

Causal Aggregation: Estimation and Inference of Causal Effects by Constraint-Based Data Fusion

Jaime Roquero Gimenez

Department of Statistics

Stanford University

Stanford, CA 94305, USA

ROQUERO@STANFORD.EDU

Dominik Rothenhäusler

Department of Statistics

Stanford University

Stanford, CA 94305, USA

RDOMINIK@STANFORD.EDU

Editor: David Jensen

Abstract

In causal inference, it is common to estimate the causal effect of a single treatment variable on an outcome. However, practitioners may also be interested in the effect of simultaneous interventions on multiple covariates of a fixed target variable. We propose a novel method that allows to estimate the effect of joint interventions using data from different experiments in which only very few variables are manipulated. If there is only little randomized data or no randomized data at all, one can use observational data sets if certain parental sets are known or instrumental variables are available. If the joint causal effect is linear, the proposed method can be used for estimation and inference of joint causal effects, and we characterize conditions for identifiability. In the overidentified case, we indicate how to leverage all the available causal information across multiple data sets to efficiently estimate the causal effects. If the dimension of the covariate vector is large, we may only have a few samples in each data set. Under a sparsity assumption, we derive an estimator of the causal effects in this high-dimensional scenario. In addition, we show how to deal with the case where a lack of experimental constraints prevents direct estimation of the causal effects. When the joint causal effects are non-linear, we characterize conditions under which identifiability holds, and propose a non-linear causal aggregation methodology for experimental data sets similar to the gradient boosting algorithm where in each iteration we combine weak learners trained on different datasets using only unconfounded samples. We demonstrate the effectiveness of the proposed method on simulated and semi-synthetic data.

Keywords: causal inference, structural equation models, data fusion, randomized experiments.

1. Introduction

Causal inference is a centerpiece of scientific research, with applications ranging from the social sciences to biology. Often, the goal in causal inference is to estimate the effect of one single variable on one outcome: randomizing that variable and evaluating its effect on the outcome is one way to do so. Randomization provides a gold standard procedure for identifying causal effects related to that variable, as the intervention removes any spurious

association with the response due to unmeasured factors. Sometimes it is also of interest to estimate the effect of joint interventions on multiple variables on an outcome. Ideally, if all the variables are jointly randomized, one can precisely reconstruct the global causal mechanism, including potential interactions between covariates. In practice, however, we only have access to multiple individual experiments—also called environments—where just a few variables are simultaneously manipulated, which only provide partial information about such mechanism. We formulate a procedure for aggregating the knowledge obtained from several experiments that reconstructs a complex global causal model, capturing the causal effect of multiple covariates on the response as if they were all simultaneously manipulated. Our framework for aggregating causal information also allows the use of Instrumental Variables (IV) and covariate adjustment as building blocks. The approach is motivated by the following: during the last decade, internet companies have massively adopted a new experimental framework to improve their web-based products: the WebLab. Any website has a myriad of design choices built in it that affect the customer behavior, such as the ranking of articles in a newsfeed, the location of an ad within a webpage, etc. These companies have the possibility of running randomized A/B experiments where customers are redirected to slightly different versions of the website. Actionable insights may be obtained by evaluating downstream metrics that reflect the effect of a particular change in the website. Estimating the combined effect of several changes would require simultaneous randomization of many parameters, which may not be feasible in practice as the user experience would vary too much. Therefore it may be of interest to understand how insights from different individual experiments, each manipulating a small number of parameters, may be aggregated into a single causal model. The following examples illustrate the purpose of our method.

Example 1 Assume that variables are related via the structural causal model (Wright, 1921; Bollen, 1989) presented in Figure 1. The disturbance variables ϵ_i are jointly independent




Figure 1: Simple Linear Structural Equation Model: observational setting. Graphical representation of the linear SEM, along with the corresponding set of equations and regression coefficients. The disturbance variables ϵ_i are jointly independent standard Gaussian.

standard Gaussian. We observe samples from the covariates X_1, X_2 and the response Y . H is an unobserved variable that jointly affects X_1, X_2 and Y . We assume we do not know such structure, and we want to identify the vector $\beta^0 = (\beta_1^0, \beta_2^0)$ that defines the linear structural equation of Y with respect to (X_1, X_2) . This vector defines the causal effect on Y of a joint intervention that acts on both X_1 and X_2 (Pearl et al., 2009). Estimating




Figure 2: Experimental environments. Top: environment where X_1 is randomized. Graphically, the dependency on H is removed. Only the structural equation defining X_1 is modified, and the regression coefficients are still biased. Bottom: environment where X_2 is randomized. In this case, regressing on X_2 leads to the direct causal effect. In each case, the disturbance variables ϵ_i are jointly independent standard Gaussian.

β^0 from observational samples alone is not possible. A quick computation shows that the coefficient $\beta_{(X_1, X_2)}^{OLS}$ from ordinary least squares (OLS) regression of Y on $\{X_1, X_2\}$ based on observational data is biased (cf. Figure 1). Individually regressing the response on any of the covariates X_1 or X_2 (that we denote $\beta_{X_1}^{OLS}, \beta_{X_2}^{OLS}$ respectively) also produces biased results. This is expected, as the latent variable H is simultaneously affecting the covariates and the response. β^0 can be estimated by OLS in an environment where we fully randomize all covariates. However, suppose we can at most randomize one variable at a time, leading to two different environments. We represent those environments in Figure 2, along with the regression coefficients. Graphically, randomization removes the effect of the latent variable on the randomized covariate, but unfortunately the regression coefficients $\beta_{(X_1, X_2)}^{OLS}$ are still biased in both environments. Regressing Y on individual covariates does not necessarily lead to the corresponding coordinate of β^0 either, even when those covariates are the ones randomized. The “total causal effect” of X_1 on Y (Pearl et al., 2009), evaluated by regressing Y only on X_1 in the environment where X_1 is randomized, is equal to 2 because it takes into account the effect of X_1 on Y mediated by X_2 . This situation does not occur in the environment where X_2 is randomized, where $\beta_{X_2}^{OLS} = \beta_2^0$. However we assume we do not know the graphical structure of the causal model so we do not know a priori whether the effect of X_1 on Y is mediated by X_2 . Without such knowledge, can we aggregate information from experimentation on individual covariates in different environments? Based on partial experimentation on covariates and other types of causal information, we will discuss how to reconstruct β^0 —the “causal” coefficient vector, as if we were simultaneously randomizing all covariates.

The problem of causal discovery focuses on learning the causal structure, such as identifying the subset of covariates with non-zero coefficients in β^0 . Aggregating “causal information” from different environments requires additional knowledge about the environment, besides the data samples. In the above example, we know which covariates are randomized in each environment. But other situations may provide us with other types of “causal information” that we can aggregate to the experimental data. Instrumental Variable (IV) methods is one such situation where causal information is embedded in the requirements for a variable to be an instrument. More generally, knowledge about the structural causal model can also be aggregated to the previously mentioned methods. In particular, knowing the parental set of a covariate is another type of causal knowledge that we can leverage.

Example 2 *We observe an additional variable I within the linear SEM from the observational environment. Several assumptions are needed for I to be an instrument. Most importantly, I must be exogenous (has no parents in the graph) and does not directly affect the response Y (exclusion restriction). The last condition needed asks that I is relevant, that is, it affects the covariates in the model. However, we may not know in practice which covariates are direct descendants of I . Also, in usual IV regression we need at least as many instruments as there are confounded covariates. Can we leverage and combine incomplete causal knowledge derived from IV methods with experimental data? Our causal aggregation methodology does not require a priori to know which of X_1 or X_2 are affected by I . In particular, our method recovers β^0 by combining the IV based information from the model in Figure (3) and the information from the environment where X_2 is randomized.*




Figure 3: Instrumental variables as causal information. A new observed variable is added to the graph: it corresponds to an instrument that leads to the second orthogonality constraint (2). The noise contribution ϵ_i are jointly independent standard Gaussian.

Leveraging causal constraints arising from environments. Example 1 shows that recovering the causal model that defines Y is a non-trivial task if the ability to intervene on the system is limited to manipulating small subsets of covariates, even in simple linear models. Fortunately, it is possible to extract from each experimental data set a set of constraints that partially identify β^0 . Combining several of such constraints helps us identify β^0 and derive a procedure for constructing a consistent estimator. The intuition is that randomizing a covariate introduces exogenous randomness that is independent from all the other remaining elements in the system. For example, randomizing X_1 in Example 1 removes the confounding effect induced by the latent variable H when regressing Y on X_1 . This external manipulation

of the covariate modifies the structural equation defining it (cf. top Figure 2). Denoting by (Y^e, \mathbf{X}^e) for $e \in \{1, 2\}$ the random variables generated from the above model where X_1^1 (resp. X_2^2) have been randomized, we get the following system of equations in β that is solved by β^0 :

$$\begin{cases} \mathbb{E}[X_1^1(Y^1 - \mathbf{X}^{1,T}\beta)] = 0 \\ \mathbb{E}[X_2^2(Y^2 - \mathbf{X}^{2,T}\beta)] = 0 \end{cases} \quad (1)$$

under the assumption of a linear model for $Y|\mathbf{X}$ parametrized by $\beta = (\beta_1, \beta_2)$, where $\mathbf{X} = (X_1, X_2)$. These constraints reflect the independence between the randomized covariate and the residual term of the regression under the correct parameter value. Each equation imposes a ‘‘causal constraint’’ on the vector β . Given enough such constraints, β^0 is identified and the estimator $\hat{\beta}$ obtained as the solution to the empirical counterpart of the above system of equations is a consistent estimator of β^0 . Causal constraints may originate from other assumptions on the data. If I is an instrument for X_1 in the example above, then the following orthogonality constraint

$$\mathbb{E}[I(Y - \mathbf{X}^T\beta)] = 0 \quad (2)$$

is satisfied whenever $\beta = \beta^0$. Other constraints can be constructed based on additional knowledge of the structural equations, in particular whenever we know the parental set of a given variable in the context of graphical representation of causal models. More generally, the structural equation of Y may include non-linear terms in \mathbf{X} that include interaction terms between the covariates, represented by a function $f^0(\mathbf{X})$. Randomizing different subsets of covariates across different environments also leads to a joint system of equations similar to the system (1) over the several environments: we estimate f^0 by constraining an estimator \hat{f} so that the residuals $Y - \hat{f}(\mathbf{X})$ and the randomized covariates are orthogonal.

Intervening on subsets of X . In this paper, we are interested in quantifying the total causal effect of *joint* interventions on X_1, \dots, X_p , which means that the variables X_1, \dots, X_p are set to a value at the same time. This is different from intervening just on one single variable. For example, in Figure 2, intervening on X_1 changes the distribution of X_2 which then subsequently changes the distribution of Y . If on the other hand one intervenes on both X_1 and X_2 simultaneously, then randomization of X_2 destroys the causal pathway $X_1 \rightarrow X_2 \rightarrow Y$ and thus changes propagate through the system differently as in the case where only X_1 is randomized. If one is interested in the effect of an intervention on a subset $X_S \subseteq \{X_1, \dots, X_p\}$, the methods of this paper still applies since one can simply set $\tilde{X} = X_S$ and apply the methods below for the subset of variables.

Our Contribution. We define a general procedure for aggregating causal information from different environments where we leverage, in addition to the data samples, our knowledge of how the environment e is generated. We present in Section 3 a formal description of environments and how they relate via experimental manipulations. Starting with a linearity assumption in Section 4, we select *constraint-inducing* variables R based on such knowledge that define constraints of the form:

$$0 = \mathbb{E}[R(Y - \mathbf{X}^T\beta^0)]$$

that must be satisfied by the true structural parameter β^0 . In practice, constraints originate from randomized experiments, from the existence of an instrumental variable, but also if we have additional knowledge of the structure of the generative process. For example, if we know the parental set of a variable we can obtain another constraint via regression adjustment. These constraints are obtained with data from different sources that partially share some structure, in particular the structural parameter β^0 . The general idea then is to construct estimators that simultaneously satisfy all the available constraints. We propose simple estimators that are asymptotically unbiased and normally distributed, provided that we have as many constraints as covariates that are not linearly dependent, under mild additional regularity assumptions such as finite moments. We also provide conditions for identifiability of β^0 based on these constraints. Finally, we analyze the case where the linear system is over-determined. Our solution relies on the method of moments (MM) theory, and we show how to optimally weight the constraints to obtain an asymptotically efficient estimator of β^0 . Aggregating additional constraints obtained from new environments always reduces the asymptotic variance of our estimator.

A major challenge arises in the high-dimensional case, where there are not enough constraints or samples to accurately use the main methodology from the low-dimensional linear case. This is important whenever there is a very large number of experiments, but within each data set only a few samples are available. Another use case is when the system is under-determined, in the sense that we have fewer constraints than the number of covariates: multiple values for the estimator can simultaneously satisfy all constraints, and β^0 is not identifiable. Additional assumptions such as sparsity in β^0 can help: we propose in Section 5 a regularized estimator similar to Causal Dantzig (Rothenhäusler et al., 2019) that provably recovers β^0 under additional regularity assumptions. Whenever we have a large number of experiments, and the sample size in each experiment is large compared to the logarithm of the number of covariates, our method leads to a consistent and data-efficient procedure to estimate the causal parameter. In the purely under-determined case, additional constraints are required on the structure of the causal model for our method to work. Since these restrictions are relatively strong, in this high-dimensional case we recommend reducing the number of covariates through a pre-screening step before using regularized causal aggregation.

Building on the linear aggregation framework, we develop in Section 6 a non-linear causal aggregation procedure following a boosting approach to construct estimators from a simple class of base learners. This procedure opens the possibility to learn complex non-linear causal models with potential interactions between covariates. We assume that the noise is additive but not necessarily independent of covariates. Using unconfounded (randomized) covariates, we construct individual base learners within each environment, and then combine these using a linear aggregation step as a sub-routine in the general boosting update. We assume that the randomized covariates are known in each environment. In addition to this flexible procedure, we characterize necessary conditions on the environments to identify the true non-linear causal structural equation of the response.

2. Related Work

Learning a causal structure from multiple data sources has a long history in the literature. Cooper (1997) developed an algorithm for causal discovery based on independence tests. Tian

and Pearl (2001) combine a constraint based approach with background knowledge inferred from analyzing interventional data. Sachs et al. (2005) use a score-based algorithm that searches through the space of directed acyclic graphs. Eaton and Murphy (2007) use Bayesian inference to learn the causal structure from interventions with unknown effects. Eberhardt et al. (2010) describe a method to learn both the graph and the causal relations between a set of variables in presence of confounding based on experiments by first estimating a total effect matrix and then inferring direct effects from it. Hauser and Bühlmann (2012) present a modification of greedy equivalence search to learn the causal structure from interventional data. Similarly, Hyttinen et al. (2012) discuss identifiability and present a search algorithm for learning linear cyclic models in presence of latent variables. Mooij and Heskes (2013) learn cyclic causal models from equilibrium data collected under different experimental and observational contexts. Mooij et al. (2020) propose an approach that allows different types of interventions and can be seen as a unification of several existing discovery methods. In contrast to these methods, we do not aim to reconstruct the causal graph or the overall structure but leverage constraints on causal effects.

Relatively recently, invariance principles have been exploited to estimate causal effects based on several data sets (Peters et al., 2016; Heinze-Deml et al., 2018; Magliacane et al., 2018; Pfister et al., 2019; Rothenhäusler et al., 2019). Our approach is similar in the sense that we can use data from different environments. In some of the work, it is possible to add background knowledge in the form of logical constraints to the optimization procedure. (Hyttinen et al., 2014; Magliacane et al., 2018). In contrast, we allow to incorporate information about parental sets and experimental data not as logical constraints but as gradient information.

Related to our work is do-calculus and the data-fusion framework (Bareinboim and Pearl, 2013a,b; Pearl and Bareinboim, 2014; Bareinboim and Pearl, 2016). In this line of work, the authors present a powerful nonparametric framework to combine data sets to estimate a causal query, given the directed acyclic graph. Our approach is different in the sense that we do not assume that the graph is known but restrict the functional complexity of the structural equations. This distinction can be important in practice, since the graph is unknown in many cases. While the graph can be estimated using structure learning algorithms, such algorithms will usually make some errors, which propagate to the estimation step. In general, it is challenging to provide confidence intervals that take into account both the uncertainty in graph estimation and the uncertainty due to the estimation step. The proposed estimator allows to skip the graph estimation step which results in straightforward uncertainty quantification. In cases where the graph is known or can be estimated with high accuracy, the output of do-calculus and the data fusion framework can often be used in conjunction with the proposed method. This is explored in Section 7.1.3, where we demonstrate that causal constraints from the data fusion framework and other background knowledge can be used to improve precision.

In the potential outcome framework, there has been some recent work to combine data sets for causal inference. Athey et al. (2016) use surrogates to estimate long-term outcomes. Kallus et al. (2018) use limited experimental data to remove the confounding on larger observational data under a linearity assumption. Yang and Ding (2020) consider a setting where the researcher is given a small unconfounded validation data set and a large confounded data set.

In applications, it is common to use two-stage least squares with multiple instruments, see for example Mogstad et al. (2019). We add to this literature by allowing for a wider range of causal constraints not necessarily based on linear models. Non-linear instrumental variable methods (Newey and Powell, 2003; Carrasco et al., 2007; Darolles et al., 2011; Singh et al., 2019) provide a flexible framework for estimating non-linear causal effects based on exogenous instruments. Our non-linear aggregation method instead focuses on combining causal information based on different experimental datasets. A boosting version of non-linear instrumental variables has been recently proposed (Bakhitov and Singh, 2021) which is closer to our work. Again, these methods are devised for observational datasets with instruments, as opposed to considering the more general problem of leveraging different causal identification strategies across multiple data sets.

3. Setting and Notation

We assume that samples originate from different sources that we call environments, each characterized by their own data generating distribution. Let \mathcal{E} denote the set of environments. We collect data from $E = |\mathcal{E}| \geq 2$ environments, where for each $e \in \mathcal{E}$ we have n_e i.i.d. samples $(\mathbf{X}_i^e, Y_i^e)_{1 \leq i \leq n_e}$, and let $n = \sum_e n_e$ the total number of samples. The p -dimensional random vector $\mathbf{X}^e \in \mathbb{R}^p$ (we denote by bold letters multivariate random variables) corresponds to the covariate vector in environment e , and Y^e corresponds to the response variable. When needed, for simplicity we denote the response via a $p + 1$ -th indexed covariate $X_{p+1}^e := Y^e$. The environment e is therefore characterized by $\mathbb{P}^e = \mathcal{D}(\mathbf{X}^e, Y^e)$ where $\mathcal{D}(U)$ denotes the distribution of U . In particular, we assume that we collect across environments the same real-valued covariates indexed by $[p] := \{1, \dots, p\}$, so that $\{\mathbb{P}^e, e \in \mathcal{E}\} \subset \mathcal{P}(\mathbb{R}^p \times \mathbb{R})$ where $\mathcal{P}(\mathcal{X})$ denotes the set of probability distributions over the space \mathcal{X} . Additionally, we assume that there are $p' \geq 1$ unobserved variables $\mathbf{H} = (H_l)_{l \in [p']}$ that jointly affect the covariates and the response. In order to define how the different \mathbb{P}^e are related, we start with an observational base distribution $\mathbb{P}^0 \in \mathcal{P}(\mathbb{R}^p \times \mathbb{R})$. We assume that samples from \mathbb{P}^0 are generated by an acyclic linear Structural Equation Model (SEM) \mathcal{M}^0 with latent variables (Bollen, 1989; Pearl et al., 2009):

$$\mathcal{M}^0 : \begin{cases} X_j \leftarrow \sum_{k \in \text{pa}_0(j)} a_{jk} X_k + \sum_{l \in [p']} c_{jl} H_l + \epsilon_j & \text{for all } 1 \leq j \leq p \\ Y \leftarrow \sum_{k \in [p]} \beta_k^0 X_k + \sum_{l \in [p']} d_l H_l + \epsilon_Y \end{cases} \quad (3)$$

where $(a_{jk})_{1 \leq j \leq p, 1 \leq k \leq p+1} \in \mathbb{R}^{p \times (p+1)}$, $\beta^0 \in \mathbb{R}^p$ and $\text{pa}_0(j) := \{k : a_{jk} \neq 0\} \subset [p + 1]$. For concreteness, the effect of latent variables is assumed to be linear with coefficients $c_{jl} \in \mathbb{R}$, $d_l \in \mathbb{R}$, although our theory actually does not require linear latent effects on the observed variables. We associate to \mathcal{M}^0 a directed graph $G^0 = (V^0, E^0)$ where the set of vertices $V^0 = [p + 1]$ corresponds to the observed variables in the SEM, and a directed edge $(j, k) \in E^0$ iff $k \in \text{pa}_0(j)$. The subset $\text{pa}_0(j)$ corresponds to the *parent nodes* of j in G^0 , which we assume is a Directed Acyclic Graph (DAG). This translates into constraints on the sets of coefficients $\{a_{jk}\}_{jk}, \beta^0$ that define equation 3. We allow the response variable Y to be a parent node of any covariate. The random variables $\epsilon := \{\epsilon_j\}_{j \in [p]} \cup \{\epsilon_Y\}$ —also referred as disturbance terms—are assumed to be centered, and have finite second moments. Furthermore, the disturbance terms are assumed to be independent of \mathbf{H} and jointly independent. For

simplicity, we may also denote $\epsilon_{p+1} := \epsilon_Y$ whenever we refer to sets $\{\epsilon_j\}_{j \in J}$ with $J \subset [p+1]$. We can complete the graph G^0 by adding the vertices $\mathbf{H} = (H_l)_{1 \leq l \leq p'}$ to the set of nodes and define an extended graph $\bar{G}^0 := (\bar{V}^0, \bar{E}^0)$, where $\bar{V}^0 = V^0 \cup \{\mathbf{H}\}$ and \bar{E}^0 contains all edges in E^0 plus edges between a component H_l and nodes of variables in (\mathbf{X}, Y) whenever any latent variable H_l has an effect on that variable. That is, unless we have explicit indication of the contrary, we assume that all covariates are potentially affected by the latent factors (i.e. $c_{jl} \neq 0, d_l \neq 0$). The graph G^0 encodes the input-output relations between observed variables given by the structural equations in \mathcal{M}^0 . The structural model implicitly assumes that latent variables are not affected by observed variables. Figure 4 is an example of a most general \bar{G}^0 under our model. The joint distribution \mathbb{P}^0 over (\mathbf{X}, Y) is properly defined given distributions $\mathcal{D}(\epsilon), \mathcal{D}(\mathbf{H})$ and the structural equations: we reformulate the model \mathcal{M}^0 as a linear system of equations.

$$\mathcal{M}^0 : \begin{pmatrix} \mathbf{X} \\ Y \end{pmatrix} \leftarrow M^0 \begin{pmatrix} \mathbf{X} \\ Y \end{pmatrix} + M_H \mathbf{H} + \epsilon \tag{4}$$

where $M^0 \in \mathbb{R}^{(p+1) \times (p+1)}$ is a matrix that contains the structural parameters a_{jk}, β_k^0 between observed covariates and $M_H \in \mathbb{R}^{(p+1) \times p'}$ is the matrix containing coefficients c_{jl}, d_l . As G^0 is acyclic, the matrix $(I_{p+1} - M^0)$ is invertible. This guarantees that the distribution over the observed variables is well-defined. As discussed in Section 1, estimation of β^0 based only on




Figure 4: Example of a SEM model: The SEM contains latent variables that have a confounding effect on the covariate-response relationship, and the response is potentially a parent of other covariates.

samples from \mathbb{P}^0 —for example, using least squares regression—might be subject to bias due to the presence of the latent variables. We want to impose the weakest possible assumptions on how the distributions \mathbb{P}^e are generated, but that still allow us to recover β^0 , which we assume remains invariant across environments. We now introduce causal models that encode the different types of interventions that lead to different environments. Importantly, in practice we only need to know which covariates receive the interventions. Knowledge of the graph structure of \bar{G}^0 is not needed, although we will show later how such information may be helpful in certain cases. However, we will in general estimate parameters that are related to those variables that are perturbed: having a flexible model for representing perturbations for as many variables as possible is therefore crucial in our framework.

