Analysis of partial identification of treatment effects has developed and been applied over the past twenty years, beginning with Manski (1990) and continuing through our present work. See Manski (2003 and 2007) for textbook expositions. Some applications include Manski and Nagin (1998), Manski and Pepper (2000), Pepper (2000), Blundell *et al.* (2007), and Gundersen, Kreider, and Pepper (2011).

The basic insight of partial identification analysis is that identification need not be an all or nothing undertaking. Available data and credible assumptions may lead to partial conclusions. Some may find this ambiguity frustrating and be tempted to impose strong assumptions in order to provide definitive answers. We caution against such a reaction.

Imposing strong but untenable assumptions cannot not truly resolve problems of inference on the deterrent effects of capital punishment. The NRC Panel on Research on Deterrent and Incapacitative Effects recognized this when it concluded that “research on this topic is not likely to produce findings that will or should have much influence on policymakers” (Blumstein, Cohen, and Nagin, 1978, p.63). The lesson of past research is that researchers and policymakers must cope with ambiguity.

To demonstrate partial identification analysis in a relatively simple setting, this paper considers the problem of drawing inferences on the deterrent effects of death penalty statutes using data from repeated cross sections of states.¹ We focus on the years following the 1972 Supreme Court case *Furman vs. Georgia*, which resulted in a de facto moratorium on the application of the death penalty, and the 1976 case *Gregg vs. Georgia*, which ruled that the death penalty could be applied subject to certain criteria. We examine the effect of death penalty statutes on the national homicide rate in two years: 1975, the last full year of the federal moratorium on death penalty, and 1977, the first full year after the moratorium was lifted. In 1975 the death penalty was illegal throughout the country, and in 1977 thirty-two states had legal death penalty statutes. For each state and each year, we observe the homicide rate and whether the death penalty is legal.

Table 1 displays the homicide rate per 100,000 residents in 1975 and 1977 in the states that did and did not legalize the death penalty after the Gregg decision. The former are the “treated” states and the latter are the “untreated” ones. Here and throughout the paper, we include the District of Columbia and regard it as equivalent to a state. When computing averages across states, we

¹ We use data provided by Justin Wolfers at <http://bpp.wharton.upenn.edu/jwolfers/DeathPenalty.shtml>. This archive reports annual state level data from 1930 to 2004. See Donohue and Wolfers (2005) for a detailed description.

weight each state by its population. The thirty-two states with legal death penalty statutes in 1977 contained seventy percent of the total population.

Table 1: Homicide Rates per 100,000 Residents by Year and Treatment Status in 1977

Year	Group		Total
	Untreated	Treated	
1975	8.0	10.3	9.6
1977	6.9	9.7	8.8
Total	7.5	10.0	9.2

The data in the table may be used to compute three simple estimates of the effect of death penalty statutes on the national homicide rate. A “before-and-after” analysis compares homicide rates in the treated states in 1975 and 1977, yielding the estimate -0.6 ($9.7 - 10.3$). Contemporaneous comparison of 1977 homicide rates in the treated and untreated states yields the estimate 2.8 ($9.7 - 6.9$). The difference-in-difference (DID) estimate compares the time-series changes in homicide rates in the treated and untreated states, yielding the estimate 0.5 [$(9.7 - 10.3) - (6.9 - 8.0)$].²

These three estimates yield different empirical findings. Given certain assumptions, each appropriately measures the effect on death penalty statutes on the national homicide rate. However,

² DID estimates have been reported in numerous policy analyses, including evaluations of the death penalty. For example, Dezhbakhsh and Shepherd (2006) used the federal moratorium of the 1970s as a “judicial experiment.” More broadly, they used data from 1960–2000 to compare murder rates immediately before and after changes in death penalty laws. They concluded that the death penalty has a substantial deterrent effect on homicides. Examining the same questions using the same data, Donohue and Wolfers (2005) concluded that there is no evidence that the death penalty deters homicides.

the assumptions that justify this interpretation differ across estimates. Moreover, one may think that none of the requisite assumptions is credible.

In Section 2, we formally define the empirical question and the selection problem. Section 3 presents two polar approaches to inference. At one pole, deterrent effects are point-identified under strong assumptions needed to justify the three estimates described above. At the other, they are partially identified using only the data and an a priori bound on the maximum possible value of the mean counterfactual homicide rate. No other assumptions are imposed to address the selection problem.

Section 4 applies new results on partial identification developed in our companion technical paper (Manski and Pepper, 2011). These results exploit the variation of homicide rates and death penalty status in repeated cross sections of states to explore middle ground assumptions. By successively adding stronger assumptions and determining their identifying power, our analysis makes transparent how assumptions shape inferences about the effect of capital punishment on homicide.

The assumptions that we consider restrict the potential variation of treatment response or treatment effects over time and/or across states. Under the weakest assumptions, there is considerable ambiguity: we cannot rule out the possibility that the death penalty substantially increases or decreases the mean homicide rate across states. This ambiguity is reduced by imposing stronger assumptions, but inferences are highly sensitive. Given the available data, imposing certain assumptions implies that the death penalty increases homicide but other assumptions imply that the death penalty deters it.

This paper does not provide measures of statistical precision when presenting our findings on average treatment effects. That is, we view the states in 1975 and 1977 as constituting the population of interest, rather than as realizations from some sampling process. One reason we do this is expositional. We want to focus attention on the identification problem arising from the unobservability of counterfactual outcomes. Statistical precision of estimates is a second-order concern relative to this problem.

Another reason is that measurement of statistical precision requires specification of a sampling process that generates the data, but we are unsure what type of sampling process would be reasonable to assume. Existing methods for computing confidence intervals in partial identification analysis assume that the data are a random sample drawn from an infinite population; see, for example, Imbens and Manski (2004) and Chernozhukov, Hong, and Tamer (2007). This sampling assumption does not seem natural when considering states as units of observation.

2. Average Treatment Effects and the Selection Problem

We consider the problem of learning the effect of death penalty statutes on the national homicide rate. This is the population-wide average treatment effect (ATE)

$$\text{ATE}_d \equiv E[Y_d(1)] - E[Y_d(0)]. \quad (1)$$

Here there are two mutually exclusive treatments: treatment $t = 1$ denotes a state sanctions regime that includes the presence of a death penalty statute and $t = 0$ denotes one without such a statute. Defining the treatment to be the presence or absence of a death penalty statute is a representation of actual sanctions policy intended to simplify analysis. While this comparison addresses a well-defined question, the resulting analysis cannot reveal the mechanisms by which the statute impacts crime. One might like to differentiate treatments by the specifics of the death penalty statute enacted, the way it is implemented, and by the nature of the non-capital sanctions that a state has in place.

