
Sharp Bounds for Generalized Causal Sensitivity Analysis

Dennis Frauen, Valentyn Melnychuk & Stefan Feuerriegel
LMU Munich
Munich Center for Machine Learning
{frauen,melnychuk,feuerriegel}@lmu.de

Abstract

Causal inference from observational data is crucial for many disciplines such as medicine and economics. However, sharp bounds for causal effects under relaxations of the unconfoundedness assumption (causal sensitivity analysis) are subject to ongoing research. So far, works with sharp bounds are restricted to fairly simple settings (e.g., a single binary treatment). In this paper, we propose a unified framework for causal sensitivity analysis under unobserved confounding in various settings. For this, we propose a flexible generalization of the marginal sensitivity model (MSM) and then derive sharp bounds for a large class of causal effects. This includes (conditional) average treatment effects, effects for mediation analysis and path analysis, and distributional effects. Furthermore, our sensitivity model is applicable to discrete, continuous, and time-varying treatments. It allows us to interpret the partial identification problem under unobserved confounding as a distribution shift in the latent confounders while evaluating the causal effect of interest. In the special case of a single binary treatment, our bounds for (conditional) average treatment effects coincide with recent optimality results for causal sensitivity analysis. Finally, we propose a scalable algorithm to estimate our sharp bounds from observational data.

1 Introduction

Causal effects are crucial for decision-making in many disciplines, such as marketing [71], medicine [77], and economics [1]. For example, physicians need to know the treatment effects to personalize medical care, and governments are interested in the causal effects of policies on infection rates during a pandemic. In many such applications, randomized experiments are costly or infeasible, because of which causal effects must be estimated from observational data [59].

Estimating causal effects from observational data may lead to bias due to the existence of confounders, i.e., variables that affect both treatment and outcome [55]. A remedy is to observe all confounders and thus to assume unconfoundedness (e.g., as in [16, 63, 64]). However, in many practical applications, the assumption of unconfoundedness is violated. For example, electronic health records do not capture a patient’s ethnic background, which is a known confounder in medicine [52]. In such cases, the causal effect is not identified from observational data, and unbiased estimation is thus impossible [27].

A popular way to perform causal inference in the presence of unobserved confounders is *causal sensitivity analysis*. Causal sensitivity analysis aims to derive bounds on the causal effect of interest under relaxations of the unconfoundedness assumption. Here, the strength of unobserved confounding is typically controlled by some sensitivity parameter Γ , which also determines the tightness of the bounds. In practice, one chooses Γ through domain knowledge [14, 33, 82] or data-driven heuristics [29]. Then, one derives that the causal effect of interest lies in some informative region. For example,

a suitable may enable – despite unobserved confounding – to infer the sign of the treatment effect, which is often sufficient for consequential decision-making [35].

A common model for causal sensitivity analysis is the marginal sensitivity model (MSM) [31, 35, 36, 67]. A benefit of the MSM is that it does not impose any kind of parametric assumptions on the data-generating process. However, the standard MSM is only applicable in settings where the treatment is a single binary variable. Different extensions have been proposed for continuous treatments [10, 32, 45] and for time-varying treatments and confounders [10]. However, these works are restricted to specific settings, while a unified framework for causal sensitivity analysis is still missing.

Contributions: ¹ In this paper, we propose a generalized marginal sensitivity model (GMSM). Our GMSM provides a unified framework for causal sensitivity analysis under unobserved confounding in various settings with

multiple discrete, continuous, and time-varying treatments. Crucially, our GMSM and bounds are applicable to our GMSM includes existing models such as those in [67] and [32], as special cases. As a result, our GMSM enables a unified approach to deriving sharp bounds for a large class of causal effects. To do so, we bound a distribution shift in the unobserved confounders while performing an intervention on the treatments. As a result, we obtain sharp bounds for various causal effects, including (conditional) average treatment effects but also effects for mediation analysis and path analysis (see Fig. 1) and for distributional effects, which have not been studied under an MSM-type sensitivity analysis previously. We also show that, for binary treatments, our bounds coincide with recent optimality results for (conditional) average treatment effects under the MSM. Finally, we propose a scalable algorithm to estimate our sharp bounds from observational data and perform extensive computational experiments to show the validity of our bounds empirically.

2 Related work

In the following, we review related work on causal sensitivity analysis. For a more general review of partial identification and treatment effect estimation under unconfoundedness, we refer to Appendix A.

Causal sensitivity analysis: Causal sensitivity analysis dates back to a study from 1959 showing that unobserved confounders cannot explain away the causal effect of smoking on cancer [14]. This was formalized by introducing sensitivity models that yield bounds on the causal effect of interest under some restriction on the amount of confounding. Previous works introduced a variety of sensitivity models that make different assumptions about the data-generating process and the confounding mechanism. Examples include sensitivity models based on parametric assumptions [30, 62], difference between potential outcome distributions [60, 70], and Rosenbaum's sensitivity model that uses randomization tests [61].

Marginal sensitivity model (MSM): The marginal sensitivity model (MSM) [67] is a common model for sensitivity analysis aimed at settings with binary treatments. Many methods were proposed to estimate bounds from observational data under the MSM. Examples of such methods include linear fractional programs [82] and machine learning such as including kernel-based methods [35] and deep neural networks [31]. Recently, Dorn et al. [18] and Jin et al. [33] showed that estimates from the previous methods are too conservative bounds and, as a remedy, derived closed-form solutions for sharp bounds under the MSM. This was also extended to semiparametric inference [19, 53]. However, all previous methods are limited to binary treatments and (conditional) average treatment effects.

Different extensions for the MSM have been proposed. Jesson et al. [32], Bonvini et al. [10], and Marmarelis et al. [45] developed extensions for continuous treatments. Furthermore, Bonvini et al. [10] also extended the MSM to time-varying treatments and confounders. However, both works are not applicable to settings beyond the estimation of (conditional) average treatment effects, such as

¹Code is available at <https://github.com/DennisFrauen/SharpCausalSensitivity>.

mediation analysis and path analysis or distributional effects. In contrast, we propose a generalized marginal sensitivity model that is compatible with binary, continuous, and time-varying treatments for a variety of causal effects.

Research gap: To the best of our knowledge, no existing causal sensitivity analysis based on the MSM provides a unified framework for deriving bounds for binary, continuous, and time-varying treatments. Furthermore, MSM-based causal sensitivity analysis is restricted to (conditional) average treatment effects and is thus not applicable for complex settings such as effects for mediation analysis and path analysis or distributional effects.

3 Generalized marginal sensitivity model (GMSM)

We first formally define a general setting for causal sensitivity analysis that includes mediation and path analysis (Sec. 3.1 and Sec. 3.2). We then propose our generalized marginal sensitivity model (GMSM) in Sec. 3.3 and compare it with existing sensitivity models from the literature.

Notation: We write random variables as capital letters and their realizations in lowercase. Bold letters \mathbf{X} or \mathbf{x} represent (random) vectors. If $\mathbf{X} = (X_1, \dots, X_k)$ is a sequence of random variables of length k , we denote $\mathbf{X}_k = (X_1, \dots, X_k)$ for $1 \leq k \leq k$. We denote probability distributions over \mathbf{X} as $P^{\mathbf{X}}$ where required. The probability mass function for a discrete \mathbf{X} is denoted as $P(\mathbf{x}) = P(\mathbf{X} = \mathbf{x})$. If \mathbf{X} is continuous $P(\mathbf{x})$ is the probability density function. We denote $P(\mathbf{x})$ as the corresponding probability mass/density function not evaluated at a specific \mathbf{x} . Similarly, we write conditional probability mass functions/density functions as $P(\mathbf{y} | \mathbf{x}) = P(\mathbf{Y} = \mathbf{y} | \mathbf{X} = \mathbf{x})$ and conditional expectations as $E[\mathbf{Y} | \mathbf{x}] = E[\mathbf{Y} | \mathbf{X} = \mathbf{x}] = \int \mathbf{y} P(\mathbf{y} | \mathbf{x}) d\mathbf{y}$. We denote $P(\mathbf{y} | \mathbf{do}(\mathbf{X} = \mathbf{x}))$ as the probability mass function/density function of \mathbf{Y} after performing the do-intervention $\mathbf{do}(\mathbf{X} = \mathbf{x})$ [55]. Finally, we define $\mathbf{X}_{\mathbf{Y}} = \text{pa}(\mathbf{Y}) \setminus \mathbf{X}$, where $\text{pa}(\mathbf{Y})$ denotes the parents of \mathbf{Y} in a given causal graph.

3.1 Problem setup for generalized causal sensitivity analysis

We formalize causal sensitivity analysis based on Pearl's structural causal model framework [55].

Definition 1. A structural causal model (SCMM) is a tuple $\langle V; U; F; P^U \rangle$, where $V = (V_1, \dots, V_k)$ are observed endogenous variables, $U = (U_1, \dots, U_l)$ are unobserved exogenous variables determined outside of the model, $F = \{f_{V_1}, \dots, f_{V_k}\}$ is a set of functions so that each f_i maps a set of parents $\text{pa}(V_i) \subseteq V \cup U$ to V_i , and P^U is a probability distribution on U .

Every SCMM induces unique directed graph \mathcal{G}_M on $V \cup U$ by drawing a directed edge from U_1 to V_2 if $U_1 \in \text{pa}(V_2)$. In this paper, we assume that \mathcal{G}_M is acyclic, i.e., does not contain any directed cycle. Then M induces a unique joint probability distribution $P^{V \cup U}$ on $V \cup U$. Furthermore M induces unique interventional distributions $P^{V \cup U}_{\mathbf{do}(A = a)}$ when performing the do-intervention $\mathbf{do}(A = a)$ for some observed treatments $\mathbf{A} \subseteq V$ [4, 55].

We consider settings where we observe four distinct types of endogenous variables $\mathbf{X}; \mathbf{A}; \mathbf{M}; \mathbf{Y}$: observed confounders $\mathbf{X} \subseteq V \subseteq \mathbb{R}^{d_x}$, treatments $\mathbf{A} \subseteq V \subseteq \mathbb{R}^{d_a}$, discrete mediators $\mathbf{M} = (M_1, \dots, M_k)$ with $M_i \in \mathcal{N}$ for $1 \leq i \leq k$, and an outcome $\mathbf{Y} \subseteq \mathbb{R}$. In a medical setting, \mathbf{X} might be patient characteristics (gender, age, medical history, etc.), \mathbf{A} a medical treatment, \mathbf{M} a change in diet, and \mathbf{Y} a variable indicating a health outcome. Possible causal graphs are shown in Fig. 1. We assume w.l.o.g. that (M_1, \dots, M_k) are ordered causally, i.e. M_{i-1} are parents of M_i for each $i \in \{2, \dots, k\}$. Given treatment interventions $\mathbf{a}_{+1} = (a_1, \dots, a_{+1})$, we are interested in causal effects of the form

$$Q(\mathbf{x}; \mathbf{a}_{+1}; \mathbf{M}) = \sum_{\mathbf{m} \in \mathcal{N}^k} D \cdot P^{\mathbf{Y}}(\mathbf{y} | \mathbf{x}; \mathbf{m}; \mathbf{do}(\mathbf{A} = \mathbf{a}_{+1})) \prod_{i=1}^k P(m_i | \mathbf{x}; m_{i-1}; \mathbf{do}(\mathbf{A} = \mathbf{a}_i)); \quad (1)$$

where D is some functional that maps the density $P^{\mathbf{Y}}(\mathbf{y} | \mathbf{x}; \mathbf{m}; \mathbf{do}(\mathbf{A} = \mathbf{a}_{+1}))$ to a scalar value and we sum over all possible realizations of \mathbf{M} . We are also interested in averaged causal effects

²Note that we consider graphs over both observed and unobserved variables. In the causal inference literature, these are also called augmented causal graphs [9].

$\int_{\mathcal{X}} Q(x; a_{+1}; M) P(x) dx$ or differences (Appendix D). The effect $Q(x; a_{+1}; M)$ generalizes many common effects across different causal inference settings.

Example 1 (CATE, $\delta = 0$). If $D = E[\cdot]$ is the expectation functional, Eq(1) reduces to $Q(x; a; M) = E[Y | x; do(A = a)]$. When A is continuous, this is known as the conditional dose-response function [7, 32]. For binary treatments $A \in \{0, 1\}$, the query $Q(x; 1; M) - Q(x; 0; M)$ is known as the conditional average treatment effect (CATE) [16, 78], and its averaged version as the average treatment effect (ATE) [64, 69].

Example 2 (Mediation analysis, $\delta = 1$). If $D = E[\cdot]$ and $M_1 = M$ is a single mediator, Eq(1) reduces to $Q(x; a_1; M) = \int_{\mathcal{M}} E[Y | x; m; do(A = a_2)] P(m | x; do(A = a_1))$. If $A \in \{0, 1\}$ is binary and the $M \rightarrow Y$ relationship is unconfounded, we obtain the following causal effects studied in mediation analysis [56]: $Q(x; (a_1 = 0; a_2 = 1); M) - Q(x; (a_1 = 0; a_2 = 0); M)$ is the (conditional) natural direct effect (NDE), and $Q(x; (a_1 = 1; a_2 = 0); M) - Q(x; (a_1 = 0; a_2 = 0); M)$ is the (conditional) natural indirect effect (NIE).