New environments arise when subsets of covariates are manipulated. We use superscripts $e \in \mathcal{E}$ to indicate the elements of the structural equation that are environment-dependent.

For environment e , we index by $\phi(e) \subset [p]$ the subset of randomized covariates, which may be an empty set. We can then write the corresponding structural equations that define the distribution \mathbb{P}^e :

$$\mathcal{M}_{\phi(e)}^e : \begin{cases} X_j^e \leftarrow \sum_{k \in \text{pa}_e(j)} a_{jk}^e X_k^e + \sum_{l \in [p']} c_{jl}^e H_l^e + \epsilon_j^e & \forall j \notin \phi(e) \\ X_j^e \leftarrow \epsilon_j^e & \forall j \in \phi(e) \\ Y^e \leftarrow \sum_{k \in [p]} \beta_k^0 X_k^e + \sum_{l \in [p']} d_l^e H_l^e + \epsilon_Y^e \end{cases} \quad (5)$$

In particular $(p+1) \notin \phi(e)$ for all $e \in \mathcal{E}$, i.e. no interventions allowed on the response variable. We now define a graph $G^e = (V^e, E^e)$ as above that encodes the model \mathcal{M}^e : we have that $V^e = V^0$, and we assume that the edges are such that $\text{pa}_e(j) \subset \text{pa}_0(j)$ for all $j \in [p+1]$, which implies that G^e is also a DAG. The randomization intervention on covariates $\mathbf{X}_{\phi(e)}^e$ implies that the corresponding nodes have no incoming edges. That is, $E^e = E^0 \setminus \{(j, k) \in E^0 \mid k \in \phi(e)\}$. The set of coefficients $\{a_{jk}^e\}_{jk}$ may change provided that no new dependencies between covariates are created in the structural equations. Variables ϵ^e satisfy the same assumptions as in the observational environment: they are centered, have finite second moments, are independent of \mathbf{H}^e and jointly independent. We also allow for arbitrary changes in the coefficients c_{jl}^e and d_l^e that are non-zero, as well as the latent variable distribution $\mathcal{D}(\mathbf{H}^e)$ unless otherwise indicated. The completed graph $\bar{G}^e := (\bar{V}^e, \bar{E}^e)$ still has $\bar{V}^e = \bar{V}^0$, but $\bar{E}^e = E^e \setminus \{(H_l, j)\}_{j \in \phi(e)}$: randomization removes edges connecting the latent confounders with the randomized covariates.

In conclusion, the assumptions on how this intervention mechanism acts on the system can be summarized as constraints on the extended graphs \bar{G}^e . We additionally assume that the different distributions $\{\mathbb{P}^e\}_{e \in \mathcal{E}}$ have the same support.

This model for generating environment distributions \mathbb{P}^e is closely related to “surgical interventions” in Pearl et al. (2009), also called “ideal interventions” (Spirtes et al., 2000) or “structural interventions” (Eberhardt and Scheines, 2007). Such interventions replace those structural equations of the intervened variables $j \in \phi(e)$ by an independently generated ϵ_j^e , leaving the rest unchanged. That is a strong assumption: manipulating variables could lead to “spill-over” effects on other variables whose disturbance distribution is shifted or structural equations perturbed. In contrast, our definition of \mathbb{P}^e allows for changes in the structural equations of the other covariates as well as the joint distribution of the latent variables \mathbf{H} and disturbance variables ϵ : we only impose invariance assumptions on β^0 and on the parental sets $\text{pa}_e(j)$ for all e that are contained in $\text{pa}_0(j)$.

As long as we assume that β^0 is invariant, we can define additional models for generating environments. We introduce in the Appendix A a model where environments are generated by additive shifts in the distribution $\mathcal{D}(\epsilon)$ of the disturbance terms. These are a special case of so-called “parametric” interventions (Eberhardt and Scheines, 2007) or “soft interventions” (Eaton and Murphy, 2007). Randomizing covariates is a direct experimental intervention: the validity of the resulting constraints is thus verified by the scientist’s intervention on the data generating procedure. On the other hand, the effects of soft interventions are sometimes not as easily verifiable, as shifts in latent factors are often not under the direct control of the scientist. Thus we need to trade-off this “weaker” causal knowledge with stronger assumptions. In particular we require the stability of the structural equations (i.e. the coefficients a_{jk}) across environments, as well as the base distributions $\mathcal{D}(\mathbf{H})$, $\mathcal{D}(\epsilon)$. To summarize, we can

aggregate causal information originating from a flexible collection of models that represent environment heterogeneity, where side knowledge on the “causal” mechanism that generates the environment trades-off with our assumption on how stable such mechanism is across environments.

We additionally present in Section 6 a non-linear extension of the above SEM. This increased flexibility restricts the types of causal information that we can aggregate to those arising from a randomized experiment only, as we can no longer leverage instrumental variables or covariate adjustment whenever the parental set of a variable is known. This model not only allows non-linear response structural equations, but can also include interaction terms between covariates. Identifying these effects is possible when an environment simultaneously randomizes the interacting covariates: the causal information contained in environments arising from other types of interventions—such as additive shifts—is harder to aggregate.

We use the following notation throughout the paper. For a vector $\mathbf{u} \in \mathbb{R}^p$, we write for any subset $J \subset [p]$, $\mathbf{u}_J := (u_j)_{j \in J}$. We denote by $|J|$ the cardinality of set J . For any $q \in [1, \infty]$, let $\|\mathbf{u}\|_q$ be the q -norm of \mathbf{u} , and given a positive semi-definite symmetric matrix \mathbf{M} , let $\|\mathbf{u}\|_{\mathbf{M}}^2 := \mathbf{u}^T \mathbf{M} \mathbf{u}$. We denote by $e_j \in \mathbb{R}^p$ the j -th coordinate vector, a one-hot vector where the j -th coordinate is equal to 1.

4. Causal Aggregation in the Linear Case

Assuming a linear structural equation model leads to an intuitive method for aggregating causal information where constraints are derived from each environment and we build an estimator to simultaneously satisfy these constraints. We start by characterizing constraints in Section 4.1 and then formulate our aggregation procedure in Section 4.2.

4.1 Linear Constraints

We can identify the true vector β^0 in our linear SEM model \mathcal{M}^0 by aggregating multiple sources of information about the causal structure. As illustrated in Example 1, randomization of a covariate leads to a linear constraint that should be satisfied by any estimator $\hat{\beta}$ of β^0 . Linear constraints are also obtained if instrumental variable (IV) assumptions hold for some specific variable. Broadly speaking, under the linear SEM model the residual term $Y - \beta^T \mathbf{X}$ is equal to ϵ_Y only for the true value of $\beta = \beta^0$. Based on the assumption at hand, we formulate a linear equality that must be satisfied by the true parameter β^0 , which is generally a consequence of the independence between ϵ_Y and other variables in the model.

4.1.1 INSTRUMENTAL VARIABLES

We obtain an orthogonality constraint whenever we have access to an instrument I . Instrumental Variables (IV) methods (Wright, 1928; Heckman, 1990; Angrist and Imbens, 1995; Angrist et al., 1996) are based on several assumptions that we translate in graphical terms within our linear SEM. For those environments e where I is available, \bar{G}^e has an additional node I . The instrument is *exogenous*, meaning that the node has no parents: there is no directed edge from H nor any other node to I . It is *relevant* and satisfies the *exclusion restriction*: these assumptions correspond to I having only covariate nodes as potential children but not the response Y - and at least one such child. For the true β^0 , the following

constraint holds:

$$0 = \mathbb{E}[I^e(Y^e - \beta^{0,T} \mathbf{X}^e)] \quad (6)$$

The IV constraint is a direct consequence of the independence between the instrument and the response disturbance term. An instrument does not need to be measured in every single environment: the causal information does not rely on comparing different environments unlike other constraints later described. In practice, we use the sample average to build the constraint that must be satisfied by the estimator:

$$0 = \frac{1}{n_e} \sum_{i \in [n_e]} I_i^e Y_i^e - \left(\frac{1}{n_e} \sum_{i \in [n_e]} I_i^e \mathbf{X}_i^e \right) \beta$$

Estimating β^0 via IV methods is usually done by solving a generalized method of moments (MM) problem based on constraints originating from several instruments. A general treatment of IV methods can be found in reference textbooks (Hall, 2005) which provide necessary conditions for identification of β^0 . In particular, there must be at least as many instruments as covariates to build an estimator $\hat{\beta}$ by solving the potentially overidentified system of equations.

4.1.2 EXPERIMENTAL DATA FROM RANDOMIZATION

A more direct method for obtaining an orthogonality constraint is via randomization of a covariate. Following our definition of a model $\mathcal{M}_{\phi(e)}^e$ for environment e where the subset $\phi(e) \subset [p]$ indexes the covariates that are randomized, the following equation holds:

$$0 = \mathbb{E}[X_j^e(Y^e - \beta^{0,T} \mathbf{X}^e)] \quad \forall j \in \phi(e) \quad (7)$$

Again, the equation above is a direct consequence of the independence between the randomized covariate and ϵ_Y^e . Although data from experimental sources is expensive and limited, randomization provides a solid guarantee that the assumption leading to the orthogonality condition for identifying β^0 holds. Again, we use sample averages when building the constraint:

$$0 = \frac{1}{n_e} \sum_{i \in [n_e]} X_{j,i}^e Y_i^e - \left(\frac{1}{n_e} \sum_{i \in [n_e]} X_{j,i}^e \mathbf{X}_i^e \right)^T \beta$$

As indicated in Example 1, full simultaneous randomization of the covariates leads to a linear system of equations with β^0 as the unique solution. Our proposed method allows to individually treat the constraint from each randomized covariate and thus be able to aggregate them across environments.

4.1.3 REGRESSION ADJUSTMENT

Causal inference based on graphical models heavily relies on conditional independence statements between variables that are encoded by the DAG. Pearl's do-calculus identifies causal effects in a causal DAG by transforming conditional statements based on intervened distributions into conditional statements on the observational distribution (Pearl et al., 2009). A fundamental assumption is the knowledge of the DAG, which often times needs to be estimated in practice. Errors in estimating the graph can drastically change the conclusions

on causal effects, which is one motivation for developing this aggregation framework which circumvents estimating \bar{G}^0 . However, we may have partial knowledge of the graph structure, which can be incorporated into our causal aggregation methodology in the form of constraints. If, for a given variable X_j , we assume that we know its parental set in \bar{G}^0 and that the latent confounders are not in such parental set, then based on the “adjustment for directed causes” property (Pearl et al., 2009, Theorem 3.2.2) we have the following constraint for variable j :

$$0 = \mathbb{E}[X_j^e(Y^e - \beta^{0,T} \mathbf{X}^e) | \mathbf{X}_{pa_0(j)}^e] \quad (8)$$

In essence, we use a conditional independence property based on the fact that the distribution \mathbb{P}^e factorizes in \bar{G}^0 , but as we can not condition on the unobserved confounder we need to additionally assume that it is not in the parental set of the node in \bar{G}^0 , which is thus the same as the parental set in G^0 . In practice, we consider the residual variable derived from regressing X_j^e on its parents to derive the linear constraint:

$$0 = \frac{1}{n_e} \sum_{i \in [n_e]} (X_{j,i}^e - \hat{X}_{j,i}^e) Y_i^e - \left(\frac{1}{n_e} \sum_{i \in [n_e]} (X_{j,i}^e - \hat{X}_{j,i}^e) \mathbf{X}_i^e \right)^T \beta,$$

where $\hat{X}_{j,i}^e := \sum_{k \in pa_e(j)} \hat{a}_{jk} \mathbf{X}_{k,i}^e$ and $\hat{\gamma}_j := (\hat{a}_{jk})_{k \in pa_e(j)}$ is the regression coefficient of X_j^e on $\mathbf{X}_{pa(j),i}^e$. Importantly, to simplify asymptotic deductions below, we assume that the estimator $\hat{\gamma}$ is computed on a different data set than the one used for constructing the orthogonality constraint (for notation simplicity we will use γ to refer to the regression coefficient of a covariate on its parental set). In our framework, this is not very stringent, as different environments may provide enough samples to do this. Any two environments where we know that the structural equation of X_j has not changed—this precludes any environment with X_j randomized—can be used for estimating γ and the orthogonality constraint independently (hence the notation without environment subscript).

4.1.4 ADDITIVE INTERVENTIONS ACROSS ENVIRONMENTS

We can derive constraints based on how different environments *relate to each other*. In contrast with all previously mentioned examples of “causal information”, we can derive orthogonality constraints based on the *inner product invariance* (Rothenhäusler et al., 2019) under β^0 for pairs of environments that are generated via additive interventions with respect to the observational base distribution \mathbb{P}^0 . Assume that in environment e covariates X_j for $j \in \psi(e)$ have an additive intervention given by the model $\mathcal{M}_{\psi(e)}^e$ as defined in Appendix A. The distributional shift induced by the additive intervention leaves the expression $\mathbb{E}[X_j^e Y^e - X_j^e \mathbf{X}^{e,T} \beta^0]$ invariant across environments. The intuition behind this approach is that the covariance between X^e and the residuals $Y^e - \mathbf{X}^{e,T} \beta^0$ is a measure of the strength of confounding. Under certain assumptions, the strength of confounding is invariant across settings, which can be leveraged for statistical inference. We thus have the following orthogonality constraint for variable X_j for $j \in \psi(e)$ by merging data from the base distribution and the intervened one:

$$0 = \mathbb{E}[X_j^e Y^e - X_j^0 Y^0 - (X_j^e \mathbf{X}^{e,T} - X_j^0 \mathbf{X}^{0,T}) \beta^0] \quad (9)$$

This orthogonality constraint is crucial for the Causal Dantzig (Rothenhäusler et al., 2019). We leave it as an additional way of constraining the parameter vector that leads to consistent

estimators of β^0 , but for readability it will not be included in finer analysis of the asymptotic behavior. We summarize in the following proposition the set of orthogonality constraints that we derived in the previous sections.

Proposition 1 *The causal vector β^0 satisfies the linear constraints as defined above via either instrumental variables in eq. (6), randomization in eq. (7), regression adjustment in eq. (8), or inner product invariance in eq. (9). In all these cases the constraint in β can be summarized*

$$\mathbf{g}^T \beta = z \tag{10}$$

where $\mathbf{g} \in \mathbb{R}^p$, $z \in \mathbb{R}$ are some specific transformations of population-level cross-covariances obtained at each environment.

As we see, different orthogonality constraints are derived from different types of prior causal information. We can potentially have more constraints than strictly necessary to estimate β^0 , and we may be willing to discard those that rest on weaker foundations. In practice, the assumptions leading to constraints are not on an equal footing. Data obtained via covariate manipulation and the subsequent orthogonality constraint has a better standing than a constraint generated by assumptions on additive shifts in the covariance structure, which are less verifiable in practice. Conditioning on the parental set assumes an accurate knowledge of the (potentially estimated) graph, whereas there may be settings where one can be confident that the exogeneity and exclusion restriction assumptions in IV hold.

4.2 Aggregating Linear Constraints under Just-Identification

One can recover β^0 by exclusively using one type of the orthogonality constraints previously defined. For example, if all the covariates in the model are randomized, then OLS is unbiased. Whenever there are as many instruments as covariates (and a full-rank condition holds) then usual IV methods apply. Finally, Causal Dantzig (Rothenhäusler et al., 2019) provably recovers β^0 whenever for every covariate there is an environment where the given covariate has an additive intervention. The objective is now to combine constraints arising from multiple data sets into one single estimator. Let \mathcal{C} be the set of constraints, denoted by $\mathbf{g}^{c,T} \beta = z^c$ for each $c \in \mathcal{C}$ as in equation (10). We aggregate these via a linear matrix equality by first defining the vector \mathbf{Z} and matrix \mathbf{G} as follows:

$$\mathbf{Z} := (z^c)_{c \in \mathcal{C}} \in \mathbb{R}^{|\mathcal{C}|} \quad \text{and} \quad \mathbf{G} := (\mathbf{g}^{c,T})_{c \in \mathcal{C}} \in \mathbb{R}^{|\mathcal{C}| \times p}$$

The vector β^0 is then a solution to the linear system:

$$\mathbf{Z} = \mathbf{G}\beta \tag{11}$$

Without any prior assumptions on the SEM \mathcal{M}^0 , a necessary condition for identifying β^0 is to have at least as many constraints as covariates, i.e. $|\mathcal{C}| \geq p$, otherwise the system (11) has multiple solutions. A sufficient condition so that β^0 is the unique solution to equation (11) above can be stated in purely mathematical terms. This is similar to the identifiability result in Hyttinen et al. (2012) which we extend to any type of constraint, not only those based on randomization, although we constrain G^0 to be a DAG. With additional information about the constraints, we can formulate more practical necessary conditions. Consider instrumental

variables in equation (6), experimental data through randomization from an interventional environment e that follows the causal model $\mathcal{M}_{\phi(e)}^e$ in equation (7), or constraints from regression adjustment via the population regression vector of a covariate on its parental set in eq. (8). All these constraints feature a known specific random variable denoted R^c that captures the prior knowledge about the causal structure of the data, which we refer to as the *constraint-inducing* variable. We have respectively that $R^c = X_j^{ec}$ for a randomized covariate $j \in \phi(e_c)$, or $R^c = I^{ec}$ for an instrument I^{ec} whenever available, and R^c is the residual term of regressing a covariate on its known parental set (using a different data set for estimating the regression adjustment and for estimating the orthogonality constraint). Additionally, each constraint-inducing variable relates to a covariate: either the randomized covariate itself, any covariate that is correlated to the instrument, or the covariate that is regressed on its parental set. A sufficient condition for identifiability of β^0 is then that each covariate has a distinct constraint related to it. This is again similar to Hyttinen et al. (2012) where a necessary condition is to have an environment where each variable is randomized, and Y is never intervened on.

Proposition 2 β^0 is identified if there is a subset of p linearly independent constraints in \mathbf{G} in (11). If we only consider constraints derived from IV, randomization or regression adjustment, then it suffices to have at least one distinct constraint related to each covariate.

Based on the empirical counterparts $\hat{\mathbf{Z}}, \hat{\mathbf{G}}$ of \mathbf{Z}, \mathbf{G} , we look for estimators $\hat{\beta}$ of β^0 that make the two sides of the equation (11) as close as possible:

$$\hat{\mathbf{Z}} \approx \hat{\mathbf{G}} \hat{\beta} \quad (12)$$

This naïve approach only works in very limited situations, and we describe here the just-identified case. Whenever $|\mathcal{C}| = p$ and the corresponding square matrix \mathbf{G} is invertible, β^0 is identified. Additionally, if $\hat{\mathbf{G}}$ is invertible, then the unique solution to equation (12) is given by the estimator

$$\hat{\beta} := \hat{\mathbf{G}}^{-1} \hat{\mathbf{Z}} \quad (13)$$

which is consistent if $\hat{\mathbf{G}} \rightarrow \mathbf{G}$ and $\hat{\mathbf{Z}} \rightarrow \mathbf{Z}$ by continuity of the matrix inverse and product. Consistency of $\hat{\mathbf{G}}, \hat{\mathbf{Z}}$ holds by the law of large numbers as soon as $n_e \rightarrow \infty$ in every environment. In the following we analyze the asymptotic behavior of this estimator. We assume that sample sizes grow at the same rate, i.e. $n_e/n \rightarrow \rho_e$ for some $\rho_e \in (0, 1)$. We discard constraints based on inner-product invariance for readability. Denote by $e_c \in \mathcal{E}$ the environment where constraint c is generated, we can write:

$$\mathbf{Z} := (\mathbb{E}[R^c Y^{e_c}])_{c \in \mathcal{C}} \in \mathbb{R}^{|\mathcal{C}|} \quad \text{and} \quad \mathbf{G} := (\mathbb{E}[R^c \mathbf{X}^{e_c, T}])_{c \in \mathcal{C}} \in \mathbb{R}^{|\mathcal{C}| \times p}$$

and the corresponding empirical counterparts

$$\begin{aligned} \hat{\mathbf{Z}} &:= ((1/n_{e_c}) \sum_{i \in [n_{e_c}]} R_i^c Y_i^{e_c})_{c \in \mathcal{C}} \in \mathbb{R}^{|\mathcal{C}|} \\ \hat{\mathbf{G}} &:= ((1/n_{e_c}) \sum_{i \in [n_{e_c}]} R_i^c \mathbf{X}_i^{e_c, T})_{c \in \mathcal{C}} \in \mathbb{R}^{|\mathcal{C}| \times p} \end{aligned} \quad (14)$$

Standard regularity conditions on the moments of the variables R, \mathbf{X}, Y also lead to asymptotically valid confidence intervals for $\hat{\beta}$. Whenever a constraint is obtained by adjusting for the parental set we need that the estimates of γ are obtained from a data set independent that the one used for constructing the constraints. Under these assumptions, we show that the estimator is asymptotically normally distributed:

$$\sqrt{n}(\hat{\beta} - \beta^0) \xrightarrow[n \rightarrow +\infty]{d} \mathcal{N}(\mathbf{0}; \Sigma)$$

where Σ is a positive definite matrix that we can consistently estimate by some $\hat{\Sigma}$. Therefore we can form asymptotically valid confidence intervals for β_j^0 by

$$I_j = [\hat{\beta}_j \pm q_{\mathcal{N}}^{1-\alpha/2} \sqrt{\hat{\Sigma}_{jj}}] \quad (15)$$

where $q_{\mathcal{N}}^\alpha$ is the α quantile of the standard Gaussian distribution, which has exact asymptotic coverage.

$$\mathbb{P}(\beta_j^0 \in I_j) \xrightarrow{+\infty} 1 - \alpha$$

Let $\sigma_e^2 := \text{Var}(\epsilon_Y^e)$. We summarize these statements and give an expression for the asymptotic covariance in the following proposition.

