





Survivor-complier effects in the presence of selection on treatment, with application to a study of prompt ICU admission

Edward H. Kennedy, Steve Harris & Luke J. Keele

To cite this article: Edward H. Kennedy, Steve Harris & Luke J. Keele (2018): Survivor-complier effects in the presence of selection on treatment, with application to a study of prompt ICU admission, Journal of the American Statistical Association, DOI: [10.1080/01621459.2018.1469990](https://doi.org/10.1080/01621459.2018.1469990)

To link to this article: <https://doi.org/10.1080/01621459.2018.1469990>

 View supplementary material [↗](#)

 Accepted author version posted online: 03 May 2018.

 Submit your article to this journal [↗](#)

 View Crossmark data [↗](#)

Survivor-complier effects in the presence of selection on treatment, with application to a study of prompt ICU admission

Abstract

Pre-treatment selection or censoring ('selection on treatment') can occur when two treatment levels are compared ignoring the third option of neither treatment, in 'censoring by death' settings where treatment is only defined for those who survive long enough to receive it, or in general in studies where the treatment is only defined for a subset of the population. Unfortunately, the standard instrumental variable (IV) estimand is not defined in the presence of such selection, so we consider estimating a new survivor-complier causal effect. Although this effect is generally not identified under standard IV assumptions, it is possible to construct sharp bounds. We derive these bounds and give a corresponding data-driven sensitivity analysis, along with nonparametric yet efficient estimation methods. Importantly, our approach allows for high-dimensional confounding adjustment, and valid inference even after employing machine learning. Incorporating covariates can tighten bounds dramatically, especially when they are strong predictors of the selection process. We apply the methods in a UK cohort study of critical care patients to examine the mortality effects of prompt admission to the intensive care unit, using ICU bed availability as an instrument.

Keywords: censoring by death, influence function, instrumental variables, observational study, partial identification, semiparametric theory.

1 Introduction

1.1 Time to ICU Admission

In UK hospitals, nurses can call for a critical care assessment to judge whether deteriorating patients on general hospital wards should be transferred to the intensive care unit (ICU) for a higher level of critical care. When a critical care assessment is ordered, staff from the ICU assess whether the patient would improve under ICU care. Patients can at this time be admitted to the ICU, but they may not be directly transferred to the ICU. If ICU bed space is limited, the patient is admitted but often must wait to be transferred to the ICU. Clinical guidelines in the UK indicate that a patient should be transferred to the ICU within four hours (The Faculty of Intensive Care Medicine & Intensive Care Society 2013).

One open clinical question is whether delayed admission to the ICU is harmful. More specifically: does admission to the ICU in less than four hours contribute to decreased mortality? This question is made more complex by two factors. First, time to admission is undefined for patients that are not admitted to the ICU at the time of assessment. That is, we do not observe a time to transfer for any patient that is not admitted. As such, important information is censored for part of the study population. Second, the mix of patients makes naive comparisons of mortality rates for patients quickly transferred to the ICU to those with delayed admission deceiving. It is very likely that patients with short wait times tend to be sicker than patients with delayed transfer. If that is the case, prompt transfer to the ICU would be associated with higher not lower mortality rates. However, this association arises because of the characteristics of those patients who are promptly

transferred, rather than because of their prompt ICU care.

A carefully designed randomized experiment could overcome these difficulties. If we were to randomize all admitted to patients to different wait times for admission, we could identify the causal effect among a specific sub-population. However, ethical constraints make an experiment of this type infeasible. Alternatively we might seek to find a natural experiment. A natural experiment is some naturally occurring circumstance that results in haphazard if not as-if random assignment of the treatment. An instrument is one type of natural experiment. In our case, an instrumental variable (IV) would give a haphazard nudge towards being admitted to the ICU, but would only affect outcomes indirectly through ICU admission and by no other means.

However, standard IV methods only identify the effect of being admitted to the ICU, while we seek to identify the causal effect of delayed transfer to the ICU. Since wait times are undefined or censored for those patients not admitted to the ICU, standard IV methods fail. Therefore we develop methods for estimating the effect of delayed transfer, among the principal stratum of patients who would always be admitted to the ICU regardless of bed availability. Unfortunately this effect is not point-identified under standard IV assumptions, but we can nonetheless construct sharp covariate-adjusted bounds and corresponding flexible estimation methods. Combining IV methods with partial identification approaches has a long history, to which our work contributes (Manski 1990; Manski 1996; Balke & Pearl 1997; Kitagawa 2009; Siddique 2013; Mealli & Pacini 2013). Next, we describe the data in greater detail.

1.2 Data & Descriptive Statistics

The data we use is from a prospective cohort study of general ward patients who were referred to critical care in the ICU in 48 National Health Service (NHS) hospitals between 1 November 2010 and 31 December 2011 (the SPOTlight study) (Harris *et al.* 2018). The data record information for every general ward patient that was assessed for possible admission to the ICU including the decision to admit a patient to ICU care and the time from assessment to actual transfer into the ICU. Importantly, data collection also included a measure of ICU occupancy rates at the specific time of the patient assessment. This measure of ICU bed availability at time of assessment serves as an instrument in our study. The data collection also included a number of baseline covariates such as age, septic diagnosis (0/1), peri-arrest (0/1), and measures of physiology. These measures of physiology include the Intensive Care National Audit & Research Centre (ICNARC) physiology score, the NHS National Early Warning Score (NEWS) which measures whether respiratory rate, oxygen saturations, temperature, systolic blood pressure, pulse rate, a level of consciousness vary from the norm, and the Sequential Organ Failure Assessment (SOFA) score which ranges from 0 to 24, with higher scores indicating a greater degree of organ failure. The data also record the patient's existing level of care at assessment and recommended level of care after assessment using the UK Critical Care Minimum Dataset (CCMDS) levels of care. These levels are 0 and 1 for normal ward care, 2 for care within a high dependency unit, and 3 for care with intensive care unit. Finally, the data include indicators for whether it was the weekend, out of hours (between 7 PM and 7 AM), or the months from November to February. The primary endpoint is 28-day mortality.

In total, the data contain information on 15158 patients on general hospital wards that were assessed for admission to the ICU. From this total, we exclude 2141 patients due to the presence of a treatment limitation order which excluded the possibility of transfer to the ICU. We further exclude six additional patients that have missing data on the availability of beds in the ICU at the time of assessment, giving us a study population of 13011 patients that were assessed for possible transfer to the ICU. Of these 13011 patients, 27% were admitted to the ICU with a median time to transfer of 2 hours. However, 8.3% of patients that were initially admitted were never transferred to the ICU. Figure 1 plots the distribution of wait times for transfer to the ICU, where it is clear that while many patients are transferred within four hours, the time for many patients significantly exceeds that threshold.

1.3 An Instrumental Variables Approach

As we noted above, we address the question of ICU wait times within an IV framework. Specifically, we use ICU bed availability as an instrument for admission to the ICU. That is, if many beds are available at the time of assessment, this should serve as a haphazard encouragement for ICU admittance. Under a set of causal identification assumptions, IVs allow for the identification of causal effects subject to a form of unobserved confounding that is likely in our application (Angrist *et al.* 1996). See Keele *et al.* (2016) for an in-depth study of bed availability as an instrument for ICU admission.

Any use of IVs requires careful assessment of whether the identification assumptions are plausible. However, standard application of IV methods in our application will fail

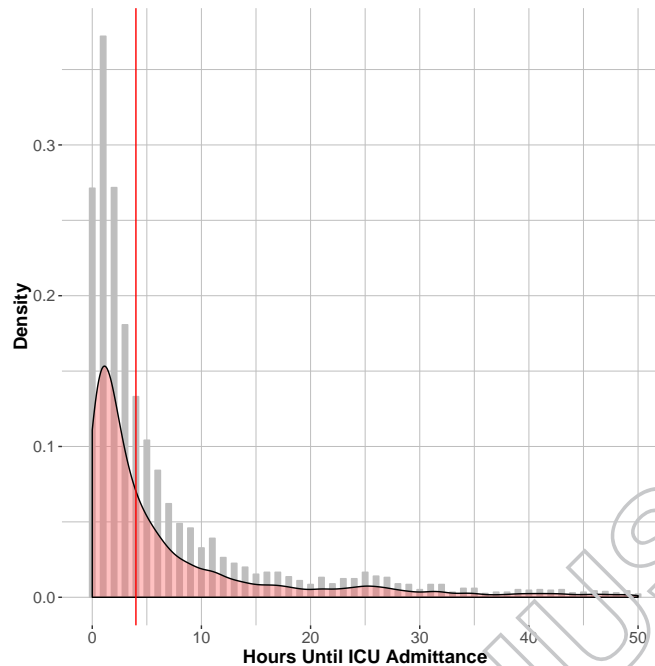


Figure 1: Distribution of ICU wait times. Vertical red line indicates four hour mark. Three observations with wait times in excess of 50 hours removed from the data.

due to selection bias. As we noted above, the waiting time for transfer to the ICU is only observed when a patient is admitted to the ICU. Thus for patients that are assessed but not admitted to critical care, wait times and the counterfactual mortality outcome under wait times are undefined. In our application, selection arises through this censoring, since the treatment may be undefined due to selection that occurs prior to treatment.