In general, Eq(1) includes so-called path-specific effects if the relationship between mediators and outcome is unconfounded [2, 15]. For example, by setting $\delta = 1$ and $a_k = 0$ for all $2 \leq k \leq K + 1$, we obtain the indirect effect that is passed through the mediator sequence (M_1, \dots, M_K) . Path-specific effects are important in various applications, including algorithmic fairness, where the aim is to mitigate effects through paths that are considered unfair [50, 51]. Furthermore, we use a quantile instead of using the mean in all the examples above. This results in distributional versions of CATE, NDE, NIE, and path-specific effects. For example, if the outcome distribution is skewed or contains outliers, practitioners might prefer the median or other quantiles over the mean due to their robustness properties [17].

If all confounders between treatment and mediators and outcome were observed, we could (under additional assumptions) identify the causal effect $Q(x; a_{+1}; M)$ from the observational distribution P^V on V by replacing all do-operations with conditional probabilities according to the backdoor criterion [55]. However, under unobserved confounding, identification is not possible because SCMs M^1 and M^0 exist, which induce the same observational distribution P^V , but which result in different causal effects $Q(x; a_{+1}; M^1) \neq Q(x; a_{+1}; M^0)$ [54, 74].

3.2 Causal sensitivity analysis

In the following, we formalize causal sensitivity analysis as maximizing/minimizing the causal effect $Q(x; a_{+1}; M)$ over all SCMs M that are compatible with a pre-defined sensitivity model. Similar approaches have been used for partial identification with instrumental variables [40, 54] and testing identification of counterfactuals [74, 75], where all SCMs are considered that are compatible with the observational data. However, a sensitivity model must additionally restrict the joint distribution of both observed and unobserved variables in order to allow for informative bounds on the causal effect.

Definition 2 (Sensitivity model) Let P^V denote the distribution of the observed variables $(X; A; M; Y)$. A sensitivity model S is a tuple $(U; P)$, where $U = (U_W)_{W \in \mathcal{W}}$ are unobserved confounders between A and $W \in \{M_1, \dots, M_K; Y\}$, respectively (see Fig. 1) and a family of joint probability distributions on $V \times U$, such that, for all $P \in \mathcal{P}$, it holds that $\int P(v; u) du = P^V(v)$.

We denote the set of all SCMs compatible with S (i.e., that respect the causal graph, induce a distribution $P \in \mathcal{P}$, and do not contain additional confounders $C \setminus S$) (see Appendix E).

Our definition of sensitivity models excludes unobserved confounding between mediators and the outcome. This ensures that the causal effect from Eq(1) can be interpreted as a path-specific effect [2, 56]. We refer to Sec. 6 for a detailed discussion.

Using the definitions above, we can define causal sensitivity analysis as the following partial identification problem: we aim to obtain bounds $Q^-(x; a_{+1}; S) \leq Q^+(x; a_{+1}; S)$ so that

$$Q^-(x; a_{+1}; S) = \sup_{M \in \mathcal{M}(S)} Q(x; a_{+1}; M) \text{ and } Q^+(x; a_{+1}; S) = \inf_{M \in \mathcal{M}(S)} Q(x; a_{+1}; M): \quad (2)$$

$Q^-(x; a_{+1}; S)$ is the maximal causal effect that can be achieved by any SCM that is compatible with the sensitivity model S (and vice versa for $Q^+(x; a_{+1}; S)$). Hence, if the sensitivity model is valid, i.e., contains the ground-truth distribution over observed variables and unobserved confounders, we know that the ground-truth causal effect must be contained in the interval $[Q^-(x; a_{+1}; S); Q^+(x; a_{+1}; S)]$. Bounds for average causal effects and effect differences follow immediately (see Appendix D).

3.3 Generalized marginal sensitivity model (GMSM)

We introduce now the GMSM. We begin by providing the general definition and then show that this extends various marginal sensitivity models from existing literature.

Definition 3 (GMSM). The generalized marginal sensitivity model (GMSM) is a sensitivity model $(U; P)$, where P contains all P that satisfy the following sensitivity constraint: For each $W \in \mathcal{M}_1; \dots; \mathcal{M}_T; Y$, there exist bounds $s_W(a; x) \leq 1 \leq s_W^+(a; x)$, so that, for all u_W, x , and a

$$s_W(a; x) = \frac{P(U_W = u_W \mid x; a)}{P(U_W = u_W \mid x; do(A = a))} \leq s_W^+(a; x) \quad (3)$$

The GMSM bounds the distribution shift in the unobserved confounders when performing the intervention $do(A = a)$ instead of conditioning on $A = a$. This is a restriction on the strength of the effect the unobserved confounders can have on the treatment. If U_W has no effect on A , Eq. (3) holds with $s_W(a; x) = s_W^+(a; x) = 1$. Hence, the further $s_W(a; x)$ and $s_W^+(a; x)$ deviate from 1, the larger is the effect from U_W on A that the GMSM allows for. We will often use a weighted GMSM, which expresses the bounds in terms of a sensitivity parameter.

Definition 4 (Weighted GMSM) A weighted GMSM is a GMSM where $s_W(a; x)$ and $s_W^+(a; x)$ can be written as $s_W(a; x) = \frac{1}{(1 - w)q_W(a; x) + w}$ and $s_W^+(a; x) = \frac{1}{(1 - w^+)q_W(a; x) + w^+}$ for a sensitivity parameter $w \in [0, 1]$ and a weight function $q_W(a; x) \in [0, 1]$ for all x and a .

In a weighted GMSM, the sensitivity parameter w captures the overall restriction on the unobserved confounding strength across individuals. If $w = 1$, no unobserved confounding is allowed, and unconfoundedness holds. For $w < 1$, the restriction is relaxed completely, and arbitrary confounding strength is allowed. The weight function $q_W(a; x) \in [0, 1]$ offers several advantages. It allows to further restrict confounding for individuals with treatment and covariates x . As $q_W(a; x) \rightarrow 1$, confounding is more strongly restricted until unconfoundedness is reached. This is helpful in applications where prior knowledge about the confounding structure is available. As an example, consider a medical setting where we want to estimate the effect of new drug treatments on the risk of developing a certain disease, but we suspect that the data is confounded by the individual's genetic risk for the disease, which is not measured in the observational data. However, we also might know that genetic risk for the disease is not relevant for a specific combination of age and gender.

Comparison with the MSM and its extensions: In the following, we show that the weighted GMSM extends popular sensitivity models from the literature. Proofs are in Appendix C. Note that existing sensitivity models do not consider settings with mediators (i.e. $U = U_Y$), we thus can write $w = w_Y$ for the sensitivity parameter $w(a; x) = w_Y(a; x)$ for the weight function, and $U = U_Y$ for the unobserved confounders. For binary treatments $A \in \{0, 1\}$, the marginal sensitivity model (MSM) [67] is defined via $\frac{1 - \frac{p(x)}{1 - p(x)}}{1 - \frac{p(x; u)}{1 - p(x; u)}}$, where $p(x) = P(A = 1 \mid x)$ denotes the observed propensity score and $p(x; u) = P(A = 1 \mid x; u)$ denotes the full propensity score. For continuous treatments A , the continuous marginal sensitivity model (CMSM) [32] is defined via $\frac{1 - \frac{P(a \mid x; u)}{P(a \mid x)}}{1 - \frac{P(a \mid x; u; u)}{P(a \mid x; u)}}$. For longitudinal settings with time-varying observed confounders $X = (X_1; \dots; X_T)$, unobserved confounders $U = (U_1; \dots; U_T)$, treatments $A = (A_1; \dots; A_T)$, we define the longitudinal marginal sensitivity model (LMSM) via $\frac{1 - \prod_{t=1}^T \frac{P(a_t \mid x_T; u_t; a_{t-1})}{P(a_t \mid x_T; a_{t-1})}}{1 - \prod_{t=1}^T \frac{P(a_t \mid x_T; u_t; a_{t-1})}{P(a_t \mid x_T; a_{t-1})}}$.

Lemma 1. The MSM, CMSM and LMSM are special cases of the weighted GMSM by choosing the weight functions $w(a; x) = P(a \mid x)$ (MSM) and $w(a; x) = 0$ (CMSM, LMSM).

Lemma 1 provides a new interpretation of the MSM: if the probability $P(a \mid x)$ is large, the MSM restricts the confounding strength because most of the randomness in the treatment is already explained by the observed confounders. We can extend this approach to arbitrary discrete treatments within the weighted GMSM framework. The CMSM and LMSM apply the same confounding restriction for all a and x .

4 Bounding causal effects under the GMSM

In this section, we derive sharp bounds under our GMSM. We first quantify the maximal shift in conditional distributions (Sec. 4.1). This allows us to derive an algorithm to compute explicit bounds (Sec. 4.2). Finally, we show how these bounds can be estimated from finite data (Sec. 4.3).

4.1 Shifting interventional distributions

In the following, we provide some intuition before stating the main result. Let us consider a simple setting with treatment A , a single unobserved confounder $U \in \mathbb{R}$, outcome Y , and GSM bounds $s_Y^-(a)$ and $s_Y^+(a)$ (see Fig. 2, left). Any SCMM that is compatible with S describes the relationship between A , U , and Y via a functional assignment $f_Y = f_Y(A; U)$. For a fixed treatment a , we denote $f_Y(a; \cdot)$ as f_a . We are interested in the interventional density $P(y \mid \text{do}(A = a)) = \int P(y \mid a; u) P(u) du = f_{a\#} P^U(y)$, where $f_{a\#} P^U$ denotes the push-forward distribution induced by f_a on P^U . However, we only have access to the observational density $P(y \mid a) = \int P(y \mid a; u) P(u \mid a) du = f_{a\#} P^{U \mid a}(y)$. Hence, for a fixed functional assignment f , we can quantify the discrepancy between $P(y \mid \text{do}(A = a))$ and $P(y \mid a)$ via the distribution shift between P^U and $P^{U \mid a}$.

Fig. 2 shows a toy example where $P^{U \mid a}$ is the uniform distribution on $[0, 1]$, and the functional assignment f_a is the inverse standard normal CDF Φ^{-1} , so that $P(y \mid a)$ is the standard normal probability density. We now want to “right-shift” the interventional density $P(y \mid \text{do}(A = a))$ as much as possible, so that $F(y \mid \text{do}(A = a)) = F(y \mid a)$ for the CDFs. To achieve this, the distribution P^U must put more probability mass on the right-hand side of the unit interval $[0, 1]$ as compared to $P^{U \mid a}$ (see Fig. 2, left). Then, the functional assignment will also push more probability mass to the right of $P(y \mid \text{do}(A = a))$ as compared to $P(y \mid a)$ (see Fig. 2, right). However, the GSM bounds the distribution shift between P^U and $P^{U \mid a}$ via $1 = s_Y^+(a) \int P(u) du - 1 = s_Y^-(a) \int P(u \mid a) du$ (Eq. (3) using that $P(u) = P(u \mid \text{do}(A = a))$). Intuitively, the maximal possible right shift under our GSM.

GSM should occur by choosing P^U so that the bounds are attained, i.e. $\int_{\mathbb{R}} P^U(u) (1 = s_Y^+(a) \mathbb{1}(u < c_Y^+) + 1(u > c_Y^+)) (1 = s_Y^-(a) \mathbb{1}(u < c_Y^-) + 1(u > c_Y^-)) du = 1$. The corresponding “right-shifted” interventional density $P_+(y \mid \text{do}(A = a))$ is defined via the push-forward (see Fig. 2).

It turns out that the density $P_+(y \mid \text{do}(A = a))$ is the maximally right-shifted interventional distribution that can be obtained under any SCM that is compatible with the GSM. Furthermore, the above arguments can be generalized beyond the toy example from Fig. 2.

Theorem 1. Let S be a GSM with bounds $s_W = s_W(a; x)$ and $s_W^+ = s_W^+(a; x)$ for $W \in \mathcal{M}_1; \dots; \mathcal{M}_M; Y$. We define $c_W^+ = \frac{(1 - s_W) s_W^+}{s_W^+ - s_W}$. If $W \in \mathbb{R}$ is continuous, we define the probability density function

$$P_+(w \mid x; m_W; a) = \begin{cases} (1 = s_W^+) P(w \mid x; m_W; a); & \text{if } F(w) < c_W^+; \\ (1 = s_W) P(w \mid x; m_W; a); & \text{if } F(w) > c_W^+; \end{cases} \quad (4)$$

where $F(\cdot)$ is the CDF corresponding to $P(\cdot \mid x; m_W; a)$. If $W \in \mathbb{N}$ is discrete, we define the probability mass function

$$P_+(w \mid x; m_W; a) = \begin{cases} (1 = s_W^+) P(w \mid x; m_W; a); & \text{if } F(w) < c_W^+; \\ (1 = s_W) P(w \mid x; m_W; a); & \text{if } F(w - 1) > c_W^+; \\ (1 = s_W^+) c_W^+ - F(w - 1) + (1 = s_W) F(w) - c_W^+; & \text{else.} \end{cases} \quad (5)$$

Let $F_+(\cdot)$ denote the conditional CDF corresponding to $P_+(\cdot \mid x; m_W; a)$ and let $F(\cdot)$ be the conditional CDF of $P(\cdot \mid x; m_W; a)$, which is defined by swapping signs (s_W^+ and s_W). For any SCMM \mathcal{M} , denote the CDF corresponding to $P_{\mathcal{M}}(\cdot \mid x; m_W; \text{do}(A = a))$ as $F_{\mathcal{M}}(\cdot)$. Then, for all w ,

$$F_+(w) = \inf_{\mathcal{M} \in \mathcal{S}} F_{\mathcal{M}}(w) \quad \text{and} \quad F(w) = \sup_{\mathcal{M} \in \mathcal{S}} F_{\mathcal{M}}(w); \quad (6)$$

Assume now that $P(u_W \mid x; \text{do}(A = a)) = P(u_W \mid x)$. Then, the bounds are sharp (i.e., equality holds in Eq.(6)) whenever A is discrete and it holds that $1 = s_W^+ P(a \mid x)$, or A is continuous.