Proposition 3 *Consider the setting described above. Assume that for all c , all variables $R^c, \mathbf{X}^{e_c}, Y^{e_c}$ have finite fourth moments. We then have that*

$$\sqrt{n}(\hat{\beta} - \beta^0) \xrightarrow[n \rightarrow +\infty]{d} \mathcal{N}(\mathbf{0}; \Sigma) \quad (16)$$

where $\Sigma = \mathbf{G}^{-1} \text{Diag}_{c \in [p]} \left(\frac{\sigma_{e_c}^2}{\rho_{e_c}} \text{Var}(R^c) \right) \mathbf{G}^{-1,T}$. A consistent estimator $\hat{\Sigma}$ is obtained via

$$\hat{\Sigma} = \hat{\mathbf{G}}^{-1} \text{Diag}_{c \in [p]} \left(\frac{n \hat{\sigma}_{e_c,2}^{e_c}}{n_{e_c}} \widehat{\text{Var}}(R^c) \right) \hat{\mathbf{G}}^{-1,T}$$

where $\hat{\sigma}_{e_c,2}^{e_c} = \frac{1}{n_{e_c}} \sum_{i \in [n_{e_c}]} (Y_i^{e_c} - \hat{\beta}^T \mathbf{X}_i^{e_c})^2$.

This is a first asymptotic result for analyzing the convergence of an aggregation estimator to the true parameter β^0 based on multiple environments. For completeness, we now develop linear aggregation extensions to the over-identified case. Provided the model is correctly specified, additional constraints improve the efficiency of the estimator that we propose, based on a method of moments framework. In practice, some constraints potentially arise in several environments, but only one is kept in the just-identified case—if we assume we know the parental set of a covariate that is invariant across environments, then we have access to several regression adjustment constraints.

4.3 Linear Aggregation in the Over-Identified Case with More Constraints than Covariates

The number of available constraints may surpass the dimension of the covariate vector. Provided that the model is well specified in the sense that the constraints are compatible with

β^0 as solution, additional constraints improve the efficiency of the estimator by decreasing the asymptotic variance. In the ideal setting where in one environment all covariates are randomized, it may be tempting to discard data from other environments as direct least squares regression generates an unbiased estimator relying only on that environment. Other methods can also provide consistent estimators of the causal effects based on data from a single environment. The do-calculus is a method that, provided the underlying causal graph is properly estimated, is capable of assessing whether in a given environment the causal effects are identifiable. In such case, if identifiability holds, then the causal vector β^0 can be estimated via a regression over a judiciously chosen set of covariates. We show that, in any of these scenarios, we can still benefit from aggregating additional constraints from other environments into a single estimator. In this overidentified setting, $\hat{\mathbf{G}}$ is no longer a potentially invertible square matrix, so $\hat{\beta}$ can not be simply derived by inverting a matrix. We instead estimate it via the method of moments estimator (MM), we refer to Hall (2005) for a general treatment of the subject. The standard MM conditions hold in our setting, where β^0 is identified if \mathbf{G} is of rank p , and the estimators that we propose are consistent and asymptotically normal. The challenge is to construct one with minimal asymptotic variance, which requires adapting the usual MM framework to the multiple environment setting. In particular, to obtain the efficient two-step efficient estimator (Hansen, 1982) we need to compute the covariance of the vector of constraints based on samples from different environments. Finally we show that incorporating additional constraints always lead to an improvement of the asymptotic variance.

We reformulate the multiple environment framework (cf. equation 5) to represent individual observations across environments as elements of a same space, sampled from a unique common distribution. We use an environment-indicator variable E and then aggregate samples across all environments, the following collection of variables represents an individual observation:

$$(E_i, Y_i^{E_i}, \mathbf{X}_i^{E_i}, \{\tilde{R}_i^c\}_{c \in \mathcal{C}})$$

The i -th observation originates from environment E_i , where $E_i \sim \text{Multinomial}((\rho_e)_{e \in \mathcal{E}})$. Conditionally on the environment E_i , the values of $Y_i^{E_i}, \mathbf{X}_i^{E_i}$ correspond to the actual observations obtained in environment E_i . Additionally, each constraint in $c \in \mathcal{C}$ is obtained in a specific environment e_c . The aggregated i -th observation concatenates all $\tilde{R}_i^c = R_i^c \mathbb{1}_{E_i=e_c}$ for each $c \in \mathcal{C}$, which are equal to 0 whenever the constraint c is not based on environment E_i . We again exclude the constraints arising from shift interventions. Samples across environments are thus merged into one data set such that samples can be considered as i.i.d. for $i \in [n]$. The c -th moment condition now becomes

$$0 = \mathbb{E}[\tilde{R}_i^c(Y_i^{E_i} - \mathbf{X}_i^{E_i}\beta)] = \mathbb{E}[R_i^c(Y_i^{E_i} - \mathbf{X}_i^{E_i}\beta)\mathbb{1}_{E_i=e_c}] \quad (17)$$

and therefore whenever E_i is the environment where R_i^c is obtained, we get that the above equality holds for β^0 by construction of our orthogonality constraints. In particular, after stacking the $|\mathcal{C}| = q$ constraints (for simplicity $\mathcal{C} = \{c_1, \dots, c_q\}$), we can reformulate the

system of equations in terms of \mathbf{Z} and \mathbf{G} as defined in the previous section:

$$0 = \mathbb{E} \begin{bmatrix} R_i^{c_1} (Y_i^{E_i} - \mathbf{X}_i^{E_i} \beta) \mathbb{1}_{E_i=e_{c_1}} \\ \vdots \\ R_i^{c_q} (Y_i^{E_i} - \mathbf{X}_i^{E_i} \beta) \mathbb{1}_{E_i=e_{c_q}} \end{bmatrix} = \mathbf{D}(\mathbf{Z} - \mathbf{G}\beta)$$

where $\mathbf{D} = \text{Diag}((\rho_{e_c})_{c \in \mathcal{C}})$. This shows that if \mathbf{G} has rank p then the solution to the system above is unique, equal to β^0 . Given a positive definite weighting matrix $\mathbf{W} \in \mathbb{R}^{q \times q}$, the method of moments estimator of β^0 is the minimizer of the following loss with a closed-form solution:

$$\hat{\beta}^{MM}(\mathbf{W}) := \arg \min_{\beta \in \Theta} \|\hat{\mathbf{D}}(\hat{\mathbf{Z}} - \hat{\mathbf{G}}\beta)\|_{\mathbf{W}}^2 \quad (18)$$

for the empirical values of $\hat{\mathbf{Z}}$, $\hat{\mathbf{G}}$ and $\hat{\mathbf{D}} := \text{Diag}((\frac{n_{e_c}}{n})_{c \in \mathcal{C}})$. Under standard regularity assumptions on the moments of the variables, provided that \mathbf{G} is of rank p , the estimator $\hat{\beta}^{MM}$ is consistent for any choice of weighting matrix. The choice of \mathbf{W} characterizes the asymptotic covariance of the estimator. Hansen (1982) proves that, to obtain the estimator with minimal asymptotic variance among all choices of weighting matrices, the optimal choice is given by $\mathbf{W} = \mathbf{S}^{-1}$, where

$$\mathbf{S} := \text{Cov} \left((\tilde{R}^c (Y^E - \mathbf{X}^E \beta^0))_{c \in \mathcal{C}} \right) \in \mathbb{R}^{q \times q} \quad (19)$$

Given that the residual term and the constraint-inducing variable R^c are independent and that any two constraint-inducing variables are pairwise independent, we get that the constraints are uncorrelated and thus obtain the following simplification of \mathbf{S} into a diagonal matrix:

$$\mathbf{S} = \text{Diag} \left(\rho_{e_c} \text{Var}(R^c) \sigma_{e_c}^2 \right)_{c \in \mathcal{C}} \quad (20)$$

\mathbf{S} can only be estimated from data, but in order to do so one needs to estimate the residual variances. These, in turn, are given by $\sigma_{e_c}^2 = \text{Var}(Y^{e_c} - \mathbf{X}^{e_c} \beta^0)$ which depend on the true unknown parameter β^0 . Hansen (1982) proposes a two-step estimator where an inefficient, consistent MM estimator $\hat{\beta}^{MM}(I_q)$ is obtained by setting $\mathbf{W} = I_q$ in equation (18). $\hat{\mathbf{S}}$ is then derived based on $\hat{\beta}^{MM}(I_q)$, which then is used to construct a new weighting matrix for the final efficient MM estimator:

- Compute a first consistent estimator $\tilde{\beta} := \hat{\beta}^{MM}(I_q)$.
- Compute a consistent estimator $\hat{\mathbf{S}} := \text{Diag} \left(\rho_{e_c} \widehat{\text{Var}}(R^c) \hat{\sigma}^{e_c,2} \right)_{c \in \mathcal{C}}$ of the weighting matrix, where $\hat{\sigma}^{e_c,2} = \frac{1}{n_{e_c}} \sum_{i \in [n_{e_c}]} (Y_i^{e_c} - \tilde{\beta}^T \mathbf{X}_i^{e_c})^2$.
- Return the two-step estimator $\hat{\beta}^{MM*} = \hat{\beta}^{MM}(\hat{\mathbf{S}}^{-1})$

If the data are i.i.d. within each environment with finite second moments, and $\tilde{\beta}$ is consistent, then $\hat{\sigma}^{e_c,2}$ and $\widehat{\text{Var}}(R^c)$ are consistent. Thus, $\hat{\mathbf{S}}$ is consistent for \mathbf{S} and $\hat{\beta}^{MM*}$ has optimal asymptotic variance. We summarize the above statements in the following proposition:

Proposition 4 (Hall (2005) Chapter 2) *Assume that β^0 satisfies the constraints defined by equation (11), that \mathbf{G} is rank p , and that data samples are i.i.d. within each environment, with finite second moments so that any MM estimator is consistent. Additionally, assume that the plug-in estimator $\hat{\mathbf{S}}$ of \mathbf{S} is consistent. The two-step MM estimator $\hat{\beta}^{MM*}$ based on our set of orthogonality constraints satisfies the following asymptotic limit*

$$\sqrt{n}(\hat{\beta}^{MM*} - \beta^0) \rightarrow \mathcal{N}(\mathbf{0}; \Sigma)$$

where convergence is in distribution, and where we get the optimal asymptotic covariance among all choices of weighting matrices: $\Sigma = (\mathbf{G}^T \text{Diag}(\frac{\rho_{ec}}{\text{Var}(R^c)\sigma_{\epsilon_c}^2})_{c \in \mathcal{C}} \mathbf{G})^{-1}$

We recover here the asymptotic covariance from Proposition 3 whenever \mathbf{G} is a square, invertible matrix. Also, this result automatically shows that adding constraints can not hurt the performance of the estimator. Given any subset of p constraints, we can construct a MM estimator of the form (13) by setting to 0 the appropriate elements of \mathbf{W} in (18). Therefore the asymptotic efficiency of $\hat{\beta}^{MM*}$ implies that the new constraints, properly weighted by \mathbf{S} , improve over the just-identified case. Incorporating all such information leads to more efficient estimators: all causal information helps. Conversely, in this over-identified setting one can wonder whether certain inconsistent constraints may be hurting the estimator performance: we refer to Hall (2005) for further discussions on tests to detect such issues.

The following two sections address the high-dimensional and non-linear cases respectively, building upon the notation and theory presented thus far. We first present in the following section the high-dimensional case where additional assumptions are required to identify β^0 , and we propose an estimator that favors sparse solutions $\hat{\beta}$ and converges to β^0 . Within this framework, under additional assumptions, we can recover β^0 even in the under-identified case. Unfortunately these are not always satisfied, in which case a pre-screening step may be used to bypass this issue. The linear aggregation procedure in the just-identified case becomes a central sub-routine in the non-linear causal aggregation framework, which is addressed in the subsequent section. The underlying intuition is the same: based on the available additional knowledge about how environments are generated, we build estimators by enforcing constraints that represent the orthogonality between an unconfounded variable and the residuals.

5. High-Dimensional Linear Aggregation under Sparsity Assumptions

In high-dimensional settings we may not have enough constraints to construct an estimator via the methods presented above. Even with enough samples per environment, β^0 may not be identifiable based only on the orthogonality constraints. Conversely, we may have access to a large number of environments, enough for identifying β^0 at the population level, but containing very few samples in each. In the following section, we develop an estimator for such settings. Under additional assumptions on the structure of β^0 , it may be possible to aggregate such causal information and obtain a reasonable estimator of β^0 . Regularization methods based on leveraging the geometry induced by the ℓ_1 -norm are helpful to overcome this issue under the assumption that the actual vector β^0 is sparse. Among these very popular techniques we mention the Lasso penalty for linear regression (Tibshirani, 1996) and basis pursuit (Chen et al., 2001). These methods were subsequently adapted to address multiple

other problems based on sparsity assumptions (eg. precision matrix estimation in Friedman et al. (2008)), deal with additional structure in the regressors (eg. when the order of the covariates matters as in the fused lasso in Tibshirani et al. (2005), where differences between consecutive coefficients are penalized) or expanded with additional penalties (eg. Zou and Hastie (2005) combine Lasso and ℓ_2 -norm penalties). Whenever we have experimental data for every covariate—even if very few samples per experiment—, we derive in Section 5.1 an estimator based on these regularization techniques by directly solving a constrained risk minimization problem. We discuss this in more detail after Proposition 6. Alternatively, in Section 5.2 we propose running a two-step procedure where a subset of covariates is first selected based on the Lasso regression, and then the proposed aggregation procedure is run on the subset of covariates.

5.1 Estimation by Constrained Optimization

We present an estimator that mirrors the formulation of the Dantzig Selector (Candes et al., 2007) designed for the problem of high-dimensional linear regression $y = \mathbf{x}^T \beta^0 + \epsilon$, where the dimension of \mathbf{x} is larger than the number of available observations. The Dantzig Selector is the solution to the following problem:

$$\text{minimize } \|\beta\|_1 \quad \text{subject to } \frac{1}{n} \|\mathbf{X}^T(Y - \mathbf{X}\beta)\|_\infty \leq \lambda$$

where $Y = (y_1, \dots, y_n)$, and $\mathbf{X} = (x_1, \dots, x_n)^T$ is the design matrix. Candes et al. (2007) derived a probabilistic upper bound for the ℓ_2 error of the estimator under some conditions on \mathbf{X} . Additional work by Bickel et al. (2009), Ye and Zhang (2010), among others lead to sharper bounds. We now present our ℓ_1 -norm minimization based causal aggregation estimator, and we then follow Ye and Zhang (2010) to derive upper bounds on the ℓ_q loss. We follow the exposition in Rothenhäusler et al. (2019) and adapt it to our setting, leaving all proofs to the Appendix. Our alternative formulation promotes sparsity by solving an ℓ_1 -norm minimization problem. We denote it by $\hat{\beta}(\lambda)$ where $\lambda > 0$ is a hyper-parameter, and the definition mirrors the above minimization problem:

$$\text{minimize } \|\beta\|_1 \quad \text{subject to } \|\hat{\mathbf{Z}} - \hat{\mathbf{G}}\beta\|_\infty \leq \lambda \quad (21)$$

where $\hat{\mathbf{Z}}$ and $\hat{\mathbf{G}}$ are defined as in equation (14). This is a convex problem, in particular a linear programming problem: Let B^λ be the feasible set to (21) and define Γ^λ as the feasible set to the linear programming problem

$$\begin{aligned} & \text{minimize } \mathbf{c}^T \gamma \\ & \text{subject to } \mathbf{A}\gamma \leq \mathbf{b} \text{ and } \gamma \geq 0, \text{ where} \\ & \mathbf{A} = \begin{pmatrix} -\hat{\mathbf{G}} & \hat{\mathbf{G}} \\ \hat{\mathbf{G}} & -\hat{\mathbf{G}} \end{pmatrix}; \mathbf{b} = \begin{pmatrix} -\hat{\mathbf{Z}} \\ \hat{\mathbf{Z}} \end{pmatrix} + \begin{pmatrix} \lambda \\ \lambda \\ \vdots \\ \lambda \end{pmatrix}; \text{ and } \mathbf{c} = \begin{pmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{pmatrix}. \end{aligned}$$

Then, as shown in Lemma 5 in Rothenhäusler et al. (2019), $B^\lambda = \{\gamma_{1:p} - \gamma_{(p+1):2p} : \gamma \in \Gamma^\lambda\}$. Assimilating $\gamma_{1:p}$ to the vector β_+ , the positive part of the β coefficients –analogously for $\gamma_{p+1:2p}$ and β_- , we get the equivalence.

Higher values of λ relax the constraint based on the causal orthogonality constraints, leading to solutions $\hat{\beta}(\lambda)$ with smaller ℓ_1 norm. Because of the geometry of the ℓ_1 -norm unit ball, sparser solutions are favored. We can prove bounds on the ℓ_q loss of our estimator with high probability. To this end, we will impose assumptions on $\hat{\mathbf{G}}$. Following Ye and Zhang (2010), we define the cone invertibility factor (CIF) as follows: for $1 \leq q \leq +\infty$, $J \subset [p]$, and $J^c := [p] \setminus J$, define

$$\text{CIF}_q(J, \mathbf{M}) := \inf \left\{ \frac{|J|^{1/q} \|\mathbf{M}\mathbf{u}\|_\infty}{\|\mathbf{u}\|_q}; \mathbf{u} \in \mathbb{R}^p \setminus \{0\}; \|\mathbf{u}_{J^c}\|_1 \leq \|\mathbf{u}_J\|_1 \right\}$$

where $\mathbf{M} \in \mathbb{R}^{\tilde{p} \times p}$ is not necessarily a square matrix, and we set by convention $|J|^{1/\infty} = 1$. Ye and Zhang (2010) show how this quantity plays a similar role and relates to the sparse eigenvalue condition, where generally the matrix \mathbf{M} is an estimate of the covariance matrix of \mathbf{X} . Intuitively, estimation in high dimensions is difficult as $\hat{\mathbf{G}}$ is not invertible. However, under a sparsity assumption on the coefficients, one only needs $\hat{\mathbf{G}}$ to be invertible on the set of sparse vectors. The cone invertibility factor (CIF) captures whether $\hat{\mathbf{G}}$ is non-invertible on the set of sparse vectors. The proof technique then proceeds by showing that the cone invertibility factor for $\hat{\mathbf{G}}$ is close to the cone invertibility factor for the population matrix \mathbf{G} , which is assumed to be invertible. As we will show, estimation of β^0 is possible in this setting and we can control the ℓ_q error of the estimator $\hat{\beta}(\lambda)$. This derives from an upper bound for $\|\hat{\beta}(\lambda) - \beta^0\|_q$ by a ratio with the CIF in the denominator as a critical quantity that must be positive.

One may ask whether under a sparsity assumption it is actually necessary to have one constraint per covariate, or whether it is sufficient to have far fewer constraints than covariates. The following example gives a negative answer to this question. Intuitively speaking, sparsity assumptions allows us to get away with few observations per constraint, but we still need at least as many constraints as covariates. This issue might be mitigated by pre-screening, which will be discussed further below. More specifically, Figure 5 provides an example where ℓ_1 -regularization leads to the wrong solution, even for $n \rightarrow \infty$. Consider two separate models where the only difference is in the structural equation of X_2 which has no effect on Y . The true causal parameter is $\beta^0 = (1, 0)$. Assuming X_1 is intervened on both settings, we obtain an estimator solving equation (21) on each separate environment, both satisfying their corresponding orthogonality constraints. Even though β^0 also satisfies the orthogonality constraints in both cases, the estimator on the right-hand side model is a solution with smaller ℓ_1 norm than β^0 , and thus is not consistent. The minimal ℓ_1 norm solution to the linear constraint for the left-hand side model is however the true β^0 .

We now refer back to the two scenarios briefly described at the beginning of this section. Consider environments generated via randomization. The population level matrix \mathbf{G} is invertible following Proposition 2 whenever we have access to a very large number of environments, such that for any covariate there is an environment where it is randomized. In this case, even if we only have few samples per environment, our high-dimensional causal aggregation procedure will be of practical use as the CIF is positive. On the other hand, whenever the number of constraints is small, we will assume that the entries in the connectivity matrix between covariates are small enough so that the matrix \mathbf{G} is invertible on the set of sparse vectors with same support as β^0 , leading to a positive CIF. This assumption is not




Figure 5: Regularization by shrinking the ℓ_1 norm of the coefficient vector requires additional structural assumptions. In this example, both $\beta^{(1)} = (1, 0)$ and $\beta^{(2)} = (0, .5)$ satisfy $\mathbf{G}\beta = \mathbf{Z}$ in their respective environments. Since $\|\beta^{(1)}\|_1 > \|\beta^{(2)}\|_1$, the proposed method would estimate $\beta = \beta^{(2)}$ in the population case for the right-hand model, which is not the correct solution. This issue can be mitigated by pre-screening, as we will discuss below.

verifiable in practice, and therefore we recommend pre-screening in these scenarios which we discuss in Section 5.2.

We thus now assume that the CIF value is positive, and in particular does not decrease too fast in the high-dimensional regime where p, n_e simultaneously grow. We formalize our main result in the following proposition.

Proposition 5 *Denote by $S^0 := \{j : \beta_j^0 \neq 0\}$ the active set of covariates. Assume that X_j^e are σ_X^2 sub-Gaussian, ϵ_Y^e are σ_E^2 sub-Gaussian, and that R^c are σ_C^2 sub-Gaussian for all $e \in \mathcal{E}, j \in [p], c \in \mathcal{C}$ and some fixed $\sigma_X^2, \sigma_E^2, \sigma_C^2 > 0$. Additionally, assume that*

$$\frac{1}{\text{CIF}_q(S^0, \mathbf{G})} \sqrt{\frac{\log p}{\min_{e \in \mathcal{E}} n_e}} \xrightarrow{\{n_e\}_{e, p} \rightarrow +\infty} 0$$

so that, in particular, $\text{CIF}_q(S^0, \mathbf{G}) > 0$. There exists a constant $K > 0$, that depends only on $\sigma_X^2, \sigma_E^2, \sigma_C^2 > 0$ and K_0 , another universal constant, such that, for the following choice of λ :

$$\lambda := K \sqrt{\frac{\log p}{\min_{e \in \mathcal{E}} n_e}}$$

we get

$$\mathbb{P} \left(\|\hat{\beta}(\lambda) - \beta^0\|_q \leq \frac{K|S^0|^{1/q}}{\text{CIF}_q(S^0, \mathbf{G})} \sqrt{\frac{\log p}{\min_{e \in \mathcal{E}} n_e}} \right) \xrightarrow{\{n_e\}_{e, p} \rightarrow +\infty} 1$$

Our assumptions require in particular that the dimension p does not grow too fast compared to the number of samples in the environments:

$$\frac{\log p}{\min_{e \in \mathcal{E}} n_e} \xrightarrow{\{n_e\}_{e, p} \rightarrow +\infty} 0$$

The estimator $\hat{\beta}(\lambda)$ is sparse and we denote $\hat{S}(\lambda) := \{j : \hat{\beta}(\lambda) \neq 0\}$ its active set. Our estimation procedure leads to a feature selection procedure as the following result holds under an additional assumption on the minimum value of the non-zero coordinates of β^0 , called beta-min assumption Bühlmann et al. (2013). This condition is required for support recovery as it provides the required separation between the non-zero coordinates of β^0 and the null vector with respect to the CIF-derived upper bound of the ℓ_q loss. Note that support recovery is different from model selection consistency, which usually needs much stronger assumptions (Zhao and Yu, 2006).