The outcome $Y_d(1)$ denotes the homicide rate if a state were to have a death penalty statute, $Y_d(0)$ denotes the analogous outcome if the state were not to have a death penalty statute, and d indicates whether the year is 1975 or 1977 ($= 0$ if 1975, $= 1$ if 1977). ATE_d expresses how the national homicide rate in year d would differ if all states were to have a death penalty statute versus what would occur if all states had a death penalty moratorium.

We also consider inference on the effect of death penalty statutes on homicide in groups of states with specified observed characteristics. Let X denote these characteristics. Then the objective is to learn the group-specific average treatment effect $ATE_{dx} \equiv E[Y_d(1)|X] - E[Y_d(0)|X]$.

Notice that for each state j and year d , there are two potential outcomes, $Y_{jd}(1)$ and $Y_{jd}(0)$. The outcome $Y_{jd}(1)$ is counterfactual for all states that did not have a death penalty statute in year d , while $Y_{jd}(0)$ is counterfactual for all states that did have a death penalty. The observed murder rate is $Y_{jd} = Y_{jd}(1) \cdot Z_{jd} + Y_{jd}(0) \cdot (1 - Z_{jd})$, where $Z_{jd} = 1$ denotes that state j has a death penalty statute in year d and $Z_{jd} = 0$ otherwise.

The fact that the data only reveal one of the two mutually exclusive outcomes constitutes the selection problem. Using the Law of Iterated Expectations, the implications for identification of $E[Y_d(1)]$ can be seen by writing this quantity as

$$E[Y_d(1)] = E[Y_d(1)|Z_d = 1]P(Z_d = 1) + E[Y_d(1)|Z_d = 0]P(Z_d = 0). \quad (2)$$

Each observation in the sample reveals $Y_{jd}(Z_{jd})$, Z_{jd} , and X_{jd} . (We will sometimes write $Y_{jd} \equiv Y_{jd}(Z_{jd})$ for short.) Hence, the sampling process identifies the selection probability $P(Z_d = 1)$, the censoring probability $P(Z_d = 0)$, and the mean of $Y_d(1)$ in states with the death penalty, $E[Y_d(1) | Z_d = 1]$. For example, in 1977 we have $E[Y_1(1)|Z_1 = 1] = 9.7$, $P(Z_1 = 1) = 0.70$, $P(Z_1 = 0) = 0.30$. However, the sampling process does not reveal the mean of $Y_d(1)$ in states without the death penalty, $E[Y_d(1)|Z_d = 0]$. Thus, $E[Y_d(1)]$ only partially identified by the data alone.

3. Polar Approaches to Inference

How might we proceed? This section considers polar approaches. One pole makes assumptions strong enough to point-identify ATEs. The other only assumes an upper bound on the homicide rate, the result being wide bounds on ATEs.

3.1. Assumptions that Point-identify Average Treatment Effects

At one extreme, researchers may impose assumptions strong enough to point-identify average treatment effects. This has been the norm in the literature, with researchers applying a variety of assumptions. In this section we give assumptions under which the three simple estimates mentioned in the Introduction identify the effect of death penalty statutes on the national homicide rate.

3.1.1. Random Treatment Selection

A common assumption is that the realized treatments Z_{jd} are statistically independent of potential outcomes, as they would be in a classical randomized experiment. This implies that $E[Y_1(1)] = E[Y_1(1)|Z_d = 1]$ and $E[Y_1(0)] = E[Y_1(0)|Z_d = 0]$. The ATE in 1977 is point-identified under this assumption because $E[Y_1(1)|Z_d = 1]$ and $E[Y_1(0)|Z_d = 0]$ are the observed mean homicide rates in states that do and do not have the death penalty.

Combining this assumption with the data implies that the death penalty increases the mean homicide rate by 2.8 per 100,000 (i.e., 9.7 – 6.9) in 1977. Without additional assumptions, the ATE in 1975 is not identified. The reason is that, with the federal moratorium in place, no states had the death penalty during that time period.

The random-selection assumption is credible in randomized experiments, but it is not generally credible in observational studies where treatments (i.e., death penalty statutes) are self-selected. A particular concern is that states may adopt death penalty statutes in part based upon their beliefs about the deterrent effect of such statutes.

3.1.2. Time-Invariant Treatment Response

Let $X_j = 1$ if state j is in the treated group (i.e., $Z_{j1} = 1$) and $X_j = 0$ if it is in the untreated group ($Z_{j1} = 0$). Suppose that, within the group of treated states, mean treatment response is the same in 1975 and 1977. Thus, $E[Y_1(\cdot)|X = 1] = E[Y_0(\cdot)|X = 1]$. The data for 1975 and the fact that no state had a death penalty statute that year imply that

$$E[Y_0(0)|X = 1] = E[Y_0(0)|X = 1, Z = 0] = E(Y_0|X = 1, Z = 0).$$

The data for 1977 and the fact that only treated states had a death penalty statute that year imply that

$$E[Y_1(1)|X = 1] = E[Y_1(1)|X = 1, Z = 1] = E(Y_1|X = 1, Z = 1).$$

Combining this with the assumption that mean treatment response does not vary over time for the treated group implies that the time-invariant *effect of treatment on the treated* (ETT) is

$$ETT \equiv ATE_{X=1} = E[Y_1(1)|X = 1] - E[Y_0(0)|X = 1] = E(Y_1|X = 1, Z = 1) - E(Y_0|X = 1, Z = 0).$$

The right-hand side is the estimate of deterrence given by before-and-after analysis of the treated states. The empirical finding with the data in Table 1 is -0.6 ($9.7 - 10.3$).

Before-and-after analysis of the treated states only reveals the average treatment effect within this group of states. One cannot perform an analogous analysis for the untreated states because they did not have a death penalty statute in either year. One can interpret the before-and-after estimate

as giving the effect of death penalty statutes on the national homicide rate if one thinks it credible to assume that the effect of treatment on the treated equals the effect of treatment on the untreated. However, the data in Table 1 make this interpretation suspect. The data show that the homicide rate in the untreated states fell from 8.0 in 1975 to 6.9 in 1977. Thus, mean treatment response varied with time in the untreated states.

3.1.3. Linear Homogeneous Treatment Response

A third assumption that point-identifies the ATE begins by posing a model of linear homogeneous treatment response. Let

$$Y_{jd}(t) = \alpha_j + \beta \cdot d + \gamma \cdot t + \delta_{jd}. \quad (3)$$

The parameter γ measures the effect on the homicide rate of having the death penalty. This effect is assumed to be homogeneous across states j and dates d . The date-specific intercept $\beta \cdot d$ shifts the response function additively by date. Observe that this intercept does not vary with the treatment t or across states j . Similarly, the state-specific intercept α_j allows the response function to differ additively by state. The unobserved random variable δ_{jd} varies across states and periods.

Evaluated at realized values of treatments and outcomes, the model yields

$$Y_{jd} = \alpha_j + \beta \cdot d + \gamma \cdot Z_{jd} + \delta_{jd}. \quad (4)$$

The conventional practice is to impose distributional assumptions that point-identify γ . Here is one assumption that achieves this objective.