Proof. See Appendix B. □

The result in Theorem 1 does not depend on the distribution or dimensionality of unobserved confounders, and it does not depend on any specific SCM. As such, our sharp bounds are applicable to a wide class of causal effects, without restricting assumptions on the confounding structure beyond the sensitivity constraint of the GSM.

In case S is a weighted GSM with sensitivity parameters s_W for all $W \in \{M_1, \dots, M_k; Y\}$, the quantiles c_W^+ and c_W^- are of the particular simple form $c_W^+ = s_W/(1+s_W)$ and $c_W^- = 1/(1+s_W)$. Furthermore, the discrete sharpness condition simplifies to $c_W^+ q_W(a; x) + c_W^- P(a; j; x)$. In particular, the bounds are sharp whenever we choose a weighting function that satisfies $P(u_W = j; x; do(A = a)) = P(u_W = j; x)$. The assumption $P(u_W = j; x; do(A = a)) = P(u_W = j; x)$ excludes the time-varying case (LMSM). Deriving sharp bounds for the LMSM is an interesting direction for future research.

4.2 Bounding causal effects

We now leverage Theorem 1 to obtain explicit solutions for the partial identification problem from Eq. (2) with monotone D (see Appendix B). This includes expectation and distributional effects.

Corollary 1 (Bounds without mediators) If $\tau = 0$ and D is monotone, we obtain sharp bounds

$$Q^+(x; a; S) \geq D \left(P_+^Y(j; x; a) \right) \quad \text{and} \quad Q^-(x; a; S) \leq D \left(P_-^Y(j; x; a) \right); \quad (7)$$

and sharpness holds under the same conditions as in Theorem 1.

Proof. See Appendix B. □

Algorithm 1: Causal sensitivity analysis with mediators

```

Input : Causal query  $Q(x; a_{-k+1}; M)$ , GSM  $S$  with  $s_W^+$  and  $s_W^-$ .
Output: Upper bound  $Q^+(x; a_{-k+1}; S)$ 
// Outcome bound
 $c_W^+ = \frac{(1-s_W^-)s_W^+}{s_W^+ + s_W^-}$  for  $W \in \{M_1, \dots, M_k; Y\}$ 
 $Q_{-k+1}^+(m_{\cdot}) \geq D \left( P_+^Y(j; x; m_{\cdot}; a_{-k+1}) \right)$  for  $m_{\cdot} \in \text{supp}(M_{\cdot})$ 
// Adjusting for confounding in mediators
for  $i \in \{1, \dots, k\}$  do
  for  $m_{i-1} \in \text{supp}(M_{i-1})$  do
    Permutation map in ascending order of
     $Q_{i+1}^+(m_{i-1}; (m_i))$ 
     $P_{-}(m_i) = \begin{cases} P(M_i = m_j; x; m_{i-1}; a_i) & \text{if } P_{-}(m_i) < c_{M_i}^+ \\ (1-s_{M_i}^+)P(m_i; j; x; m_{i-1}; a_i) & \text{if } P_{-}(m_i) > c_{M_i}^+ \end{cases}$ 
     $P_{+}(m_i) = \begin{cases} P(M_i = m_j; x; m_{i-1}; a_i) & \text{if } P_{+}(m_i) < c_{M_i}^+ \\ (1-s_{M_i}^+)P(m_i; j; x; m_{i-1}; a_i) & \text{if } P_{+}(m_i) > c_{M_i}^+ \end{cases}$ 
     $Q_i^+(m_{i-1}) = \begin{cases} c_{M_i}^+ P_{-}(m_i) & \text{if } P_{-}(m_i) < c_{M_i}^+ \\ c_{M_i}^+ P_{+}(m_i) & \text{if } P_{+}(m_i) > c_{M_i}^+ \\ P_{-}(m_i) & \text{else.} \end{cases}$ 
  end
end
 $Q^+(x; a_{-k+1}; S) = Q_1^+$ 

```

If D is the expectation functional, A is binary, and Y is continuous, this coincides with the optimality result from Dorn and Guo [18]. Hence, Corollary 1 generalizes the result from Dorn and Guo [18] to distributional effects and arbitrary treatments (e.g., categorical, continuous, or time-varying). In their paper, Dorn and Guo proved the sharpness of the bounds by using the Neyman-Pearson Lemma. In contrast, we take the more principled approach outlined in Sec. 4.1, which is applicable to more general settings.

In the following, we consider settings with mediators, i.e., we aim to derive bounds for the causal effect in Eq(1). The idea is to first obtain outcome bounds $D \left(P_+^Y(j; x; m_{\cdot}; a_{-k+1}) \right)$ conditioned on all possible values $m_{\cdot} \in \text{supp}(M_{\cdot})$ in the support of the mediator M_{\cdot} . Without unobserved confounding between treatments and mediators, we can obtain the upper bound $Q^+(x; a_{-k+1}; S)$ from Eq.(1) by replacing $D \left(P_+^Y(j; x; m_{\cdot}; do(A = a_{-k+1})) \right)$

with $D \left(P_+^Y(j; x; m_{\cdot}; a_{-k+1}) \right)$. With unobserved confounding, we need to additionally take the distribution shift in the mediators into account. To maximize the causal effect, the shifted mediator distribution should put more probability mass on values for which $D \left(P_+^Y(j; x; m_{\cdot}; a_{-k+1}) \right)$ is large. Hence, we can apply the maximal right-shift from Theorem 1 to the mediator distributions in Eq.(1), but where the values are permuted to order $D \left(P_+^Y(j; x; m_{\cdot}; a_{-k+1}) \right)$ in ascending order.

We provide the details on our iterative procedure to compute the upper bound $Q^+(x; a_{-k+1}; S)$ in Algorithm 1. The lower bound $Q^-(x; a_{-k+1}; S)$ can be computed analogously by swapping signs in c_W^+ and s_W^- .

Corollary 2. Under the Assumptions of Theorem 1, Algorithm 1 returns the sharp bounds from Eq. (2).

Proof. See Appendix B. □

4.3 Empirical bounds via importance sampling

In practice, we only have access to an empirical distribution of sample size n instead of the full observational distribution P^Y . We thus obtain estimates $\hat{P}(w | j, x; m_w; a)$ of the conditional density or probability mass function $P(w | j, x; m_w; a)$ for all $W \in \{M_1, \dots, M_k; Y\}$. If W is discrete, this reduces to a standard (multi-class) classification problem, and the estimated class probabilities can be plugged into Algorithm 1. If $W = Y$ is continuous, we can use arbitrary conditional density estimators to obtain $\hat{P}(y | j, x; m; a)$. We propose an importance sampling approach to estimate $D P_+^Y(j, x; m; a)$ assuming that we can sample from our estimated density $\hat{P}(y | j, x; m; a)$ (see Appendix F for details and derivations). If D is the expectation functional, our estimator is

$$D P_+^Y(j, x; m; a) = \frac{1}{\delta_v^k} \sum_{i=1}^k P_{bc}^c y_i + \frac{1}{\delta_v^k} \sum_{i=bc+1}^k P_{bc}^k y_i; \quad \text{where } (y_i)_{i=1}^k \sim \hat{P}(y | j, x; m; a) \text{ is sorted.} \quad (8)$$

We also provide importance sampling estimators for distributional effects in Appendix F. We can also obtain empirical confidence intervals by using the same bootstrap procedure as described in [32].

Implementation: We use feed-forward neural networks with softmax activation function to estimate discrete probability mass functions. For densities, we use conditional normalizing flows [73] (neural spline flows [21]), which are universal density approximators and allow for sampling. We perform training using the Adam optimizer [41]. We also perform extensive hyperparameter tuning in our experiments. Implementation and hyperparameter tuning details are in Appendix G.

5 Experiments

Baselines: Most existing methods focus on sensitivity analysis for binary treatments under the MSM [18, 19, 31, 35, 53]. In this setting, our bounds coincide with existing optimality results [18]. For mediation analysis, we are to the best of our knowledge the first to propose bounds for MSM-based sensitivity analysis. This is why we refrain from benchmarking against baselines in our experiments and only show the validity of our bounds.

Synthetic data: We perform extensive experiments using synthetic data from various causal inference settings to evaluate the validity of our bounds. Synthetic data are commonly used to evaluate causal inference methods as they ensure that the causal ground truth is available [6, 16, 76]. Here, we generate six different synthetic datasets with 50,000 samples in the following manner: We first construct three SCMs with binary treatments for the causal graphs in Fig. 1, that is, (i) no mediators, (ii) a single mediator M , and (iii) two mediators M_1 and M_2 . We then construct three more SCMs for settings (i)–(iii) with continuous treatments. Details are in Appendix H.

To demonstrate the validity of our bounds, we aim to show that they contain the oracle causal effect whenever the sensitivity constraints are satisfied. We evaluate two versions of our GSM: the MSM for binary treatments and the CMSM for continuous treatments (see Sec. 3.3). Using oracle knowledge from the SCMs, we estimate the respective density ratio from

Figure 3: Results for the binary treatment setting. Settings (i)–(iii) are ordered from left to right. The top row shows the oracle sensitivity parameter w (depending on x), and the bottom row shows the bounds (details are in Appendix H).

We thus demonstrate that our bounds contain the oracle causal effect whenever $\gamma \geq \gamma_w$, i.e., whenever we choose sensitivity parameters at least as large as the oracle γ_w .

The results for binary treatments are shown in Fig. 3 and for continuous treatments in Fig. 4. For binary treatments, we evaluate the causal effect $\tau(x; a_{+1}; M)$ for $x \in \{-1, 1\}$ with three (arbitrary) treatment combinations corresponding to the three different settings:

- (i) $a_1 = 1$, (ii) $a_2 = (1; 0)$,
- and (iii) $a_3 = (1; 0; 0)$.

For continuous treatments, we also use three (arbitrary) treatment combinations: (i) $a_1 = 0:6$, (ii) $a_2 = (0:9; 0:5)$, and (iii) $a_3 = (0:2; 0:4; 0:5)$.

Results for additional treatment combinations are in Appendix J. Evidently, the oracle causal effect is contained within our bounds. Hence, the results confirm the validity of our bounds.

Figure 4: Results for continuous treatment setting. Settings (i)–(iii) are ordered from left to right. The top row shows the oracle sensitivity parameter γ_w (depending on x), and the bottom row shows the bounds from Jesson et al.

[32] with our sharp bounds under the CMSM and examine whether we can improve on these by using a different weight function (Def. 4). For this purpose, we modify the SCM from setting (i) in the continuous treatment setting so that there is no unobserved confounding for individuals with $x = 1$ (i.e., $\gamma = 0$). Details are in Appendix H. Table 1 reports the bounds from Jesson et al., our bounds under the CMSM, and our bounds under a weighted CMSM with weight function $w(x) = 1(x > 0)$. The two main findings are: (i) Using the weight function, we can leverage prior knowledge about the confounding structure to obtain even tighter bounds. (ii) Our bounds are much faster to compute (using 5,000 samples). The latter is because we derived closed-form solutions for our bounds, while the method proposed by Jesson et al. is an approximation that uses grid search.

Table 1: Experiment with weighted CMSM

Bounds	Metric/ γ_w		Interval length						Coverage			Time (sec.)	
			1:2		1:5		2		1:2	1:5	2		
Jesson et al. [32] (CMSM)	0:33	0:00	0:74	0:01	1:27	0:01	1	1	1	13748	2:02		
Our bounds (CMSM)	0:33	0:00	0:74	0:01	1:25	0:01	1	1	1	639	0:02		
Our bounds (weighted CMSM)	0:17	0:01	0:37	0:02	0:63	0:03	0:6	1	1	042	0:05		

Reported: Average standard deviation over random seeds (best in bold).

Real-world data: We demonstrate our bounds using an example with real-world data. We consider a setting from the COVID-19 pandemic where mobility (captured through telephone movement) was monitored to obtain a leading predictor of case growth [57]. Details regarding the data (publicly available) and our analysis are in Appendix I. Here, mobility is the mediator, case growth is the outcome, and stay-home order (ban of gatherings with more than 5 people) is the treatment. We are interested in the natural directed effect (NDE, see Example 2) of a stay-home order on the case growth. We suspect that unobserved confounders between treatment and mediator might exist (e.g., adherence, etc. of the population). Hence we perform a causal sensitivity analysis where we vary γ_M and plot the bounds in Fig. 5. We observe that the estimated effect under unconfoundedness ($\gamma_M = 1$) is negative, i.e., the stay-home order decreases case growth. For $\gamma_M < 1$, we obtain bounds around this estimand. For $\gamma_M < 6$ the bounds are negative, which means that under moderate unobserved confounding, it seems likely that the NDE is nonzero, in line with prior evidence [57].