Proposition 6 *Assume that the conditions of the result above hold with $q = \infty$, and that*

$$\min_{j \in S^0} |\beta_j^0| > \frac{K}{\text{CIF}_\infty(S^0, \mathbf{G})} \sqrt{\frac{\log p}{\min_{e \in \mathcal{E}} n_e}}$$

We then have

$$\mathbb{P}(S^0 \subset \hat{S}(\lambda)) \xrightarrow[\{n_e\}_e, p \rightarrow +\infty]{} 1$$

In conclusion, our regularized estimator recovers the support of β^0 and converges in probability to β^0 under the ℓ_q norm under the assumption that the CIF value is positive. This holds whenever we have a very large number of environments where every covariate is randomized in at least one environment, even if we have few samples per environment. In cases where the number of randomized constraints is too small this assumption may not hold in practice. In this setting, we recommend the pre-screening estimation procedure described below.

5.2 Estimation via Pre-Screening

Given the limitations of the direct shrinkage approach, which requires few interventional samples per environment, but a large number of experiments, we propose an alternative formulation for estimation of β^0 in the high-dimensional setting based on a two-step procedure. However, we instead assume that some observational data is available: we can start with a pre-screening step that chooses a subset of covariates by running a Lasso regression of Y on X on the observational data set. Under regularity assumptions, the set \hat{S} of covariates selected by running a Lasso regression on observational data contains the Markov blanket of the response variable Y : in particular, $\mathbb{P}[\hat{S} \supseteq S^0] \rightarrow 1$ under some conditions on the non-zero regression coefficients (Bühlmann and Van De Geer, 2011, Section 2.5). This procedure assumes that we are able to choose covariates to intervene on: randomizing variables in \hat{S} in (potentially) multiple environments generates orthogonality constraints that allow us to estimate β^0 . The estimator is computed on a different dataset from the one used to pre-select covariates, which guarantees asymptotically valid confidence intervals if \hat{S} contains the Markov blanket. We report in Appendix B simulations based on synthetic data to validate our procedure based on a pre-screening step.

6. Non-Linear Causal Aggregation

Here we extend our causal aggregation procedure beyond the linear case. This methodology allows us to recover non-linear interaction terms between covariates if we have access to

environments where those covariates are simultaneously randomized. If we assume that the response variable model has no such interactions, environments where a single variable is randomized still allow for estimation of non-linear responses. Compared to previous sections, in this section we focus on causal constraints that arise through randomization.

We first define in Section 6.1 a non-linear SEM extension, and then characterize in Section 6.2 sufficient conditions on the set of environments that allow us to identify the non-linear function of the response variable structural equation. Finally, in Section 6.3 we propose a causal aggregation procedure inspired from the Boosting methodology.

6.1 Non-linear Structural Equation Models

Estimating causal effects is more challenging in presence of interactions, and linear approximations may not capture true causal relationships. We define a non-linear extension of the SEM in (3) via the following causal model \mathcal{M}^0 :

$$\mathcal{M}^0 : \begin{cases} X_j \leftarrow f_j(\mathbf{X}_{\text{pa}_0(j)}) + \epsilon_j & \text{for all } 1 \leq j \leq p \\ Y \leftarrow f^0(\mathbf{X}_{\text{pa}_0(p+1)}) + \epsilon_Y \end{cases} \quad (22)$$

where we define structural equations as real-valued functions f_j for $1 \leq j \leq p$ over a subset of covariates indexed by $\text{pa}_0(j) \subset [p+1]$. The structural equations from the causal model \mathcal{M}^0 given by (22) define a DAG $G^0 = (V^0, E^0)$ over the $p+1$ nodes denoted by V^0 . As in the linear case, we assume that G^0 is a DAG and that the response variable Y can belong to the parental sets of covariate nodes. Unlike the linear case, we do not make explicit the confounding effect via latent variables, instead allowing disturbance terms $\epsilon = \{\epsilon_i\}_{i \in [p+1]}$ to be dependent. Therefore, due to such confounding the observational distribution \mathbb{P}^0 factorizes in an extended graph $\bar{G}^0 = (V^0, \bar{E}^0)$ with additional edges in \bar{E}^0 . Conditional independence statements based only on the structure of G^0 may not hold if a latent variable simultaneously influences several nodes. The function f^0 represents the causal effect of the covariates on the response, which takes as argument a subset $S^0 := \text{pa}_0(p+1) \subset [p]$ of covariates and potentially contains non-linear effects of covariates on the response as well as interaction effects between covariates. We assume that all functions $(f_j)_j, f^0$ are square integrable under any environment distribution \mathbb{P}^e . The function f^0 represents the unconfounded relationship between covariates and response: under an interventional distribution \mathbb{P}^e where $\mathbf{X}_{S^0}^e$ are randomized, we assume that $\mathbb{E}[\epsilon_Y^e | \mathbf{X}_{S^0}^e] = 0$. The model for generating distributions \mathbb{P}^e from interventions on covariates indexed by $\phi \subset [p]$ is defined analogously to (5) by:

$$\mathcal{M}_{\phi(e)}^e : \begin{cases} X_j^e \leftarrow f_j^e(\mathbf{X}_{\text{pa}_e(j)}^e) + \epsilon_j^e & \forall j \notin \phi(e) \\ X_j^e \leftarrow \epsilon_j^e & \forall j \in \phi(e) \\ Y^e \leftarrow f^0(\mathbf{X}_{S^0}^e) + \epsilon_Y^e \end{cases} \quad (23)$$

where randomized covariate disturbance terms $(\epsilon_j^e)_{j \in \phi}$ are jointly independent, independent of non-randomized covariate ones $(\epsilon_j^e)_{j \notin \phi}$. As in the linear case, an interventional causal model $\mathcal{M}_{\phi(e)}^e$ still factors in a simplified DAG $\bar{G}^e = (V^0, \bar{E}^e)$ where incoming edges into $\mathbf{X}_{\phi(e)}$ nodes are deleted. Our proposed non-linear aggregation procedure relies on the structure of \bar{G}^e .

6.2 Identification of the Response Structural Equation

Our previous approach for aggregating information across environments does not carry on to this scenario as we can no longer stack vectors representing linear orthogonality constraints into one system of equations that aggregates the causal information derived from each environment. Similar to the linear case, estimating f^0 is possible whenever all the covariates in the parental set of Y are randomized: in environment e where $S^0 \subset \phi(e)$, we have $\mathbb{E}[Y^e | \mathbf{X}_{\phi(e)}^e] = f^0(\mathbf{X}_{S^0}^e) + \mathbb{E}[\epsilon_Y^e | \mathbf{X}_{\phi(e)}^e] = f^0(\mathbf{X}_{S^0}^e)$. Any non-parametric regression method can be used to estimate f^0 in this setting, where non-randomized covariates are ignored. However, as we indicated in the introduction, simultaneous randomization may not be feasible in practice, so that instead of one fully randomized experiment we may only have access to several datasets where different subsets of covariates are randomized. Our method for constructing an estimator of f^0 still relies on orthogonality constraints that we obtain environment-wise and then aggregate into a single estimator \hat{f} . In the example above, denoting by $\sigma(U)$ the Borel σ -algebra generated by the random variable U , we get that $f^0(\mathbf{X}_{S^0}^e)$ is characterized as the orthogonal projection of Y^e on the subspace $L_2(\sigma(\mathbf{X}_{\phi(e)}^e))$ of square-integrable $\sigma(\mathbf{X}_{\phi(e)}^e)$ -measurable random variables: for any square-integrable Borel function h , we have

$$\mathbb{E}[h(\mathbf{X}_{\phi(e)}^e)(Y^e - f^0(\mathbf{X}_{S^0}^e))] = 0 \quad (24)$$

Conversely, this projection does not capture the true f^0 if some covariates in S^0 are confounded: the L_2 projection by conditioning over \mathbf{X}_{S^0} leads to a biased estimate of f^0 . However, under additional assumptions on f^0 and the set of causal models $(\mathcal{M}_{\phi(e)}^e)_e$, exact recovery of f^0 is still possible based on multiple environments, where only a few covariates in S^0 are simultaneously randomized in each environment. Our first result shows that, given a collection of environments $(\mathcal{M}_{\phi(e)}^e)_e$, assuming f^0 can be decomposed as follows:

$$f^0(\mathbf{x}_{S^0}) := \sum_{e \in \mathcal{E}} f_e(\mathbf{x}_{\phi(e) \cap S^0})$$

then we can identify f^0 . We do not know S^0 in practice, and thus we assume that f^0 is a sum of functions representing each an interaction term between subsets of covariates in S^0 which are simultaneously randomized in a given environment:

$$f^0(\mathbf{x}_{S^0}) = \sum_{e \in \mathcal{E}} f_e(\mathbf{x}_{\phi(e)}) \quad (25)$$

Let \mathcal{F} be the set of square-integrable Borel functions for all \mathbb{P}^e , and $\mathcal{F}_{\mathcal{E}} \subset \mathcal{F}$ the set of functions f that can be decomposed as above (25), where $f_e : \mathbb{R}^{|\phi(e)|} \rightarrow \mathbb{R}$ and $f_e \in \mathcal{F}$ for all $e \in \mathcal{E}$.

Proposition 7 *Assume that the distribution \mathbb{P}^e defining environment $e \in \mathcal{E}$ is generated through interventions following model (23), and that all \mathbb{P}^e have the same support. Then there exists at most one function $\bar{f} \in \mathcal{F}_{\mathcal{E}}$ such that, for all $e \in \mathcal{E}$, for all $h \in \mathcal{F}$, we have*

$$0 = \mathbb{E}[h(\mathbf{X}_{\phi(e)}^e)(Y^e - \bar{f}(\mathbf{X}^e))] \quad (26)$$

Additionally, f^0 satisfies equation (26) above over the shared support of $(\mathbb{P}^e)_e$ whenever $f^0 \in \mathcal{F}_{\mathcal{E}}$.

This effectively corresponds to projecting the response Y on a smaller subspace of random variables. However, this projection is only on the unconfounded covariates within each environment: the objective function varies between environments when fitting any regression model to estimate such projections.

6.3 Causal Aggregation via Boosting

Based on the available environments, we propose a causal aggregation estimator $\hat{f} \in \mathcal{F}_{\mathcal{E}}$ of f^0 following decomposition in equation (25):

$$\hat{f} = \sum_{e \in \mathcal{E}} \hat{f}_e(\mathbf{x}_{\phi(e)})$$

by individually estimating each term \hat{f}_e via samples from the corresponding environment—specifically, the randomized covariates of those samples. Our ability to recover f^0 is therefore limited by the overall “causal information” derived from each environment: the more covariates are simultaneously randomized in some experiment, the better our procedure will be at capturing potential interactions. The identification result can inspire a naive fitting procedure that closely resembles the backfitting algorithm for additive models (Breiman and Friedman, 1985), but uses samples from environment e to estimate the different additive terms \hat{f}_e . We first initialize all the additive term estimators \hat{f}_e (for example, setting these to 0). We then randomly pick an environment e , compute the residual term given by

$$R^e := Y^e - \sum_{\bar{e} \neq e} \hat{f}_{\bar{e}}(\mathbf{X}_{\phi(\bar{e})}^e)$$

and update \hat{f}_e based on the orthogonality constraint for that environment where h is any square-integrable random variable:

$$0 = \mathbb{E} \left[h(\mathbf{X}_{\phi(e)}^e) (R^e - \hat{f}_e(\mathbf{X}_{\phi(e)}^e)) \right] \tag{27}$$

This in turn corresponds to choosing \hat{f}_e as the minimizer of

$$\hat{f}_e = \arg \min_{f_e} \mathbb{E} \left[(R^e - f_e(\mathbf{X}_{\phi(e)}^e))^2 \right] \tag{28}$$

We then loop until some convergence criterion is attained, for example, taking the supremum over e of an ℓ_{∞} loss on the difference between the two last most recent updates of \hat{f}_e . We describe this procedure in Algorithm (1). The connection between this procedure and the proof of the identification result in Proposition 7 is the following. Assume the sequential training of the component \hat{f}_e in \hat{f} starts with the environment index e whose input variables are the most downstream in the DAG derived from \mathcal{M}^0 . Then this estimator is set to the true population f_e when fitted on the environment e if the model is well specified—similar to the observation leading to (24) in the fully randomized case. Iterating this procedure we recover f^0 by sequentially cancelling out the contribution of each f_e in Y^e when computing (28). Unfortunately, this procedure would require knowing the graph G^0 , and other challenges arise when using this method that showcase the differences with respect to the original backfitting

algorithm. In particular, the loss (28) is minimized on different datasets. Therefore, the sequence of loss values for subsequent iterations of our estimator on one given environment is not decreasing. After updating any given term in decomposition (25) based on minimizing the loss over the corresponding environment, the loss evaluated at another environment may increase instead.

Algorithm 1: Non-linear Causal Aggregation via Backfitting

Initialize $\hat{f}_e = 0$, set $\delta_0 > 0$ convergence threshold, $\delta = 2 * \delta_0$ current update gap.
while $\delta > \delta_0$ **do**
 Sample uniformly: $e \sim \mathcal{U}(|\mathcal{E}|)$
 Compute the residual: $R^e \leftarrow Y^e - \sum_{\bar{e} \neq e} \hat{f}_{\bar{e}}(\mathbf{X}_{\phi(\bar{e})}^e)$
 Compute estimator over environment: $\hat{g}_e \leftarrow \arg \min_{f_e \in \mathcal{F}_{\mathcal{E}}} \mathbb{E} \left[(R^e - f_e(\mathbf{X}_{\phi(e)}^e))^2 \right]$
 Update gap: $\delta \leftarrow \sup_e \sup_{\mathbf{X}^e} |\hat{f}_e(\mathbf{X}^e) - \hat{g}_e(\mathbf{X}^e)|$
 Update estimator component: $\hat{f}_e \leftarrow \hat{g}_e$
return $\hat{f} = \sum_{e \in \mathcal{E}} \hat{f}_e$

This algorithm then suffers from unstable training, in particular when fitting models with small sample sizes. These issues have the same underlying source: exactly enforcing the orthogonality constraint on one environment to update one component can in turn make the orthogonality conditions fail on the remaining components.

We propose a procedure that addresses this issue, called non-linear causal aggregation via boosting. Boosting is a greedy algorithm where weak learners are aggregated to form a regression estimator or a classifier. Weak learners correspond to individual simple estimators derived from some base procedure such as regression or decision trees. Initial boosting formulations for classification sequentially fit weak learners to samples re-weighted in such way that the mis-classified ones by previously generated weak learners get higher weights in the next iteration. The procedure thus increasingly focuses on the “hard” training samples (Freund, 1995; Freund and Schapire, 1997; Schapire, 1990). The final output is then a linear weighted average of the weak learners. In Friedman et al. (2000), boosting is reformulated as a sequential additive modeling procedure where fitted weak learners correspond to the negative gradient of the classification loss in the function space. Weak learners are then sequentially added to the current function estimator. This characterization was first observed in Breiman et al. (1998); Breiman (1999), but using this additive modeling perspective boosting methods were subsequently developed beyond classification problems, in particular for regression problems (Friedman, 2001). At each iteration a weak learner is trained on top of the current iteration of the full estimator: using the square loss, this corresponds to training the weak learner on the residual term obtained based on the current regression model (Bühlmann and Yu, 2003). Within our framework, instead of replacing altogether the individual component \hat{f}_e with an updated fitted model trained on environment e alone, we incrementally update all components \hat{f}_e simultaneously by adding weighted function estimates \hat{h}_e defined over the same set of covariates $\mathbf{x}_{\phi(e)}$ as the corresponding \hat{f}_e . This can also be understood as aggregating “weakly informative datasets”, as each boosting step only takes into account for each sample those covariates that are unconfounded. Those \hat{h}_e are first estimated independently: within each environment, we use the current aggregate function \hat{f}

to generate the environment residuals

$$R^e := Y^e - \sum_{\tilde{e}} \hat{f}_{\tilde{e}}(\mathbf{X}_{\phi(\tilde{e})}^e) = R^e - \hat{f}(\mathbf{X}^e)$$

Then, independently for each environment, we estimate \hat{h}_e by solving a minimization problem as in (28) based on samples from the corresponding environment e . We can view this in light of the gradient boosting methodology (Friedman et al., 2001). We similarly fit a set of candidate learners to the residuals derived from a quadratic loss function. However, individual components \hat{h}_e use only information pertaining to the randomized covariates in each individual data set. Recent work has developed the boosting framework for non-linear instrumental variables regression (Bakhitov and Singh, 2021). Instead of using instruments from a single data set, our procedure deals with heterogeneous datasets with different distributions generated via different experimental settings. The second step in this procedure scales the candidate $(\hat{h}_e)_{e \in \mathcal{E}}$ by weights constructed to simultaneously satisfy all orthogonality constraints. That is, we update each component as follows:

$$\hat{f}_e \leftarrow \hat{f}_e + \hat{\alpha}_e \hat{h}_e$$

where $(\hat{\alpha}_e)_{e \in \mathcal{E}}$ are chosen to satisfy the constraints (27) for the corresponding aggregated $\sum_e \hat{f}_e + \hat{\alpha}_e \hat{h}_e$:

$$0 = \mathbb{E} \left[\hat{h}_e(\mathbf{X}_{\phi(e)}^e) \left(Y^e - \left(\sum_{\tilde{e}} \hat{f}_{\tilde{e}}(\mathbf{X}_{\phi(\tilde{e})}^e) + \hat{\alpha}_e \hat{h}_{\tilde{e}}(\mathbf{X}_{\phi(\tilde{e})}^e) \right) \right) \right] \quad \forall e \in \mathcal{E}$$

One choice of loss function to implement this weight parameter search leads to the convex minimization problem:

$$(\hat{\alpha}_e)_{e \in \mathcal{E}} := \arg \min_{\alpha_e \in \mathbb{R}} \sum_{e \in \mathcal{E}} \left| \mathbb{E} \left[\left(R^e - \sum_{\tilde{e} \in \mathcal{E}} \alpha_{\tilde{e}} \hat{h}_{\tilde{e}}(\mathbf{X}_{\phi(\tilde{e})}^e) \right) \hat{h}_e(\mathbf{X}_{\phi(e)}^e) \right] \right| + \nu \|\alpha\|_2^2 \quad (29)$$

where we regularize the weight vector α with a quadratic penalty scaled by $\nu \geq 0$ that we empirically observed improves the training behavior. This step is key to guarantee the stability of the procedure. Although individual \hat{h}_e are fitted through individual environment e samples, the global re-weighting over the joint data across environments prevents that individual updates by a single \hat{h}_e perturb the orthogonality constraints on environments $\tilde{e} \neq e$. Finally, we empirically observe that shrinking the updates by a learning rate $1 > \eta > 0$ improves convergence, so our final proposed update is given by:

$$\hat{f} \leftarrow \hat{f} + \eta \sum_{\tilde{e} \in \mathcal{E}} \hat{\alpha}_{\tilde{e}} \hat{h}_{\tilde{e}}$$

We again stop the training procedure as soon as the new term $\sum_{\tilde{e} \in \mathcal{E}} \hat{\alpha}_{\tilde{e}} \hat{h}_{\tilde{e}}$ no longer substantially updates the aggregated \hat{f} , and hyper-parameters such as the learning rate and the penalty weight ν can be chosen by keeping a separate validation data set within each environment and evaluating the unpenalized orthogonality constraint loss in the minimization problem (29).

Algorithm 2: Causal Aggregation Boosting

Initialize $\hat{f}_e = 0$, $\hat{f} = \sum_e \hat{f}_e$, set η learning rate, ν penalty weight, set $\delta_0 > 0$
convergence threshold, $\delta = 2 * \delta_0$ current update gap.
while $\delta > \delta_0$ **do**
 for $e \in \mathcal{E}$ **do**
 Compute residual: $R^e \leftarrow Y^e - \hat{f}(\mathbf{X}^e)$
 Compute individual component over environment:
 $\hat{h}_e \leftarrow \arg \min_{h_e \in \mathcal{F}_{\mathcal{E}}} \mathbb{E} \left[(R^e - h_e(\mathbf{X}_{\phi(e)}^e))^2 \right]$
 Enforce orthogonality constraints:
 $(\hat{\alpha}_e)_{e \in \mathcal{E}} := \arg \min_{\alpha_e \in \mathbb{R}} \sum_{e \in \mathcal{E}} \left| \mathbb{E} \left[(R^e - \sum_{\bar{e} \in \mathcal{E}} \alpha_{\bar{e}} \hat{h}_{\bar{e}}(\mathbf{X}_{\phi(\bar{e})}^e)) \hat{h}_e(\mathbf{X}_{\phi(e)}^e) \right] \right| + \nu \|\alpha\|_2^2$
 Update gap: $\delta \leftarrow \sup_{e \in \mathcal{E}} \sup_{\mathbf{X}^e} \left| \sum_{\bar{e} \in \mathcal{E}} \hat{\alpha}_{\bar{e}} \hat{h}_{\bar{e}}(\mathbf{X}_{\phi(\bar{e})}^e) \right|$
 Update model: $\hat{f} \leftarrow \hat{f} + \eta \sum_{\bar{e} \in \mathcal{E}} \hat{\alpha}_{\bar{e}} \hat{h}_{\bar{e}}$
 return \hat{f}

We summarize the procedure in Algorithm (2). The proposed procedure leverages the same causal aggregation principle as in the linear aggregation framework, but as a subroutine of the fitting procedure. Our procedure can be understood as gradient descent on the space of functions $\mathcal{F}_{\mathcal{E}}$, where at each step we linearize the space of functions by computing the environment-wise direction of maximal descent, and then fit the best global linear approximation before taking the gradient step. The crucial point is that the environment-wise components are fitted based on the unconfounded covariates of each sample, as the variables $\mathbf{X}_{\phi(e)}^e$ used to estimate \hat{h}_e in environment e are precisely those that are randomized.

7. Numerical Simulations

We run simulations based on synthetic and semi-synthetic data. We start by validating in Section 7.1 the main results of linear causal aggregation in the just-identified case, its extension to the over-identified case and compare it to do calculus. We then validate non-linear causal aggregation via boosting in Section 7.2. Finally, we then generate in Section 7.3 a semi-synthetic data set based on a gene perturbation experiment where we know the true causal relationship between \mathbf{X} and Y by design and use it to validate our causal aggregation method.