As in Section 3.1.2, let $X_j = 1$ if state j is in the treated group (i.e., $Z_{j1} = 1$) and $X_j = 0$ if it is in the untreated group ($Z_{j1} = 0$). Assume that, for each date d , $E(\delta_d|X, Z_d) = 0$ and $E(\alpha|X, Z_d) = E(\alpha|X)$. Then

$$E(Y_d|X, Z_d) = E(\alpha|X) + \beta \cdot d + \gamma \cdot Z_d.$$

It follows that

$$\begin{aligned} \gamma = & [E(Y_1|X = 1, Z_1 = 1) - E(Y_0|X = 1, Z_0 = 0)] - \\ & [E(Y_1|X = 0, Z_1 = 0) - E(Y_0|X = 0, Z_0 = 0)]. \end{aligned} \quad (5)$$

Thus, the ATE is point identified. The right-hand side of (5) is the DID form. Given the model and the data on homicide rates and executions summarized in Table 1, we find that the death penalty statute increases the homicide rate in every state and date by 0.5.

The problem with this approach to identification is, again, credibility. The linear homogeneous response model in equation (3) is generally difficult to justify, as policies are typically thought to have heterogeneous effects (Manski, 1990; Moffitt, 2005). The distributional assumptions used above in conjunction with the linear model are also hard to justify. Instead, researchers often apply instrumental-variable assumptions asserting that potential outcomes are mean-independent of some observed covariate that is statistically associated with the realized

treatment. Finding instrumental variables that satisfy this condition, however, has proven to be difficult in studies of the death penalty (Donohue and Wolfers, 2005).

3.2. Partial Identification Assuming Bounded Outcomes

Imposing assumptions strong enough to yield a definitive finding is alluring, but strong assumptions may be inaccurate, yielding flawed and conflicting conclusions. Rather than attempt to point-identify the ATE, partial-identification analysis does not impose the strong assumptions that have been used in the literature. Instead, we make weaker assumptions that yield bounds on the deterrent effect of the death penalty. Given the conflicting findings in the literature and the methodological challenges in addressing the selection problem, deriving bounds under assumptions that may be credible seems an important step forward.

A natural starting point is to ask what the data alone reveal about the ATE. Recall that

$$E[Y_d(1)] = E[Y_d(1) | Z_d = 1] \cdot P(Z_d = 1) + E[Y_d(1) | Z_d = 0] \cdot P(Z_d = 0).$$

The selection problem arises because the data do not reveal the homicide rate if a death penalty statute were in place for states where there was no such statute. However, we do know that the homicide rate per 100,000 residents logically cannot be larger than 100,000. Thus, $Y_d(1) \in [0, 100,000]$ and, hence, $E[Y_d(1) = 1 | Z_d = 0] \in [0, 100,000]$. To put a more reasonable upper bound on this counterfactual mean outcome, note that across all states and both years (1975 and 1977), the

observed homicide rate always was in the range $[0.8, 32.8]$. Thus, it seems reasonable to assume that $E[Y_d(1)|Z_d = 0] \in [0, 35]$.

Using this upper bound, it follows that

$$E[Y_d(1)] \in \{E[Y_d(1) | Z_d = 1] \cdot P(Z_d = 1) + 0 \cdot P(Z_d = 0), \\ E[Y_d(1) | Z_d = 1] \cdot P(Z_d = 1) + 35 \cdot P(Z_d = 0)\} \quad (6)$$

Observe that the width of this bound increases with the censoring probability $P(Z_d = 0)$. Thus, if a large fraction of states adopt death penalty statutes, the width of the bound on $E[Y_d(1)]$ is relatively narrow. In that case, the data do not reveal much about the distribution of $Y_d(0)$, so the analogous bound on $E[Y_d(0)]$ is wide.

Consider, for example, drawing inferences on mean potential outcomes during the moratorium, when no states had the death penalty. Hence, $P(Z_0 = 0) = 1$. In this case, the data alone provide no information on the mean outcome if all states were to adopt a death penalty, but they point-identify the mean outcome if all states were to not have a death penalty.

The sharp bound on the ATE can be found by taking the appropriate difference between the lower (upper) bound on $E[Y_d(1)]$ and the upper (lower) bound on $E[Y_d(0)]$ (Manski, 1990). Given the restriction that the counterfactual mean outcomes lie in the interval $[0, 35]$, the width of the bound on the ATE necessarily equals 35. It follows that, in the absence of additional assumptions, the data cannot reveal the sign of the effect of the death penalty on the murder rate.

Table 2 displays the bounds on the ATE for 1975 and 1977. The data show that in 1977, seventy percent of the population resided in states which legalized the death penalty. In this year,

the population weighted murder rate was 9.7 in states with the death penalty and 6.9 in states without it. Thus, $P(Z_1 = 1) = 0.70$, $E[Y_1(1) | Z_1 = 1] = 9.7$, and $E[Y_1(0) | Z_1 = 0] = 6.9$. Thus, evaluation of the bound in (6) shows that $E[Y_1(1)]$ must be in the interval $[6.8, 17.3]$ and $E[Y_1(0)]$ must be in the interval $[2.1, 26.6]$. These bounds on mean potential murder rates imply that the ATE must be in the interval $[-19.8, 15.2]$.

Importantly, these bounds are not a confidence interval – they do not express statistical imprecision created by sampling variability. Rather, the bounds express the ambiguity created by the selection problem. Assuming only that counterfactual mean potential murder rates cannot exceed 35, the data reveal that the ATE lies in the interval $[-19.8, 15.2]$. Recall that the ATE under the random-selection assumption was 2.8, with before-and-after analysis was -0.6 , and the DID estimate was 0.5.

Table 2 also derives bounds for 1975, a year in which no states had the death penalty. In this year, $P(Z_0 = 1) = 0$, so the data are uninformative about what the mean homicide rate would be if all states had a death penalty statute. However, the data point-identify $E[Y_0(0)] = 9.6$. Thus, for 1975, we find that the ATE must be in the interval $[-9.6, 25.4]$. While the bounds for 1975 are different than those for 1977, both have a width of 35 and both include zero.

Table 2: Partial Identification of the ATE Under the Bounded Outcomes Assumption

	1975	1977
Probability of a Death Penalty Statute: $P(Z_d = 1)$	0	0.7
Murder Rate in States with the Death Penalty: $E[Y_d(1) Z_d = 1]$	N. A.	9.7
Murder Rate in States without the Death Penalty: $E[Y_d(0) Z_d = 0]$	9.6	6.9
Bounds:		
$E[Y_d(1)]$	[0, 35]	[6.8 , 17.3]
$E[Y_d(0)]$	9.6	[2.1 , 26.6]
ATE_d	[-9.6, 25.4]	[-19.8 ,15.2]

4. Middle-Ground Assumptions

The analysis of Section 3.2 made no assumptions that relate criminal behavior in 1975 and 1977. Nor did it make assumptions that relate criminal behavior in states that did and did not enact a death penalty in 1977. Deterrent effects were permitted to vary across years and states.