Some recent papers have considered settings where the outcome can be undefined but treatment is always observed, i.e., it is assumed that the selection or censoring event (e.g., death) can only occur *after* treatment (Imai 2008; Yang & Small 2016). However, this is substantially different from our application, where the selection is based on treatment status; as we show, selection on treatment leads to different causal estimands and requires

new methods. More relevant is the work by Swanson *et al.* (2015) and Ertefaie *et al.* (2016), who first discovered and discussed the fundamental selection-on-treatment problem with IVs. Our work adds to these papers in several critical respects. First, Swanson *et al.* (2015) importantly pointed out the problem and discussed the biases of standard IV analyses, but did not present any solutions. Ertefaie *et al.* (2016) followed with a sensitivity analysis approach for the same estimand we consider in this paper, but their approach requires numerous sensitivity parameters, does not yield provably sharp bounds, and does not accommodate covariates (even low-dimensional). In contrast, our approach can exploit potentially high-dimensional covariate information to make bounds as sharp as possible (and weaken identifying assumptions), and otherwise is indexed by only two interpretable sensitivity parameters. Importantly we use influence function theory to develop nonparametric estimators that can converge at fast parametric rates, even when based on flexible machine learning methods; our approach can similarly be extended to other partial identification problems involving high-dimensional covariates, which we feel is a critical but understudied area.

This article is organized as follows. In Section 2 we outline notation, define the causal estimand, and derive identification conditions. In Section 3 we outline a set of sharp nonparametric bounds for the causal estimand of interest. We also derive a data-driven sensitivity analysis that uses the bounds to allow for more informative inferences, and propose new estimators that allow for high-dimensional covariate adjustment. We then apply the methods in the aforementioned UK cohort study to examine mortality effects of prompt admission to the ICU. In Section 4 we use simulations to explore the conditions

under which the bounds are most informative. Section 5 concludes.

2 Notation, Causal Estimand, & Identification

2.1 Notation

We suppose we observe an independent and identically distributed sample $(\mathbf{O}_1, \dots, \mathbf{O}_n)$ from distribution \mathbb{P} where $\mathbf{O} = (\mathbf{X}, Z, S, A, Y)$ with \mathbf{X} denoting covariate information, Z a binary instrument, S an indicator of the selection event, A a binary treatment that is only defined when $S = 1$, and Y some outcome of interest (which may also only be defined when $S = 1$). In our example, Z is an indicator of bed availability (with $Z = 1$ indicating that there are more than the average number of beds available), and S indicates whether a patient was accepted for ICU admission. Treatment $A = 1$ indicates a wait time less than 4 hours before actually being transferred into the ICU. For all patients not admitted to the ICU (i.e., with $S = 0$), wait time until transfer into the ICU is undefined, as are counterfactual mortality outcomes under different waiting times. In our application, Y is a binary indicator of death; however all results can also be extended to arbitrary bounded outcomes.

A directed acyclic graph showing the causal structure is given in Figure 2, with arrows between variables indicating causal relationships. The graph shows baseline covariates \mathbf{X} , instrument Z , selection event S (the box enclosing S indicates that subsequent variables are only defined conditional on $S = 1$), treatment A , and outcome Y , along with any unmeasured confounders U . If only time-ordering is taken into account, then all arrows would

be present (excluding those that allow the future to cause the past); however the relationships underlying the gray dotted arrows are assumed absent by standard IV assumptions, as discussed in the next subsection.




Figure 2: Directed acyclic graph showing covariates \mathbf{X} , instrument Z , selection event S , treatment or exposure A , outcome Y , and unmeasured variables U . The exposure A (and possibly outcome Y) are only defined when $S = 1$. Gray dotted arrows indicate relationships that are assumed absent in a standard instrumental variable analysis.

Although we illustrate the data structure in graphical terms in Figure 2, we use potential outcome notation to formalize assumptions and characterize causal effects (Rubin 1974). Specifically we let $(S^{(z)}, A^{(z)}, Y^{(z)})$ denote selection, treatment, and outcome had the instrument been set to $Z = z$, and similarly we let $(Y^{(z,a)}, Y^{(a)})$ denote outcomes had the instrument and/or treatment been set to $Z = z$ and/or $A = a$. Our goal is to learn about the distribution of the effect $Y^{(a=1)} - Y^{(a=0)}$, but this contrast only makes sense for some units in the population, as we will make more explicit in the next subsection.

Finally, we note that selection on treatment can also arise when the goal is to compare only two levels of a multivalued treatment (Swanson *et al.* 2015; Ertefaie *et al.* 2016). In this setting, $S = 1$ would indicate that one of the two treatment levels of interest was received (e.g., statin type), while $S = 0$ would indicate receipt of an alternative treatment

(e.g., no statins at all); here A would indicate which of the two treatments of interest was received, and would be undefined when $S = 0$. This kind of selection bias also occurs when patients die during follow-up, after encouragement by the instrument but prior to the measurement of treatment. Then the selection event S is survival, so that $S = 0$ means a patient died; in this case subsequent treatment and outcomes are undefined, and it often does not make sense to consider what values might have been had the patient survived (Tchetgen Tchetgen 2014).

2.2 Causal Estimand

In this subsection we make explicit our target of inference, i.e., formally define which causal contrast we aim to identify and estimate. Since the treatment of interest A is only defined when $S = 1$ (e.g., how long a patient waits to be transferred to the ICU is only defined for patients who are admitted to be transferred to the ICU), and since setting $S = 1$ for all patients might not be a meaningful counterfactual (e.g., some patients might be too healthy or too sick to ever be admitted to the ICU), it is reasonable to restrict inference to the subgroup of patients who would always be selected, regardless of instrument value, i.e., those patients with

$$S^{(z=1)} = S^{(z=0)} = 1$$

with probability one. These are the patients who would always be admitted to the ICU, regardless of bed availability. Note that this is really the only subgroup for which we can discuss treatment effects, since treatment is not defined if $S = 0$. This is an example of a principal stratum since it is defined based on a cross-classification of joint counterfactuals

(Frangakis & Rubin 2002). This particular kind of principal stratum commonly arises in censoring-by-death settings (Zhang & Rubin 2003; Tchetgen Tchetgen 2014) where counterfactuals are defined based on treatment rather than instrument values, as well as in vaccine effect studies where interest centers on effects among sicker patients who would always be infected regardless of treatment status (Hudgens *et al.* 2003).

One could also argue that the distribution of $Y^{(a=1)} - Y^{(a=0)}$ should be explored in a further refined subgroup, in particular only within those always-selected or survivor patients who would actually respond to encouragement by the instrument, i.e., those with

$$A^{(z=1)} > A^{(z=0)}.$$

These units are commonly termed compliers in the IV literature (Angrist *et al.* 1996). Under appropriate assumptions, these are patients whose wait time would depend on bed availability. That is, these patients would only wait less than 4 hours if there were beds available at the time of critical care assessment. There is an ongoing debate about the merits of exploring causal effects within such subgroups (Imbens 2014; Swanson & Hernán 2014). In this paper we take a pluralistic perspective, viewing complier effects as meaningful but only one piece of the puzzle, which can be combined with inference about other subgroups. Our focus on them is driven by both subject matter concerns in our ICU application and because they are a popular target of inference in IV studies.

As such, in our analysis we focus on the causal effect among those always-selected patients who would respond to IV encouragement, i.e.,

$$\psi = \mathbb{E}\left(Y^{(a=1)} - Y^{(a=0)} \mid S^{(z=1)} = S^{(z=0)} = 1, A^{(z=1)} > A^{(z=0)}\right).$$

We call this parameter the survivor-complier average treatment effect, or SCATE. It represents the mortality risk difference due to waiting less than four hours ($A = 1$) versus more than four hours ($A = 0$), among patients who would always be admitted to the ICU but whose wait time would depend on bed availability. Negative values of this estimate indicate that shorter wait times reduce mortality in this subgroup.