6 Discussion

Assumptions: As common in the causal inference literature, our results rely on assumptions on the data-generating process that must be justified by domain knowledge. Our main assumption is that we exclude confounders between mediators and the outcome in our analysis (see Def. 2). The reason

for this assumption is that it allows us to interpret the causal query from (1) as a path-specific effect. Path-specific effects are defined as so-called nested counterfactuals that lie in the third layer of Pearl's causal hierarchy [55]. Using the assumption from Theorem 1, they can be reduced to the query from Eq. (1), which lies in layer 2 of Pearl's hierarchy, i.e., only depends on interventions [15]. Our sensitivity analysis then bridges the gap from layer 2 to layer 1 (observational data). Approaches for relaxing this assumption, e.g., by combining our results with a sensitivity analysis from layer 3 to layer 2, are left for future work.

Sensitivity parameter and weighting function: Both the sensitivity parameter τ_W and the weighting function $q(a; x)$ incorporate domain knowledge about the unobserved confounding and need to be chosen by the practitioner accordingly.

The sensitivity parameter τ_W controls the strength of unobserved confounding. In practice, one typically chooses τ_W by domain knowledge or data-driven heuristics [35, 28]. One approach is to obtain the smallest τ_W so that the corresponding partially identified interval includes 0. Then,

Figure 5: Estimated upper/lower bound can be interpreted as a level of “causal uncertainty”, for the NDE on real-world data. Re-quantifying the smallest violation of unconfoundedness reported: mean and standard deviation that would explain away the causal effect [31, 33].

over 10 runs. The weight function $q(a; x)$ offers additional opportunities to incorporate domain knowledge about the confounding structure to obtain tighter bounds. Consider an observational study on the effect of smoking on cancer risk, confounded by certain unobserved genes. For example, it may be known that genes that act as unobserved confounders do not affect the cancer risk for a certain population with certain covariates x , which allows us to set $q(a; x) = 1$. Note that we can always set $q(a; x) = P(a | x)$ (discrete treatments) or $q(a; x) = 0$ (continuous treatments) if no domain knowledge is available, leading to established sensitivity models from the literature. For example, practitioners may use our bounds for the MSM in a mediation analysis setting without ever explicitly using the GSM formulation via a weighting function.

Other sensitivity models: In this paper, we provide bounds for MSM-type sensitivity models. Recently, other types of sensitivity models have been proposed in the literature, which may provide less conservative bounds in situations where the data-generating process does not follow an MSM. Examples include ϵ -sensitivity models [34], L_2 -sensitivity models [81], curvature sensitivity models [48] and the τ -MSM [45]. Extending our results to these sensitivity models may be another possible direction for future work.

Efficient estimation: Our main results (Theorem 1, Corollary 1, 2) and identifiability results, i.e., hold in the limit of infinite data. We did not provide results on efficient estimation. Therefore, future work may consider extending our approach to incorporate semiparametric efficiency theory [37].

Conclusion: We proposed a flexible generalization of the MSM and derived sharp bounds for a variety of different causal inference settings and data types (e.g., continuous, categorical). Our work provides practitioners with a unified framework for causal sensitivity analysis. This enables reliable causal inferences from observational data in the presence of unobserved confounders, thus promoting safe decision-making.

Acknowledgments

SF acknowledges funding from the Swiss National Science Foundation (SNSF) via Grant 186932.

References

- [1] Joshua D. Angrist. “Lifetime earnings and the vietnam era draft lotter: Evidence from social security administrative records”. In: *The American Economic Review* 80.3 (1990), pp. 313–336.
- [2] Chen Avin, Ilya Shpitser, and Judea Pearl. “Identifiability of path-specific effects”. In: *JCAI*. 2005.
- [3] Vahid Balazadeh, Vasilis Syrgkanis, and Rahul G. Krishnan. “Partial identification of treatment effects with implicit generative models”. In: *NeurIPS 2022*.
- [4] Elias Bareinboim et al. “On Pearl’s hierarchy and the foundations of causal inference”. In: *Probabilistic and Causal Inference: The Works of Judea Pearl* (2022), pp. 507–556.
- [5] Andrew Bennett, Nathan Kallus, and Tobias Schnabel. “Deep generalized method of moments for instrumental variable analysis”. In: *NeurIPS 2019*.
- [6] Jeroen Berrevoets et al. “To Impute or not to Impute? Missing data in treatment effect estimation”. In: *AISTATS2023*.
- [7] Ioana Bica, James Jordon, and Mihaela van der Schaar. “Estimating the effects of continuous-valued interventions using generative adversarial networks”. In: *NeurIPS 2020*.
- [8] Ioana Bica et al. “Estimating counterfactual treatment outcomes over time through adversarially balanced representations”. In: *ICLR*. 2020.
- [9] Stephan Bongers et al. “Foundations of structural causal models with cycles and latent variables”. In: *The Annals of Statistics* 49.5 (2021), pp. 2885–2915.
- [10] Matteo Bonvini et al. “Sensitivity analysis for marginal structural models”. *arXiv preprint arXiv:2210.04681* (2022).
- [11] Kan Chen, Bingkai Wang, and Dylan S. Small. “A differential effect approach to partial identification of treatment effects”. In: *arXiv preprint arXiv:2303.06332* (2023).
- [12] Victor Chernozhukov, Ivan Fernández-Val, and Blaise Melly. “Inference on counterfactual distributions”. In: *Econometrica* 81.6 (2013), pp. 2205–2268.
- [13] Victor Chernozhukov et al. “Double/debiased machine learning for treatment and structural parameters”. In: *The Econometrics Journal* 21.1 (2018), pp. C1–C68.
- [14] James Cornfield et al. “Smoking and lung cancer: Recent evidence and a discussion of some questions”. In: *Journal of the National Cancer Institute* 22.1 (1959), pp. 173–203.
- [15] Juan D. Correa, Sanghack Lee, and Elias Bareinboim. “Nested counterfactual identification from arbitrary surrogate experiments”. In: *NeurIPS 2021*.
- [16] Alicia Curth and Mihaela van der Schaar. “Nonparametric estimation of heterogeneous treatment effects: From theory to learning algorithms”. In: *AISTATS2021*.
- [17] David L. Donoho and Peter J. Huber. “The notion of breakdown point”. In: *Festschrift for Erich L. Lehman* (1983), pp. 157–184.
- [18] Jacob Dorn and Kevin Guo. “Sharp sensitivity analysis for inverse propensity weighting via quantile balancing”. In: *Journal of the American Statistical Association* (2022).
- [19] Jacob Dorn, Kevin Guo, and Nathan Kallus. “Doubly-valid/ doubly-sharp sensitivity analysis for causal inference with unmeasured confounding”. *arXiv preprint arXiv:2112.11449* (2022).
- [20] Guilherme Duarte et al. “An automated approach to causal inference in discrete settings”. In: *Journal of the American Statistical Association* (2023).
- [21] Conor Durkan et al. “Neural splines”. In: *NeurIPS 2019*.
- [22] Helmut Farbmacher et al. “Causal mediation analysis with double machine learning”. In: *Econometrics Journal* 25.2 (2022), pp. 277–300.
- [23] Dennis Frauen and Stefan Feuerriegel. “Estimating individual treatment effects under unobserved confounding using binary instruments”. In: *ICLR*. 2023.
- [24] Dennis Frauen et al. “Estimating average causal effects from patient trajectories”. In: *ASA An*: 2023.

- [25] Florian Gunsilius. “A path-sampling method to partially identify causal effects in instrumental variable models”. In: arXiv preprint arXiv:1910.09502 (2020).
- [26] Wenshuo Guo et al. “Partial identification with noisy covariates: A robust optimization approach”. In: CLeaR 2022.
- [27] Jason Hartford et al. “Deep IV: A flexible approach for counterfactual prediction”. In: ICML. 2017.
- [28] Tobias Hatt, Daniel Tschernutter, and Stefan Feuerriegel. “Generalizing off-policy learning under sample selection bias”. In: IJAI. 2022.
- [29] Jesse Y. Hsu and Dylan S. Small. “Calibrating sensitivity analyses to observed covariates in observational studies”. In: Biometrics 69.4 (2013), pp. 803–811.
- [30] Guido W. Imbens. “Sensitivity to exogeneity assumptions in program evaluation”. In: American Economic Review 93.2 (2003), pp. 128–132.
- [31] Andrew Jesson et al. “Quantifying ignorance in individual-level causal-effect estimates under hidden confounding”. In: ICML. 2021.
- [32] Andrew Jesson et al. “Scalable sensitivity and uncertainty analysis for causal-effect estimates of continuous-valued interventions”. In: NeurIPS 2022.
- [33] Ying Jin, Zhimei Ren, and Emmanuel J. Candès. “Sensitivity analysis of individual treatment effects: A robust conformal inference approach”. In: Proceedings of the National Academy of Sciences (PNAS) 120.6 (2023).
- [34] Ying Jin, Zhimei Ren, and Zhengyuan Zhou. “Sensitivity analysis under the sensitivity models: A distributional robustness perspective”. In: arXiv preprint arXiv:2203.04373 (2022).
- [35] Nathan Kallus, Xiaojie Mao, and Angela Zhou. “Interval estimation of individual-level causal effects under unobserved confounding”. In: AISTATS 2019.
- [36] Nathan Kallus and Angela Zhou. “Confounding-robust policy improvement”. In: NeurIPS 2018.
- [37] Edward H. Kennedy. “Semiparametric doubly robust targeted double machine learning: A review”. In: arXiv preprint (2022).
- [38] Edward H. Kennedy. “Towards optimal doubly robust estimation of heterogeneous causal effects”. In: arXiv preprint (2022).
- [39] Edward H. Kennedy, Sivaraman Balakrishnan, and Larry Wasserman. “Semiparametric counterfactual density estimation”. In: Biometrika (2023).
- [40] Niki Kilbertus, Matt J. Kusner, and Ricardo Silva. “A class of algorithms for general instrumental variable models”. In: NeurIPS 2020.
- [41] Diederik P. Kingma and Jimmy Ba. “Adam: A method for stochastic optimization”. In: ICLR. 2015.
- [42] Sören R. Künzel et al. “Metalearners for estimating heterogeneous treatment effects using machine learning”. In: Proceedings of the National Academy of Sciences (PNAS) 116 (2019), pp. 4156–4165.
- [43] Bryan Lim, Ahmed M. Alaa, and Mihaela van der Schaar. “Forecasting treatment responses over time using recurrent marginal structural networks”. In: NeurIPS 2018.
- [44] Charles F. Manski. “Nonparametric bounds on treatment effects”. In: The American Economic Review 80.2 (1990), pp. 319–323.
- [45] Myrl G. Marmarelis et al. “Partial identification of dose responses with hidden confounders”. In: UAI. 2023.
- [46] Valentyn Melnychuk, Dennis Frauen, and Stefan Feuerriegel. “Causal transformer for estimating counterfactual outcomes”. In: ICML. 2022.
- [47] Valentyn Melnychuk, Dennis Frauen, and Stefan Feuerriegel. “Normalizing flows for interventional density estimation”. In: ICML. 2023.
- [48] Valentyn Melnychuk, Dennis Frauen, and Stefan Feuerriegel. “Partial counterfactual identification of continuous outcomes with a curvature sensitivity model”. In: NeurIPS 2023.
- [49] Krikamol Muandet et al. “Counterfactual mean embeddings”. In: Journal of Machine Learning Research 22 (2021), pp. 1–71.
- [50] Razieh Nabi, Daniel Malinsky, and Ilya Shpitser. “Learning optimal fair policies”. In: ICML. 2019.
- [51] Razieh Nabi and Ilya Shpitser. “Fair inference on outcomes”. In: IJAI. 2018.