One may reasonably believe that there is some commonality in criminal behavior across years and states. However, it is not credible to assume as much commonality as the linear homogeneous model, which supposes that deterrent effects are the same in every year and every state. This leads

us to consider “middle-ground” assumptions that presume some commonality across years or states, but not homogeneity. In particular, we apply new analysis of partial identification with repeated cross-sections developed in Manski and Pepper (2011).³

To begin, Section 4.1 permits treatment effects to vary across states but assumes that they do not vary across years. Section 4.2 adds an assumption that date-specific intercepts do not vary across specified groups of states. Section 4.3 assumes the existence of bounded instrumental variables, which bound the variation of average treatment effects across groups of states. An appendix explains the algorithms used to compute our empirical findings.

Our intent is not to endorse any particular assumption. It is rather to demonstrate how the the conclusions drawn depends on the assumptions imposed, thus providing a menu of possibilities to readers of research on deterrence.

4.1. Date-Invariant Treatment Effects

We begin by assuming that the deterrent effect of the death penalty is the same in 1975 and 1977, in the formal sense that $ATE_1 = ATE_0$. Then the date-invariant ATE must lie in the intersection of the two date-specific intervals shown in Table 2, these being $[-9.6, 25.4]$ and $[-19.8, 15.2]$. The result is $[-9.6, 15.2]$.

While there remains much ambiguity about the deterrent effect of the death penalty, the assumption of date-invariant treatment effects has identifying power. It reduces the width of the

³ Other authors have studied inference using assumptions that relax the linear homogeneous model. See Athey and Imbens (2006), Chernozhukov, Fernandez-Val, Hahn and Newey (2010) and Evdokimov (2010). Their assumptions and analyses differ considerably from what we present here.

bound on the ATE from 35 to 25. This bound does not allow us to the sign of the ATE, but it does rule out claims that the death penalty reduces the mean murder rate by more than -9.6 per 100,000 or increases it by more than 15.2.

The analysis below builds on this basic finding. Section 4.1.1 introduces new notation and uses it to re-derive the basic finding. The payoff from introducing the new notation is that it provides the basis for consideration of further assumptions. Section 4.1.2 studies the additional identifying power of placing a priori bounds on time-series variation in mean response levels. Section 4.1.3 shows the further identifying power of placing a tighter priori bound on counterfactual mean response levels than the $[0, 35]$ bound assumed heretofore.

4.1.1. Basic Analysis

As earlier, suppose that one observes cross-sections of all the states including the District of Columbia in 1975 ($d = 0$) and 1977 ($d = 1$). Assume that mean treatment response at date d has the form

$$E[Y_1(t)] = E[Y_0(t)] + \beta. \quad (7)$$

Here β is a date-specific intercept that distinguishes mean response at dates 0 and 1. Equation (7) permits mean treatment response levels to vary across dates, but it assumes that the average treatment effect is invariant across dates. Specifically, $E[Y_1(1)] - E[Y_1(0)] = E[Y_0(1)] - E[Y_0(0)]$.

To shorten the notation, let $E_t \equiv E[Y_0(t)]$. Then (7) is equivalent to

$$E[Y_0(t)] = E_t, \quad (8a)$$

$$E[Y_1(t)] = E_t + \beta. \quad (8b)$$

Equations (8a)–(8b) have identifying power because they reduce the number of unknown mean potential outcomes by one. Without the assumption, we do not know the four quantities $E[Y_d(t)]$, $d = 0, 1$; $t = 0, 1$. With the assumption, we do not know the three quantities (E_0, E_1, β) .

To obtain the identifying power of the assumption, first consider each date-treatment pair (d, t) separately and obtain the identification region for $E[Y_d(t)]$ using only the assumption of bounded outcomes, as was done in Section 3.2. Let this interval be called $[L_d(t), U_d(t)]$. Combining this with (8a)–(8b), the available information is

$$L_0(t) \leq E_t \leq U_0(t), \quad t = 0, 1; \quad (9a)$$

$$L_1(t) \leq E_t + \beta \leq U_1(t), \quad t = 0, 1. \quad (9b)$$

Thus, the feasible values of the three unknowns (β, E_0, E_1) are all the triples that satisfy the four inequalities given in (9a)–(9b).

Table 2 gives the values of $[L_d(t), U_d(t)]$, $d = 0, 1$; $t = 0, 1$. Inserting these values in (9a)–(9b) yields these findings for (E_0, E_1, β) : $\beta \in [-7.5, 17.0]$, $E_0 = 9.6$, and $E_1 \in [0, 24.8]$. Hence, the date-invariant ATE lies in the interval $[-9.6, 15.2]$, as shown earlier by a more direct argument.

4.1.2. Bounding Time-Series Variation in Mean Response Levels

The analysis of Section 4.1.1 placed no a priori restrictions on time-series variation in treatment response levels between 1975 and 1977. Some states could have become much more prone to homicide over this period while others could have become much less prone to homicide. We only assumed that the overall deterrent effect of the death penalty remains stable over time.

Our objective was to learn about the ATE, but we also found that assumption (7) and the data implied a bound on the variation in mean response levels between 1975 and 1977, namely $\beta \in [-7.5, 17.0]$. Thus, mean potential homicide rates may have decreased by as much as 7.5 per 100,000 or increased by as much as 17.0 per 100,000 over the three-year period.


One might not think it credible that such large variations in mean potential homicide rates could have occurred over such a short time period. One might be willing to assume that β must lie in some narrower interval than $[-7.5, 17.0]$. Such an assumption may imply a narrower bound on the ATE.

Consider, for example, the assumption that β lies in the interval $[-5.0, 3.0]$. One might motivate this assumption by the fact that the largest state-specific observed decrease in the homicide rate between 1975 and 1977 was 5.0 and the largest observed increase was 2.9. If one uses $[-5.0, 3.0]$ as an a priori bound on β , application of (8) implies that the ATE lies in the interval $[-5.8, 12.7]$, narrowing the bound $[-9.6, 15.2]$ derived earlier.

Alternatively, consider the much stronger assumption that $\beta = 0$. This assumption permits individual states to experience time-series variation in their proneness to homicide, but it supposes

that there is no national trend. Combining this assumption with (8) implies that the ATE lies in the interval $[-2.8, 7.7]$.