2.3 Identification Conditions

Next, we summarize the standard IV assumptions that we rely on in this paper, and discuss why the SCATE ψ is not point-identified under them. We also give an expression for ψ that will be important for constructing bounds and interpretable sensitivity analyses. Since ψ is expressed in terms of only partially observed potential outcomes, we will need causal assumptions to be able to learn anything about it. Instead of tailoring our assumptions so as to ensure point identification (i.e., so that ψ can be uniquely expressed as a functional of the observed data distribution \mathbb{P}), we instead see what we can learn about the SCATE ψ under standard IV assumptions. Specifically we assume the following:

Assumption 1 (Consistency).

$$S = ZS^{(z=1)} + (1 - Z)S^{(z=0)},$$

$$A = ZA^{(z=1)} + (1 - Z)A^{(z=0)} \text{ (if } S = 1\text{),}$$

$$Y = ZY^{(z=1)} + (1 - Z)Y^{(z=0)} \text{ (if } S = 1\text{).}$$

Assumption 2 (Positivity). $0 < \mathbb{P}(Z = 1 \mid \mathbf{X}) < 1$.

Assumption 3 (Instrumentation). $\mathbb{P}(A^{(z=1)} > A^{(z=0)}, S^{(z=1)} = S^{(z=0)} = 1) \geq \epsilon > 0$.

Assumption 4 (Unconfoundedness). $Z \perp\!\!\!\perp (S^{(z)}, A^{(z)}, Y^{(z)}) \mid \mathbf{X}$ for $z = 0, 1$.

Assumption 5 (Exclusion). $\mathbb{E}(Y^{(z,a)} - Y^{(a)} \mid S^{(z=1)} = S^{(z=0)} = 1) = 0$ for $z = 0, 1$.

Assumption 6 (Monotonicity Restriction).

$$\mathbb{P}(S^{(z=1)} \geq S^{(z=0)}) = 1,$$

$$\mathbb{P}(A^{(z=1)} \geq A^{(z=0)} \mid S^{(z=1)} = S^{(z=0)} = 1) = 1.$$

Analogous of Assumptions 1–6 are all standard in the IV literature, although monotonicity is sometimes replaced by effect homogeneity or no-interaction assumptions (Robins 1994; Hernán & Robins 2006; Tan 2010). Under consistency, we assume that any variation in how patients receive care in the ICU corresponds to the same potential outcomes. The positivity assumption implies that, regardless of covariate value, each patient has some chance of receiving each instrument value. Under instrumentation, the survivor-complier subgroup must be nonempty. We can partly assess this assumption from the data. When the number of ICU beds available is higher than average, patients are admitted to the ICU 61% of the time, and they waited less than four hours 42% of the time. When the number of beds available was lower than average, patients were admitted to the ICU 53% of the time but only 34% waited less than four hours. Unconfoundedness says that (within strata of observed covariates) the instrument is as good as randomized, since it is conditionally independent of potential outcomes. In the data, bed availability in general tells us little about patient characteristics or risk severity. If the exclusion restriction holds among the always-selected, the instrument has no direct effect on outcomes and thus can only affect outcomes through the treatment. We judge the exclusion restriction plausible, since bed occupancy in the ICU is measured when the patient is assessed for

ICU admission. To violate the exclusion restriction, there would have to be some aspects of bed availability at the time of assessment that contributes directly to patient mortality. Finally, the monotonicity assumption implies that there are no units who would always do the opposite of what the instrument encourages them to do, either with respect to selection or treatment; in other words this assumption rules out the possibility of ‘defiers’. The monotonicity assumption would be violated if the hospital staff assessing patients for ICU admission encouraged prompt transfer to the ICU when there were few beds available, and discouraged prompt transfer when there were many beds available. Such behavior seems unlikely. See Angrist *et al.* (1996) and Hernán & Robins (2006) for a more in depth discussion of Assumptions 1–6, including ways in which they could be violated.

With (Z, S, A) all binary, there are nine possible principal strata, three of which are ruled out by monotonicity; this is illustrated in Table 1. Monotonicity implies that any patient who would be admitted with few beds available would also be admitted with more beds available, and among those patients who would always be admitted to the ICU, having more beds available would not delay transfer times relative to having fewer beds available. Thus, in this application, monotonicity is probably a reasonable assumption.

Under Assumptions 1–6 the SCATE parameter ψ is not identified. The following proposition states this, and is proved in the Appendix (with all other results).

Proposition 1. *The survivor-complier effect ψ is not identified under Assumptions 1–6, i.e., there exist data-generating processes with $\psi_1 \neq \psi_2$ that imply the same observed probability distribution \mathbb{P} for (\mathbf{X}, Z, S, A, Y) .*

In other words, because the same observed data distribution \mathbb{P} can be consistent with

Table 1: Principal strata defined by (Z, S, A) and strata excluded by assumption. Strata excluded by monotonicity are shaded in gray. Primary stratum of interest is in row seven.

$S^{(z=0)}$	$S^{(z=1)}$	$A^{(z=0)}$	$A^{(z=1)}$
0	0	.	.
0	1	.	0
0	1	.	1
1	0	0	.
1	0	1	.
1	1	0	0
1	1	0	1
1	1	1	0
1	1	1	1

different values of ψ , even knowing \mathbb{P} without error (e.g., with infinite sample size) is not enough to learn the true value of ψ with certainty. As pointed out for example by Yang & Small (2016) and Ertefaie *et al.* (2016) in related identification problems, this is essentially a result of the fact that the conditional distribution of Y given (Z, S, A) is a mixture of potential outcome distributions from different principal strata.

Although the SCATE ψ is not point-identified, we show in the next section that it is still possible to construct informative bounds and sensitivity analysis. Before that, though, we first define a fundamental expression for the ψ , indicating how it relates to the proportion of survivor-compliers and an intention-to-treat effect. Again, proofs of all results are given in the Appendix.

Proposition 2. *Under Assumptions 1, 3, 5, and 6, we can write the SCATE as*

$$\psi = (\beta/\alpha)\mathbb{P}(S^{(z=1)} = S^{(z=0)} = 1)$$

where

$$\alpha = \mathbb{P}(S^{(z=1)} = S^{(z=0)} = 1, A^{(z=1)} > A^{(z=0)})$$

is the survivor-complier proportion, and

$$\beta = \mathbb{E}(Y^{(z=1)} - Y^{(z=0)} \mid S^{(z=1)} = S^{(z=0)} = 1)$$

is the survivor intention-to-treat effect (i.e., effect of Z among always-selected).

The parameter α is the survivor-complier proportion, which is a measure of the strength of the instrument, and is also important since it measures the size of the non-identified population in which the effect is being estimated. In our application, this is the proportion of patients that wait less than four hours to gain access to the ICU because they were encouraged by bed availability in the ICU. The intention-to-treat effect, β , is the change in mortality caused by bed availability among the sub-population that would always be admitted to the ICU in less than four hours. The term β , in addition to providing an upper-bound on complier effects, does not require monotonicity of treatment (though in our setting its bounds will require monotonicity of selection). The quantities (α, β) may be of interest in their own right, but they also play a crucial role in constructing bounds and sensitivity analyses for the SCATE ψ , since they are the unidentified parts of ψ .

3 Bounds, Sensitivity Analysis, & Estimation

Next, we derive bounds on the SCATE ψ , and we use these bounds to construct optimally informative and interpretable sensitivity analysis parameters; we then detail flexible yet efficient estimation methods.

3.1 Bounds

Although the SCATE ψ is not point-identified, we show here that it is possible to construct potentially informative bounds. Bounds are often used to derive inferences under weaker conditions than those needed for point identification (Grilli & Mealli 2008; Schwartz *et al.* 2012; Zhang *et al.* 2008). Importantly, our analysis exploits covariate information to sharpen bounds as much as possible; then we use these worst-case bounds to aid a more informative sensitivity analysis.

For notational simplicity, we denote

$$\begin{aligned}\pi_z(\mathbf{x}) &= \mathbb{P}(Z = z \mid \mathbf{X} = \mathbf{x}), \\ \theta_z(a \mid \mathbf{x}) &= \mathbb{P}(A = a \mid \mathbf{X} = \mathbf{x}, Z = z), \\ \lambda_z(\mathbf{x}) &= \mathbb{E}(S \mid \mathbf{X} = \mathbf{x}, Z = z), \\ \mu_z(\mathbf{x}) &= \mathbb{E}(YS \mid \mathbf{X} = \mathbf{x}, Z = z)\end{aligned}$$

as important nuisance functions that will be used throughout. Note that since $A \in \{0, 1, \cdot\}$, we have $1 - \theta_z(0 \mid \mathbf{x}) = \theta_z(1 \mid \mathbf{x}) + \theta_z(\cdot \mid \mathbf{x})$ where $A = \cdot$ if $S = 0$. Also, as pointed out earlier, we focus on the case where Y is binary, although our results equally apply to arbitrary outcomes as long as they are bounded.