- [52] Ziad Obermeyer et al. "Dissecting racial bias in an algorithm used to manage the health of populations". In: *Science* 366.6404 (2019), pp. 447–453.
- [53] Miruna Oprescu et al. "B-learner: Quasi-oracle bounds on heterogeneous causal effects under hidden confounding". *IntCML*. 2023.
- [54] Kirtan Padh et al. "Stochastic causal programming for bounding treatment effects". *ICLR* 2023.
- [55] Judea Pearl. *Causality*. New York City: Cambridge University Press, 2009.
- [56] Judea Pearl. "Interpretation and identification of causal mediation". *Psychological methods* 19.4 (2014), pp. 459–481.
- [57] Joel Persson, Jurriaan F. Parie, and Stefan Feuerriegel. "Monitoring the COVID-19 epidemic with nationwide telecommunication data". *Proceedings of the National Academy of Sciences of the United States of America* 118.26 (2021).
- [58] Danilo Jimenez Rezende and Shakir Mohamed. "Variational inference with normalizing flows". In: *ICML*. 2015.
- [59] James M. Robins, Miguel A. Hernán, and Babette Brumback. "Marginal structural models and causal inference in epidemiology". *Epidemiology* 11.5 (2000), pp. 550–560.
- [60] James M. Robins, Andrea Rotnitzky, and Daniel o. Scharfstein. "Sensitivity analysis for selection bias an unmeasured confounding in missing data and causal inference models". In: *Statistical Models in Epidemiology, the Environment, and Clinical Trials* 16 (2000), pp. 1–94.
- [61] Paul R. Rosenbaum. "Sensitivity analysis for certain permutation inferences in matched observational studies". In: *Biometrika* 74.1 (1987), pp. 13–26.
- [62] Paul R. Rosenbaum and Donald B. Rubin. "Assessing sensitivity to an unobserved binary covariate in an observational Study with binary outcome". *Journal of the Royal Statistical Society: Series B* 45.2 (1983), pp. 212–218.
- [63] Uri Shalit, Fredrik D. Johansson, and David Sontag. "Estimating individual treatment effect: Generalization bounds and algorithms". *ICML*. 2017.
- [64] Claudia Shi, David M. Blei, and Victor Veitch. "Adapting neural networks for the estimation of treatment effects". In: *NeurIPS* 2019.
- [65] Rahul Singh, Maneesh Sahani, and Arthur Gretton. "Kernel instrumental variable regression". In: *NeurIPS* 2019.
- [66] Vasilis Syrgkanis et al. "Machine learning estimation of heterogeneous treatment effects with instruments". In: *NeurIPS* 2019.
- [67] Zhiqiang Tan. "A distributional approach for causal inference using propensity scores". In: *Journal of the American Statistical Association* 101.476 (2006), pp. 1619–1637.
- [68] Eric J. Tchetgen Tchetgen and Ilya Shpitser. "Semiparametric theory for causal mediation analysis: Efficiency bounds, multiple robustness, and sensitivity analysis". *Annals of Statistics* 40.3 (2012), pp. 1816–1845.
- [69] Mark J. van der Laan and Donald B. Rubin. "Targeted maximum likelihood learning". *The International Journal of Biostatistics* 3.1 (2006).
- [70] Stijn Vansteelandt et al. "Ignorance and uncertainty regions as inferential tools in a sensitivity analysis". In: *Statistica Sinica* 16 (2006), pp. 953–979.
- [71] Hal R. Varian. "Causal inference in economics and marketing". *Proceedings of the National Academy of Sciences (PNAS)* 113.27 (2016), pp. 7310–7315.
- [72] Stefan Wager and Susan Athey. "Estimation and inference of heterogeneous treatment effects using random forests". In: *Journal of the American Statistical Association* 113.523 (2018), pp. 1228–1242.
- [73] Christina Winkler et al. "Learning likelihoods with conditional normalizing flows". *arXiv preprint arXiv:1912.00042* (2019).
- [74] Kevin Xia, Yushu Pan, and Elias Bareinboim. "Neural causal models for counterfactual identification and estimation". *IntCLR*. 2023.
- [75] Kevin Xia et al. "The causal-neural connection: Expressiveness, learnability, and inference". In: *NeurIPS* 2021.
- [76] Liyuan Xu et al. "Learning deep features in instrumental variable regression". *ICLR*. 2021.
- [77] Azam M. Yazdani and Eric Boerwinkle. "Causal inference in the age of decision medicine". In: *Journal of Data Mining in Genomics & Proteomics* 16.1 (2015).

- [78] Jinsung Yoon, James Jordon, and Mihaela van der Schaar. “GANITE: Estimation of individualized treatment effects using generative adversarial nets”. *ICLR*. 2018.
- [79] Marco Zaffalon, Alessandro Antonucci, and Rafael Cabañas. “Structural causal models are (solvable by) credal networks”. In *International Conference on Probabilistic Graphical Models 2020*.
- [80] Junzhe Zhang, Jin Tian, and Elias Bareinboim. “Partial counterfactual identification from observational and experimental data”. *ICML*. 2022.
- [81] Yao Zhang and Qingyuan Zhao. “Bounds and semiparametric inference and L^2 sensitivity analysis for observational studies”. *arXiv preprint arXiv:2211.04697* (2022).
- [82] Qingyuan Zhao, Dylan S. Small, and Bhaswar B. Bhattacharya. “Sensitivity analysis for inverse probability weighting estimators via the percentile bootstrap”. *Journal of the Royal Statistical Society: Series B* 81.4 (2019), pp. 735–761.

A Extended related work

A.1 Partial identification under unobserved confounding

There are various works for partial identification (i.e., bounding causal effects) under unobserved confounding that do not impose sensitivity models. In the following, we discuss the difference between this literature stream, which we call causal partial identification (CPA), to causal sensitivity analysis (CSA).

The main difference between CSA and CPA is that CSA imposes sensitivity models, that is, assumptions on the strength of unobserved confounding, which is controlled by a sensitivity parameter γ . Methods for CPA do not impose sensitivity constraints but instead impose other assumptions. In practice, CSA can be used to test the robustness of causal effect estimates to violations of the unconfoundedness assumption (by varying γ) while CPA may be applicable in situations where no domain knowledge about the confounding strength is available.

Approaches for CPA can be roughly described by two categories that describe the type of assumptions that are imposed to derive informative bounds.

1. Additional variables: In order to achieve informative bounds without restricting the strength of unobserved confounding, some works impose assumptions on the data-generating process by postulating the existence of additional variables. One example is instrumental variables (IVs), i.e., variables, which only have a direct effect on treatment variables but not on outcomes. Under certain assumptions, IVs render bounds for causal effects informative without assumptions on the underlying confounding structure [3, 25, 40, 54]. Other examples include leaky mediation [3, 54], differential effects [11], noisy proxy settings [26].
2. Discrete SCMs: Another stream of literature derives informative bounds in discrete SCMs. Examples include [20, 75, 79, 80].

Comparison with our bounds: None of the methods above is applicable in the causal inference settings we consider (e.g., continuous treatments or outcomes, no IVs available, etc.). In contrast, a well-known CPA result in our setting is the so-called no-assumptions bound [44] that leverages discrete treatments. In the following, we consider the standard setting for CATE estimation without mediators ($\gamma = 0$) and a binary treatment $A \in \{0, 1\}$. We show that we obtain the no-assumptions bound with our sensitivity analysis for the query $E[Y_j | x; \text{do}(A = a)]$ for $a \in \{0, 1\}$, i.e. when lifting the sensitivity constraint. Let $\mathcal{P}_1; \mathcal{P}_2$ denote the support of the conditional density $f(y_j | x; a)$ with c.d.f. F . Using the MSM we obtain

$$Q^+(x; a; S) = \int_{F^{-1}(\frac{a}{1+a})}^1 y (1 - a) P(a | x) + a P(y | x; a) dy \quad (9)$$

$$+ \int_{F^{-1}(\frac{a}{1+a})}^{\mathcal{P}_2} y ((1 - a) P(a | x) + a) P(y | x; a) dy \quad (10)$$

$$\int_{\mathcal{P}_1} P(a | x) E[Y_j | x; a] + (1 - P(a | x)) \mathcal{P}_2; \quad (11)$$

which corresponds to the result in [44]. The result for the lower bound $Q^-(x; a; S)$ follows analogously.

A.2 Estimation of causal effects under unconfoundedness

Under certain additional assumptions, unconfoundedness makes it possible to point-identify causal effects from the observational data, so that the causal inference problem reduces to a purely statistical estimation problem. Various methods for estimating point-identified causal effects under unconfoundedness have been proposed that make use of machine learning and/or (semiparametric) statistical theory. Examples include methods for conditional average treatment effects [13, 16, 38, 42, 63, 72, 78], average treatment effects [24, 64, 69], instrumental variables [5, 23, 27, 65, 66, 76], time-varying data [8, 43, 46], mediation analysis [68, 22], and distributional effects [12, 39, 47, 49]. Note that all the methods above are biased if the unconfoundedness assumption is violated, which outlines the need for our causal sensitivity analysis.

B Proofs of the GSM bounds

B.1 Proof of Theorem 1

Proof. We give the proof for continuous $W \in \mathbb{R}$ (see Eq.(4) in the main paper). The derivation for discrete $W \in \mathbb{N}$ (see Eq.(5) in the main paper) follows with the same arguments and the normalization constraint $\int_{\mathcal{W}} P_+(w | x; m_W; a) = 1$. We first show the inequality

$$F_+(w) \geq \inf_{M \in \mathcal{C}(S)} F_M(w) \quad (12)$$

and then provide sharpness results by showing

$$F_+(w) = \inf_{M \in \mathcal{C}(S)} F_M(w) \quad (13)$$

under the sharpness conditions of Theorem 1. The result follows analogously.

Validity of the bounds (12): We provide a proof by contradiction. To do so, we assume that there exists an SCM $M \in \mathcal{C}(S)$ and $w \in \mathbb{R}$ so that

$$F_+(w) > F_M(w): \quad (14)$$

By the definition of $F_+(w)$, there must exist a set $W_1 \subseteq \mathbb{R}$ with $P_+(w | x; m_W; a) > P(w_1 | x; m_W; \text{do}(A = a))$ for all $w_1 \in W_1$;

$$P_+(w_1 | x; m_W; a) > P(w_1 | x; m_W; \text{do}(A = a)) \quad \text{for all } w_1 \in W_1; \quad (15)$$

or a set $W_2 \subseteq \mathbb{R}$ with $P_+(w_2 | x; m_W; a) < P(w_2 | x; m_W; \text{do}(A = a))$ for all $w_2 \in W_2$;

$$P_+(w_2 | x; m_W; a) < P(w_2 | x; m_W; \text{do}(A = a)) \quad \text{for all } w_2 \in W_2; \quad (16)$$

as otherwise $P(w | x; m_W; \text{do}(A = a))$ would not integrate to 1. Let $W = f(X; M_W; A; U_W)$ be the functional assignment M and let $U_1 = f_{x; m_W; a}^{-1}(W_1) \subseteq \mathbb{R}^d$ and $U_2 = f_{x; m_W; a}^{-1}(W_2) \subseteq \mathbb{R}^d$ denote the preimages of W_1 and W_2 under $f_{x; m_W; a}$ in the confounding space.

We can again write $P_+(w | x; m_W; a)$ as a push forward

$$P_+(w | x; m_W; a) = \int_{\mathbb{R}^d} P_+^{U_W | x; a}(w) P(u_W | x; a) du_W \quad (17)$$

for some density $P_+^{U_W | x; a}$ on the confounding space. By the definition of $P_+(w | x; m_W; a)$ and Eq. (28), we obtain

$$P_+(u_1 | x; a) = \frac{1}{s_W^+} P(u_1 | x; a) \quad \text{and} \quad P_+(u_2 | x; a) = \frac{1}{s_W} P(u_2 | x; a) \quad (18)$$

for all $u_1 \in U_1$ and $u_2 \in U_2$. Due to the definition of U_1 and U_2 , it follows that there exist $u_1 \in U_1$ and $u_2 \in U_2$, so that

$$\frac{P(u_1 | x; a)}{P(u_1 | x; \text{do}(A = a))} > \frac{P(u_1 | x; a)}{P_+(u_1 | x; a)} = s_W^+ \quad (19)$$

and

$$\frac{P(u_2 | x; a)}{P(u_2 | x; \text{do}(A = a))} < \frac{P(u_2 | x; a)}{P_+(u_2 | x; a)} = s_W. \quad (20)$$

Both Eq.(19) and Eq.(20) are contradictions to the GSM constraint (Eq. 1) of the main paper. Hence, $M \notin \mathcal{C}(S)$.

Sharpness of the bounds (13): We show that under the sharpness conditions from Theorem 1, there exists an SCM $M \in \mathcal{C}(S)$ with induced interventional density $P_+(w | x; m_W; a)$ for all w . The construction of M is similar to that of our motivational toy example in Sec. 4.1 of the main paper. We first define an (interventional) probability density for the unobserved confounder $U \in \mathbb{R}^d$ given X via

$$P(u_W | x; \text{do}(A = a)) = 1(0 \leq u_W^{(1)} \leq c_W^+) (1 = s_W^+) + 1(1 - u_W^{(1)} < c_W^+) (1 = s_W); \quad (21)$$

where $u_W^{(1)}$ denotes the first coordinate of u_W . $P(u_W | x; m_W)$ is a properly normalized density with support $[0, 1]^d$ because

$$\int_{\mathbb{R}^d} P(u_W | x; \text{do}(A = a)) du_W = \frac{c_W^+}{s_W^+} + \frac{1 - c_W^+}{s_W} = \frac{1}{s_W^+} \frac{s_W}{s_W} + \frac{s_W^+}{s_W^+} \frac{1}{s_W} = 1; \quad (22)$$

where we used the definition $\alpha_W^+ = \frac{(1 - s_W) s_W^+}{s_W^+ - s_W}$.

We now define the probability density for the unobserved confounder $U_W \in \mathbb{R}^d$ given X, M_W , and treatment A as the uniform density on $[0, 1]^d$, i.e.,

$$P(u_W \mid j; x; a) = \mathbb{1}(0 \leq u_W \leq 1); \quad (23)$$

where the indicator function is defined coordinate-wise.

We now need to show that there always exists an SCM which induces the densities in Eq.(21) and Eq.(22) and that satisfies the sensitivity constraints. For this purpose, we use the assumptions to write the interventional density for discrete A as

$$P(u_W \mid j; x; \text{do}(A = a)) = P(u_W \mid j; x) = P(a \mid j; x) + (1 - P(a \mid j; x))P(u_W \mid j; x; A \neq a); \quad (24)$$

Hence, our definitions for $P(u_W \mid j; x; \text{do}(A))$ and $P(u_W \mid j; x; a)$ are valid whenever

$$0 \leq P(u_W \mid j; x; A \neq a) = \frac{P(u_W \mid j; x; \text{do}(A = a)) - P(a \mid j; x)}{1 - P(a \mid j; x)}; \quad (25)$$

which follows from the sharpness condition $s_W^+ \geq P(a \mid j; x)$. Sharpness for continuous A follows because we can approximate A with a sequence of discrete treatments $(A_n)_{n \in \mathbb{N}}$ and it holds that $P(A_n \mid j; x) \rightarrow 0$ (as long as $P(a \mid j; x)$ is non-degenerate). Finally, it holds that

$$s_W \leq \frac{P(U_W = u_W \mid j; x; a)}{P(U_W = u_W \mid j; x; \text{do}(A = a))} \leq s_W^+ \quad \text{for all } u_W \in [0, 1]^d; \quad (26)$$

so that M respects the sensitivity constraint of the GMSM.