7.1 Aggregation in the linear case

7.1.1 AGGREGATION IN THE JUST-IDENTIFIED CASE

We build on Example 1 to illustrate a simple case where we aggregate causal information from a diverse set of environments, with different types of constraints built within each environment. We add new covariates as represented in Figure 6, and generate samples




Figure 6: SEM for observational environment: Experimental environments are obtained by randomizing covariates in the above SEM, for which their incoming edges are removed.

according to a SEM \mathcal{M}^0 as in (3). The structural equations defining the model are as follows:

$$\begin{aligned}
 H &\leftarrow \epsilon_H \\
 X_1 &\leftarrow 2H + \epsilon_1 \\
 X_2 &\leftarrow X_1 + H + \epsilon_2 \\
 X_3 &\leftarrow -X_1 + 2X_2 + \epsilon_3 \\
 X_4 &\leftarrow X_1 + X_3 + \epsilon_4 \\
 X_5 &\leftarrow 2X_2 + X_4 - Y + \epsilon_5 \\
 Y &\leftarrow X_2 + 2X_4 + H + \epsilon_Y
 \end{aligned} \tag{30}$$

This model contains latent factors that simultaneously affect the covariates and the response. The causal vector is given by $\beta^0 = (0, 1, 0, 2, 0)$, but the presence of a latent variable H biases the least squares estimates of Y on \mathbf{X} . Additionally, the vector ϵ is sampled from a standard Gaussian distribution (where for simplicity we include ϵ_H in the vector ϵ previously defined). We consider four different environments. First, samples are collected from an observational environment e_1 where an instrument I is available for X_1 , so that its structural equation becomes $X_1 \leftarrow 2H + I + \epsilon_1$. An experimental environment e_2 is generated by intervening on X_3 and X_5 . This intervention is translated into a SEM with the same equations above except for X_3, X_5 that follow $X_i \leftarrow \epsilon_i$. We generate samples for the third environment e_3 from an interventional data set where we assume we intervene on X_2 and also we know the parental set of X_4 . These interventional environments follow the the causal model $\{\mathcal{M}_{\phi(e)}^e\}_e$ where $\phi(e_2) = \{3, 5\}$ and $\phi(e_3) = \{2\}$ (cf. equation 5). Randomization in the experimental datasets removes spurious correlations due to the latent variable H and truncate the direct contribution of parental variables of intervened covariates. Last, we generate an environment e_4 where all covariates are randomized and thus the latent variable H has no longer any confounding effect, i.e. $\phi(e_4) = \{1, 2, 3, 4, 5\}$.

We run several experiments and evaluate the performance when using different subsets of constraints derived from the environments above. The first criterion we look at is whether




Figure 7: Graphical representation of experimental environments: In environment e_2 (left), X_3 and X_5 are randomized. In environment e_3 (right) only X_2 is randomized. Nodes in red correspond to covariates that are associated to an orthogonality constraint.

Table 1: We construct confidence intervals with .95 nominal coverage based on the SEM in Figure 6. We report the actual coverage of those confidence intervals for increasing sample sizes, over 500 repetitions. Bold values indicate that the target coverage is achieved.

	Estimator	$n = 50$	$n = 100$	$n = 200$	$n = 500$	$n = 1000$
Experiment A	Causal Aggregation	0.98 ± 0.02	0.97 ± 0.02	0.96 ± 0.02	0.95 ± 0.02	0.96 ± 0.02
	Pooled data OLS	0.35 ± 0.04	0.23 ± 0.04	0.14 ± 0.03	0.04 ± 0.02	0.01 ± 0.01

the actual coverage from the confidence intervals matches the nominal pre-specified level. We compute confidence intervals with nominal coverage 0.95 for each coordinate and report actual coverage for increasing sample sizes $n \in \{50, 100, 200, 500, 1000\}$, averaging across 500 repetitions. In those cases where several methods attain the correct coverage, we compare them based on the mean length of the individual coordinate-wise confidence intervals. We begin comparing our causal aggregation estimator $\hat{\beta}$ presented above and simple OLS, and we later add other methods with competitive performance.

Our first experiment (referred to as experiment A) relies on constraints from environments e_1 , e_2 and e_3 that are directly derived from the constraint inducing variables. Each variable has exactly one such constraint across all environments, this corresponds to the just-identified case. We report in Table 1 the coverage of the confidence intervals derived in equation (15) for different values of sample size, based on 500 repetitions of the simulation. As expected, our causal aggregation estimator has proper coverage following our theoretical results, and OLS, being inconsistent, does not achieve the targeted coverage.

7.1.2 SIMULATIONS IN THE OVER-IDENTIFIED CASE

We define additional experiments built upon the previous example based on the SEM (30). We already assumed the parental set for X_4 is known, but previously we only constructed the corresponding constraint in e_3 , leaving information on the side. Based on our assumptions on

Table 2: Extending Table 1 to experiments B, C and D. As expected, the confidence intervals derived from our causal aggregation method always achieve the target coverage.

	Estimator	$n = 50$	$n = 100$	$n = 200$	$n = 500$	$n = 1000$
Experiment B	Causal Aggregation	0.97 ± 0.02	0.96 ± 0.02	0.96 ± 0.02	0.96 ± 0.02	0.96 ± 0.02
	Pooled Data OLS	0.36 ± 0.04	0.23 ± 0.04	0.13 ± 0.03	0.05 ± 0.02	0.01 ± 0.01
Experiment C	Causal Aggregation	0.94 ± 0.02	0.94 ± 0.02	0.96 ± 0.02	0.95 ± 0.02	0.95 ± 0.02
	Pooled Data OLS	0.95 ± 0.02	0.95 ± 0.02	0.95 ± 0.02	0.95 ± 0.02	0.95 ± 0.02
Experiment D	Causal Aggregation	0.95 ± 0.02	0.94 ± 0.02	0.94 ± 0.02	0.95 ± 0.02	0.95 ± 0.02
	Pooled Data OLS	0.39 ± 0.04	0.20 ± 0.04	0.06 ± 0.03	0.01 ± 0.01	0.00 ± 0.00

the causal model $\mathcal{M}_{\phi(e)}^e$, the parental set for X_4 remains unchanged in those environments where X_4 is not intervened on, and thus we have access to two additional constraints. We define experiment B by including those, obtained in environments e_1 and e_2 , on top of those already available in experiment A. Experiment C is defined though the randomization constraints from environment e_4 only. Last, experiment D combines all constraints from experiments B and C.

We first check the actual coverage of causal aggregation and OLS and report the results in Table 2 where each row corresponds to a different experiment. Again, causal aggregation achieves the target coverage in all experiments, but OLS only does so in experiment C, where full randomization removes any confounding due to the latent variable and thus the OLS estimator becomes consistent and the derived confidence intervals attain the nominal coverage. Actually, OLS and our causal aggregation method are the same in the just-identified setting of experiment C. We also notice a slight over-coverage in certain settings: the achieved coverage may be higher than the target. We conjecture that this is due to the two-step nature of the regression adjustment constraints: the estimated variance of the constraints may be slightly upwardly biased due to the adjustment step previous to the construction of the orthogonality constraint, which leads to wider-than-expected confidence intervals in finite samples.

Having checked that the confidence intervals from our causal aggregation method achieve the nominal coverage, we now turn to comparing size across different experiments. Figure 8 shows the main point of these simulations: increasing the number of constraints improves the asymptotic efficiency of the causal aggregation estimator. Experiments A and C correspond to two just-identified cases, and additional constraints define experiments B and D respectively, which are thus over-identified cases where we use the two-step MM aggregation method. In both cases the confidence intervals become tighter: experiment B shows an improvement over A, and experiment D improves over C and B. Even if causal aggregation (and OLS in experiment C) are able to properly estimate β^0 , our method leverages additional information to improve efficiency.

7.1.3 SIMULATIONS WITH ADDITIONAL DO-CALCULUS CONSTRAINTS

Within environment e_3 , we report the performance of another standard, well-known estimation procedure of the causal effects: regression after applying do-calculus on a known causal graph.




Figure 8: Box plots of the mean length of the confidence intervals across 500 repetitions for a range of sample sizes in different experimental settings. The causal aggregation estimator is consistent and the confidence intervals achieve the target coverage. Our method can additionally leverage information from every environment that generates a valid constraint to tighten the confidence interval length.

Table 3: We construct confidence intervals with .95 nominal coverage and report the actual coverage of those confidence intervals for increasing sample sizes, over 500 repetitions. Bold values indicate that the target coverage is achieved.

	$n = 50$	$n = 100$	$n = 200$	$n = 500$	$n = 1000$
Do-calculus	0.94 ± 0.02	0.95 ± 0.02	0.95 ± 0.01	0.96 ± 0.02	0.95 ± 0.02
Causal Aggregation	0.94 ± 0.03	0.96 ± 0.02	0.95 ± 0.01	0.97 ± 0.01	0.98 ± 0.01

To showcase the advantage of our methodology that leverages all the available data from each environment, we assume that the causal graphs in each environment are *known* to the do-calculus method. We have applied the causal transportability formula (a generalization of the do-calculus as implemented in the R-package `causaleffect`) to the environments $\{e_1, e_2, e_3\}$. Among the environments $\{e_1, e_2, e_3\}$ (or combinations of environments), only e_3 leads to identifiable causal effects on Y of intervening on all X_1, X_2, X_3, X_4, X_5 . These are obtained by regressing Y on the set of variables X_1, X_2, X_3, X_4 and reporting the coefficients of X_2, X_4 . Therefore, based on samples from e_3 only, we can estimate β^0 based on two constraints derived from the do-calculus approach and the perfect knowledge of the causal graph structure. We omit any analysis with samples from the fully randomized environment e_4 , which could be used in an equivalent manner by both methods.

The causal transportability formula does not recommend using any samples from e_1 and e_2 . This is potentially due to the fact that the approach is non-parametric. Samples from e_1 and e_2 help improving the efficiency of the estimator whenever we use causal aggregation. Indeed, the previous causal aggregation constraints in those environments can be updated with the knowledge derived from the causal graph: certain coefficients in β^0 are equal to 0 regardless of the environment. We can therefore restrict the problem to a regression of Y over X_2, X_4 and build a causal aggregation estimator combining the two do-calculus constraints with the updated ones from environments e_1 and e_2 . We verify in Table 3 that both methods achieve the target coverage as expected. We then compare the length of the confidence intervals in Figure 9, showing again an improvement in efficiency as we include additional constraints from multiple environments with causal aggregation.




Figure 9: Box plots of the mean length of the confidence intervals across 500 repetitions for a range of sample sizes.

In conclusion, even in situations where OLS and do-calculus lead to consistent estimators of the causal effects, we can benefit from collecting data from other environments with confounding where those previous methods may fail to work: if we can extract partial information based on constraints derived from valid assumptions, we can improve the efficiency of our estimators.

7.2 Non-Linear Aggregation

Consider the same SEM used for simulations in the low dimensional case, whose DAG structure is given in Figure 6, but where the structural equation of Y is given by a complex non-linear function. We also modify structural equations of \mathbf{X} to validate the fact that causal boosting performs well when allowing flexible non-linear covariate structural equations.

In particular, we consider the following SEM:

$$\begin{aligned}
 H &\leftarrow \epsilon_H \\
 X_1 &\leftarrow 2H + \epsilon_1 \\
 X_2 &\leftarrow X_1 + H + \epsilon_2 \\
 X_3 &\leftarrow -2X_1X_2 + \epsilon_3 \\
 X_4 &\leftarrow \log(1 + |X_1|) + X_3 + \epsilon_4 \\
 X_5 &\leftarrow 2X_2 + X_4 - Y + \epsilon_5 \\
 Y &\leftarrow f^0(X_1, X_2, X_3, X_4) + 2H + \epsilon_Y
 \end{aligned}$$

where all the disturbance terms $\epsilon_H, (\epsilon_j)_{j \in [5]}, \epsilon_Y$ are jointly independent, standard Gaussian. This model satisfies the assumptions on our non-linear SEM framework (cf. equation 22). We make confounding explicit by introduced the latent factor H that makes the aggregated residual $H + \epsilon_Y$ of the response structural equation non independent of the covariates—as some of these covariates are also affected by H . However, whenever we randomize the covariates in the parental set $S^0 = \{1, 2, 3, 4\}$ of the response Y , we get that the residual $H + \epsilon_Y$ is centered conditionally on the randomized covariates. We run several simulations where we analyze the performance of the proposed causal aggregation boosting algorithm (2) for different choices of function f^0 and sets of environments \mathcal{E} . Our simulations show the correctness of our method whenever the model is well specified in the sense that for any interaction term in f^0 there is an environment in \mathcal{E} where all the input covariates to the interaction term are randomized.

We fit the model on training data and put aside a test set for evaluating the performance of our method. We report the L_2 loss of our estimator \hat{f} given by an oracle with access to f^0 :

$$\mathbb{E}[(f^0(\mathbf{X}^e) - \hat{f}(\mathbf{X}^e))^2]$$

where the expectation is taken over the test set samples \mathbf{X}^e averaged across all environments.

We use two different response models in our simulations, denoted by f_I^0, f_{II}^0 . We first use a piece-wise constant function given by:

$$f_I^0(x_1, x_2, x_3, x_4) = \mathbb{1}_{x_1 > 0} + \mathbb{1}_{x_2 > 0} - 2\mathbb{1}_{x_2 > 0, x_3 > -1} + 2\mathbb{1}_{x_1 < 0, x_4 < -1} + 3\mathbb{1}_{x_1 < 0, x_2 < 1, x_3 < -1}$$

We choose as the base model to fit the environment-wise \hat{h}_e a decision tree: f_I^0 should be simple enough that the causal aggregation boosting algorithm recovers it efficiently. If we instead use another response function that is no longer piece-wise constant, then the number of samples required to achieve the same level of accuracy increases, and we run simulations for the following loss function f_{II}^0 that additively combines a linear component and a piece-wise constant component:

$$f_{II}^0(x_1, x_2, x_3, x_4) = 2x_1 - 2x_2 - 2\mathbb{1}_{x_2 > 0, x_3 > -1} + 2\mathbb{1}_{x_1 < 0, x_4 < -1} + 3\mathbb{1}_{x_1 < 0, x_2 < 1, x_3 < -1}$$

Let a set of two environments $\mathcal{E}^0 = \{e_1, e_2\}$ given by $\phi(e_1) = \{1, 2, 3\}$ and $\phi(e_2) = \{1, 4\}$, so that $f_I^0, f_{II}^0 \in \mathcal{F}_{\mathcal{E}^0}$. We can therefore recover these target functions using such set of environments as our data source. Conversely, given that the set of environments \mathcal{E} constrains the class of functions $\mathcal{F}_{\mathcal{E}}$ our estimator \hat{f} belongs to, we analyze how the performance degrades as the approximation error increases when limiting the interventions across environments.

Simulation	Sets of randomized covariates
Simulation A	$\{1, 2, 3\}, \{1, 4\}$
Simulation B	$\{1, 2, 3\}$
Simulation C	$\{1\}, \{2, 3\}$
Simulation D	$\{1\}, \{2\}, \{3\}$

Table 4: Four different experimental settings: for each simulation, a set of subsets of indices represents the subsets of simultaneously randomized covariates within an environment, across the set of environments.

We define four simulations each corresponding to a different set of environments in Table 4. For each simulation we characterize the set of environments as a set of subsets of indices, each representing one environment. Within such environment, the indices represent the simultaneously randomized covariates. Simulation A corresponds to collecting data from \mathcal{E}^0 above. Simulations B, C and D correspond to fitting our causal boosting procedure in cases where not all the interacting covariates have an environment where they are simultaneously randomized, leading to an approximation bias that increases as we limit the amount of simultaneous randomization. We report the results of fitting our causal boosting procedure to data generated with the response function given by f_I^0 in Figure 10, and as expected we observe that causal boosting correctly recovers the non-linear function: the L_2 loss decreases towards 0 with sample size. However, the loss reaches a plateau in those other simulations where not all the interactions in the target function have corresponding randomized data.

We now turn to analyzing the performance of causal boosting with respect to alternative methods. We focus on the set of environments defined in Simulation A, so that there is no approximation error. A first naive baseline corresponds to fitting a non-parametric regression method to the pooled data across environments, similar to the OLS baseline previously used in the linear case. We choose random forests for that purpose, given that the base estimator in the boosting procedure is a decision tree. We also compare the boosting procedure to the causal backfitting procedure defined in Algorithm 1. Both these procedures will perform poorly: the naive pooled method will indeed fail to recover the true non-linear response function. The backfitting procedure will suffer from instability during training, and will often fail to converge. We define a simpler one-pass backfitting procedure as an alternate baseline that does recover the response function, albeit with an additional key assumption similar to the one in the above section. If we assume that the causal graph is *known*, then we can implement a single-pass backfitting procedure that has competitive performance. The individual components of the estimator \hat{f} following equation 25 are fitted sequentially, starting by those components whose variables are most downstream in the DAG.

We report the result of the simulations in Figure 11, comparing the causal aggregation boosting procedure versus the competing methods, where the response function is either f_I^0 or f_{II}^0 and \mathcal{E}^0 is the set of environments. Again, our aggregation procedure recovers a good estimate of the true causal function for both $f_I^0, f_{II}^0 \in \mathcal{F}_{\mathcal{E}^0}$. However, for a given sample size, the recovery loss for f_I^0 is smaller than that of f_{II}^0 , as expected from the inclusion of linear components in the latter. Random forests fitted on the pooled data perform poorly. This is




Figure 10: Causal boosting loss under different simulation settings. Only Simulation A has environments with properly randomized data such that there is no approximation error.

expected, as this method is oblivious to the confounding. Training the causal backfitting procedure is unstable and does not converge towards the true non-linear response, and this issue is even more problematic for f_{II}^0 . Finally, the single pass backfitting performs reasonably well but requires larger sample sizes to achieve the same performance as the causal boosting procedure.

In conclusion, our causal boosting procedure performs very well. The per-environment model fitting, coupled with the linearization step that re-weights the components of the estimator by drawing from our linear causal aggregation theory is a competitive method for estimating non-linear causal responses from partially randomized data from multiple environments.

7.3 Semi-Synthetic Data Set

We finally validate our procedure on semi-synthetic data set that we create from a single-cell RNA sequencing (scRNA-seq) data set published in Gasperini et al. (2019). This data set contains mRNA counts of 207,324 individual cells measuring the expression of around 10,000 genes. Each of those genes is in a neighborhood of several potential DNA regulatory elements—a section of non-protein coding DNA that modulates the expression of a gene. This experiment in particular focuses on enhancers, regulatory elements that promote gene expression. The goal of the experiment is to find true biological regulatory associations between candidate enhancers pre-selected based on their chemical characteristics, and genes that they potentially up-regulate. This particular data set is obtained by perturbing cells via CRISPR interference (Jinek et al., 2012). This technology allows to directly intervene in




Figure 11: Comparative performance of different aggregation methods for two separate target functions

the gene regulation mechanism of by targeting a section of the genome with a CRISPR/Cas9 molecule attached to a specifically designed guide RNA (gRNA). That gRNA determines the locations where CRISPR/Cas9 introduces DNA mutations that perturb gene regulation. In this data set the targeted DNA sections are the candidate enhancers, therefore genes that were actually regulated by a CRISPR-intervened candidate enhancer region are differentially down-regulated.

We focus on a particular gene called PRKCB, with 10 selected candidate enhancers. Cells are then intervened on a random subset of those enhancers: that is, each cell receives a random combination of CRISPR molecules with gRNAs that target the 10 enhancers. In addition to measuring the response Y corresponding to PRKCB gene expression, the data set has a vector of binary variables (X_1, \dots, X_{10}) for each cell indicating the presence of CRISPR molecules, corresponding to the treatment covariates. Additionally, one needs to control for technical factors W such as temperature and batch ID that simultaneously affect the treatment variables and response, usually introducing those as covariates in a regression model.

Table 5: Our aggregation MM estimator $\hat{\beta}^{MM}$ on the perturbed samples closely matches the original regression estimates $\hat{\beta}^{OLS}$ based on the original unperturbed data set (first two columns). However, the confidence intervals are much larger, owing to the additional noise introduced in the perturbed data set by the latent factor. The OLS estimates $\hat{\beta}^{OLS}(e)$ on any single perturbed data set $e = 1, 2, 3$ leads to biased estimates of the regression vector due to the additional synthetic confounding we introduce. Thus, some of the OLS estimates are far from the estimates on the original, unperturbed data set. On the other hand, the estimates $\hat{\beta}^{MM}$ are close to the estimates from the original data set.

Targeted Candidate Enhancer	Original data set		Perturbed datasets		
	$\hat{\beta}^{OLS}$	$\hat{\beta}^{MM}$	$\hat{\beta}^{OLS}(1)$	$\hat{\beta}^{OLS}(2)$	$\hat{\beta}^{OLS}(3)$
1822 top	0.00 ± 0.04	0.29 ± 0.29	0.05 ± 0.07	0.67 ± 0.07	0.35 ± 0.07
1856 top	0.01 ± 0.04	-0.04 ± 0.31	0.05 ± 0.08	0.76 ± 0.07	0.35 ± 0.07
1857 top	0.06 ± 0.05	0.02 ± 0.32	1.09 ± 0.09	0.04 ± 0.09	0.60 ± 0.08
1863 second	0.03 ± 0.04	-0.29 ± 0.26	0.56 ± 0.07	-0.04 ± 0.08	0.39 ± 0.07
1863 top	0.00 ± 0.04	-0.24 ± 0.25	0.58 ± 0.07	-0.03 ± 0.07	0.30 ± 0.06
1865 top	-0.10 ± 0.04	0.00 ± 0.13	0.30 ± 0.06	0.41 ± 0.06	-0.11 ± 0.06
1866 top	-0.50 ± 0.04	-0.46 ± 0.13	-0.02 ± 0.06	0.08 ± 0.06	-0.51 ± 0.07
1866 second	-0.46 ± 0.05	-0.42 ± 0.17	0.29 ± 0.08	0.53 ± 0.08	-0.43 ± 0.09
1867 top	-0.16 ± 0.05	-0.13 ± 0.15	0.47 ± 0.08	0.78 ± 0.08	-0.12 ± 0.08
1897 top	-0.03 ± 0.05	0.17 ± 0.17	0.72 ± 0.08	0.76 ± 0.07	-0.02 ± 0.08

For the purpose of our semi-synthetic data example we assume a linear model for the response, which is regressed on the binary vector of covariates to identify the actual regulatory elements of the gene conditionally on the technical factors, though other models can better capture these effects Gasperini et al. (2019); Katsevich and Roeder (2020). We obtain an estimate of the effects of perturbing each candidate enhancer via the regression vector. We now perturb the original data set and create several environments with synthetic confounding that leads to biases in estimation at each individual environment. We then recover the original regression vector through our aggregation procedure. We generate three environments $\mathcal{E} = \{1, 2, 3\}$ by randomly partitioning the initial data set samples in three smaller datasets, and for each partition we pick a subset of covariates S_e , $e = 1, 2, 3$. We then introduce in each environment e a confounding latent factor H that simultaneously affects the subset of selected covariates $(X_i)_{i \in S_e}$ and the response Y . We also assume that no covariate is selected in all three new environments. We represent in Figure 12 graphical models corresponding to two potential environments. Recovering the original regression estimate based on just one environment is no longer possible, as we assume that in practice H is not observed. However, if we assume that for each environment we know the subset of unconfounded covariates, we can then construct orthogonality constraints for these, and use the aggregation estimator to recover an estimate that is closer to the original one based on the unperturbed data set. In




Figure 12: Semi-synthetic environments with added confounders: We split the initial data set into three environments and sample three subsets of covariates. In each environment we add a confounding term between the corresponding set of covariates and the response. No covariate is perturbed in all three environments. Graphical models above represent two potential environments.

practice, given one of the environments $e \in \mathcal{E}$, we perturb samples as follows:

$$\begin{aligned} H &\leftarrow \text{Ber}(0.5) \\ \tilde{X}_i &\leftarrow X_i + H \quad \forall i \in S_e \\ \tilde{Y} &\leftarrow Y - 4 * (H - 1) \end{aligned}$$

leaving the other covariates unchanged. The data set covariates are coded via the set of enhancers that are targeted by the CRISPR gRNA. In particular, for $e = 1$ we perturb the set of covariates {1857 top, 1863 second, 1863 top, 1865 top, 1866 top, 1866 second, 1867 top, 1897 top}, for $e = 2$ we perturb the set of covariates {1822 top, 1856 top, 1865 top, 1866 top, 1866 second, 1867 top, 1897 top} and for $e = 3$ we perturb the set of covariates {1822 top, 1856 top, 1857 top, 1863 second, 1863 top}. The OLS estimator $\hat{\beta}^{OLS}(e)$ on any environment $e \in \{1, 2, 3\}$ is thus biased. However, aggregated estimator $\hat{\beta}^{MM}$ provides us with estimates that broadly match those of the original OLS estimator, albeit with wider confidence intervals due to the increased variance because of the latent factor. In particular, confidence intervals for each covariate overlap between these two methods, which is not the case for OLS estimators built on just one environment.