In the next theorem we provide bounds (α_ℓ, α_u) on the survivor-complier proportion α and bounds (β_ℓ, β_u) on the survivor intention-to-treat effect β . We will use these bounds to construct bounds and a sensitivity analysis for the SCATE ψ .

Theorem 1. *Under Assumptions 1–6, the survivor-complier proportion is bounded as*

$$\mathbb{E}\{\{\theta_0(0 \mid \mathbf{X}) - \theta_1(0 \mid \mathbf{X})\}_+\} \leq \alpha \leq \mathbb{E}\{\theta_1(1 \mid \mathbf{X}) - \theta_0(1 \mid \mathbf{X})\},$$

and the survivor intention-to-treat effect is bounded as

$$\frac{\mathbb{E}[\{\mu_1(\mathbf{X}) + \lambda_0(\mathbf{X}) - \lambda_1(\mathbf{X})\}_+ - \mu_0(\mathbf{X})]}{\mathbb{E}\{\lambda_0(\mathbf{X})\}} \leq \beta \leq \frac{\mathbb{E}[\{\mu_1(\mathbf{X}) \wedge \lambda_0(\mathbf{X})\} - \mu_0(\mathbf{X})]}{\mathbb{E}\{\lambda_0(\mathbf{X})\}}.$$

The intuition behind the bound on α is as follows. Under Assumptions 1–6, the quantity $\theta_1(1 | \mathbf{X}) - \theta_0(1 | \mathbf{X}) = \mathbb{P}(A^{(z=1)} = 1 | \mathbf{X}) - \mathbb{P}(A^{(z=0)} = 1 | \mathbf{X})$ is the difference in proportions of units exposed to the treatment (less than a 4 hour wait) when encouraged versus not among units with covariates \mathbf{X} . This equals the proportion with covariates \mathbf{X} who either are survivor-compliers or are selection-compliers ($S^{(z=1)} > S^{(z=0)}$) that take treatment when encouraged, which is the upper bound on the proportion of survivor-compliers. Averaging the covariate-specific upper bound gives an upper bound on the marginal proportion of survivor-compliers. Similar logic shows that $\theta_0(0 | \mathbf{X}) - \theta_1(0 | \mathbf{X}) = \mathbb{P}(A^{(z=1)} \neq 0 | \mathbf{X}) - \mathbb{P}(A^{(z=0)} \neq 0 | \mathbf{X})$ equals the covariate-specific proportion of survivor-compliers minus the proportion of selection-compliers that take control when encouraged, which is clearly a lower bound. Note that when there is no selection (i.e., $S = 1$ with probability one), the lower and upper bounds collapse to $\mathbb{E}(A | \mathbf{X}, Z = 1) - \mathbb{E}(A | \mathbf{X}, Z = 0)$, which is equivalent to the denominator of the usual Wald estimator.

The bounds on β are a covariate-adjusted version of those previously given by Hudgens *et al.* (2003) and Zhang & Rubin (2003), since β is the survivor average causal effect of encouragement by the IV. As with α , when there is no selection, the bounds on β collapse to the numerator of the usual Wald estimand given by $\mathbb{E}(Y | \mathbf{X}, Z = 1) - \mathbb{E}(Y | \mathbf{X}, Z = 0)$, since in that case we have $\lambda_z(\mathbf{X}) = 1$. In fact the lower bound β_ℓ can collapse to a scaled version of the Wald numerator even in the presence of selection, whenever there is no differential selection for those encouraged or not, i.e., when $\mathbb{P}\{\lambda_1(\mathbf{X}) = \lambda_0(\mathbf{X})\} = 1$ or

equivalently $S \perp\!\!\!\perp Z \mid \mathbf{X}$.

Using the bounds for α and β , we can construct bounds on the SCATE ψ , since the other term in the numerator is identified. These bounds are given in the next theorem.

Theorem 2. *Let (α_ℓ, α_u) and (β_ℓ, β_u) denote the lower and upper bounds for α and β , respectively, given in Theorem 1. Under Assumptions 1–6, the SCATE is bounded as*

$$\frac{\beta_\ell \mathbb{E}\{\lambda_0(\mathbf{X})\}}{(\alpha_u - \alpha_\ell)\mathbb{1}(\beta_\ell > 0) + \alpha_\ell} \leq \psi \leq \frac{\beta_u \mathbb{E}\{\lambda_0(\mathbf{X})\}}{(\alpha_\ell - \alpha_u)\mathbb{1}(\beta_u > 0) + \alpha_u}.$$

As with the bounds on α and β , the bounds (ψ_ℓ, ψ_u) on the SCATE collapse to the usual Wald estimator when there is no selection. However, even in the presence of selection, the bounds operate in a fashion analogous to the Wald estimator. In each case, we have a ratio estimator based on the bounds for α and β . The numerators are bounds on β , which is the analog of the intention-to-treat effect, and the denominators are bounds on $\alpha/\mathbb{P}(S^{(z=1)} = S^{(z=0)} = 1)$, which is the proportion of compliers among the survivors. The correction factor $(\alpha_\ell - \alpha_u)$ picks out the appropriate bound on α depending on the sign for the bound on β . When $S \perp\!\!\!\perp Z \mid \mathbf{X}$, the width of the bounds is completely determined by the non-identification of the proportion of survivor-compliers α .

While the width of the bounds depends on the use of covariates, not all covariates are equally useful in narrowing the bounds. Covariates that are strong predictors of selection S are especially important. In the extreme case where $S = g(\mathbf{X})$ for some g (i.e., covariates can perfectly predict selection), then the bounds collapse to a single point. A proof of this fact is given in the Appendix. In intermediate cases, where the covariates for S are good but not perfect predictors, the bounds can still be much tighter than those not incorporating covariates; this is illustrated via simulations in Section 4. These results show how covariates

can be critical to improving the bounds. They should also help guide data collection in that researchers should focus on covariates that are strong predictors of selection.

Apart from the cases mentioned above where covariates are perfectly predictive, the only way to achieve point identification of the SCATE ψ is to add stronger assumptions. For example, if one assumed that the third stratum in Table 1 was empty, i.e., that $\mathbb{P}(S^{(z=1)} > S^{(z=0)}, A^{(z=1)} = 1 \mid \mathbf{X}) = 0$, then α would be identified with its upper bound $\alpha = \alpha_u$. Conversely, if one assumed the second stratum in Table 1 was empty, i.e., that $\mathbb{P}(S^{(z=1)} > S^{(z=0)}, A^{(z=1)} = 0 \mid \mathbf{X}) = 0$, then α would be identified with its lower bound $\alpha = \alpha_\ell$. Since β is a survivor average effect of IV encouragement, it could be identified by adding assumptions considered by previous authors, such as parametric modeling assumptions (Zhang *et al.* 2009), structural assumptions about pre-treatment covariates (Ding *et al.* 2011), or no unmeasured confounding assumptions given baseline or post-IV covariates (Tchetgen Tchetgen 2014). In our view, additional identifying assumptions of this type are likely implausible in the ICU application.

3.2 Sensitivity Analysis

In some cases, the bounds given in Theorem 2 will be sufficiently narrow so as to be informative and practically meaningful. In other cases, however, they may be wide and thus unsatisfactory in that they only give a worst-case snapshot of possible SCATE values.

In the latter case, we propose that it is more informative to display all possible values of ψ as a function of sensitivity parameters, not just the two most extreme worst-case possibilities. This is in the same spirit as Ichino *et al.* (2008), for example, since the goal

is to use the data to inform more specific inferences about ψ .

In our setting, there are two natural sensitivity parameters: the survivor-complier proportion α and the survivor intention-to-treat effect β . We can vary these sensitivity parameters by interpolating between their upper and lower bounds given in Theorem 1. Therefore we consider the parameter

$$\psi^*(\boldsymbol{\delta}) = \left\{ \frac{\delta_2 \beta_u + (1 - \delta_2) \beta_\ell}{\delta_1 \alpha_u + (1 - \delta_1) \alpha_\ell} \right\} \mathbb{E}\{\lambda_0(\mathbf{X})\}$$

for $(\delta_1, \delta_2) = \boldsymbol{\delta} \in [0, 1]^2$. Note that larger values of δ_1 correspond to larger values of α , and larger values of δ_2 correspond to larger values of β , with $\delta_1 = 1$ recovering α_u and $\delta_1 = 0$ recovering α_ℓ , and similarly for β . And of course we have $\psi_\ell \leq \psi^*(\boldsymbol{\delta}) \leq \psi_u$ for any $\boldsymbol{\delta}$, where (ψ_ℓ, ψ_u) are the bounds on the SCATE given in Theorem 2. However, consideration of $\psi^*(\boldsymbol{\delta})$ allows for more nuance than just (ψ_ℓ, ψ_u) , since $\psi^*(\boldsymbol{\delta})$ allows us to incorporate prior knowledge about the sizes of the survivor-complier proportion α or the intention-to-treat effect β .