Let now $P(w \mid j; x; m_W; a)$ denote the observational density of W given X, M_W , and A with corresponding cumulative distribution function (CDF) given by $F_W(w)$. To complete our construction of M , we define functional assignment $W = f_W(X; M_W; A; U_W)$ via the inverse CDF

$$f_W(x; m_W; a; u_W) = F_W^{-1}(u_W^{(1)}); \quad (27)$$

By denoting $f_W(x; m_W; a; \cdot)$ as $f_{x; m_W; a}$, we can write the observational distribution under M as the push forward

$$f_{x; m_W; a; \#} P^{U_W \mid j; x; a}(w) = P(w \mid j; x; m_W; a) \quad (28)$$

due to Eq.(22) and Eq.(27). Note that we used here that W is not a parent of M_W (see Fig. 1). Hence, $M \in \mathcal{C}(S)$ is compatible with the sensitivity model S . Furthermore, the induced interventional distribution can be written as the push-forward

$$f_{x; m_W; a; \#} P^{U_W \mid j; x; \text{do}(A = a)}(w) = P_+(w \mid j; x; m_W; a) \quad (29)$$

because of Eq. (21). □

B.2 Proof of Corollary 1

Here, we formally restate Corollary 1 for monotone functionals. For two probability densities $P^1(y)$ and $P^0(y)$, we denote $P^1 \geq P^0$ if $F^1 \leq F^0$ holds almost surely for the corresponding CDFs.

Definition 5. A functional D is called monotone if $D(P^1) \geq D(P^0)$ whenever $P^1 \geq P^0$.

Intuitively, a monotone functional increases if applied on a distribution that is further right-shifted. Note that both the expectation functional $E(P) = \int y P(y) dy$ and the quantile functionals $D(P) = F^{-1}(\cdot)$ for $\cdot \in [0, 1]$ are monotone.

Corollary 3 (Restatement) If $M = \cdot$; and D is monotone, we obtain sharp bounds

$$Q^+(x; a; S) = D(P_+^Y(\cdot \mid j; x; a)) \quad \text{and} \quad Q^-(x; a; S) = D(P_-^Y(\cdot \mid j; x; a)); \quad (30)$$

Proof. Follows directly from Theorem 1 for $W = Y$. □

B.3 Proof of Corollary 2

Proof. We derive Algorithm 1 for $Q^+(x; a_{+1}; S)$. The case for $Q^-(x; a_{+1}; S)$ follows analogously. Recall that we want to maximize the causal effect

$$Q(x; a; M) = \prod_m^X D P^Y(j | x; m; do(A = a_{+1})) \prod_{i=1}^Y P(m_i | j; x; m_{i-1}; do(A = a_i)); \quad (31)$$

over all possible SCM $M \in \mathcal{C}(S)$ that are compatible with the GMSM. By using the assumption (no unobserved confounding between mediators and outcome), we can write M in terms of functional assignments defined via $W = f^W(X; M_W; A; U_W)$ and induced (interventional) distributions $P^{U_W | j^X}$ in the following way:

$$Q(x; a_{+1}; M) = \prod_m^X D f_{x; m_{\cdot}; a_{+1} \#}^Y P^{U_Y | j^X}(\cdot) \prod_{i=1}^Y f_{x; m_{i-1}; a_i \#}^{M_i} P^{U_{M_i} | j^X}(m_i); \quad (32)$$

Hence, the optimization problem reduces to maximizing (32) over all functional assignments $f_{x; m_{\cdot}; a_{+1} \#}^W$ and distributions $P^{U_W | j^X}$ that are compatible with S . Note that the terms in the product do not depend on each other or the term in the sum. Thus, by rearranging the suprema and products, we can equivalently perform the following iterative procedure: First, we initialize

$$Q_{+1}^+(x; m_{\cdot}; a_{+1}; S) = \sup_{M2C(S)} D f_{x; m_{\cdot}; a_{+1} \#}^Y P^{U_Y | j^X}(\cdot) \quad (33)$$

and then define

$$Q_i^+(x; m_{i-1}; a_{+1}; S) = \sup_{M2C(S)} \prod_{m_i}^X Q_{i+1}^+(x; m_{i-1}; m_i; a_{+1}) f_{x; m_{i-1}; a_i \#}^{M_i} P^{U_Y | j^X}(m_i) \quad (34)$$

for all $i \in \{2, \dots, 1g\}$, which results in the sharp upper bound

$$Q^+(x; a_{+1}; S) = Q_1^+(x; a_{+1}; S); \quad (35)$$

For Eq. (33) and monotonicity, we can directly apply Theorem 1 and obtain

$$Q_{+1}^+(x; m_{\cdot}; a_{+1}; S) = D P_+^Y(j | x; m_{\cdot}; a_{+1}); \quad (36)$$

For Eq.(34), we need to find an induced distribution $f_{x; m_{i-1}; a_i \#}^{M_i} P^{U_Y | j^X}$ on M_i that is compatible with S and puts most probability mass on where $Q_{i+1}^+(x; m_{i-1}; m_i; a_{+1})$ is large. Hence, we can apply the discrete version of Theorem 1 with $\pi = \pi(M_i)$, where $\pi: \text{supp}(M_i) \rightarrow \text{supp}(M_i)$ is a permutation map so that $Q_{i+1}^+(x; m_{i-1}; \pi(m_i)); a_{+1}$ is ordered in ascending order. The corresponding update step is shown in Algorithm 1.

□

C Special cases of the GSM

In this section, we prove Lemma 1, i.e., we show that all sensitivity models introduced in Sec. 3.3 of the main paper are special cases of our (weighted) GSM. Recall that we consider settings without mediators (i.e., $M = \emptyset$), and write $\gamma = \gamma(a; x)$ for the sensitivity parameter and $q(a; x) = q_\gamma(a; x)$ for the weight function, and $U = U_\gamma$ for the unobserved confounders. In this case, the weighted GSM is defined via the confounding restriction

$$\frac{1}{(1 - \gamma)q(a; x) + \gamma} = \frac{P(U = u | x; a)}{P(U = u | x; do(A = a))} = \frac{1}{(1 - \gamma)q(a; x) + \gamma}; \quad (37)$$

C.1 Marginal sensitivity model (MSM):

The MSM [67] for binary treatment $A = \{0, 1\}$ is defined via

$$\frac{1}{1 - \gamma(x)} = \frac{1}{1 - \gamma(x; u)}; \quad (38)$$

where $\gamma(x) = P(A = 1 | x)$ denotes the observed propensity score and $\gamma(x; u) = P(A = 1 | x; u)$ denotes the full propensity score. By rearranging the terms, we obtain

$$\frac{1}{(1 - \gamma)P(a | x) + \gamma} = \frac{P(a | x; u)}{P(a | x)} = \frac{1}{(1 - \gamma)P(a | x) + \gamma}; \quad (39)$$

for $A = \{0, 1\}$. Furthermore, by Bayes' theorem, it follows that

$$\frac{P(a | x; u)}{P(a | x)} = \frac{P(u | x; a)P(a | x)}{P(u | x)P(a | x)} = \frac{P(u | x; a)}{P(u | x)} = \frac{P(u | x; a)}{P(u | x; do(A = a))}; \quad (40)$$

which implies that the MSM is a weighted GSM with weight function $q(a; x) = P(a | x)$.

Comparison to the MSM definition using potential outcomes: Some papers define the MSM in terms of potential outcomes (Y_1, Y_0) instead of unobserved confounders [31, 67]. Here, Y_a denotes the potential outcome under the treatment intervention $do(A = a)$. In the following, we show that this is equivalent to the definition we use in our paper.

Lemma 2. The following two statements are equivalent:

1. There exists an unobserved confounder U that satisfies $Y_1, Y_0 \perp\!\!\!\perp A | X; U$ so that it holds that $\frac{P(u | x; a)}{P(u | x)} \leq s^+$.
2. It holds that $\frac{P(Y_1, Y_0 | x; a)}{P(Y_1, Y_0 | x)} \leq s^+$.

Proof. The second statement follows directly from the first one by definition $\gamma = \gamma(Y_1, Y_0)$. For the other direction, we proceed via proof by contradiction. Assume there exists a γ that violates the second statement, say w.l.o.g. $\frac{P(Y_1, Y_0 | x; a)}{P(Y_1, Y_0 | x)} > s^+$. We can use the independence condition from the first statement to write

$$P(Y_1, Y_0 | x; a) = \int_Z P(Y_1, Y_0 | x; u; a)P(u | x; a) du = \int_Z P(Y_1, Y_0 | x; u)P(u | x; a) du \quad (41)$$

Furthermore, we have that

$$P(Y_1, Y_0 | x) = \int_Z P(Y_1, Y_0 | x; u)P(u | x) du; \quad (42)$$

It follows that

$$s^+ < \frac{\int_Z P(Y_1, Y_0 | x; u)P(u | x; a) du}{\int_Z P(Y_1, Y_0 | x; u)P(u | x) du}; \quad (43)$$

Hence, there exists a u such that $s^+ < \frac{P(Y_1, Y_0 | x; u)P(u | x; a)}{P(Y_1, Y_0 | x; u)P(u | x)} = \frac{P(u | x; a)}{P(u | x)}$, which is a contradiction to the sensitivity constraint of the first statement. \square

C.2 Continuous marginal sensitivity model (CMSM):

For continuous treatments $A \in \mathbb{R}^d$, the continuous marginal sensitivity model (CMSM) [32] is defined via

$$\frac{1}{P(a; x)} \frac{P(a; x; u)}{P(a; x)} : \quad (44)$$

With the same arguments as in (40), it follows that the CMSM is a weighted GMSM with weight function $q(a; x) = 0$.

C.3 Longitudinal marginal sensitivity model (LMSM):

For longitudinal settings with time-varying observed confounders $X_T = (X_1; \dots; X_T)$, unobserved confounders $U_T = (U_1; \dots; U_T)$, treatments $A = A_T = (A_1; \dots; A_T)$, we define the longitudinal marginal sensitivity model (LMSM) via

$$\frac{1}{\prod_{t=1}^T P(a_t; x_T; a_{t-1})} \prod_{t=1}^T \frac{P(a_t; x_T; u_t; a_{t-1})}{P(a_t; x_T; a_{t-1})} : \quad (45)$$

It holds that

$$P(u_T; x_T; a_T) = \prod_{t=1}^{Q_T} \frac{P(u_t; x_T; a_t)}{\prod_{t=1}^{Q_T} P(u_t; x_T; a_t)} = \prod_{t=1}^T \frac{P(u_t; x_T; a_t)}{P(u_{t-1}; x_T; a_{t-1})} \quad (46)$$

$$\stackrel{(\cdot)}{=} \prod_{t=1}^T \frac{P(a_t; x_T; u_t; a_{t-1})}{P(a_t; x_T; a_{t-1})} \prod_{t=1}^T \frac{P(u_t; x_T; a_{t-1})}{P(u_{t-1}; x_T; a_{t-1})} \quad (47)$$

$$= \prod_{t=1}^T \frac{P(a_t; x_T; u_t; a_{t-1})}{P(a_t; x_T; a_{t-1})} \prod_{t=1}^T P(u_t; x_T; a_{t-1}; u_{t-1}) \quad (48)$$

$$= \prod_{t=1}^T \frac{P(a_t; x_T; u_t; a_{t-1})}{P(a_t; x_T; a_{t-1})} P(u_T; x_T; \text{do}(A_T = a_T)); \quad (49)$$

where (·) follows by applying Bayes' theorem on $P(u_t; x_T; a_t)$. Hence, the LMSM is a weighted GMSM with weight function $q(a_T; x_T) = 0$. Note that we can also define LMSMs with different weight functions via

$$\frac{1}{(1 - \epsilon) q(a_T; x_T) + \epsilon} \prod_{t=1}^T \frac{P(a_t; x_T; u_t; a_{t-1})}{P(a_t; x_T; a_{t-1})} \frac{1}{(1 - \epsilon) q(a_T; x_T) + \epsilon} : \quad (50)$$

D Bounds for average causal effects and differences

Here, we show that we can use our sharp bounds to obtain sharp bounds for causal effect averages and differences. We state the results for the upper bound

$$Q^+(x; a_{+1}; S) = \sup_{M_2 \in \mathcal{C}(S)} Q(x; a_{+1}; M) \quad (51)$$

All definitions and bounds for the lower bound $Q^-(x; a_{+1}; S)$ can be obtained by swapping the signs.