Readers of research on deterrence may vary in their beliefs on the credible range of values for β . To enable readers to bring to bear their own beliefs and determine the implications for inference on the ATE, the solid lines in Figure 1 display the bound on the ATE as a function of β . The figure shows how a priori restrictions on β reduce ambiguity about deterrence. A person who believes that $\beta \leq -3$ can conclude that the ATE is positive; that is, the death penalty increases the expected homicide rate. In contrast, someone who believes that $\beta \geq 8$ can conclude that the ATE is negative; that is, the death penalty deters crime. Someone who thinks that β may lie in the interval $(-3, 8)$ cannot identify the sign of the ATE.



4.1.3. Tighter Bounds on Counterfactual Mean Response Levels

We have thus far placed only a very weak bound on the counterfactual homicide rates $E[Y_d(t) | Z_d \neq t]$, supposing that they must lie in the range $[0, 35]$. Of the 102 state-specific homicide rates observed to occur in 1975 and 1977, the central ninety percent fall in the interval $[2, 15]$. Suppose that one uses this interval as a bound on $E[Y_d(t) | Z_d \neq t]$ rather than the earlier bound $[0, 35]$. Then application of (8) implies that $\beta \in [-6.1, 3.0]$ and $ATE \in [-5.2, 5.4]$.

The dashed lines in Figure 1 display how the bound on the ATE varies with β . Assuming the tighter bound on mean counterfactual outcomes substantially narrows the bounds on the ATE relative to those reported in Section 4.1.2. Whereas the assumption $\beta = 0$ earlier implied that $ATE \in [-2.8, 7.7]$, it now implies that $ATE \in [-2.2, 1.7]$. A person who believes that $\beta < -2$ can now conclude that the ATE is positive, while one who believes that $\beta \geq 2$ can conclude that the ATE is negative.

4.2. Date-Invariant Treatment Effects and Covariate-Invariant Date Intercepts

In Section 4.1, β was the mean difference in potential murder rates between 1975 and 1977, the mean being computed across all states in the nation. Let X be a covariate that separates states into K distinct groups, each group containing at least one state. We now combine the assumption of date-invariant treatment effects with the assumption that groups of states with different values of X have the same date intercepts. A stronger version of the assumption was made in the linear

homogeneous model of Section 3.1, where it was assumed that states with different (covariate, realized treatment) values share the same mean date intercepts.

Let $E_{t|X} \equiv E[Y_0(t)|X]$. Assume that

$$E[Y_0(t)|X] = E_{t|X}, \quad (10a)$$

$$E[Y_1(t)|X] = E_{t|X} + \beta. \quad (10b).$$

This repeats assumption (7), now conditional on X , and also assumes that β does not vary with X .

Repeating the earlier derivation, but now conditional on X , let $[L_d(t|x), U_d(t|x)]$ be the bound on $E[Y_d(t)|X]$ obtained using only the assumption that outcomes are bounded in the range $[0, 35]$.

Combining this with (10a)–(10b), the available information is

$$L_0(t|X) \leq E_{t|X} \leq U_0(t|X), \quad t = 0, 1; \text{ all } X \quad (11a)$$

$$L_1(t|X) \leq E_{t|X} + \beta \leq U_1(t|X), \quad t = 0, 1; \text{ all } X \quad (11b).$$

Thus, the feasible values of the $(2K + 1)$ unknowns $(\beta, E_{0|X}, E_{1|X}, \text{ all } X)$ satisfy the $4K$ inequalities given in (11a)–(11b).

Given that β does not vary with X or t , adding covariates provides additional identifying information. When $K = 1$, as in Section 4.1, there are three unknowns that satisfy four inequalities. When $K = 2$, five unknowns satisfy eight inequalities. When $K = 4$, nine unknowns satisfy sixteen inequalities. And so on.

To illustrate, we evaluate the ATE with two definitions of X . First, X indicates whether a state does or does not have a death penalty statute in 1977; that is, whether it is treated or untreated. Second, we let X indicate the location of a state in one of four mutually exclusive and exhaustive census regions. The derivation of findings in the latter case is not particularly revealing, so we omit the details. However, the former case is simple and yields an interesting analytical result. Hence, we give the derivation first before examining the empirical findings.

4.2.1. Treatment Group as the Covariate

When the covariate differentiates treated and untreated states, the effect of treatment on the treated (ETT) is the DID estimate and the effect of treatment on the untreated (ETU) is partially identified. To see this, let $X_j = 1$ if $Z_{j1} = 1$ and $X_j = 0$ if $Z_{j1} = 0$. The former are the treated states and the latter are the untreated ones. With this definition of X , the eight inequalities (11a)–(11b) become

$$\begin{aligned}
 E_{0|0} &= E(Y_0|X=0), & E_{0|0} + \beta &= E(Y_1|X=0), \\
 0 &\leq E_{1|0} \leq 35, & 0 &\leq E_{1|0} + \beta \leq 35, \\
 E_{0|1} &= E(Y_0|X=1), & 0 &\leq E_{0|1} + \beta \leq 35, \\
 0 &\leq E_{1|1} \leq 35, & E_{1|1} + \beta &= E(Y_1|X=1).
 \end{aligned}$$

Given these bounds, the equalities in the first row point-identify the date intercept, with

$$\beta = E(Y_1|X=0) - E(Y_0|X=0).$$

In our application, $\beta = -1.1$. Recall that, using only assumption (7) rather than (10), we could only conclude that $\beta \in [-7.5, 17.0]$.

With knowledge of β and the full set of inequalities, it is straightforward to assess identification of the four mean response values. We find that

$$\begin{aligned} E_{0|0} &= E(Y_0|X=0), & E_{0|1} &= E(Y_0|X=1), \\ E_{1|0} &\in [\max(0, -\beta), \min(35, 35 - \beta)]; & E_{1|1} &= E(Y_1|X=1) - \beta. \end{aligned}$$

Thus, the effect of treatment on the treated is point-identified, with

$$ETT \equiv E_{1|1} - E_{0|1} = [E(Y_1|X=1) - E(Y_0|X=1)] - [E(Y_1|X=0) - E(Y_0|X=0)].$$

This is the DID form obtained earlier using the stronger assumptions of the linear homogeneous model.⁴

The effect of treatment on the untreated is partially identified. The data and assumptions reveal that

$$ETU \equiv E_{1|0} - E_{0|0} \in [\max(0, -\beta) - E(Y_0|X=0), \min(35, 35 - \beta) - E(Y_0|X=0)].$$

⁴ Manski and Pepper (2011) show that the DID form arises under a yet weaker assumption where the date-invariance restriction only applies to the $Y(0)$ response function. In this case, the ETT is point-identified at date $d = 1$ but not at $d = 0$.

4.2.2. Empirical Findings

Table 3 displays the empirical findings for the ETT, ETU and ATE under many of the assumptions considered thus far. Two conclusions warrant attention.

First, while incorporating covariates clearly tightens the estimated bounds on the ATE, there remains much ambiguity. For example, when census region is used as a covariate, the bound on the ATE shrinks from $[-9.6, 15.2]$ to $[-9.0, 10.1]$. Thus, while the ATE bound shrinks by nearly six points, we still cannot determine whether the death penalty increases or decrease the national homicide rate.