Although we index $\psi^*(\boldsymbol{\delta})$ with the parameter $\boldsymbol{\delta}$ to simplify notation and theoretical analysis, in practice (e.g., when plotting results) it will be more useful to index estimates by the implied values of (α, β) , i.e., by $\alpha^* = \delta_1 \alpha_u + (1 - \delta_1) \alpha_\ell$ and $\beta^* = \delta_2 \beta_u + (1 - \delta_2) \beta_\ell$. More specifically, we index $\psi^*(\boldsymbol{\delta})$ by $\boldsymbol{\delta}$ to reflect the fact that the implied values (α^*, β^*) are not arbitrarily chosen, but instead depend on the bounds (α_ℓ, α_u) and (β_ℓ, β_u) implied by \mathbb{P} . This will be important when accounting for uncertainty in estimating $\psi^*(\boldsymbol{\delta})$ due to the fact that we have estimated the bounds for (α, β) .

Note that what we propose is different from a standard sensitivity analysis, which typically relies on fixed arbitrary values of sensitivity parameters. In contrast, our sensitivity

parameters are informed by the data. In other words, we do not need to consider the entire range of values $\alpha^* \in (0, 1]$ and $\beta \in [-1, 1]$, instead we can consider the narrower range of values consistent with the observed data. This will typically lead to a more meaningful and interpretable sensitivity analysis.

3.3 Estimation & Inference

Next, we construct efficient nonparametric estimators for the parameter $\psi^*(\boldsymbol{\delta})$ and detail their asymptotic properties, showing in particular how to construct confidence intervals. Our approach to estimation makes use of semi- and non-parametric efficiency theory and influence functions (Bickel *et al.* 1993; van der Laan & Robins 2003; Tsiatis 2006). Influence functions are essential for a number of reasons. First, the influence function with the smallest variance (called the efficient influence function) determines the efficiency bound for estimating a parameter, thus giving a benchmark against which estimators can be measured and indicating the difficulty of the problem. Second, influence functions can be used to construct estimators with very favorable properties, such as double robustness or general second-order bias. For this reason, influence function-based estimators can attain fast parametric rates of convergence, even while allowing flexible nonparametric estimation of complex high-dimensional nuisance functions.

First, we outline additional notation. We let \mathbb{P}_n denote the empirical measure so that for example sample averages can be written as $n^{-1} \sum_i f(\mathbf{O}_i) = \mathbb{P}_n(f) = \mathbb{P}_n\{f(\mathbf{X})\}$. Further, for any arbitrary variable $T \subset \mathbf{O}$, we let

$$\phi_z(T) = \frac{\mathbb{1}(Z = z)}{\pi_z(\mathbf{X})} \left\{ T - \mathbb{E}(T \mid \mathbf{X}, Z = z) \right\} + \mathbb{E}(T \mid \mathbf{X}, Z = z)$$

denote the uncentered component of the efficient influence function for the parameter $\mathbb{E}\{\phi_z(T)\} = \mathbb{E}\{\mathbb{E}(T \mid \mathbf{X}, Z = z)\}$. Note that, in addition to z and T , the quantity $\phi_z(T)$ also depends on the variables (\mathbf{X}, Z) as well as the nuisance functions $\pi_z(\mathbf{x})$ and $\mathbb{E}(T \mid \mathbf{X} = \mathbf{x}, Z = z)$; for now we keep this dependence implicit for notational simplicity. Similarly we let

$$\hat{\phi}_z(T) = \frac{\mathbb{1}(Z = z)}{\hat{\pi}_z(\mathbf{X})} \left\{ T - \hat{\mathbb{E}}(T \mid \mathbf{X}, Z = z) \right\} + \hat{\mathbb{E}}(T \mid \mathbf{X}, Z = z)$$

denote the estimated version of $\phi_z(T)$, based on nuisance estimators $\hat{\pi}_z(\mathbf{X})$ and $\hat{\mathbb{E}}(T \mid \mathbf{X}, Z = z)$. Note that we are largely agnostic about which particular methods are used to construct these nuisance estimators, leaving this up to the analyst. Our theory allows for either parametric or nonparametric estimators, with asymptotic normality and root- n rates requiring certain convergence rate conditions, as detailed in Theorem 3. The functions ϕ_z and $\hat{\phi}_z$ will play an important role in constructing efficient estimators, analyzing their asymptotic behavior, and computing confidence intervals.

Before discussing our proposed estimators, we need to point out two regularity conditions required for such estimators to be well-behaved. These conditions are required since the bounds $(\alpha_\ell, \beta_\ell, \beta_u)$ from the previous subsection are nonsmooth, i.e., α_ℓ and β_ℓ involve maxima, and β_u involves a minimum. In particular, without such conditions the parameter $\psi^*(\boldsymbol{\delta})$ would not admit an influence function, and regular estimators would not exist. Informally, regular estimators are nicely behaved estimators whose asymptotic distribution is invariant to small shifts in the data-generating process; for more details see van der Vaart (2000), Tsiatis (2006), and Hirano & Porter (2012). The two regularity conditions we need are as follows.

Condition 1. $\mathbb{P}\{\theta_1(0 | \mathbf{X}) = \theta_0(0 | \mathbf{X})\} = 0$.

Condition 2. $\mathbb{P}[\mu_1(\mathbf{X}) \in \{\lambda_0(\mathbf{X}), \lambda_0(\mathbf{X}) - \lambda_1(\mathbf{X})\}] = 0$.

Conditions 1–2 are similar in spirit to so-called non-exceptional law conditions in optimal treatment regime estimation (Robins 2004; van der Laan & Luedtke 2014). However, our conditions may be more plausible since the quantities they restrict are not conditional effects of the treatment A , which might reasonably be zero in practice.

Since under Assumptions 1–6 we have $\theta_z(0 | \mathbf{X}) = \mathbb{P}(A^{(z)} = 0 | \mathbf{X})$, Condition 1 says that there are almost no covariate strata in which the proportion of survivor-compliers equals the proportion of units in the second stratum in Table 1 (i.e., compliers with respect to selection, who take control when selected). In the ICU application this means that, for almost all covariate strata, the proportion of patients who would always be admitted to the ICU regardless of bed availability, but whose wait time would depend on availability (i.e., survivor-compliers), cannot equal the proportion of patients who would only be admitted to the ICU with more beds available, and who would wait more than 4 hours if admitted. This appears to be a plausible assumption in the ICU application; it would be surprising if these particular strata happened to be exactly the same size.

Condition 2 is similar, though perhaps slightly less intuitive. Since under Assumptions 1–6 we have $\mu_z(\mathbf{X}) = \mathbb{E}(Y^{(z)} = 1 | \mathbf{X}, S^{(z)} = 1)\lambda_z(\mathbf{X})$ and $\lambda_z(\mathbf{X}) = \mathbb{P}(S^{(z)} = 1 | \mathbf{X})$, Condition 2 says that there are almost no covariate strata in which the mean outcome under treatment for those selected under encouragement from the IV equals either (a) the ratio $\mathbb{P}(S^{(z=0)} = 1 | \mathbf{X})/\mathbb{P}(S^{(z=1)} = 1 | \mathbf{X})$ of those who would be selected under no encouragement versus encouragement, or (b) this ratio minus one. Note that we must have

$\lambda_1(\mathbf{X}) \geq \lambda_0(\mathbf{X})$ under monotonicity, so the only way the latter part of Condition 2 could be violated is if for some non-negligible strata there are only always- and never-takers with respect to selection, and in these exact strata the outcome for the always-takers happens to be always zero. As with Condition 1, this is a somewhat contrived scenario that we would generally not expect to encounter in practice.