8. Discussion

We have introduced a method for aggregating causal information across data sets, with the goal of estimating the effect of simultaneous interventions. The method is based on causal constraints, which arise from experimental manipulations and background knowledge. On observational data, instrumental variables and knowledge about parental sets can be used to define causal constraints. In the low-dimensional case, we discuss a two-stage procedure that allows for asymptotically efficient estimation and inference of causal effects. In the high-dimensional case we provide an ℓ_1 -regularized estimator and derive finite sample bounds. These finite sample bounds rely on a cone invertibility factor, which play a similar role as the sparse eigenvalue condition in high-dimensional linear regression. As our high-dimensional theory indicates, the method might be of use whenever a very large number of experiments are

available, but only very few samples per experiment are observed. Whenever few covariates are randomized, we recommend instead a pre-screening step on observational data to reduce the dimensionality of the problem. In addition, we provide a non-linear version of causal aggregation of experimental data that flexibly estimates interactions between covariates and non-linearities in the response. This non-linear method uses the linear aggregation step as a sub-routine when training the model following a boosting-like procedure, using at every step only those covariates in each sample that are unconfounded. On synthetic and semi-synthetic data sets we show that the proposed method outperforms naive methods that do not take into account the special structure induced by the causal constraints. Most of the causal constraints individually make use of data from only one environment. However, as in the case of cross-product invariance (9), a causal constraint can leverage data from several environments. Looking ahead, it would be interesting to explore whether novel constraints can be derived that make use of data from different environments simultaneously.

References

- Joshua D. Angrist and Guido W. Imbens. Identification and estimation of local average treatment effects. Technical report, National Bureau of Economic Research, 1995.
- Joshua D. Angrist, Guido W. Imbens, and Donald B. Rubin. Identification of causal effects using instrumental variables. *Journal of the American Statistical Association*, 91(434):444–455, 1996.
- Susan Athey, Raj Chetty, Guido Imbens, and Hyunseung Kang. Estimating treatment effects using multiple surrogates: The role of the surrogate score and the surrogate index. *arXiv preprint arXiv:1603.09326*, 2016.
- Edvard Bakhitov and Amandeep Singh. Causal gradient boosting: Boosted instrumental variable regression. *arXiv preprint arXiv:2101.06078*, 2021.
- Elias Bareinboim and Judea Pearl. A general algorithm for deciding transportability of experimental results. *Journal of Causal Inference*, 1(1):107–134, 2013a.
- Elias Bareinboim and Judea Pearl. Meta-transportability of causal effects: A formal approach. In *Proceedings of the 16th International Conference on Artificial Intelligence and Statistics*, pages 135–143, 2013b.
- Elias Bareinboim and Judea Pearl. Causal inference and the data-fusion problem. *Proceedings of the National Academy of Sciences*, 113(27):7345–7352, 2016.
- Peter J Bickel, Ya’acov Ritov, Alexandre B Tsybakov, et al. Simultaneous analysis of lasso and dantzig selector. *The Annals of Statistics*, 37(4):1705–1732, 2009.
- Kenneth A. Bollen. *Structural equations with latent variables*. Wiley Series in Probability and Statistics, New Jersey, 1989.
- Leo Breiman. Prediction games and arcing algorithms. *Neural Computation*, 11(7):1493–1517, 1999.
- Leo Breiman and Jerome H. Friedman. Estimating optimal transformations for multiple regression and correlation. *Journal of the American Statistical Association*, 80(391):580–598, 1985.
- Leo Breiman et al. Arcing classifier (with discussion and a rejoinder by the author). *The Annals of Statistics*, 26(3):801–849, 1998.
- Peter Bühlmann and Sara Van De Geer. *Statistics for high-dimensional data: methods, theory and applications*. Springer Science & Business Media, 2011.
- Peter Bühlmann and Bin Yu. Boosting with the l_2 loss: regression and classification. *Journal of the American Statistical Association*, 98(462):324–339, 2003.
- Peter Bühlmann et al. Statistical significance in high-dimensional linear models. *Bernoulli*, 19(4):1212–1242, 2013.

- Emmanuel Candes, Terence Tao, et al. The dantzig selector: Statistical estimation when p is much larger than n . *The Annals of Statistics*, 35(6):2313–2351, 2007.
- Marine Carrasco, Jean-Pierre Florens, and Eric Renault. Linear inverse problems in structural econometrics estimation based on spectral decomposition and regularization. *Handbook of econometrics*, 6:5633–5751, 2007.
- Scott Shaobing Chen, David L. Donoho, and Michael A. Saunders. Atomic decomposition by basis pursuit. *SIAM review*, 43(1):129–159, 2001.
- Gregory Cooper. A simple constraint-based algorithm for efficiently mining observational databases for causal relationships. *Data Mining and Knowledge Discovery*, 1(2):203–224, 1997.
- Serge Darolles, Yanqin Fan, Jean-Pierre Florens, and Eric Renault. Nonparametric instrumental regression. *Econometrica*, 79(5):1541–1565, 2011.
- Daniel Eaton and Kevin Murphy. Exact bayesian structure learning from uncertain interventions. In *Proceedings of the 11th International Conference on Artificial Intelligence and Statistics*, pages 107–114, 2007.
- Frederick Eberhardt and Richard Scheines. Interventions and causal inference. *Philosophy of Science*, 74(5):981–995, 2007.
- Frederick Eberhardt, Patrik O. Hoyer, and Richard Scheines. Combining experiments to discover linear cyclic models with latent variables. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 185–192, 2010.
- Yoav Freund. Boosting a weak learning algorithm by majority. *Information and Computation*, 121(2):256–285, 1995.
- Yoav Freund and Robert E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55(1):119–139, 1997.
- Jerome Friedman, Trevor Hastie, Robert Tibshirani, et al. Additive logistic regression: a statistical view of boosting (with discussion and a rejoinder by the authors). *Annals of Statistics*, 28(2):337–407, 2000.
- Jerome Friedman, Trevor Hastie, Robert Tibshirani, et al. *The elements of statistical learning*, volume 1. Springer Series in Statistics, New York, 2001.
- Jerome H. Friedman. Greedy function approximation: a gradient boosting machine. *Annals of Statistics*, pages 1189–1232, 2001.
- Jerome H. Friedman, Trevor Hastie, and Robert Tibshirani. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441, 2008.

- Molly Gasperini, Andrew J Hill, José L McFaline-Figueroa, Beth Martin, Seungsoo Kim, Melissa D Zhang, Dana Jackson, Anh Leith, Jacob Schreiber, William S Noble, et al. A genome-wide framework for mapping gene regulation via cellular genetic screens. *Cell*, 176(1-2):377–390, 2019.
- Alastair R Hall. *Generalized method of moments*. Oxford University Press, 2005.
- Lars Peter Hansen. Large sample properties of generalized method of moments estimators. *Econometrica: Journal of the Econometric Society*, 50(4):1029–1054, 1982.
- Alain Hauser and Peter Bühlmann. Characterization and greedy learning of interventional markov equivalence classes of directed acyclic graphs. *Journal of Machine Learning Research*, 13(79):2409–2464, 2012.
- James Heckman. Varieties of selection bias. *The American Economic Review*, 80(2):313–318, 1990.
- Christina Heinze-Deml, Jonas Peters, and Nicolai Meinshausen. Invariant causal prediction for nonlinear models. *Journal of Causal Inference*, 6(2), 2018.
- Antti Hyttinen, Frederick Eberhardt, and Patrik O. Hoyer. Learning linear cyclic causal models with latent variables. *The Journal of Machine Learning Research*, 13(109):3387–3439, 2012.
- Antti Hyttinen, Frederick Eberhardt, and Matti Järvisalo. Constraint-based causal discovery: Conflict resolution with answer set programming. In *UAI*, pages 340–349, 2014.
- Martin Jinek, Krzysztof Chylinski, Ines Fonfara, Michael Hauer, Jennifer A. Doudna, and Emmanuelle Charpentier. A programmable dual-rna-guided dna endonuclease in adaptive bacterial immunity. *Science*, 337(6096):816–821, 2012.
- Nathan Kallus, Aahlad Manas Puli, and Uri Shalit. Removing hidden confounding by experimental grounding. *Advances in Neural Information Processing Systems*, 31:10888–10897, 2018.
- Eugene Katsevich and Kathryn Roeder. Conditional resampling improves sensitivity and specificity of single cell crispr regulatory screens. *bioRxiv preprint bioRxiv:2020.08.13.250092*, 2020.
- Sara Magliacane, Thijs Van Ommen, Tom Claassen, Stephan Bongers, Philip Versteeg, and Joris M Mooij. Domain adaptation by using causal inference to predict invariant conditional distributions. *Advances in neural information processing systems*, 31, 2018.
- Magne Mogstad, Alexander Torgovitsky, and Christopher Walters. Identification of causal effects with multiple instruments: Problems and some solutions. *NBER Working Paper*, 2019.
- Joris M. Mooij and Tom Heskes. Cyclic causal discovery from continuous equilibrium data. In *Proceedings of the Twenty-Ninth Conference on Uncertainty in Artificial Intelligence*, pages 431–439, 2013.

- Joris M. Mooij, Sara Magliacane, and Tom Claassen. Joint causal inference from multiple contexts. *Journal of Machine Learning Research*, 21(99):1–108, 2020.
- Whitney K. Newey and James L. Powell. Instrumental variable estimation of nonparametric models. *Econometrica*, 71(5):1565–1578, 2003.
- Judea Pearl and Elias Bareinboim. External validity: From do-calculus to transportability across populations. *Statistical Science*, pages 579–595, 2014.
- Judea Pearl and Rina Dechter. Identifying independencies in causal graphs with feedback. *arXiv preprint arXiv:1302.3595*, 2013.
- Judea Pearl et al. Causal inference in statistics: An overview. *Statistics Surveys*, 3:96–146, 2009.
- Jonas Peters, Peter Bühlmann, and Nicolai Meinshausen. Causal inference by using invariant prediction: identification and confidence intervals. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 78(5):947–1012, 2016.
- Niklas Pfister, Peter Bühlmann, and Jonas Peters. Invariant causal prediction for sequential data. *Journal of the American Statistical Association*, 114(527):1264–1276, 2019.
- Dominik Rothenhäusler, Peter Bühlmann, Nicolai Meinshausen, et al. Causal dantzig: fast inference in linear structural equation models with hidden variables under additive interventions. *The Annals of Statistics*, 47(3):1688–1722, 2019.
- Karen Sachs, Omar Perez, Dana Pe’er, Douglas Lauffenburger, and Garry Nolan. Causal protein-signaling networks derived from multiparameter single-cell data. *Science*, 308(5721):523–529, 2005.
- Robert E Schapire. The strength of weak learnability. *Machine Learning*, 5(2):197–227, 1990.
- Rahul Singh, Maneesh Sahani, and Arthur Gretton. Kernel instrumental variable regression. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- Peter Spirtes, Clark N. Glymour, Richard Scheines, and David Heckerman. *Causation, prediction, and search*. MIT Press, 2000.
- Jin Tian and Judea Pearl. Causal discovery from changes. In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence*, 2001.
- Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 58(1):267–288, 1996.
- Robert Tibshirani, Michael Saunders, Saharon Rosset, Ji Zhu, and Keith Knight. Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(1):91–108, 2005.
- Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge University Press, 2018.

Philip G. Wright. *Tariff on animal and vegetable oils*. Macmillan Company, New York, 1928.

Sewall Wright. Correlation and causation. *J. agric. Res.*, 20:557–580, 1921.

Shu Yang and Peng Ding. Combining multiple observational data sources to estimate causal effects. *Journal of the American Statistical Association*, 115(531):1540–1554, 2020.

Fei Ye and Cun-Hui Zhang. Rate minimaxity of the lasso and dantzig selector for the lq loss in lr balls. *Journal of Machine Learning Research*, 11(114):3519–3540, 2010.

Peng Zhao and Bin Yu. On model selection consistency of lasso. *The Journal of Machine Learning Research*, 7:2541–2563, 2006.

Hui Zou and Trevor Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320, 2005.

Appendix A. SEM Extension to Additive Shifts

In Section 3 we defined how new environments are generated by randomization of subsets of covariates. In addition to this intervention mechanism, we consider additive noise interventions where the distribution of the disturbance variable changes across environments by an additive shift that is independent of the base distribution. Assume that within environment e there is a set $\psi(e) \subset [p]$ of covariates that have an additive intervention. With respect to the base model \mathcal{M}^0 , the following structural equations are modified.

$$\mathcal{M}_{\psi(e)}^e : \begin{cases} X_j^e \leftarrow \sum_{k \in pa_0(j)} a_{jk} X_k^e + \epsilon_j + \delta_j^e & \forall j \in \psi(e) \\ Y^e \leftarrow \sum_{k \in [p]} \beta_k^0 X_k^e + \epsilon_Y \\ \{\epsilon_j\}_{j \in [p+1]} \perp\!\!\!\perp \{\delta_j^e\}_{j \in [p]} \end{cases} \quad (31)$$

where $\delta_j^e = 0$ for all $j \notin \psi(e)$. Let us give a justification for this model of environments with a practical example. Experimentation by randomization provides a very concrete way of perturbing a system, by direct manipulation of a covariate that modifies the structural mechanism that generates it. A concrete example of this is a gene knock-out experiment via CRISPR-Cas9: the targeted genes are no longer expressed, shifting the expression of other downstream genes in the regulation pathway. Additionally, environments may differ by some change in the overall environment. For example, different cell lines may have different baseline expressions of some subset of genes that shift the overall gene expression distribution. In practice, we want to assume that we know the subset of variables that are affected by this background shift, and that they precede in the causal mechanism any other variable that is experimentally manipulated.

Appendix B. High-Dimensional Aggregation Simulations: Pre-screening Before Collecting Experimental Data

Whenever the number of covariates is too large, obtaining orthogonality constraints for each covariate may become prohibitive. In practice we may be allowed to choose a subset

Table 6: Results of simulation from Appendix B. We estimate the regression coefficient on a subset of covariates that are pre-selected with a Lasso regression on an initial observational data set. We split the set of selected covariates into two, and build two experimental environments where covariates from each split are randomized in the corresponding environment. We then generate orthogonality constraints and run our aggregation procedure, as well as the pooled OLS. We additionally report the average number of selected covariates by Lasso.

Estimator	Sample Size	$n = 50$	$n = 100$	$n = 200$	$n = 500$	$n = 1000$
	Selected MB vars.	2.72 ± 0.26	2.98 ± 0.26	3.74 ± 0.16	4.70 ± 0.27	5.3 ± 0.12
Causal Aggregation	Coverage	0.99 ± 0.01	0.99 ± 0.01	0.97 ± 0.04	0.97 ± 0.05	0.95 ± 0.06
	Average length	136.6 ± 171.9	3.84 ± 2.04	0.51 ± 0.03	0.32 ± 0.01	0.12 ± 0.01
Pooled data OLS	Coverage	0.80 ± 0.11	0.81 ± 0.11	0.75 ± 0.12	0.72 ± 0.12	0.58 ± 0.14
	Average length	0.43 ± 0.04	0.31 ± 0.02	0.19 ± 0.02	0.13 ± 0.01	0.05 ± 0.01

of covariates to randomize based on an initial observational data set. Assuming that the connectivity matrix is sparse, we can run a pre-selection step to select a few variables, then based on an experimental environment construct the appropriate orthogonality constraints. We then assume the direct effects of the discarded covariates is 0, and thus we construct an estimator in the low-dimensional framework. However, to construct valid confidence intervals for such estimator based on pre-selected covariates we need to avoid using the same samples for model selection and inference. Our description above does indeed use two different sets of samples, where (easily accessible) observational data is used for covariate selection and then experimental data is obtained at a second stage. Otherwise we would split the data set to solve this issue.

Consider the following SEM where the covariate dimension is 200, where the connectivity matrix has a sparse structure given by the graph in Figure 13. The only covariate with a non-zero direct effect on Y is X_{99} . The values of the entries in the connectivity matrix are sampled from a Gaussian distribution centered in 1 and with variance 0.5. We run a Lasso regression (Tibshirani, 1996) on a standardized observational data set to select a subset of the covariates. Under some conditions on the coefficient sizes, as the sample size increases the selected set of covariates contains the the Markov blanket of Y (Bühlmann and Van De Geer, 2011, Section 2.5), which is equal to the subset $\{X_1, X_2, X_{99}, X_{100}, X_{199}, X_{200}\}$. In practice, this pre-selection step is adding a few random variables to the set of those that are randomized, our aggregation method works as long as the variables in the Markov blanket are selected and randomized in the experimental data. We report in our simulations the average number of Markov blanket covariates selected by Lasso (out of 6) as sample size increases.

Based on this selection step, we partition the selected covariates and for each subset we generate experimental datasets where those covariates are randomized. We do so by removing the dependence of the randomized covariates in their parents and assigning it a random standard Gaussian variable. We run the procedure for the just-identified setting based on the orthogonality constraints obtained from the experimental environments, as well as the OLS on the pooled data from the experimental environments. We run this procedure for different

sample sizes ($n \in \{100, 200, 500, 1000, 5000\}$) and report the confidence intervals coverage with nominal coverage 0.95 and average length. We report the results in Table 6. As the sample size increases, our aggregation procedure correctly estimates the sparse regression coefficient. For small sample sizes, the confidence intervals are meaningless as with few samples errors propagate in both the screening step and the aggregation step in the just-identified case, which is too imprecise with small sample sizes as previously seen. Pooled OLS coverage is well below the nominal value.




Figure 13: SEM for high-dimensional example: Observational samples are generated via the above SEM with 200 covariates, with a confounder affecting the response and two covariates.

Appendix C. Proofs

C.1 Proof of Proposition 1

Proof The linear constraint is a consequence of the independence between the residual term obtained $\epsilon_Y = Y - \mathbf{X}^T \beta^0$ and an exogenous variable. For an instrumental variable I or a randomized covariate X_j , independence arises by definition. This leads to constraints of the type:

$$\begin{aligned} \mathbb{E}[I(Y - \mathbf{X}^T \beta^0)] &= 0 \\ \mathbb{E}[X_j(Y - \mathbf{X}^T \beta^0)] &= 0 \end{aligned}$$

In a DAG, a variable is independent of its non-descendent nodes conditionally on its parental set (Pearl et al., 2009, Theorem 3.2.2). Our constraint is based on conditioning on the parental set of a covariate X_j . This parental set corresponds is obtained from the DAG \bar{G}^0 , where the observational distribution factorizes. However, in practice we can not condition on the unobserved variables, hence we must assume that the parental set of X_j in \bar{G}^0 is the same as the parental set in G^0 . When adjusting X_j for the parental set we get a random variable given by $X_j - \sum_{k \in pa_0(j)} a_{jk} X_k = \epsilon_j$ as we assumed $c_{jk} = 0$: i.e. the latent factors do not affect X_j . Additionally, given that the graph has no cycles whenever we know that the response Y is in the parental set of X_j we immediately get as a constraint the fact that the regression coefficient corresponding to X_j is 0. Therefore from now on we assume that Y is not in the parental set. The residual is not represented in the graph on its own, but we can derive the orthogonality constraint as follows. Replacing Y by $(Y - \mathbf{X}^T \beta^0) = \sum_{k \in [p]} d_k H_k + \epsilon_Y$, we get that this term is independent of the adjusted term $X_j - \hat{X}_j$ where $\hat{X}_j = \sum_{k \in pa_0(j)} a_{jk} X_k$.