We are interested in the sharp upper bound for the average causal effect

$$Q_{\text{avg}}^+(a_{+1}; S) = \sup_{M_2 \in \mathcal{C}(S)} \int_X Q(x; a_{+1}; M) dx \quad (52)$$

and the sharp upper bound for the causal effect difference

$$Q_{\text{di}}^+(x; a_{+1}^{(1)}; a_{+1}^{(2)}; S) = \sup_{M_2 \in \mathcal{C}(S)} Q(x; a_{+1}^{(1)}; M) - Q(x; a_{+1}^{(2)}; M) \quad (53)$$

Lemma 3. We can compute $Q_{\text{avg}}^+(a_{+1}; S)$ and $Q_{\text{di}}^+(x; a_{+1}^{(1)}; a_{+1}^{(2)}; S)$ from our sharp bounds $Q^+(x; a_{+1}; S)$ and $Q^-(x; a_{+1}; S)$ via

$$Q_{\text{avg}}^+(a_{+1}; S) = \int_X Q^+(x; a_{+1}; S) dx \quad (54)$$

and

$$Q_{\text{di}}^+(x; a_{+1}^{(1)}; a_{+1}^{(2)}; S) = Q^+(x; a_{+1}^{(1)}; S) - Q^-(x; a_{+1}^{(2)}; S) \quad (55)$$

Proof. The result for $Q_{\text{avg}}^+(a_{+1}; S)$ follows directly from interchanging the supremum and integral.

For $Q_{\text{di}}^+(x; a_{+1}^{(1)}; a_{+1}^{(2)}; S)$, we note that

$$Q_{\text{di}}^+(x; a_{+1}^{(1)}; a_{+1}^{(2)}; S) = \sup_{M_1 \in \mathcal{C}(S)} Q(x; a_{+1}^{(1)}; M_1) - \inf_{M_2 \in \mathcal{C}(S)} Q(x; a_{+1}^{(2)}; M_2) \quad (56)$$

$$= Q^+(x; a_{+1}^{(1)}; S) - Q^-(x; a_{+1}^{(2)}; S) \quad (57)$$

To show the equality in Eq(56), we show that, for each pair of $M_1, M_2 \in \mathcal{C}(S)$, we can find an $M \in \mathcal{C}(S)$ such that

$$Q(x; a_{+1}^{(1)}; M_1) - Q(x; a_{+1}^{(2)}; M_2) = Q(x; a_{+1}^{(1)}; M) - Q(x; a_{+1}^{(2)}; M) \quad (58)$$

We can assume w.l.o.g. that $M_1, M_2 \in \mathcal{C}(S)$ induce the same distributions $f_{M_j}^{U_j|X}$ on the confounding space. We denote the functional assignments of M_1 and M_2 as $f_{M_1}^W(x; m_W; a; u_W)$ and $f_{M_2}^W(x; m_W; a; u_W)$. We can now define a functional assignment $f_M^W(x; m_W; a; u_W)$ for M so that

$$f_M^W(;; a^{(1)};) = f_{M_1}^W(;; a^{(1)};) \quad \text{and} \quad f_M^W(;; a^{(2)};) = f_{M_2}^W(;; a^{(2)};); \quad (59)$$

which implies Eq. (58). □

E Compatibility of SCMs with sensitivity models

Here, we provide details regarding our class of SCMs that are compatible with a sensitivity model S from Def. 2.

Definition 6 (Compatibility). Let P^V denote the distribution of the observed variables $(X; A; M; Y)$ and let S be a sensitivity model with unobserved confounders U and a family of joint distributions P . Let M be an SCM with observed variables $V_M = V$ and unobserved variables $U_M = U$ that contain the unobserved confounders from the sensitivity model but also potential exogenous noise. Then, M is compatible with S , if the following three conditions are satisfied:

1. Graph compatibility: the induced causal graph by M on $V \cup U$ must coincide with our assumed causal structure (see Sec. 3.1).
2. Latent unconfoundedness:
$$P(w_j | x; m_w; do(A = a)) = \int P(w_j | x; u_w; m_w; a) P(u_w | x; m_w) du_w$$
 for all $x \in X$, $m_w \in \text{supp}(M_w)$, $a \in A$, and $W \in \{M_1; \dots; M_n; Y\}$.
3. Sensitivity constraints: $P_M^{V \cup U} \in \mathcal{P}$, where $P_M^{V \cup U}$ denotes the distribution on $V \cup U$ induced by M .

In Definition 6, the latent unconfoundedness condition ensures that all interventional densities $P(w_j | x; m_w; do(A = a))$ are point-identified under access to the unobserved confounders. In particular, this ensures that the SCM can not contain additional unobserved confounders (e.g., between mediators and outcome). Similar assumptions are standard in the causal sensitivity literature [18].

F Importance sampling estimators for finite sample bounds

In this section, we derive estimators for the outcome bound $D_{P_+^Y}(j; x; m; a)$ for continuous $Y \in \mathbb{R}$. We assume that we have already obtained an estimator $\hat{P}^Y(j; x; m; a)$ of the observational distribution $P^Y(j; x; m; a)$, and that we are able to sample $(y_i)_{i=1}^k \sim \hat{P}^Y(j; x; m; a)$ (see Appendix G for implementation details). Note that the outcome bound $D_{P_+^Y}(j; x; m; a)$ depends on the shifted distribution $P_+^Y(j; x; m; a)$ and not on the observational distribution $P^Y(j; x; m; a)$. Hence, we use an importance sampling approach to derive our estimators, which we outline in the following for the expectation functional and distributional effects. We denote the CDFs corresponding to $P_+^Y(j; x; m; a)$ and $P^Y(j; x; m; a)$ by $F_{P_+^Y(j; x; m; a)}$ and $F_{P^Y(j; x; m; a)}$, respectively.

Expectation functional: For the expectation functional, we can rewrite the outcome bound as

$$D_{P_+^Y}(j; x; m; a) = E_{Y \sim P_+^Y(j; x; m; a)}[Y] \quad (60)$$

$$= E_{Y \sim P^Y(j; x; m; a)} \left[Y \frac{P_+^Y(Y; j; x; m; a)}{P^Y(Y; j; x; m; a)} \right] \quad (61)$$

$$= E_{Y \sim P^Y(j; x; m; a)} \left[\frac{Y}{s_Y^+} \mathbb{1}_{Y \leq F_{P_+^Y(j; x; m; a)}(c_Y^+)} + \frac{Y}{s_Y} \mathbb{1}_{Y > F_{P_+^Y(j; x; m; a)}(c_Y^+)} \right] \quad (62)$$

to obtain the consistent estimator

$$\hat{D}_{P_+^Y}(j; x; m; a) = \frac{1}{k} \sum_{i=1}^k \frac{X_i^{b k c_Y^+ c}}{s_Y^+} y_i + \frac{1}{k} \sum_{i=b k c_Y^+ c+1}^k \frac{X_i^k}{s_Y} y_i; \quad (63)$$

where $(y_i)_{i=1}^k \sim \hat{P}^Y(j; x; m; a)$ are sampled from the estimated observational distribution. This corresponds to Eq. (8) in the main paper.

Computational complexity: Given trained models and a sample $(y_i)_{i=1}^k \sim \hat{P}^Y(j; x; m; a)$, our estimator from Eq. (63) has a complexity of $\mathcal{O}(k)$ as it only involves summing and quantile computation. To set this in relation to existing work, Algorithm 1 of Jesson et al. [32] has a complexity of $\mathcal{O}(kn)$, where n is a number of grid search points. This is demonstrated in Table 1, where we choose $k = n = 5000$.

Distributional effects: We now derive estimators for distributional effects, i.e., for quantile functionals D of the form

$$D_{P_+^Y}(j; x; m; a) = F_{P_+^Y(j; x; m; a)}^{-1}(\alpha) \quad (64)$$

with $\alpha \in (0; 1)$. We again use an importance sampling approach and rewrite

$$F_{P_+^Y(j; x; m; a)}(y) = E_{Y \sim P_+^Y(j; x; m; a)}[\mathbb{1}(Y \leq y)] \quad (65)$$

$$= E_{Y \sim P^Y(j; x; m; a)} \left[\mathbb{1}(Y \leq y) \frac{P_+^Y(Y; j; x; m; a)}{P^Y(Y; j; x; m; a)} \right] \quad (66)$$

$$= E_{Y \sim P^Y(j; x; m; a)} \left[\frac{\mathbb{1}(Y \leq \min\{y; F_{P_+^Y(j; x; m; a)}(c_Y^+)\})}{s_Y^+} + \frac{\mathbb{1}(F_{P_+^Y(j; x; m; a)}(c_Y^+) < Y \leq y)}{s_Y} \right] \quad (67)$$

Hence, we can sample $(g_i)_{i=1}^k \sim \hat{P}^Y(j; x; m; a)$ and obtain the consistent estimator

$$\hat{D}_{P_+^Y}(j; x; m; a) = \min_{F_{P_+^Y(j; x; m; a)}(g_i)} y_i; \quad (68)$$

where

$$\hat{F}_{P_+^Y(j; x; m; a)}(y) = \frac{1}{k} \sum_{i=1}^k \frac{X_i^{b k c_Y^+ c}}{s_Y^+} \mathbb{1}(y_i \leq y) + \frac{1}{k} \sum_{i=b k c_Y^+ c+1}^k \frac{X_i^k}{s_Y} \mathbb{1}(y_i \leq y); \quad (69)$$

G Implementation and hyperparameter tuning details

All our experimental settings feature a continuous outcome $Y \in \mathbb{R}$ and (optionally) discrete mediators $M_i \in \mathcal{M}_i$, $2 \leq i \leq N$. Hence, we need to estimate the conditional outcome density $p^Y(j; x; m_{\setminus j}; a)$ and conditional probability mass functions $p^{M_i}(j; x; m_{\setminus i}; a)$ in order to estimate our bounds with Eq.(63), Eq.(68), and Algorithm 1 from the main paper.

Conditional outcome density: We use conditional normalizing flows (CNFs) [73] for estimating the conditional density $p^Y(j; x; m_{\setminus j}; a)$. Normalizing flows (NFs) model a distribution P^Y of a target variable Y by transforming a simple base distribution P^U (e.g., standard normal) of a latent variable U through an invertible transformation $Y = f(U)$, where f denotes learnable parameters [58]. In order to estimate the conditional density $p^Y(j; x; m_{\setminus j}; a)$, we leverage CNFs, that is, we define the parameters as an output of a hyper network $\theta = g(x; m_{\setminus j}; a)$ with learnable parameters. Given a sample $\{x_i; m_{\setminus i}; a_i; y_i\}_{i=1}^n$, we learn θ by maximizing the log-likelihood

$$\ell(\theta) = \sum_{i=1}^n \log f_{g(x_i; m_{\setminus i}; a_i)} \# P^U(y_i) \quad (70)$$

$$\stackrel{(\cdot)}{=} \sum_{i=1}^n \log P^U(f_{g(x_i; m_{\setminus i}; a_i)}^{-1}(y_i)) + \log \frac{d}{dy} f_{g(x_i; m_{\setminus i}; a_i)}^{-1}(y_i) \quad ; \quad (71)$$

where $f_{g(x_i; m_{\setminus i}; a_i)} \# P^U(y_i)$ denotes the (push-forward) density induced by $f_{g(x_i; m_{\setminus i}; a_i)}$ on \mathbb{R} and (\cdot) follows from the change-of-variables theorem for invertible transformations.

In our implementation, we use neural spline flows. That is, we model the invertible transformation via a spline flow as described in [21]. We use a feed-forward neural network for the hypernetwork $g(x; m_{\setminus j}; a)$ with 2 hidden layers, ReLU activation functions, and linear output. We set the latent distribution P^U to a standard normal distribution $\mathcal{N}(0, 1)$. For training, we use the Adam optimizer [41].

Conditional probability mass functions: The estimation of the conditional probability mass function $p^{M_i}(j; x; m_{\setminus i}; a)$ is a standard (multi-class) classification problem. We use feed-forward neural networks with 3 hidden layers, ReLU activation functions, and softmax output. For training, we minimize the standard cross-entropy loss by using the Adam optimizer [41]. We use the same approach to estimate the propensity scores $p^S(j; x)$ for discrete treatments S .

Hyperparameter tuning:

We perform hyperparameter tuning for our experiments on synthetic data using grid search on a validation set. The tunable parameters and search ranges are shown in Table 2. For reproducibility purposes, we report the selected hyperparameters and values.³

Table 2: Hyperparameter tuning details.

MODEL	TUNABLE PARAMETERS	SEARCH RANGE
CNFs	Epochs	50
	Batch size	32, 64, 128
	Learning rate	0:0005, 0:001, 0:005
	Hidden layer size (hyper network)	5, 10, 20, 30
	Number of spline bins	2, 4, 8
Feed forward neural networks	Epochs	30
	Batch size	32, 64, 128
	Learning rate	0:0005, 0:001, 0:005
	Hidden layer size	5, 10, 20, 30
	Dropout probability	0, 0:1

³Code is available in the supplementary materials and at <https://github.com/DennisFrauen/SharpCausalSensitivity>.

H Experiments using synthetic data

Here we provide details regarding our experiments using synthetic data. This includes data generation, obtaining oracle sensitivity parameters, and details regarding experimental evaluation.