Second, when the treatment group is used as a covariate, the ETT is point-identified and estimated to equal 0.5, but the ATE is only partially identified and estimated to lie in the interval $[-1.9, 8.3]$. Researchers often loosely report the DID estimate as “the” effect of deterrence, without being careful to state their maintained assumptions. The ATE equals the DID estimate if one assumes that the ETU equals the ETT, as is the case with the linear homogenous response model. However, without this or another assumption that makes the ETU equal the ETT, using treatment group as a covariate does not identify the sign of the ATE.

Table 3: Treatment Effects with Date-Invariant Treatment Effects, with and without Covariate-Invariant Date Intercepts			
Assumption	ETT	ETU	ATE
Linear Homogeneous Response	0.5	0.5	0.5
Bounded Outcomes, 1977	[-25.3, 9.7]	[-6.9, 28.1]	[-19.8, 15.2]
Bounded Outcomes, 1975	[-35.0, 35.0]	[-9.6, 25.4]	[-9.6, 25.4]
Date-Invariant Treatment Effects			
No Covariate			[-9.6, 15.2]
Region as Covariate			[-9.0, 10.1]
Treatment Group as Covariate	0.5	[-6.9, 27.0]	[-1.9, 8.3]

4.3. Bounded Instrumental Variables

The model introduced in Section 4.2 assumed common time-series variation in treatment response across groups of states, but placed no restrictions on cross-sectional variation in treatment response. Traditional instrumental variables (IVs) assume that specified groups of treatment units have the same mean treatment response or the same average treatment effects. It often is difficult to motivate such sharp assumptions, but it may be easier to motivate weaker assumptions asserting that mean response or average treatment effects do not differ too much across groups. We refer to such assumptions as asserting the existence of *bounded instrumental variables*.

To demonstrate the idea, we apply it here to group-specific average treatment effects. Formally, we consider identification of the ATE when the researcher selects a non-negative constant Δ and assumes that

$$|ATE_x - ATE_{x'}| \leq \Delta \quad \text{for all } x \text{ and } x'. \quad (12)$$

This assumption bridges the gap between the linear homogenous model, which assumes that the ATE is identical across all states j and dates d , and the model in (7) which allows the ATE to vary across states.

When $\Delta = 0$, inequality (12) gives a traditional instrumental-variable assumption asserting that groups of states with different covariates have the same average treatment effect. For example, the ATE might be assumed to be the same across the two treatment groups or the four census regions considered in Section 4.2. The identifying power of this assumption was first analyzed in Manski (1990), where the bound on the overall ATE was shown to be the intersection of the bounds on the group-specific ATEs.

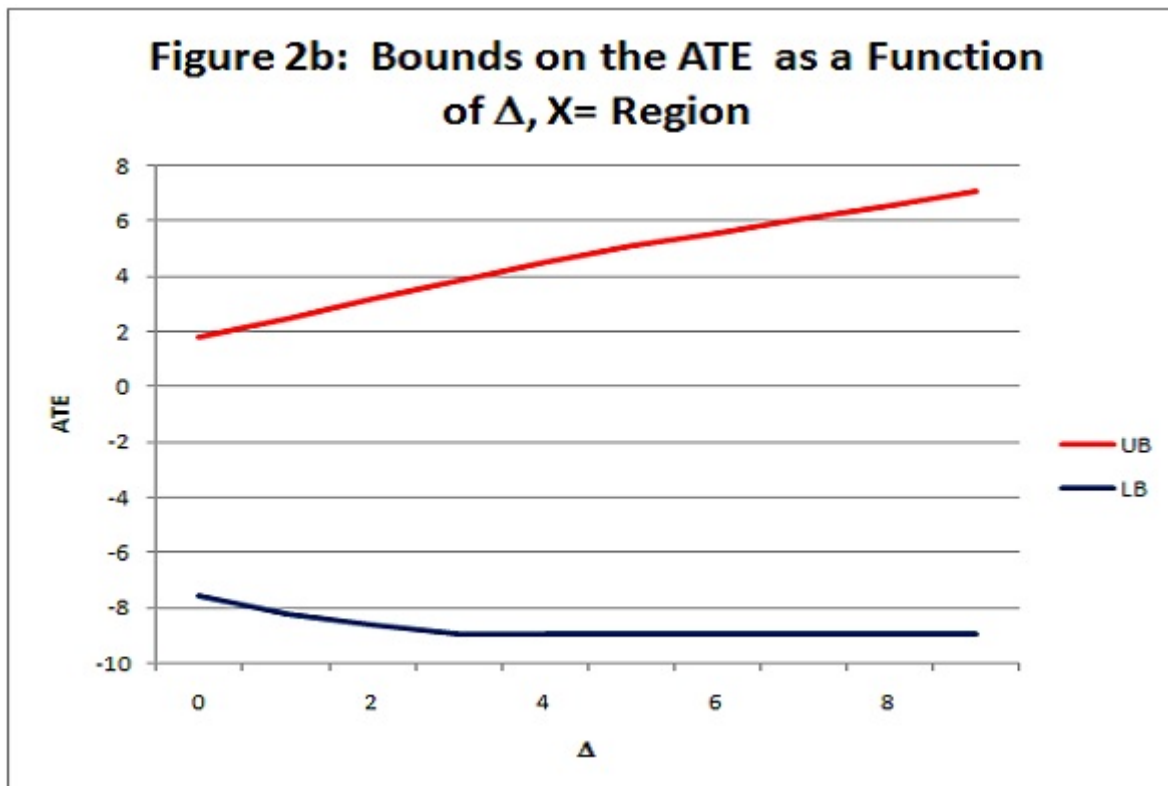
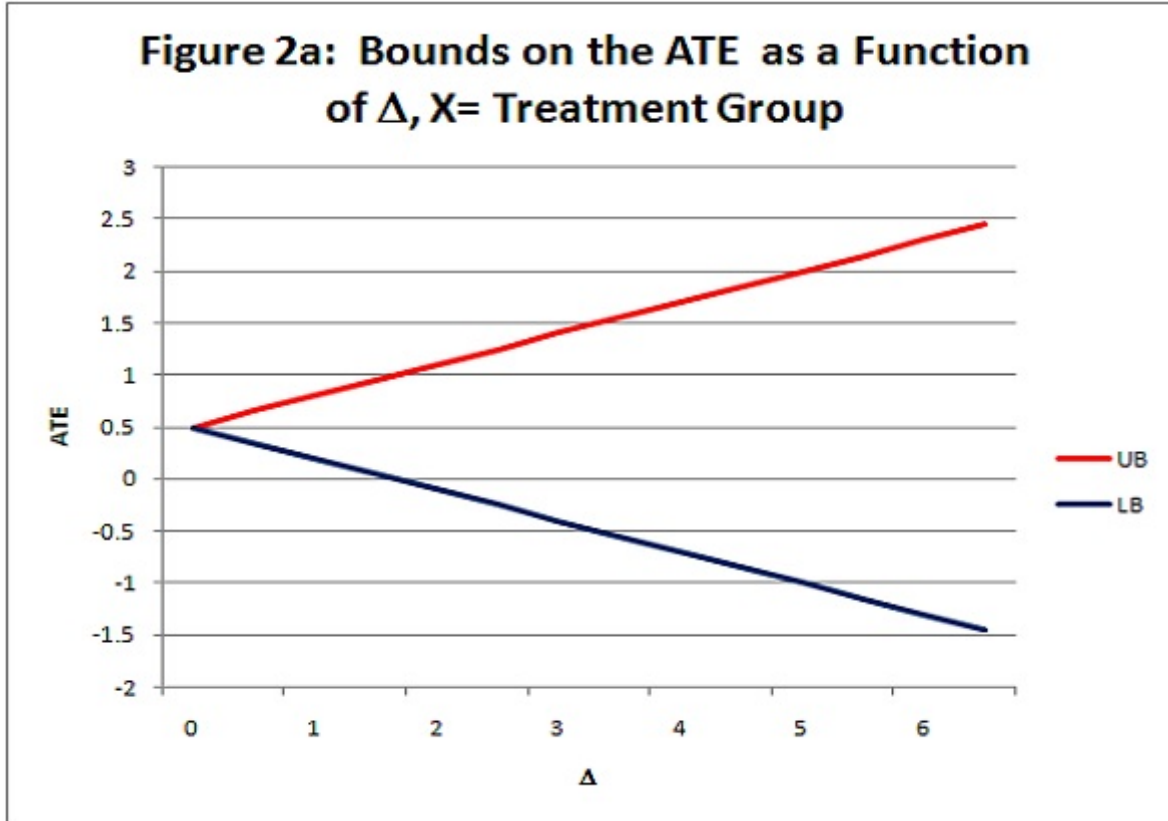
Letting $\Delta > 0$ weakens the traditional assumption by supposing that the ATE may differ across groups by no more than Δ . The larger the selected value of Δ , the weaker the assumption. To assess the sensitivity of inference to choice of Δ , Figure 2 maintains the assumptions of Sections 4.1 and 4.2, adds assumption (12), and displays the bound on the ATE as a function of Δ .