Hence we obtain the result. Finally, the last constraint derived from additive interventions (called *inner-product invariance*) is proved in Rothenhäusler et al. (2019, Proposition 1). ■

C.2 Proof of Proposition 2

Proof We can assume without loss of generality that the ordering $\{1, \dots, p\}$ is a topological ordering for the DAG G^0 . We stack constraint vectors that form \mathbf{G} following the same topological ordering in the constraint related variable. We will now show that stacking the vertical vectors in such ordering leads to an upper triangular matrix that is invertible. The idea is that each constraint is related to one covariate, and we show that the corresponding constraint inducing variable R is independent of the previous covariates in the topological ordering. Consider the j -th vector in \mathbf{G} . In the IV setting, the instrument is independent of all the non-descendant variables of X_j in the \bar{G}^0 graph. Therefore the constraint vector has its first $j - 1$ entries equal to 0. Whenever randomizing X_j , that variable is now independent of all its non-descendants in the \bar{G}^0 graph. Adjusting for direct causes (i.e. conditioning on the parental set), assuming there is no latent variable effect, we get that the residual after adjusting X_j given by $R = X_j - \hat{X}_j$ is independent of the non-descendants of X_j . Again, the constraint vector has its first $j - 1$ entries equal to 0. ■

C.3 Proof of Proposition 3

Proof For notation simplicity, we associate $\mathcal{C} = [p]$. For $c \in [p]$, let $u_c \in \mathbb{R}^p$, $v_c \in \mathbb{R}$, and let $u = (u_c)_{c \in [p]}$, $v = (v_c)_{c \in [p]}$. Consider the mapping

$$\xi : \begin{cases} \mathbb{R}^{(p+1) \times p} \rightarrow \mathbb{R}^p \\ (u, v) \mapsto \mathbf{G}^{-1}(u)\mathbf{Z}(v) \end{cases}$$

where

$$\mathbf{G}(u) = [u_c^T]_{1 \leq c \leq p} := \begin{pmatrix} u_1^T \\ \vdots \\ u_p^T \end{pmatrix} \in \mathbb{R}^{p \times p}, \quad \mathbf{Z}(v) = \begin{pmatrix} v_1 \\ \vdots \\ v_p \end{pmatrix} \in \mathbb{R}^p$$

where the notation $[u_c^T]_{1 \leq c \leq p}$ denotes the matrix obtained by stacking the row vectors u_j^T . If we consider the restriction of ξ to those vectors (u, v) such that the matrix $\mathbf{G}(u)$ is invertible, then we have that ξ is continuously differentiable at (u, v) and its derivative is given by

$$d\xi_{u,v}(\tilde{u}, \tilde{v}) = -\mathbf{G}(u)^{-1}\mathbf{G}(\tilde{u})\mathbf{G}(u)^{-1}\mathbf{Z}(v) + \mathbf{G}(u)^{-1}\mathbf{Z}(\tilde{v})$$

Therefore, via a Taylor approximation, we have that

$$\begin{aligned}
 & \xi(u', v') - \xi(u, v) - d\xi_{u,v}(u' - u, v' - v) \\
 &= \xi(u', v') - \xi(u, v) + \mathbf{G}(u)^{-1}(\mathbf{G}(u') - \mathbf{G}(u))\mathbf{G}(u)^{-1}\mathbf{Z}(v) - \mathbf{G}(u)^{-1}(\mathbf{Z}(v') - \mathbf{Z}(v)) \\
 &= \xi(u', v') - \xi(u, v) + \mathbf{G}(u)^{-1}(C(u', v') - C(u, v)) \begin{pmatrix} \mathbf{G}(u)^{-1}\mathbf{Z}(v) \\ -1 \end{pmatrix} \\
 &= o_{\substack{u' \rightarrow u \\ v' \rightarrow v}}(\|(u', v') - (u, v)\|)
 \end{aligned}$$

where

$$C(u, v) = [u_c^T, v_c]_{1 \leq c \leq p} \in \mathbb{R}^{p \times (p+1)}$$

We collect n_e samples $(\mathbf{X}_i^e, Y_i^e)_{i \in [n_e]}$ in environment e , where \mathbf{X}_i^e is the i -th vector sample of the covariates $\mathbf{X}_i^e = (X_{1,i}^e, \dots, X_{p,i}^e)$. Each constraint $c \in [p]$ is based on samples from environment $e_c \in \mathcal{E}$: for each $c \in [p]$ we collect a constraint inducing variable $(R_i^c)_{i \in [n_{e_c}]}$. Such R^c may correspond to $X_c^{e_c}$ if such covariate is randomized, or an instrument in environment e_c for covariate $X_c^{e_c}$. R^c can also be the residual variable when adjusting for the parental set of a given covariate (although as indicated in Section 4.1.3, regressing on the parental set and estimating the orthogonality constraint must be done with distinct datasets). Recall that $n = \sum_{e \in \mathcal{E}} n_e$ is the total number of samples, and that the sample sizes from different environments grow at the same rate: $\frac{n_e}{n} \rightarrow \rho_e \in (0, 1)$. The vector β^0 is characterized as the solution to the system of equations:

$$\mathbf{Z}(v) = \mathbf{G}(u)\beta \iff \begin{cases} \mathbb{E}[R^{e_1}(Y^{e_1} - \beta^{0,T} \mathbf{X}^{e_1})] = 0 \\ \dots \\ \mathbb{E}[R^{e_p}(Y^{e_p} - \beta^{0,T} \mathbf{X}^{e_p})] = 0 \end{cases}$$

so that $\beta^0 = \xi(u, v)$ where $u = (\mathbb{E}[R^c \mathbf{X}^{e_c}])_{c \in [p]}$, and $v = (\mathbb{E}[R^c Y^{e_c}])_{c \in [p]}$. Analogously, the estimator $\hat{\beta}$ is given by plugging in the previous equation the sample averages: $\hat{\beta} = \xi(\hat{u}, \hat{v})$ where

$$\begin{cases} \hat{u} = ((1/n_{e_c}) \sum_{i \in [n_{e_c}]} R_i^c \mathbf{X}_i^{e_c})_{c \in [p]} \\ \hat{v} = ((1/n_{e_c}) \sum_{i \in [n_{e_c}]} R_i^c Y_i^{e_c})_{c \in [p]} \end{cases} \quad (32)$$

We have by the strong law of large numbers that $(\hat{u}, \hat{v}) - (u, v) = O_P(1/\sqrt{n})$. Therefore we get that

$$\begin{aligned}
 & \xi(\hat{u}, \hat{v}) - \xi(u, v) + \mathbf{G}(u)^{-1}(C(\hat{u}, \hat{v}) - C(u, v)) \begin{pmatrix} \xi(u, v) \\ -1 \end{pmatrix} \\
 &= \hat{\beta} - \beta^0 + \mathbf{G}(u)^{-1}(C(\hat{u}, \hat{v}) - C(u, v)) \begin{pmatrix} \beta^0 \\ -1 \end{pmatrix} \\
 &= o_P(\|(\hat{u}, \hat{v}) - (u, v)\|)
 \end{aligned}$$

and thus

$$\sqrt{n}(\hat{\beta} - \beta^0) + \mathbf{G}(u)^{-1}\sqrt{n}(C(\hat{u}, \hat{v}) - C(u, v)) \begin{pmatrix} \beta^0 \\ -1 \end{pmatrix} = o_P(O_P(1)) = o_P(1)$$

Let $\mathbf{U}_i^c := R_i^c(\mathbf{X}_i^{e_c}, Y_i^{e_c}) \in \mathbb{R}^{p+1}$ for $c \in [p]$. We have the following convergence in distribution by the central limit theorem in environment e by combining all the constraints c based on environment e :

$$\sqrt{n_e} \text{Vect} \left[\left(\frac{1}{n_e} \sum_{i=1}^{n_e} \mathbf{U}_i^c - \mathbb{E}[\mathbf{U}^c] \right)^T \right]_{e_c=e} \xrightarrow[n \rightarrow +\infty]{d} \mathcal{N} \left(\mathbf{0}; [\text{Cov}(\mathbf{U}^c, \mathbf{U}^{\tilde{c}})]_{\substack{e_c=e \\ e_{\tilde{c}}=e}} \right)$$

Now, given the assumptions on the environment e , we have that

$$\begin{aligned} (\mathbf{U}^c)^T \begin{pmatrix} \beta^0 \\ -1 \end{pmatrix} &= R^c(\beta^{0,T} \mathbf{X}^e - Y^e) \\ &= -R^c \epsilon_Y^e \end{aligned}$$

As R^c is an instrument, a randomized covariate, or the residual from regressing the covariate on its parental set, we have R^c and ϵ_Y^e are independent as indicated in Proposition 1. Therefore, as ϵ_Y^e is centered, we get:

$$\begin{pmatrix} \beta^0 \\ -1 \end{pmatrix}^T \text{Cov}(\mathbf{U}^c, \mathbf{U}^{\tilde{c}}) \begin{pmatrix} \beta^0 \\ -1 \end{pmatrix} = \text{Cov}(R^c \epsilon_Y^e, R^{\tilde{c}} \epsilon_Y^e) = \sigma_e^2 \text{Cov}(R^c, R^{\tilde{c}})$$

where $\sigma_e^2 := \mathbb{E}[\epsilon_Y^{e,2}]$. We have by Slutsky's theorem, given that $\frac{n_e}{n} \rightarrow \rho_e$,

$$\sqrt{n} \left[\left(\frac{1}{n_e} \sum_{i=1}^{n_e} \mathbf{U}_i^c - \mathbb{E}[\mathbf{U}^c] \right)^T \begin{pmatrix} \beta^0 \\ -1 \end{pmatrix} \right]_{e_c=e} \xrightarrow[n \rightarrow +\infty]{d} \mathcal{N} \left(\mathbf{0}; \frac{\sigma_e^2}{\rho_e} \text{Cov}(R^c)_{e_c=e} \right)$$

Also, the covariance matrix $\text{Cov}(R^c)_{e_c=e}$ is diagonal and invertible, as within environment e constraint inducing variables that are either randomized covariates or instruments are jointly independent. We now concatenate the results for different environments. Given the independence of samples across environments, we have that:

$$\sqrt{n} (C(\hat{u}, \hat{v}) - C(u, v)) \begin{pmatrix} \beta^0 \\ -1 \end{pmatrix} \xrightarrow[n \rightarrow +\infty]{d} \mathcal{N} \left(\mathbf{0}; \text{Cov} \left(\frac{\sigma_{e_c}}{\sqrt{\rho_{e_c}}} R^c \right)_{c \in [p]} \right)$$

where the covariance matrix $\text{Cov} \left(\frac{\sigma_{e_c}}{\sqrt{\rho_{e_c}}} R^c \right)_{c \in [p]}$ is still diagonal. We conclude:

$$\sqrt{n} (\hat{\beta} - \beta^0) \xrightarrow[n \rightarrow +\infty]{d} \mathcal{N} \left(\mathbf{0}; \Sigma \right)$$

where $\Sigma = \mathbf{G}^{-1}(u) \text{Diag} \left(\frac{\sigma_{e_c}^2}{\rho_{e_c}} \text{Var}(R^c) \right)_{c \in [p]} \mathbf{G}^{-1,T}(u)$. ■

C.4 Proof of Proposition 5

Proof

Letting $z_\infty := \|\hat{\mathbf{Z}} - \hat{\mathbf{G}}\beta^0\|_\infty$, we show that in the event $\{z_\infty \leq \lambda\}$ our regularized estimator $\hat{\beta}(\lambda)$ satisfies

$$\|\hat{\beta}(\lambda) - \beta^0\|_q \leq \frac{|S^0|^{1/q} \|\hat{\mathbf{G}}(\hat{\beta}(\lambda) - \beta^0)\|_\infty}{\text{CIF}_q(S^0, \hat{\mathbf{G}})} \leq \frac{2|S^0|^{1/q} \lambda}{\text{CIF}_q(S^0, \hat{\mathbf{G}})} \quad (33)$$

We follow Ye and Zhang (2010) and show the following inequality:

$$\begin{aligned}
 \|\hat{\beta}(\lambda)_{(S^0)^c} - \beta_{(S^0)^c}^0\|_1 &= \|\hat{\beta}(\lambda)_{(S^0)^c}\|_1 \\
 &= \|\hat{\beta}(\lambda)\|_1 - \|\hat{\beta}(\lambda)_{S^0}\|_1 \\
 &\leq \|\beta^0\|_1 - \|\hat{\beta}(\lambda)_{S^0}\|_1 \\
 &\leq \|\beta_{S^0}^0\|_1 - \|\hat{\beta}(\lambda)_{S^0}\|_1 \\
 &\leq \|\beta_{S^0}^0 - \hat{\beta}(\lambda)_{S^0}\|_1
 \end{aligned}$$

where we used the fact that S^0 is the support of β^0 and that in the event $\{z_\infty \leq \lambda\}$ the true vector β^0 is feasible, therefore we get $\|\hat{\beta}(\lambda)\|_1 \leq \|\beta^0\|_1$. If the upper bound is equal to 0, then we get $\hat{\beta}(\lambda) = \beta^0$ as these vectors coincide over S^0 and $(S^0)^c$ and the inequality above holds. Otherwise, this shows that $\hat{\beta}(\lambda) - \beta^0$ belongs to the cone $\mathfrak{C} := \{\mathbf{u} : \|\mathbf{u}_{(S^0)^c}\|_1 \leq \|\mathbf{u}_{S^0}\|_1 \neq 0\}$ and we get the first inequality in equation (33) by definition of the CIF. Furthermore, we show that $\|\hat{\mathbf{G}}(\hat{\beta}(\lambda) - \beta^0)\|_\infty \leq 2\lambda$ in the event $\{z_\infty \leq \lambda\}$:

$$\|\hat{\mathbf{G}}(\hat{\beta}(\lambda) - \beta^0)\|_\infty = \|(\hat{\mathbf{Z}} - \hat{\mathbf{G}}\beta^0) - (\hat{\mathbf{Z}} - \hat{\mathbf{G}}\hat{\beta}(\lambda))\|_\infty \leq 2\lambda$$

as $\beta^0, \hat{\beta}(\lambda)$ are in the feasible set.

We need a high probability bound for the event $\{z_\infty \leq \lambda\}$, and we also need to control the CIF value for the empirical matrix $\hat{\mathbf{G}}$ in equation (33), which entails using concentration inequalities to control the deviation of $\hat{\mathbf{G}}$ from \mathbf{G} . Lemma 3 in Rothenhäusler et al. (2019) provides the following bound for the gap between the CIF under the estimator matrix $\hat{\mathbf{G}}$ and the CIF value under \mathbf{G} :

$$|\text{CIF}_q(S^0, \mathbf{G}) - \text{CIF}_q(S^0, \hat{\mathbf{G}})| \leq 2|S^0| \|\mathbf{G} - \hat{\mathbf{G}}\|_\infty$$

therefore in the event $\{2|S^0| \|\mathbf{G} - \hat{\mathbf{G}}\|_\infty \leq \frac{1}{2} \text{CIF}_q(S^0, \mathbf{G})\}$ the following upper bound holds:

$$\|\hat{\beta}(\lambda) - \beta^0\|_q \leq \frac{4|S^0|^{1/q} \lambda}{\text{CIF}_q(S^0, \mathbf{G})} \quad (34)$$

We now apply lemma 8 to obtain a bound with high probability for $\|\mathbf{G} - \hat{\mathbf{G}}\|_\infty$ and the event $\{z_\infty \leq \lambda\}$. Let $t = 2 \log p$ and set

$$\lambda := k\sigma_C\sigma_E \sqrt{\frac{t + \log p}{\min_{e \in \mathcal{E}} n_e}} = k\sqrt{3}\sigma_C\sigma_E \sqrt{\frac{\log p}{\min_{e \in \mathcal{E}} n_e}} \rightarrow 0$$

where the universal constant k is defined in the lemma. Given that choice of λ we get by lemma 8:

$$\mathbb{P}(z_\infty \leq \lambda) \geq 1 - \frac{2}{p^2} \quad \text{and} \quad \mathbb{P}(\|\hat{\mathbf{G}} - \mathbf{G}\|_\infty \leq \frac{\sigma_X}{\sigma_E} \lambda) \geq 1 - \frac{2}{p}$$

In addition, as by assumption we have $\frac{1}{\text{CIF}_q(S^0, \mathbf{G})} \sqrt{\frac{\log p}{\min_{e \in \mathcal{E}} n_e}} \rightarrow 0$, we get that eventually

$$\begin{aligned}
 \sqrt{\frac{\log p}{\min_{e \in \mathcal{E}} n_e}} &\leq \frac{\text{CIF}_q(S^0, \mathbf{G})}{4\sqrt{3}k|S^0|\sigma_X\sigma_C} \\
 \frac{\sigma_X}{\sigma_E} \lambda &\leq \frac{\text{CIF}_q(S^0, \mathbf{G})}{4|S^0|}
 \end{aligned}$$

The high probability bounds above then control the two events leading to inequality (33). Therefore with probability at least $1 - \frac{4}{p}$, we have that

$$\begin{aligned} \|\hat{\beta}(\lambda) - \beta^0\|_q &\leq \frac{4|S^0|^{1/q}\lambda}{\text{CIF}_q(S^0, \mathbf{G})} \\ &\leq \frac{4\sqrt{3}k\sigma_C\sigma_E|S^0|^{1/q}}{\text{CIF}_q(S^0, \mathbf{G})} \sqrt{\frac{\log p}{\min_{e \in \mathcal{E}} n_e}} \end{aligned}$$

hence the result. \blacksquare

Lemma 8 *Assume that X_j^e are σ_X^2 sub-gaussian, ϵ_Y^e are σ_E^2 sub-gaussian, and that R^c are σ_C^2 sub-gaussian for all $e \in \mathcal{E}, j \in [p], c \in \mathcal{C}$ and some fixed $\sigma_X^2, \sigma_E^2, \sigma_C^2 > 0$. There exists a universal constant $k > 0$, such that for any $t > 0$:*

$$\begin{aligned} \mathbb{P}\left(z_\infty \leq k\sigma_C\sigma_E \max\left(\frac{t + \log p}{\min_{e \in \mathcal{E}} n_e}, \sqrt{\frac{t + \log p}{\min_{e \in \mathcal{E}} n_e}}\right)\right) &\geq 1 - 2e^{-t} \\ \mathbb{P}\left(\|\hat{\mathbf{G}} - \mathbf{G}\|_\infty \leq k\sigma_C\sigma_X \max\left(\frac{t + 2 \log p}{\min_{e \in \mathcal{E}} n_e}, \sqrt{\frac{t + 2 \log p}{\min_{e \in \mathcal{E}} n_e}}\right)\right) &\geq 1 - 2e^{-t} \end{aligned}$$

Proof We will prove the result by relying on concentration inequalities for sub-gaussian and sub-exponential random variables. We use Orlicz spaces and norms since this allows us to bound the products of quantities easily, by invoking inequalities that we discuss in the following. We refer to Vershynin (2018) for further details on the use of Orlicz spaces in concentration inequalities. We also introduce universal constants, finite positive real numbers that do not depend on the other elements in the problem. The Orlicz norm of a random variable X with respect to an Orlicz function ψ is defined as

$$\|X\|_\psi := \inf \left\{ t > 0 : \mathbb{E}\left[\psi\left(\frac{X}{t}\right)\right] \leq 1 \right\}$$

The Orlicz space with respect to ψ is the space of random variables with finite Orlicz norm. Given the choices of ψ_1, ψ_2 defined below, we get that the corresponding Orlicz spaces are the families of sub-gaussian and sub-exponential random variables respectively.

$$\psi_1 : x \mapsto e^x - 1 \quad \psi_2 : x \mapsto e^{x^2} - 1$$

These two spaces are connected by the following inequality that applies for any variables X, Y :

$$\|XY\|_{\psi_1} \leq \|X\|_{\psi_2} \|Y\|_{\psi_2}$$

Hence the products of random variables we have are sub-exponential as products of sub-gaussian random variables. The following inequalities hold for an universal constant k_0 that does not depend on the random variables.

$$\begin{aligned} \|R^c \epsilon^{ec}\|_{\psi_1} &\leq \|R^c\|_{\psi_2} \|\epsilon^{ec}\|_{\psi_2} \leq k_0 \sigma_C \sigma_E \\ \|R^c X_j^{ec} - \mathbb{E}[R^c X_j^{ec}]\|_{\psi_1} &\leq \|R^c\|_{\psi_2} \|X_j^{ec}\|_{\psi_2} \leq k_0 \sigma_C \sigma_X \end{aligned}$$

where we relied on the fact that there exists a universal constant k_{00} , independent of the choice of the random variable, such that for any X we have $\|X - \mathbb{E}[X]\|_{\psi_1} \leq k_{00}\|X\|_{\psi_1}$, and, for a σ^2 sub-gaussian random variable, there is another universal constant k_{01} such that $\|X\|_{\psi_2} \leq k_{01}\sigma$. We apply Bernstein's inequality to the product $R^c \epsilon^{e_c}$ (cf. Theorem 2.8.1 in Vershynin (2018)):

$$\begin{aligned} \mathbb{P}\left(\left|\frac{1}{n_{e_c}} \sum_{i \in [n_{e_c}]} R_i^c \epsilon_i^{e_c}\right| \geq t\right) &\leq 2 \exp\left(-k_1 \min\left(\frac{(n_{e_c} t)^2}{n_{e_c} \|R^c \epsilon^{e_c}\|_{\psi_1}^2}, \frac{n_{e_c} t}{\|R^c \epsilon^{e_c}\|_{\psi_1}}\right)\right) \\ &\leq 2 \exp\left(-k_1 \min\left(\frac{n_{e_c} t^2}{k_0^2 \sigma_C^2 \sigma_E^2}, \frac{n_{e_c} t}{k_0 \sigma_C \sigma_E}\right)\right) \\ &\leq 2 \exp\left(-k_1 n_{e_c} \min\left(\frac{1}{k_0^2}, \frac{1}{k_0}\right) \min\left(\frac{t^2}{\sigma_C^2 \sigma_E^2}, \frac{t}{\sigma_C \sigma_E}\right)\right) \\ &\leq 2 \exp\left(-k_2 n_{e_c} \min\left(\frac{t^2}{\sigma_C^2 \sigma_E^2}, \frac{t}{\sigma_C \sigma_E}\right)\right) \end{aligned}$$

Therefore we get by inverting the term in the exponential bound:

$$\begin{aligned} 2 \exp(-t) &\geq \mathbb{P}\left(\left|\frac{1}{n_{e_c}} \sum_{i \in [n_{e_c}]} R_i^c \epsilon_i^{e_c}\right| \geq \sigma_C \sigma_E \max\left(\frac{t}{n_{e_c} k_2}, \sqrt{\frac{t}{n_{e_c} k_2}}\right)\right) \\ &\geq \mathbb{P}\left(\left|\frac{1}{n_{e_c}} \sum_{i \in [n_{e_c}]} R_i^c \epsilon_i^{e_c}\right| \geq \sigma_C \sigma_E \max\left(\frac{t}{n_{e_c}}, \sqrt{\frac{t}{n_{e_c}}}\right) \max\left(\frac{1}{k_2}, \sqrt{\frac{1}{k_2}}\right)\right) \\ &= \mathbb{P}\left(\left|\frac{1}{n_{e_c}} \sum_{i \in [n_{e_c}]} R_i^c \epsilon_i^{e_c}\right| \geq k_3 \sigma_C \sigma_E \max\left(\frac{t}{n_{e_c}}, \sqrt{\frac{t}{n_{e_c}}}\right)\right) \end{aligned}$$

where k_0, k_1, k_2, k_3 are universal constants. Analogously for the product $R^c X_j^{e_c}$ we get that:

$$\mathbb{P}\left(\left|\frac{1}{n_{e_c}} \sum_{i \in [n_{e_c}]} R_i^c X_{j,i}^{e_c} - \mathbb{E}[R^c X_j^{e_c}]\right| \geq k_3 \sigma_C \sigma_X \max\left(\frac{t}{n_{e_c}}, \sqrt{\frac{t}{n_{e_c}}}\right)\right) \leq 2 \exp(-t)$$