Overall data-generating process: We first describe the overall data-generating process which we use as a basis to generate data for all settings (i)-(iii) and binary/continuous treatments. We construct an SCM following the causal graph in Fig. 1 (right) from the main paper. We have an observed confounder $X \in \mathbb{R}$, a (binary or continuous) treatment $A \in \{0, 1\}$, two binary mediators M_1 and M_2 , and a continuous outcome $Y \in \mathbb{R}$. Furthermore, we consider three unobserved confounders U_M (i) confounding the A - M_1 relationship, (ii) U_{M_2} confounding the A - M_2 relationship, and (iii) U_Y confounding the A - Y relationship. Our data-generating process is inspired by synthetic experiments from previous works on causal sensitivity analysis [31, 35]. We start the data-generating process by sampling

$$X \sim \text{Uniform}[-1; 1]; \text{ and } U_{M_1}, U_{M_2}, U_Y \stackrel{\text{i.i.d.}}{\sim} \text{Bernoulli}(p = 0.5) \quad (72)$$

Depending on the setting, we either generate binary treatments $A \in \{0, 1\}$ via

$$A \sim \text{Bernoulli}(\text{sigmoid}(3x + \mu_{M_1} u_{M_1} + \mu_{M_2} u_{M_2} + \mu_Y u_Y)) \quad (73)$$

or continuous treatments $A \in \mathbb{R}$ via

$$A \sim \text{Beta}(\cdot; \cdot) \text{ with } \mu = 2 + x + \mu_{M_1}(u_{M_1} - 0.5) + \mu_{M_2}(u_{M_2} - 0.5) + \mu_Y(u_Y - 0.5); \quad (74)$$

where μ_{M_1} , μ_{M_2} , and μ_Y are parameters controlling the strength of unobserved confounding. We then generate the mediators and outcome via functional assignments

$$M_1 = f_{M_1}(X; A; U_{M_1}; \mu_{M_1}); \quad M_2 = f_{M_2}(X; A; M_1; U_{M_2}; \mu_{M_2}) \quad (75)$$

and

$$Y = f_Y(X; A; M_1; M_2; U_Y; \mu_Y); \quad (76)$$

where $U_{M_1}, U_{M_2}, U_Y \sim \mathcal{N}(0, 1)$ are standard normal distributed noise variables. The functional assignments are defined as

$$f_{M_1}(x; a; u_{M_1}; \mu_{M_1}) = 1 - a \sin(x) + (1 - a) \sin(4x) + \mu_{M_1} ((u_{M_1} - 0.5) + \mu_{M_1}) > 0 \quad (77)$$

for M_1 ,

$$f_{M_2}(x; a; m_1; u_{M_2}; \mu_{M_2}) = 1 - a m_1 \sin(x) + (1 - a) m_1 \sin(4x) \quad (78)$$

$$+ a(1 - m_1) \sin(x) - (1 - a)(1 - m_1) \sin(4x) \quad (79)$$

$$+ \mu_{M_2} ((u_{M_2} - 0.5) + \mu_{M_2}) > 0 \quad (80)$$

for M_2 , and

$$f_Y(x; a; m_1; m_2; u_Y; \mu_Y) = a m_1 m_2 \sin(x) + (1 - a) m_1 m_2 \sin(4x) \quad (81)$$

$$+ a m_1 (1 - m_2) \sin(8x) + (1 - a) m_1 (1 - m_2) \sin(x) \quad (82)$$

$$+ a(1 - m_1) m_2 \sin(x) - (1 - a)(1 - m_1) m_2 \sin(4x) \quad (83)$$

$$+ a(1 - m_1)(1 - m_2) \sin(8x) \quad (84)$$

$$+ (1 - a)(1 - m_1)(1 - m_2) \sin(x) \quad (85)$$

$$+ \mu_Y ((u_Y - 0.5) + \mu_Y) \quad (86)$$

for Y , where μ_{M_1} , μ_{M_2} , and μ_Y are parameters that control the noise level.

Settings (i)-(iii): We define the settings (i)-(iii) in Sec. 5 via specific values of the confounding parameters μ_{M_1} , μ_{M_2} , and μ_Y , and the noise parameters μ_{M_1} , μ_{M_2} , and μ_Y (see Table 3). Note that the settings are defined to mimic the causal graphs in Fig. 1 from the main paper. For example, the only unobserved confounder in setting (i) is U_M , which means that we can ignore the mediators and use our data to evaluate our bounds for settings without mediators.

Table 3: Definition of settings (i)-(iii).

	M_1	M_2	Y	M_1	M_2	Y
Setting (i), binary	0	0	1:5	0:2	0:2	2
Setting (i), continuous	0	0	1:5	0:2	0:2	1
Setting (ii)	1:5	0	1:5	1	0:2	1
Setting (iii)	1:5	1:5	1:5	0:2	0:2	1

Obtaining r_{W^*} : We provide details regarding our approach to obtain oracle sensitivity parameters r_{W^*} for all $W \in \{M_1, M_2, Y\}$. By sampling from our previously defined SCM, we can obtain Monte Carlo estimates of the GSM density ratio

$$r(u_W; x; a) = \frac{P(u_W | j; x; a)}{P(u_W | j; x; a)} = \frac{P(a | j; x; u_W)}{P(a | j; x)} \quad (87)$$

for all $u_W \in [0, 1]$, a , and x , where () follows from Bayes' theorem. We then define

$$r_W^+(x; a) = \max_{u_W \in [0, 1]} r(u_W; x; a) \quad \text{and} \quad r_W^-(x; a) = \min_{u_W \in [0, 1]} r(u_W; x; a): \quad (88)$$

For binary treatment settings, we define parameters r_W^+ and r_W^- as the sensitivity parameters that attain the density ratio bounds in the MSM from Eq. (39), i.e.

$$r_W^+(x; a) = \frac{1}{(1 - r_W^+)P(a | j; x) + r_W^+} \quad \text{and} \quad r_W^-(x; a) = \frac{1}{(1 - r_W^-)P(a | j; x) + r_W^-}: \quad (89)$$

For continuous treatment settings, we define r_W^+ and r_W^- as the sensitivity parameters that attain the density ratio bounds in the CMSM from Eq. (44), i.e.

$$r_W^+(x; a) = r_W^+ \quad \text{and} \quad r_W^-(x; a) = \frac{1}{r_W^-}: \quad (90)$$

Finally, we define r_W as the parameter corresponding to the maximum possible violation of unconfoundedness, i.e.,

$$r_W = \max\{r_W^+, r_W^-\} \quad (91)$$

By definition of r_W , our bounds should contain the oracle causal effect whenever we choose sensitivity parameters r_W for all $W \in \{M_1, M_2, Y\}$.

Weighted GSM experiment (Table 1): For our experiment in Table 1, we modify the treatment assignment from Eq. (74) in setting (i) to

$$A = \text{Beta}(\cdot; \cdot) \quad (92)$$

with

$$\cdot = 2 + x + 1(x < 0) (r_{M_1}(u_{M_1} - 0.5) + r_{M_2}(u_{M_2} - 0.5) + r_Y(u_Y - 0.5)): \quad (93)$$

Hence, unobserved confounding only affects individuals with 0. We then compare our bounds under the CMSM with our bounds under a weighted CMSM (Def. 4) with weight function $w(x) = 1(x > 0)$.

We also provide results for the bounds from Jesson et al. [32] under the CMSM. We implemented the grid search algorithm from Jesson et al. [32] and used 5000 samples for the search space. For a fair comparison, we also used 5000 samples for our importance sampling estimators. Note that the method from Jesson et al. [32] requires estimation of both the conditional outcome density $P(Y | j; x; a)$ and the conditional expectation $E(Y | j; x; a)$. For $P(Y | j; x; a)$, we use the same (normalizing flow-based) estimator as for our bounds. For $E(Y | j; x; a)$, we train a separate feed-forward neural network with linear output activation for continuous outcomes. Implementation and hyperparameter tuning are done the same way as described in Appendix G for the feed-forward neural networks.

I Experiment using real-world data

Data: We consider a setting from the COVID-19 pandemic where mobility in Switzerland (captured through telephone movement) was monitored to obtain a leading predictor of case growth. In total, 1.5 billion trips were monitored from 10 February through 26 April 2020. All data are recorded across 26 different states (cantons). For our analysis, we use an aggregated, de-identified, and pre-processed version of the data provided by Persson, Parie, and Feuerriegel [57]. The preprocessed data is publically available at <https://github.com/jopersson/covid19-mobility/blob/main/Data>. The code for our analysis is available at <https://github.com/DennisFrauen/SharpCausalSensitivity>.

We consider a binary treatment A in the form of a stay-at-home order, which bans gatherings with more than 5 people. We encode mobility as a single binary mediator M , which is 1 if the total number of trips on a specific day is larger than the median number of trips during the entire time horizon, and 0 otherwise. Our outcome is the 10-day-ahead case growth. We include the following observed variables as confounders X : the canton code (swiss member state at a subnational level), the canton population, and whether the weekday is a Monday or not. After removing the unrecorded days for each canton (due to spillover effects from other countries) and rows with missing values, we obtain a dataset with $n = 3276$ observations.

Analysis: We perform a causal sensitivity analysis for the natural directed effect (NDE) of the stay-at-home order A on the case growth Y . That is, we are interested in the part of the causal effect of A on Y that is not explained by the path $A \rightarrow M \rightarrow Y$ (i.e., through the change in mobility). The NDE in an SCMM is defined as

$$\text{NDE}(M) = \int Q(x; (a_1 = 0; a_2 = 1); M) - Q(x; (a_1 = 0; a_2 = 0); M) dx \quad (94)$$

Fig. 5 (main paper) shows causal sensitivity analysis for violations of the unconfoundedness between treatment A and mediator M . Hence, we consider a GSM for binary treatments with sensitivity parameters γ_M and $\gamma_Y = 0$. For each M , we estimate our bounds for the expectation functional and the treatment combinations $a = (0; 1)$ and $a = (0; 0)$. We then obtain bounds for the NDE as described in Appendix D.

J Additional experimental results

Here, we provide additional experimental results on synthetic data that extend the results from Sec. 5 in the main paper. We provide (i) results for additional treatment combinations and (ii) results for distributional effects. We follow the same experimental setup described in Sec. 5 (main paper) and Appendix. H.

J.1 Additional treatment combinations

Results for additional treatment combinations are shown in Fig. 6 (binary treatment settings) and Fig. 7 (continuous treatment settings). The results are similar to those in Sec. 5 in the main paper and empirically confirm the validity of our bounds. Hence, our results remain valid independently of the choice of treatment combination.

Figure 6: Results for additional treatments in the binary treatment setting. From left to right is shown: setting (ii) with $a = (0; 1)$, setting (iii) with $a = (0; 1; 0)$, and setting (iii) with $a = (0; 0; 1)$. The top row shows the oracle sensitivity parameter w (depending on x), and the bottom row shows the bounds.

J.2 Distributional effects

We also provide results for distributional effects, that is, we choose the quantile functional $D_{P^+}^Y(j; x; m; a) = F_{P^+}^{Y; j; x; m; a}(\cdot)$. Here, we consider three quantiles with $\tau = 0.7$, $\tau = 0.5$ (median), and $\tau = 0.3$. We use our importance sampling estimator derived in Appendix. F(8) to estimate our bounds. The results are shown in Fig. 8 (binary treatment) and Fig. 9 (continuous treatment) for settings (i)-(iii) from Fig. 1 in the main paper. Again, our bounds cover the underlying oracle effect in regions where the chosen sensitivity parameter w are larger than the oracle sensitivity parameters w . This also confirms empirically the validity of our bounds for distributional effects.

J.3 Semi-synthetic data

We provide additional results for the semi-synthetic IHDP data with unobserved confounding from Jesson et al. [31]. Here, we demonstrate the effectiveness of our bounds for distributional effects for decision-making. We follow Jesson et al. [31] and evaluate our bounds by measuring the error rate of associated deferral policies, which defer a prespecified fraction of test samples with the largest uncertainty to an expert. We mimic a high-stakes decision-making problem where a wrong decision to prescribe treatment is much worse than a wrong decision not to prescribe treatment. To do so, we measure performance with a weighted error rate, that penalizes wrong

Figure 7: Results for additional treatments in the continuous treatment setting. From left to right is shown: setting (i) with $\alpha = 0.9$, setting (ii) with $\alpha = (0.2; 0.4)$, and setting (iii) with $\alpha = (0.4; 0.5; 0.3)$. The top row shows the oracle sensitivity parameter (depending on x), and the bottom row shows the bounds.

treatment decisions twenty times heavier than wrong decisions not to treat. We then compare three different deferral policies: (i) the (standard) policy based on the expectation which treats whenever $E[Y|X = x; A = 1] > E[Y|X = x; A = 0]$, (ii) a more conservative quantile policy that treats whenever $Q_q[Y|X = x; A = 1] > Q_{1-q}[Y|X = x; A = 0]$ with $q = 0.4$ (Q_q denotes the q -quantile of $P(Y|X = x; A = i)$), and (iii) the same quantile policy with $q = 0.2$. Note that our bounds for policy (i) coincide with sharp bounds for binary CATE from the literature [18]. As done in [31], we report means and standard deviations of 400 random initializations of the IHDP data. The results (Fig. 10) show that the conservative quantile policies improve over the expectation policy.

Figure 8: Results for distributional effects in the binary treatment setting using the same treatments as in Fig. 3 (main paper). Settings (i)–(iii) are ordered from left to right. The top row shows the oracle sensitivity parameter ω (depending on α). Rows 2, 3, and 4 show the bounds for the quantiles of the interventional distribution with $\alpha = 0.7$, $\alpha = 0.5$, and $\alpha = 0.3$.