Now we are ready to present our proposed estimator for $\psi^*(\boldsymbol{\delta})$, which is based on estimates of the bounds on α and β . Specifically we let

$$\begin{aligned}\hat{\alpha}_\ell &= \mathbb{P}_n \left(\mathbb{1} \left\{ \hat{\theta}_0(0 | \mathbf{X}) > \hat{\theta}_1(0 | \mathbf{X}) \right\} \left[\hat{\phi}_1 \{ \mathbb{1}(A \neq 0) \} - \hat{\phi}_0 \{ \mathbb{1}(A \neq 0) \} \right] \right) \\ \hat{\alpha}_u &= \mathbb{P}_n \left[\hat{\phi}_1 \{ \mathbb{1}(A = 1) \} - \hat{\phi}_0 \{ \mathbb{1}(A = 1) \} \right] \\ \hat{\beta}_\ell &= \mathbb{P}_n \left(\mathbb{1} \left\{ \frac{\hat{\mu}_1(\mathbf{X})}{\hat{\lambda}_1(\mathbf{X}) - \hat{\lambda}_0(\mathbf{X})} > 1 \right\} \left[\hat{\phi}_1 \{ S(Y = 1) \} - \hat{\phi}_0(S) \right] - \hat{\phi}_0(SY) \right) / \mathbb{P}_n \{ \hat{\phi}_0(S) \} \\ \hat{\beta}_u &= \mathbb{P}_n \left[\mathbb{1} \left\{ \frac{\hat{\mu}_1(\mathbf{X})}{\hat{\lambda}_0(\mathbf{X})} > 1 \right\} \left\{ \hat{\phi}_0(S) - \hat{\phi}_1(SY) \right\} + \hat{\phi}_1(SY) - \hat{\phi}_0(SY) \right] / \mathbb{P}_n \{ \hat{\phi}_0(S) \}\end{aligned}$$

denote efficient influence-function-based estimators of the bounds on (α, β) , and then the corresponding estimator of $\psi^*(\boldsymbol{\delta})$ is given by

$$\hat{\psi}^*(\boldsymbol{\delta}) = \left\{ \frac{\delta_2 \hat{\beta}_u + (1 - \delta_2) \hat{\beta}_\ell}{\delta_1 \hat{\alpha}_u + (1 - \delta_1) \hat{\alpha}_\ell} \right\} \mathbb{P}_n \{ \hat{\phi}_0(S) \}.$$

In the Appendix we show that this estimator solves an estimating equation based on the efficient influence function for $\psi^*(\boldsymbol{\delta})$; this is a standard technique for constructing efficient estimators.

The next theorem gives conditions under which our proposed estimator is asymptotically normal and efficient in a nonparametric model, which puts no constraints on the probability distribution \mathbb{P} (other than, at most, nonparametric smoothness, sparsity, or other structural constraints). To ease notation we let $\theta_{za} = \theta_z(a | \mathbf{x})$ and generally suppress the dependence

of the nuisance functions on \mathbf{X} . We also let $\|f\|^2 = \int f(\mathbf{x})^2 d\mathbb{P}(\mathbf{x})$ denote the squared $L_2(\mathbb{P})$ norm.

Theorem 3. *Along with Assumptions 1–6 and Conditions 1–2, assume*

1. $\{\widehat{\psi}^*(\boldsymbol{\delta}), \widehat{\boldsymbol{\eta}}\} \xrightarrow{p} \{\psi^*(\boldsymbol{\delta}), \boldsymbol{\eta}\}$ for true value $\boldsymbol{\eta} = (\pi_z, \theta_{za}, \lambda_z, \mu_z)$.
2. The influence functions $\boldsymbol{\varphi}(\boldsymbol{\eta}) = (\varphi_\ell^{(\alpha)}, \varphi_u^{(\alpha)}, \varphi_\ell^{(\beta)}, \varphi_u^{(\beta)})$ (as defined in the Appendix) and their estimates $\boldsymbol{\varphi}(\widehat{\boldsymbol{\eta}})$ fall in a Donsker class with probability one as $n \rightarrow \infty$, and are estimated consistently in the sense that $\|\boldsymbol{\varphi}(\boldsymbol{\eta}) - \boldsymbol{\varphi}(\widehat{\boldsymbol{\eta}})\|^2 = o_{\mathbb{P}}(1)$.
3. $\|\widehat{\pi}_1 - \pi_1\| \left(\max_{z,a} \|\widehat{\theta}_{za} - \theta_{za}\| + \max_z \|\widehat{\lambda}_z - \lambda_z\| + \max_z \|\widehat{\mu}_z - \mu_z\| \right) = o_{\mathbb{P}}(1/\sqrt{n})$.
4. $\|\widehat{\eta}_j - \eta_j\| \sqrt{\mathbb{P}(|\eta_j| < |\widehat{\eta}_j - \eta_j|)} = o_{\mathbb{P}}(1/\sqrt{n})$ for $\eta_j \in \{\theta_{10} - \theta_{00}, \mu_1 - (\lambda_1 - \lambda_0), \mu_1 - \lambda_0\}$.

Then

$$\sqrt{n} \left\{ \widehat{\psi}^*(\boldsymbol{\delta}) - \psi^*(\boldsymbol{\delta}) \right\} \rightsquigarrow N \left(0, \text{var} \left[\frac{\varphi^{(\beta)}(\delta_2) - \psi^*(\boldsymbol{\delta}) \varphi^{(\alpha)}(\delta_1)}{\mathbb{E}\{\varphi^{(\alpha)}(\delta_1)\}} \right] \right)$$

where

$$\varphi^{(\alpha)}(\delta_1) = \delta_1 \varphi_u^{(\alpha)} + (1 - \delta_1) \varphi_\ell^{(\alpha)}, \quad \varphi^{(\beta)}(\delta_2) = \delta_2 \varphi_u^{(\beta)} + (1 - \delta_2) \varphi_\ell^{(\beta)}$$

for $(\varphi_\ell^{(\alpha)}, \varphi_u^{(\alpha)}, \varphi_\ell^{(\beta)}, \varphi_u^{(\beta)})$ the efficient influence functions as defined in the Appendix.

Theorem 3 is critically important. It shows that our proposed estimator is \sqrt{n} consistent and asymptotically normal even if the nuisance functions $\boldsymbol{\eta}$ are estimated at slower than parametric \sqrt{n} rates, as long as they attain the $n^{1/4}$ -type rates described in Conditions 3–4 of Theorem 3. Importantly, such $n^{1/4}$ rates are attainable under sparsity, smoothness, or other nonparametric structural constraints, whereas the \sqrt{n} rates required by maximum-likelihood-type estimators necessitate the use of correct parametric models (which can be

difficult if not impossible to specify even in the presence of a few continuous covariates). Further, 95% confidence intervals can be easily constructed as $\widehat{\psi}^*(\boldsymbol{\delta}) \pm 1.96\widehat{\sigma}(\boldsymbol{\delta})/\sqrt{n}$, using empirical estimates of the asymptotic variance $\widehat{\sigma}^2 = \mathbb{P}_n\{(\widehat{\varphi}^{(\beta)} - \widehat{\psi}^*\widehat{\varphi}^{(\alpha)})^2\}/\{\mathbb{P}_n(\widehat{\varphi}^{(\alpha)})\}^2$ (the dependence on $\boldsymbol{\delta}$ is suppressed for notational simplicity). This does not require any extra estimation or model-fitting since the estimated influence functions $\widehat{\varphi}$ are already computed to construct the estimator $\widehat{\psi}^*$. Condition 2 of Theorem 3 puts a mild restriction on the complexity of the nuisance estimators we use, however this can be completely avoided by sample splitting (Zheng & van der Laan 2010). In summary, Theorem 3 allows for efficient \sqrt{n} -rate estimation and valid inference (e.g., confidence intervals), even though our analysis is nonparametric and can incorporate complex machine learning methods.

3.4 Bounding the Effect of ICU Wait Times

We estimate the bounds for SCATE ψ both with and without covariates to understand how covariate information might narrow the bounds. Without using covariates, the upper and lower bounds on the survivor-complier proportion α are 0 and 0.08, respectively. Thus the percentage of patients that would wait less than four hours to gain access to the ICU due to bed availability is between 0 and 8%. The upper and lower bounds for β are -0.19 and 0.06. As such, bed availability among those always admitted to the ICU in less than four hours may have reduced mortality by 19% or increased mortality by 6%. The corresponding bounds for the SCATE ψ are -1 and 1 and thus are uninformative.