Figure 2a takes the covariate to be the treatment group, treated or untreated. The traditional IV assumption ($\Delta = 0$) point identifies the ATE, revealing that the death penalty increases the mean murder rate by 0.5. This holds because the IV assumption implies that $ETT = ETU$. However,

ambiguity about the ATE increases with Δ . Any value of Δ larger than one renders it impossible to sign the ATE. For example, the bounds on the ATE when $\Delta = 2$ and $\Delta = 5$ are $[-0.3, 0.9]$ and $[-1.2, 1.8]$ respectively. Still, these bounds are substantially more informative than the bound of $[-1.9, 8.3]$ reported in Table 3.

Figure 2b takes the covariate to be the census region. Setting $\Delta = 0$ assumes that all four regions have the same ATE. This implies that the overall ATE lies within the interval $[-7.6, 1.8]$. Setting $\Delta = 2$ implies that the ATE lies in the interval $[-8.6, 3.2]$. Recall that without this restriction linking the ATE across the four regions, we earlier found that the deterrent effect of the death penalty lies in the interval $[-9.0, 10.1]$.

The bound on the ATE using census region as a bounded instrumental variable can be further narrowed if one brings to bear information on the date intercept β . Figure 3 displays the bound as a function of β in the most restrictive case where the ATE is assumed to be the same across all four regions; that is, $\Delta = 0$. This assumption implies that $\beta \in [-2.4, 7.0]$. Prior knowledge of the value of β within this range substantially narrows the bounds on the ATE. In fact, the ATE is nearly point identified if a person knows the exact value of β . For example, the ATE = 0.4 if β is known to equal -1.0 , the value found when treatment group is used as the covariate. And the ATE = -1.6 if β is known to equal 1.0 . Moreover, a person who believes that β is less (greater) than -0.5 can conclude that the ATE is positive (negative). These strong conclusions, however, require sufficiently strong prior information on Δ and β . Someone who thinks that β lies in an interval that includes -0.5 or that the ATE across regions may differ (i.e., $\Delta > 0$) cannot necessarily identify the sign of the ATE.





5. Conclusion

Readers of the 1978 report of the NRC Panel on Research on Deterrent and Incapacitative Effects (Blumstein, Cohen, and Nagin, 1978) will not be surprised by the persistent problems researchers have had in providing credible inference on the deterrent effect of the death penalty. The NRC report warned the research community of the fundamental shortcomings of the data and methods, and questioned whether empirical research could provide useful information at all.

Despite these warnings, various researchers have continued to examine the same or more recent data using the same or similar methods. To yield point identification, research continues to combine data with untenable assumptions. Yet, as in 1978, the results have been found to be highly sensitive to these assumptions and no consensus has emerged. As we see it, the research has failed to provide meaningful answers.

Given that deterrence remains an important and controversial question, it seems useful to consider alternative methodological paradigms. This paper has demonstrated some of what can be learned about the deterrent effect of the death penalty under relatively weak assumptions. In particular, we have studied the identifying power of assumptions restricting variation in treatment response across places and time. The results are findings of partial identification that bound the deterrent effect of capital punishment.

By successively adding stronger identifying assumptions, the analysis makes transparent how assumptions shape inferences. If one assumes only that outcomes are bounded, one cannot identify the sign of the average treatment effect and one can only draw weak conclusions about its magnitude. Those who find it credible to make further assumptions can obtain more informative findings.

Imposing certain assumptions implies that adoption of a death penalty statute increases homicide, but other assumptions imply that the death penalty deters it. Thus, society at large can draw strong conclusions only if there is a consensus favoring particular assumptions. Without such a consensus, data on sanctions and murder rates cannot settle the debate about deterrence. However, data combined with weak assumptions can bound and focus the debate. See Manski and Nagin (1998) for similar analysis of the effect on recidivism of alternative sentencing of juvenile offenders.

To demonstrate partial identification analysis in a simple setting, this paper used only two years of data (1975 and 1977) and compared two broad treatments (the presence or absence of a death penalty statute). Future work can exploit richer data. Whereas the traditional DID framework uses only two periods of data, the approach developed in Manski and Pepper (2011) can exploit multiple periods.

Future work can also use more refined definitions of treatments. One might differentiate treatments by the specifics of the death penalty statute enacted, the way it is implemented, and by the nature of the non-capital sanctions that a state has in place. Using more detailed treatment measures may enable one to study how the specifics of sanctions regimes influence homicide.