We next estimate these bounds exploiting covariate information to weaken identifying assumptions and gain more efficiency. We model the covariate relationships (i.e., nuisance

functions η) flexibly using Super Learner (van der Laan & Rose 2011), combining logistic regression, generalized additive models, lasso, and multivariate adaptive regression splines. With covariates, the bounds on α are 0.01 and 0.07, and the bounds on β are -0.14 and 0.05. The bounds for SCATE ψ remain -1 and 1. Thus although incorporating covariate information sharpens the bounds on α and β considerably (reducing length by 23% for α and 22% for β), it does not sharpen the bounds for the ratio ψ . Next, we conduct a sensitivity analysis where we estimate ψ using the estimated bounds on α and β as sensitivity parameters. This analysis creates a grid of possible ψ estimates, and for each of these ψ estimates, we construct point-wise confidence intervals. The results are contained in Figure 3. These results are substantially more informative than the bounds on ψ alone. We observe that the preponderance of estimates for ψ indicate that shorter waiting times for the ICU reduces mortality. The sensitivity plot does not rule out the possibility that shorter waiting times increase mortality. However, the bounds clearly favor the possibility that prompt ICU care reduces mortality.

As we outlined in Section 3.1, one could invoke stronger assumptions to point identify the SCATE. For example, if one assumes that patients who would be admitted to the ICU only when beds are more available (i.e., those with $S^{(z=1)} > S^{(z=0)}$, perhaps due to less severe symptoms) would always have a longer wait time when admitted (i.e., $A^{(z=1)} = 0$), then the survivor-complier fraction is identified with its upper bound of 7.1%. If the effect of bed availability among those always admitted is further assumed to take its middle value of -4.4%, then the SCATE is point identified, indicating that shorter wait times cause a 21.4% reduction in mortality (among survivor-compliers, 95% CI: 7.6%–35.2%). Alternatively, we

could apply the Wald estimator to the $S = 1$ subpopulation. This assumes that there are no backdoor paths between S and Y . If we apply this to the data, we find that shorter wait times cause a 20.9% reduction in mortality (95% CI: -52.4%–10.6%). Interestingly, these two strategies result in nearly identical inferences. One can use Figure 3 to read off estimates under other identifying assumptions, by fixing α and β values. However, in this example it is not clear that such additional assumptions are plausible.

Alternatively, we could also use the two strategies for point identification outlined in Swanson *et al.* (2015). Both strategies rest on the assumption of observed covariates that block back door paths between either S and Y or A and Y . Under the first strategy, one assumes a set of observed covariates that block back door paths between S and Y . Here, one estimates the following set of inverse probability weights:

$$w_1 = \frac{\Pr[S = 1|Z]}{\Pr[S = 1|Z, X]}.$$

We then estimate $E[Y|A = 1] - E[Y|A = 0]$ by fitting the following linear mean model $E[Y|A] = \theta_0 + \theta_1 A$ using weighted least squares with patients weighted by the estimated values of w_1 . The second approach assumes that there are a set of observed covariates that block back door paths between and A and Y . Now, one estimates the following set of inverse probability weights in the $S = 1$ subpopulation:

$$w_2 = \frac{\Pr[A|S = 1]}{\Pr[A|Z, X, S = 1]}.$$

Again, we estimate $E[Y|A = 1] - E[Y|A = 0]$ by fitting the linear mean model $E[Y|A] = \theta_0 + \theta_1 A$ using weighted least squares but now with patients weighted by the estimated values of w_2 . We implemented both methods using the X data available. Using the first




Figure 3: Sensitivity plot for the effect of ICU waiting time on mortality. Points in the rejection region do not include zero in a 95% point-wise confidence interval. Red regions indicate a harmful effect for shorter waiting times, while black regions indicate a protective effect for a short ICU waiting time.

method, we find that shorter wait times cause a 1.8% reduction in mortality (95% CI: -5.2%–1.5%). The estimate is similar under the second method, as we find that shorter wait times cause a 2.7% reduction in mortality (95% CI: -6.0%–0.5%). However, while the sign of the effects all agree, the magnitude differs substantially, which is consistent with the sensitivity analysis.

4 Simulation

Finally, we conduct a simulation to better understand under what conditions the bounds will be informative. Here, we simulate data from the true causal model and calculate the bounds while varying the magnitude of the true values for α and ψ . The sample size for each simulation is 1000, and the simulations were repeated 1000 times for each combination of the true α and ψ parameters. In all scenarios, the instrumental variable, Z , was generated as a Bernoulli random variable with the probability of success set to 0.5. We simulated principal strata membership probabilistically, with the probability of being in the complier-survivor strata set to α , and the probability of being in the other five principal strata as $(1 - \alpha)/5$. Within the complier-survivor stratum, Y is a draw from a binomial distribution with probability $(0.5 - \psi)/2$, and is a draw from a binomial distribution with probability 0.5 for the other five principal strata. We calculated the bounds for ψ using values of 0 to 1 for α and -1 to 1 for ψ in increments of 0.01. For each simulation, we recorded the length of the bounds and whether the sign for the lower and upper bound agree. These two quantities capture how much information is in the bounds as a function of whether the instrument encourages $A = 1$ and the magnitude of the true causal effect. When A is not highly affected by the instrument and when the true causal effect is small the bounds are unlikely to be informative. However, of greater interest is the point at which the bounds can at least be informative about the sign of the true causal effect.

In the simulation, we also studied how covariates can narrow the bounds. In the first set of simulations, we did not condition on covariates. In two additional simulations, we generated S as a function of two Normally distributed covariates. In one simulation,

the covariates were more weakly informative for S as both were correlated with S at approximately 0.50; in the second set, the two covariates were more informative as both were correlated with S at approximately 0.85.


We summarize the results from the simulations in Figure 4. First, we examine the length of the bounds. In this plot, we observe that the overall length of the bounds is primarily a function of the value of α . That is, the bounds will be at their maximum length of 2 when α is approximately 0.40 or less regardless of the size of ψ . Next, we plot when the bounds are informative as to the sign of the true causal effect. Here, the results depend on the magnitude of both α and ψ . We observe that when the ψ is larger and α is less than 0.40, the bounds may be informative about the sign of the causal effect, but given the information in the two other panels the bounds are likely to be quite wide. However, once α is larger than 0.40 the information about the sign of the causal effect depends directly on the magnitude of ψ . Moreover, for larger values of α the bounds can be highly informative. For example, if ψ is 0.10, and α is 0.65, the bounds are 0.04 and 0.26. If α increases to 0.75 and ψ remains 0.10, the bounds narrow to 0.04 and 0.18. If α remains 0.65, but ψ increases to 0.20, the bounds actually widen to 0.06 and 0.42.

Figure 4 also contains results from the two sets of simulations where covariates were used to construct the bounds. First, we observe that even with the covariates that are more weakly predictive of S , the width of the bounds decreases and the region in which the sign of the bounds is informative is much larger. Across all the values for α and ψ in the simulations, the average length of the bounds without covariates was 0.91. With the weakly predictive covariates, the average length of the bounds decreases to 0.58, and


with the strongly predictive covariates, the average length of the bounds was 0.06. The simulations clearly indicate that the bounds can be highly informative, but it is essential that either values of α are fairly high or that the covariates be at least weakly predictive of S . In the ICU data, we found α to be much lower, and the covariates were not strongly associated with ICU admission. In these situations, use of the sensitivity analysis will likely be more informative.

5 Discussion


One critical clinical question of practical importance given limited health care resources is whether delayed admission to the ICU is beneficial. Answers to this question are elusive since the population referred to critical care tends to be sicker than the population that remains on a general hospital ward, and wait times for the ICU are censored for any patients not admitted to the ICU. While instruments can allow for causal inference in the presence of unobserved confounders, we proved that IV methods do not identify the usual causal estimand of interest in our application. We demonstrated that while point-identification is not possible without further assumptions, partial identification is possible. These bounds can incorporate covariate information in a flexible and data-adaptive way. Moreover, our use of influence function-based methods allows for efficient estimation and inference, even when modeling covariate relationships very flexibly using machine learning tools. Despite their usefulness, influence-function-based approaches seem not to have been used previously in partial identification problems, perhaps due to the relative complexity and nonsmoothness issues discussed in Section 3.3. Using simulations, we find that our




(a) Length of Bounds - No Covariates



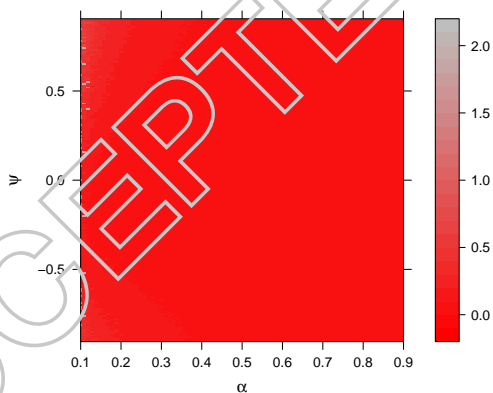
(b) Informative Sign - No Covariates




(c) Length of Bounds - Weak Covariates



(d) Informative Sign - Weak Covariates



(e) Length of Bounds - Strong Covariates



(f) Informative Sign - Strong Covariates

Figure 4: Estimated bounds using simulated data from a known causal model, across 1000 simulation for each α and ψ value. Grey areas indicate that bounds are uninformative. “Weak covariates” refers to case where covariates weakly predict $S = 1$. “Strong covariates” refers to the case where covariates strongly predict $S = 1$.

bounds can be highly informative, but it must be the case that the instrument encourages a relatively large portion of the population to be survivor-compliers, or that covariates strongly predict selection. That is, to produce bounds that are highly informative about the true causal effect, covariate selection is critical unless the effect of the instrument is very strong. Our results serve as an important guide for data collection, by indicating to investigators the importance of measuring covariates that predict selection. For example, one could utilize pilot study data to identify such covariates. Nonetheless our data-driven sensitivity analysis approach can provide a useful picture even when bounds are wide, as observed in our ICU application.

As is often the case with studies based on partial identification, we cannot offer the certainty of a point estimate and confidence interval. However, we think our proposed methods offer important evidence about the clinical question of ICU waiting times. While the bounds on the causal effect of waiting times were uninformative about the true causal effect, use of a sensitivity analysis did provide evidence that waiting less than four hours to receive critical care reduces mortality. While we cannot rule out causal effects of the opposite sign, we argue that, taken with other analyses, there is some evidence that reducing waiting times for the ICU is a clinical goal worth pursuing. Although for some causal questions the best available evidence may not take the form of a single point estimate, we have showed that advances in nonparametric efficiency theory and data-adaptive regression methods can still play a crucial role.

6 References

1. Angrist, J. D., Imbens, G. W. & Rubin, D. B. Identification of causal effects using instrumental variables. *Journal of the American Statistical Association* **91**, 444–455 (1996).
2. Balke, A. & Pearl, J. Bounds on treatment effects from studies with imperfect compliance. *Journal of the American Statistical Association* **92**, 1171–1176 (1997).
3. Bickel, P. J., Klaassen, C. A., Ritov, Y. & Wellner, J. A. *Efficient and Adaptive Estimation for Semiparametric Models* (Johns Hopkins University Press, 1993).
4. Ding, P., Geng, Z., Yan, W. & Zhou, X.-H. Identifiability and estimation of causal effects by principal stratification with outcomes truncated by death. *Journal of the American Statistical Association* **106**, 1578–1591 (2011).
5. Ertefaie, A., Small, D., Flory, J. & Hennessy, S. Selection bias when using instrumental variable methods to compare two treatments but more than two treatments are available. *The International Journal of Biostatistics* **12**, 219–232 (2016).
6. Frangakis, C. E. & Rubin, D. B. Principal stratification in causal inference. *Biometrics* **58**, 21–29 (2002).
7. Grilli, L. & Mealli, F. Nonparametric Bounds on the Causal Effect of University Studies on Job Opportunities Using Principal Stratification. *Journal of Educational and Behavioral Statistics* **33**, 111–130 (2008).

8. Harris, S. *et al.* Impact on mortality of prompt admission to critical care for deteriorating ward patients: an instrumental variable analysis using critical care bed strain. *Intensive Care Medicine* **in press** (2018).
9. Hernán, M. A. & Robins, J. M. Instruments for causal inference: an epidemiologist's dream? *Epidemiology* **17**, 360–372 (2006).
10. Hirano, K. & Porter, J. R. Impossibility results for nondifferentiable functionals. *Econometrica* **80**, 1769–1790 (2012).
11. Hudgens, M. G., Hoering, A. & Self, S. G. On the analysis of viral load endpoints in HIV vaccine trials. *Statistics in Medicine* **22**, 2281–2298 (2003).
12. Ichino, A., Mealli, F. & Nannicini, T. From Temporary Help Jobs to Permanent Employment: What Can We Learn from Matching Estimators and Their Sensitivity? *Journal of Applied Econometrics* **23**, 305–327 (2008).
13. Imai, K. Sharp bounds on the causal effects in randomized experiments with “truncation-by-death”. *Statistics & Probability Letters* **78**, 144–149 (2008).
14. Imbens, G. W. Instrumental variables: An econometrician's perspective. *Statistical Science* **29**, 323–358 (2014).
15. Keele, L. J., Harris, S., Pimentel, S. D. & Grieve, R. Stronger instruments and refined covariate balance in an observational study of the effectiveness of prompt admission to the ICU in the UK. *Unpublished manuscript* (2016).
16. Kitagawa, T. Identification region of the potential outcome distributions under instrument independence. *Centre for Microdata Methods and Practice (Cemmap)* (2009).

17. Manski, C. F. Learning about treatment effects from experiments with random assignment of treatments. *Journal of Human Resources*, 709–733 (1996).
18. Manski, C. F. Nonparametric bounds on treatment effects. *The American Economic Review* **80**, 319–323 (1990).
19. Mealli, F. & Pacini, B. Using Secondary Outcomes to Sharpen Inference in Randomized Experiments With Noncompliance. *Journal of the American Statistical Association* **108**, 1120–1131 (2013).
20. Robins, J. M. Correcting for non-compliance in randomized trials using structural nested mean models. *Communications in Statistics-Theory and methods* **23**, 2379–2412 (1994).
21. Robins, J. M. *Optimal structural nested models for optimal sequential decisions* in *Proceedings of the Second Seattle Symposium in Biostatistics* (2004), 189–326.
22. Rubin, D. B. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology* **66**, 688–701 (1974).
23. Schwartz, S., Li, F. & Reiter, J. P. Sensitivity analysis for unmeasured confounding in principal stratification settings with binary variables. *Statistics in medicine* **31**, 949–962 (2012).
24. Siddique, Z. Partially identified treatment effects under imperfect compliance: the case of domestic violence. *Journal of the American Statistical Association* **108**, 504–513 (2013).

25. Swanson, S. A. & Hernán, M. A. Think globally, act globally: An epidemiologist's perspective on instrumental variable estimation. *Statistical Science* **29**, 371–374 (2014).
26. Swanson, S. A., Robins, J. M., Miller, M. & Hernán, M. A. Selecting on treatment a pervasive form of bias in instrumental variable analyses. *American Journal of Epidemiology* **181**, 191–197 (2015).
27. Tan, Z. Marginal and nested structural models using instrumental variables. *Journal of the American Statistical Association* **105**, 157–169 (2010).
28. Tchetgen Tchetgen, E. J. Identification and estimation of survivor average causal effects. *Statistics in Medicine* **33**, 3601–3628 (2014).
29. The Faculty of Intensive Care Medicine & Intensive Care Society. Core standards for intensive care units (2013).
30. Tsiatis, A. A. *Semiparametric Theory and Missing Data* (Springer, 2006).
31. van der Laan, M. J. & Luedtke, A. R. Targeted learning of an optimal dynamic treatment and statistical inference for its mean outcome. *UC Berkeley Division of Biostatistics Working Paper Series Paper 329*, 1–96 (2014).
32. van der Laan, M. J. & Robins, J. M. *Unified Methods for Censored Longitudinal Data and Causality* (Springer, 2003).
33. van der Laan, M. J. & Rose, S. *Targeted Learning: Causal Inference for Observational and Experimental Data* (Springer, 2011).
34. van der Vaart, A. W. *Asymptotic Statistics* (Cambridge University Press, 2000).

35. Yang, F. & Small, D. S. Using post-outcome measurement information in censoring-by-death problems. *Journal of the Royal Statistical Society: Series B* **78**, 299–318 (2016).
36. Zhang, J. L. & Rubin, D. B. Estimation of causal effects via principal stratification when some outcomes are truncated by death. *Journal of Educational and Behavioral Statistics* **28**, 353–368 (2003).
37. Zhang, J. L., Rubin, D. B. & Mealli, F. in *Modelling and Evaluating Treatment Effects in Econometrics* 117–145 (Emerald Group Publishing Limited, 2008).
38. Zhang, J. L., Rubin, D. B. & Mealli, F. Likelihood-based analysis of causal effects of job-training programs using principal stratification. *Journal of the American Statistical Association* **104**, 166–176 (2009).
39. Zheng, W. & van der Laan, M. J. Asymptotic theory for cross-validated targeted maximum likelihood estimation. *UC Berkeley Division of Biostatistics Working Paper Series Paper* **273**, 1–58 (2010).