We caution that examining more refined treatments can further complicate identification. All else equal, the selection problem intensifies as one refines the definition of treatments. This occurs because the probability that a person receives a refined treatment is necessarily no larger than and typically is smaller than the probability of receiving an aggregated treatment. Refinement of treatment definition may also raise measurement issues. Whereas the presence or absence of a death penalty statute is straightforward to measure, more refined features of sanctions regimes may be less readily observed.

Finally, future work might address the problem of measuring statistical precision. We noted in the Introduction that measurement of statistical precision requires specification of a sampling process that generates the data. However, we are unsure what type of sampling process is reasonable to assume when the data are a repeated cross-section of states. Existing methods for statistical analysis in settings of partial identification assume that the data are a random sample drawn from an infinite population, but this sampling assumption does not seem natural when considering states as units of observation.

Appendix: Computation of Bounds on the ATE under Middle-Ground Assumptions

We used linear programming simplex algorithms coded in Gauss V11 to compute the feasible values of the unknown parameters in (11a)–(11b). For this application, these algorithms are equivalent to a straightforward two step computational method. The first step derives the feasible values of β as the intersection of 2K bounds, each obtained by differencing the bounds on mean response at dates 0 and 1. The second step computes the feasible values of $E_{1|X}$ as the intersection of the dates 0 and 1 bounds on $E_{1|X}$, where the date 1 bounds follow directly from (11b). $E_{0|X}$ is point identified using data from date 0.

To illustrate, consider the basic no-covariate date-invariant treatment effect model in Section 4.1. The feasible values of the three unknowns (β , E_0 , E_1) are the triples that satisfy the four inequalities given in (9a)–(9b). Equation (9a) provides initial (non-sharp) bounds on E_0 and E_1 . In particular, data from date 0 point identify E_0 and provide no information about E_1 .

Given this initial information on (E_0 , E_1), the first step is to derive feasible values of β . Differencing (9a) and (9b) yields two sets of bounds on β , one for each value of t :

$$L_1(t) - U_0(t) \leq \beta \leq U_1(t) - L_0(t), \quad t = 0, 1. \quad (13)$$

Given that β does not vary across treatments, it must lie in the intersection of the bounds for $t = 0$ and $t = 1$. Let $\{\beta_L, \beta_U\}$ denote the lower and upper intersection bounds on β .

Given these bounds on β , the second step is to update the bounds on E_1 . Equation (9) implies two distinct bounds. From (9a), we know that $L_0(1) \leq E_1 \leq U_0(1)$ and, from (9b), we have $L_1(1) -$

$\beta_U \leq E_1 \leq U_1(1) - \beta_L$. Given that E_1 does not vary across dates, it must lie in the intersection of these two bounds.

Further iterations provide no additional information about the feasible values of the parameters. Consider, for example, using the updated bounds on E_1 to further refine the bounds on β . Focusing on the potentially interesting case where the intersection bounds on E_1 are $\{L_1(1) - \beta_U, U_1(1) - \beta_L\}$, we see that

$$L_1(1) - U_1(1) + \beta_L \leq \beta \leq U_1(1) - L_1(1) + \beta_U.$$

This lower bound is smaller than β_L and the upper bound is larger than β_U . Hence, iteration yields no improvement.

We use the same two-step approach to derive bounds on the $2K + 1$ unknown parameters in date-invariant models with covariates. When evaluating the feasible values under the bounded instrumental variable model, values of β which are inconsistent with inequality (12) are classified as infeasible.

References

Athey, S. and G. Imbens (2006), “Identification and Inference in Nonlinear Difference-in Differences Models,” *Econometrica*, 74, 431-97.

Blumstein, A., J. Cohen, and D. Nagin, eds. (1978), *Deterrence and Incapacitation: Estimating the Effects of Criminal Sanctions on Crime Rates*, Washington, D.C.: National Academy Press.

Blundell, R., A. Gosling, H. Ichimura, and C. Meghir (2007), “Changes in the Distribution of Male and Female Wages Accounting for Employment Composition Using Bounds,” *Econometrica*, 75, 323-363.

Chernozhukov, V., I. Fernandez-Val, and W. K. Newey (2010), “Quantile and Average Effects in Nonseparable Panel Models,” Department of Economics, MIT.

Chernozhukov, V., H. Hong, and E. Tamer (2007), “Estimation and Confidence Regions for Parameter Sets in Econometric Models,” *Econometrica*, 75, 1243-1284.

Donohue, J. And J. Wolfers. (2005). “Uses and Abuses of Empirical Evidence in the Death Penalty Debate,” *Stanford Law Review*, 58, 791-841.

Dezhbakhsh, H. and J. Shepherd (2006), “The Deterrent Effect of Capital Punishment: Evidence from a Judicial Experiment,” *Economic Inquiry*, 14, 512-535.

Dezhbakhsh, H., P. Rubin, and J. Shepherd (2003), “Does Capital Punishment Have a Deterrent Effect? New Evidence from Post Moratorium Panel Data,” *American Law and Economic Review*, 5, 344-76.

Ehrlich, I. (1975), “The Deterrent Effect of Capital Punishment: A Question of Life and Death.” *American Economic Review*, 65, 397–417.

Evdokimov, K. (2010). “Identification and Estimation of a Nonparametric Panel Data Model with Unobserved Heterogeneity,” Department of Economics, Princeton University.

Gundersen, C., B. Kreider, and J. Pepper (2011), “The Impact of the National School Lunch Program on Child Health: A Nonparametric Bounds Analysis,” *Journal of Econometrics*, forthcoming.

Imbens, G. and C. Manski (2004), “Confidence Intervals for Partially Identified Parameters.” *Econometrica*, 72, 1845-1857.

Manski, C. (1990), “Nonparametric Bounds on Treatment Effects,” *American Economic Review Papers and Proceedings*, 80, 319–323.

Manski, C. (2003), *Partial Identification of Probability Distributions*, New York: Springer-Verlag.

Manski, C. (2007), *Identification for Prediction and Decision*, Cambridge: Harvard University Press.

Manski, C. and D. Nagin (1998), “Bounding Disagreements About Treatment Effects: A Case Study of Sentencing and Recidivism,” *Sociological Methodology*, 28, 99-137.

Manski, C. and J. Pepper (2000), “Monotone Instrumental Variables: with an Application to the Returns to Schooling,” *Econometrica*, 68, 997–1010.

Manski, C. and J. Pepper (2011), “Partial Identification of Treatment Response with Data on Repeated Cross Sections,” in preparation.

Moffitt, R. (2005), “Remarks on the Analysis of Causal Relationships in Population Research,” *Demography*, 91-108.

Pepper, J. (2000). “The Intergenerational Transmission of Welfare Receipt: A Nonparametric Bounds Analysis,” *The Review of Economics and Statistics*, 82, 472-488.