We conclude in both cases by applying an union bound. For the term z_∞ we have:

$$\begin{aligned} &\mathbb{P}\left(z_\infty \geq k_3 \sigma_C \sigma_E \max\left(\frac{t + \log p}{\min_{e \in \mathcal{E}} n_e}, \sqrt{\frac{t + \log p}{\min_{e \in \mathcal{E}} n_e}}\right)\right) \\ &= \mathbb{P}\left(\bigcup_{c \in [p]} \left\{\left|\frac{1}{n_{e_c}} \sum_{i \in [n_{e_c}]} R_i^c \epsilon_i^{e_c}\right| \geq k_3 \sigma_C \sigma_E \max\left(\frac{t + \log p}{\min_{e \in \mathcal{E}} n_e}, \sqrt{\frac{t + \log p}{\min_{e \in \mathcal{E}} n_e}}\right)\right\}\right) \\ &\leq \sum_{c \in [p]} \mathbb{P}\left(\left|\frac{1}{n_{e_c}} \sum_{i \in [n_{e_c}]} R_i^c \epsilon_i^{e_c}\right| \geq k_3 \sigma_C \sigma_E \max\left(\frac{t + \log p}{n_{e_c}}, \sqrt{\frac{t + \log p}{n_{e_c}}}\right)\right) \\ &\leq 2p \exp(-(t + \log p)) \\ &\leq 2 \exp(-t) \end{aligned}$$

For the term $\|\hat{\mathbf{G}} - \mathbf{G}\|_\infty$ we have:

$$\begin{aligned}
 & \mathbb{P}\left(\|\hat{\mathbf{G}} - \mathbf{G}\|_\infty \geq k_3 \sigma_C \sigma_X \max\left(\frac{t + 2 \log p}{\min_{e \in \mathcal{E}} n_e}, \sqrt{\frac{t + 2 \log p}{\min_{e \in \mathcal{E}} n_e}}\right)\right) \\
 &= \mathbb{P}\left(\bigcup_{\substack{c \in \mathcal{C} \\ j \in [p]}} \left\{ \left| \frac{1}{n_{e_c}} \sum_{i \in [n_{e_c}]} R_i^c X_{j,i}^{e_c} - \mathbb{E}[R^c X_j^{e_c}] \right| \geq k_3 \sigma_C \sigma_X \max\left(\frac{t + 2 \log p}{\min_{e \in \mathcal{E}} n_e}, \sqrt{\frac{t + 2 \log p}{\min_{e \in \mathcal{E}} n_e}}\right) \right\}\right) \\
 &\leq \sum_{\substack{c \in \mathcal{C} \\ j \in [p]}} \mathbb{P}\left(\left| \frac{1}{n_{e_c}} \sum_{i \in [n_{e_c}]} R_i^c X_{j,i}^{e_c} - \mathbb{E}[R^c X_j^{e_c}] \right| \geq k_3 \sigma_C \sigma_X \max\left(\frac{t + 2 \log p}{n_{e_c}}, \sqrt{\frac{t + 2 \log p}{n_{e_c}}}\right)\right) \\
 &\leq 2p |\mathcal{C}| \exp(- (t + 2 \log p)) \\
 &\leq 2 \frac{|\mathcal{C}|}{p} \exp(-t) \leq 2 \exp(-t)
 \end{aligned}$$

Therefore we get the result. ■

C.5 Proof of Proposition 6

Proof For completeness, we adapt the proof in Rothenhäusler et al. (2019) to show that

$$\lim \mathbb{P}\left(\min_{j \in S^0} |\hat{\beta}(\lambda)_j| > 0\right) \rightarrow 1.$$

In the event

$$\|\hat{\beta}(\lambda) - \beta^0\|_\infty \leq \frac{K \sigma_C \sigma_E}{\text{CIF}_\infty(S^0, \mathbf{G})} \sqrt{\frac{\log p}{\min_{e \in \mathcal{E}} n_e}}$$

the beta-min condition implies

$$\begin{aligned}
 0 &< \min_{j \in S^0} |\beta_j^0| - \frac{K \sigma_C \sigma_E}{\text{CIF}_\infty(S^0, \mathbf{G})} \sqrt{\frac{\log p}{\min_{e \in \mathcal{E}} n_e}} \\
 &\leq \min_{j \in S^0} |\hat{\beta}(\lambda)_j|.
 \end{aligned}$$

This completes the proof. ■

C.6 Proof of Proposition 7

Our proof is based on two key properties of the graphical structure of the interventional DAGs $(\bar{G}^e)_e$: they all share a same topological ordering—the topological ordering from \bar{G}^0 still holds when intervening on covariates—and the nodes of randomized covariates $\phi(e)$ in environment e have no incoming edges by definition of randomization. These two properties are independent of the response node, and therefore to simplify our proofs we consider a different graph structure by marginalizing out the response node as follows. Given the

causal models \mathcal{M}^0 and $\mathcal{M}_{\phi(e)}^e$, we define the marginal distributions of $\mathbb{P}^0, \mathbb{P}^e$ over \mathbf{X} , denoted $\mathbb{P}_X^0, \mathbb{P}_X^e$. These factorize in DAGs $\bar{G}_X^0 = ([p], \bar{E}_X^0)$ and $\bar{G}_X^e = ([p], \bar{E}_X^e)$ where the set of nodes $[p]$ represents covariate nodes. The edges in \bar{E}_X^0, \bar{E}_X^e are the same as in \bar{E}^0, \bar{E}^e for those that are not connecting Y to another node. Previous edges in (\bar{E}^0, \bar{E}^e) that connected Y to other nodes are replaced by edges in \bar{E}_X^0, \bar{E}_X^e between covariates that connect in \bar{G}^0, \bar{G}^e every parent of Y to every child of Y . Finally, graphs \bar{G}_X^e share the same previously mentioned two properties as \bar{G}^e : they all share a same topological ordering given by \bar{G}_X^0 and nodes of randomized covariates $\phi(e)$ in environment e have no incoming edges in \bar{G}_X^e .

Assume without loss of generality that the order $[p] = \{1, \dots, p\}$ is a topological order of \bar{G}_X^0 . We consider the reversed lexicographical order between two subsets $A, B \subset [p]$, $A = \{a_1, \dots, a_{n_A}\}$ and $B = \{b_1, \dots, b_{n_B}\}$, where $a_1 > a_2 > \dots > a_{n_A}$ (and analogously for B), as the total order given by

$$A \succ B \iff (\exists i | (a_j = b_j \forall j < i) \text{ and } (a_i > b_i)) \text{ or } (n_A > n_B \text{ and } \{a_1, \dots, a_{n_B}\} = B) \quad (35)$$

For a set of environments \mathcal{E} we define the set $\phi(\mathcal{E}) := \{\phi(e) \subset [p], e \in \mathcal{E}\}$ of subsets $\phi(e)$ of indices in $[p]$ that index the variables intervened in environment $e \in \mathcal{E}$. Such set has a maximal element $\max(\phi(\mathcal{E}))$. For two sets of environments $\mathcal{E}_1, \mathcal{E}_2$, we define an order $\mathcal{E}_1 \succeq \mathcal{E}_2$ by comparing their maximal elements: $\max(\phi(\mathcal{E}_1)) \succeq \max(\phi(\mathcal{E}_2))$. We now prove Proposition 7.

Proof

Uniqueness We first prove the uniqueness statement. Assume that there are two functions \bar{f}_1 and \bar{f}_2 that follow the decomposition given in equation (25) such that the orthogonality conditions given by equation (26) hold and consider the difference $f := \bar{f}_1 - \bar{f}_2$. By linearity f follows the decomposition given in equation (25). Additionally, for all $e \in \mathcal{E}$, for all square-integrable $h : \mathbb{R}^{|\phi(e)|} \rightarrow \mathbb{R}$, by subtracting the two orthogonality constraints for \bar{f}_1 and \bar{f}_2 given by equation (26), we have

$$0 = \mathbb{E}^e[h(\mathbf{X}_{\phi(e)})f(\mathbf{X})] \quad (36)$$

Applying Proposition 9 from the Appendix we get that $f = 0$ over the support of the random variables which by assumption is the same across environments, hence the uniqueness of \bar{f} .

Identifiability We finally show that f^0 satisfies equation (26) whenever $f^0 \in \mathcal{F}_{\mathcal{E}}$. We have that the structural equation defining the response variable is given by $Y = f^0(\mathbf{X}_{S^0}) + \epsilon_Y$. Then, for all $e \in \mathcal{E}$, we have that the residual ϵ_Y in environment e is independent of the intervened variables $\mathbf{X}_{\phi(e)}$ under \mathbb{P}^e . Therefore, given that the residuals are centered random variables, we get that

$$\mathbb{E}^e[h(\mathbf{X}_{\phi(e)})(Y - f^0(\mathbf{X}_{S^0}))] = 0 \quad (37)$$

This implies that $f^0 = \bar{f}$ whenever $f^0 \in \mathcal{F}_{\mathcal{E}}$. ■

Proposition 9 *Assume the model of environments defined in (23), let $f \in \mathcal{F}_{\mathcal{E}}$, i.e. it can be decomposed as follows:*

$$f(\mathbf{x}) = \sum_{e \in \mathcal{E}} f_e(\mathbf{x}_{\phi(e)}) \quad (38)$$

for some functions square-integrable $(f_e)_{e \in \mathcal{E}}$ defined over the subsets of covariates indexed by $\phi(e)$. Assume that all \mathbb{P}^e have the same support. If for all $e \in \mathcal{E}$, for all $h : \mathbb{R}^{|\phi(e)|} \rightarrow \mathbb{R}$ square-integrable we have

$$\mathbb{E}^e[h(\mathbf{X}_{\phi(e)})f(\mathbf{X})] = 0 \quad (39)$$

then $f = 0$.

Proof We prove the result by recursion: we show that, if f satisfies the conditions of the lemma for \mathcal{E} , then there exists another set of environments $\tilde{\mathcal{E}}$, such that $\mathcal{E} \succ \tilde{\mathcal{E}}$ (strictly), and f satisfies conditions (38) and (39) with $\tilde{\mathcal{E}}$. We recursively show that f must satisfy a decomposition of the type (38) with increasingly fewer variables and interactions. Given that the set of sets of subsets of $[p]$ is finite, after a finite number of steps we get that f must satisfy the decomposition (38) for $\mathcal{E} = \emptyset$ (i.e. f is constant). We then conclude that $f = 0$ using condition (39) with $h = 1$.

Consider the maximal element $\phi_0 := \max(\phi(\mathcal{E}))$, uniquely attained at e_0 , i.e. $\phi_0 = \phi(e_0)$. Let $h : \mathbb{R}^{|\phi_0|} \rightarrow \mathbb{R}$ a square-integrable function under \mathbb{P}^{e_0} such that for all $i \in \phi_0$,

$$\mathbb{E}^{e_0}[h(\mathbf{X}_{\phi_0})|\mathbf{X}_{\phi_0 \setminus \{i\}}] = 0$$

under the distribution \mathbb{P}^{e_0} from environment e_0 . By assumption (39), and the decomposition (38), we have that

$$\begin{aligned} 0 &= \mathbb{E}^{e_0}[h(\mathbf{X}_{\phi_0})f(\mathbf{X})] \\ &= \mathbb{E}^{e_0}[h(\mathbf{X}_{\phi_0}) \sum_{e \in \mathcal{E}} f_e(\mathbf{X}_{\phi(e)})] \\ &= \sum_{e: \phi_0 \succ \phi(e)} \mathbb{E}^{e_0}[h(\mathbf{X}_{\phi_0})f_e(\mathbf{X}_{\phi(e)})] + \mathbb{E}^{e_0}[h(\mathbf{X}_{\phi_0})f_{e_0}(\mathbf{X}_{\phi_0})] \end{aligned}$$

Now for all $e \neq e_0$, we have $\phi_0 \succ \phi(e)$. Let $\phi_0 =: \{a_1, \dots, a_{n_A}\}$, $\phi(e) =: \{b_1, \dots, b_{n_B}\}$. We define a subset of indices Δ_e as follows. If there exists i such that $a_j = b_j$ for all $j < i$, and $a_i > b_i$, then let $\Delta_e := \{a_1, \dots, a_{i-1}\} \cup [a_i - 1]$. Otherwise, if $n_A > n_B$ and $\{a_1, \dots, a_{n_B}\} = \phi(e)$, then $\Delta_e := \{a_1, \dots, a_{n_B}\} \cup [a_{n_B+1} - 1]$ (so it can be written as in the first case with $i = n_B + 1$). In both cases, $\phi(e) \subset \Delta_e$, $\phi_0 \setminus \{a_i\} \subset \Delta_e$ and $a_i \notin \Delta_e$. Then

$$\begin{aligned} \mathbb{E}^{e_0}[h(\mathbf{X}_{\phi_0})f_e(\mathbf{X}_{\phi(e)})] &= \mathbb{E}^{e_0}[\mathbb{E}^{e_0}[h(\mathbf{X}_{\phi_0})f_e(\mathbf{X}_{\phi(e)})|\mathbf{X}_{\Delta_e}]] \\ &= \mathbb{E}^{e_0}[\mathbb{E}^{e_0}[h(\mathbf{X}_{\phi_0})|\mathbf{X}_{\Delta_e}]f_e(\mathbf{X}_{\phi(e)})] \\ &\stackrel{(*)}{=} \mathbb{E}^{e_0}[\mathbb{E}^{e_0}[h(\mathbf{X}_{\phi_0})|\mathbf{X}_{\phi_0 \setminus \{a_i\}}]f_e(\mathbf{X}_{\phi(e)})] \\ &= 0 \end{aligned}$$

where we used the fact that by assumption on h , we have $\mathbb{E}^{e_0}[h(\mathbf{X}_{\phi_0})|\mathbf{X}_{\phi_0 \setminus \{a_i\}}] = 0$, and we later prove (*). We then get that

$$0 = \mathbb{E}^{e_0}[h(\mathbf{X}_{\phi_0})f_{e_0}(\mathbf{X}_{\phi_0})]$$

We now use Lemma 10, given that by assumption \mathbf{X}_{ϕ_0} are randomized, independently one of another under \mathbb{P}^{e_0} . We have that the following holds over the support of the variables

indexed by ϕ_0 in environment e_0 , where $g_i : \mathbb{R}^{|\phi_0|-1} \rightarrow \mathbb{R}$ are square-integrable functions:

$$f_{e_0}(\mathbf{x}_{\phi_0}) = \sum_{i \in \phi_0} g_i(\mathbf{x}_{\phi_0 \setminus \{i\}})$$

We get that f follows the following decomposition as in equation (38) over the support of \mathbb{P}^{e_0} , which by assumption is the same as that of the other environment distributions. We write:

$$\begin{aligned} f(\mathbf{x}) &= \sum_{e \in \mathcal{E} \setminus \{\phi_0\}} f_e(\mathbf{x}_{\phi(e)}) + \sum_{i \in \phi_0} g_i(\mathbf{x}_{\phi_0 \setminus \{i\}}) \\ &= \sum_{e \in \tilde{\mathcal{E}}} \tilde{f}_e(\mathbf{x}_{\phi(e)}) \end{aligned}$$

where we defined a new set of environments $\tilde{\mathcal{E}}$. The maximal element in $\tilde{\mathcal{E}}$ is smaller than ϕ_0 , as for all $i \in \phi_0$, $\phi_0 \succ \phi_0 \setminus \{i\}$, therefore $\mathcal{E} \succ \tilde{\mathcal{E}}$. Now for any $i \in \phi_0$, we can use e_0 as the environment where variables in $\phi_0 \setminus \{i\}$ are perturbed, so that condition (39) still holds. We thus get that f satisfies the same assumptions for the new $\tilde{\mathcal{E}}$. This concludes the recursion.

The remaining statement to prove is the equality (*) above, which is a consequence of the following identity:

$$\mathbb{E}^{e_0}[h(\mathbf{X}_{\phi_0}) | \mathbf{X}_{\Delta_e}] = \mathbb{E}^{e_0}[h(\mathbf{X}_{\phi_0}) | \mathbf{X}_{\phi_0 \setminus \{a_i\}}]$$

We now prove this identity holds, for which it suffices to show the following conditional independence statement holds under $\mathbb{P}_X^{e_0}$:

$$\mathbf{X}_{\phi_0} \perp\!\!\!\perp \mathbf{X}_{\Delta_e \setminus \{\phi_0 \setminus \{a_i\}\}} | \mathbf{X}_{\phi_0 \setminus \{a_i\}}$$

which is equivalent to

$$X_{a_i} \perp\!\!\!\perp \mathbf{X}_{\Delta_e \setminus \{\phi_0 \setminus \{a_i\}\}} | \mathbf{X}_{\phi_0 \setminus \{a_i\}} \quad (40)$$

To prove the last conditional independence statement, we rely on the equivalence between d-separation and separation in the moral ancestral graph (Pearl and Dechter, 2013). We use here our assumptions on how the model \mathcal{M}^{e_0} is generated. We know that $\mathbb{P}_X^{e_0}$ factorizes in the extended graph $\tilde{G}_X^{e_0}$. We need to show that X_{a_i} and $\mathbf{X}_{\Delta_e \setminus \{\phi_0 \setminus \{a_i\}\}}$ are separated by $\mathbf{X}_{\phi_0 \setminus \{a_i\}}$ in the moral ancestral graph of variables $\mathbf{X}_{\Delta_e \cup \{a_i\}}$ (union of all variables in the conditional independence statement) with respect to $\tilde{G}_X^{e_0}$. Variables indexed by ϕ_0 are randomized, hence do not have any ancestors in $\tilde{G}_X^{e_0}$. Therefore the ancestral graph of $\mathbf{X}_{\Delta_e \cup \{a_i\}}$ does not contain any additional nodes. Also, variables $\mathbf{X}_{\{a_1, \dots, a_i\}}$ do not have descendants in the ancestral graph: by topological ordering, variables $\mathbf{X}_{[a_i-1]}$ can not be descendants of $\mathbf{X}_{\{a_1, \dots, a_i\}}$. And any variable within $\mathbf{X}_{\{a_1, \dots, a_i\}}$ can not be descendant of any other node as they do not have ancestors. Therefore, variables $\mathbf{X}_{\{a_1, \dots, a_i\}}$ are isolated in the ancestral graph, and moralizing the ancestral graph does not connect these nodes to any other node. Therefore, in particular we get the separation statement between X_{a_i} and $\mathbf{X}_{\Delta_e \setminus \{\phi_0 \setminus \{a_i\}\}}$ by $\mathbf{X}_{\phi_0 \setminus \{a_i\}}$.

Importantly, the assumptions on the distributions are needed only to show that the conditional independence statement (40) holds, as this is only a property of the extended graph $\bar{G}_X^{e_0}$. The assumptions on \mathcal{M}^{e_0} impose constraints on $\bar{G}_X^{e_0}$: the choice of the structural equation functions defining the covariates as well as the joint distribution of the non-randomized disturbance variables $\{\epsilon_j\}_{j \notin \phi_0}$ is irrelevant for the validity of the proof. \blacksquare

Lemma 10 Consider a square-integrable function $g : \mathbb{R}^p \rightarrow \mathbb{R}$. Let $\mathbf{X} = (X_1, \dots, X_p)$ be independent random variables. Assume that, for any square integrable $h : \mathbb{R}^p \rightarrow \mathbb{R}$ satisfying

$$\mathbb{E}[h(\mathbf{X}) | \mathbf{X}_{-i}] = 0 \quad \forall i \in [p]$$

where $\mathbf{X}_{-i} := \{X_j, j \neq i\}$, the following holds:

$$\mathbb{E}[h(\mathbf{X})g(\mathbf{X})] = 0 \tag{41}$$

Then there exist functions $g_i : \mathbb{R}^{p-1} \rightarrow \mathbb{R}$ for $i \in [p]$ such that

$$g(\mathbf{x}) = \sum_{i=1}^p g_i(\mathbf{x}_{-i})$$

over the support of the random variables \mathbf{X} .

Proof Denote $Y := g(\mathbf{X})$. Define

$$h(\mathbf{X}) := \sum_{I \subset [p]} (-1)^{p-|I|} \mathbb{E}[Y | \mathbf{X}_I]$$

We have that, for any $i \in [p]$,

$$\begin{aligned} \mathbb{E}[h(\mathbf{X}) | \mathbf{X}_{-i}] &= \sum_{I \subset [p]} (-1)^{p-|I|} \mathbb{E}[\mathbb{E}[Y | \mathbf{X}_I] | \mathbf{X}_{-i}] \\ &= \sum_{\substack{I \subset [p] \\ i \in I}} (-1)^{p-|I|} \mathbb{E}[\mathbb{E}[Y | \mathbf{X}_I] | \mathbf{X}_{-i}] + \sum_{\substack{I \subset [p] \\ i \notin I}} (-1)^{p-|I|} \mathbb{E}[\mathbb{E}[Y | \mathbf{X}_I] | \mathbf{X}_{-i}] \\ &= \sum_{\substack{I \subset [p] \\ i \in I}} (-1)^{p-|I|} \mathbb{E}[Y | \mathbf{X}_{I \setminus \{i\}}] + \sum_{\substack{I \subset [p] \\ i \notin I}} (-1)^{p-|I|} \mathbb{E}[Y | \mathbf{X}_I] \\ &= \sum_{J \subset [p] \setminus \{i\}} (-1)^{p-|J|-1} \mathbb{E}[Y | \mathbf{X}_J] + \sum_{\substack{I \subset [p] \\ i \notin I}} (-1)^{p-|I|} \mathbb{E}[Y | \mathbf{X}_I] \\ &= 0 \end{aligned}$$

By assumption we thus get that $\mathbb{E}[h(\mathbf{X})g(\mathbf{X})] = 0$. We can decompose $h(\mathbf{X})$ as follows:

$$h(\mathbf{X}) = g(\mathbf{X}) + \sum_{\substack{I \subset [p] \\ |I| \leq p-1}} (-1)^{p-|I|} \mathbb{E}[Y | \mathbf{X}_I]$$

Therefore we get:

$$\begin{aligned}
 \mathbb{E}[h(\mathbf{X})^2] &= \mathbb{E}[h(\mathbf{X})(g(\mathbf{X}) - \sum_{\substack{I \subset [p] \\ |I| \leq p-1}} (-1)^{p-|I|} \mathbb{E}[Y|\mathbf{X}_I])] \\
 &= - \mathbb{E}[h(\mathbf{X}) \sum_{\substack{I \subset [p] \\ |I| \leq p-1}} (-1)^{p-|I|} \mathbb{E}[Y|\mathbf{X}_I]] \\
 &= - \sum_{\substack{I \subset [p] \\ |I| \leq p-1}} (-1)^{p-|I|} \mathbb{E}[h(\mathbf{X}) \mathbb{E}[Y|\mathbf{X}_I]] \\
 &= - \sum_{\substack{I \subset [p] \\ |I| \leq p-1}} (-1)^{p-|I|} \mathbb{E}[\mathbb{E}[h(\mathbf{X})|\mathbf{X}_I] \mathbb{E}[Y|\mathbf{X}_I]] \\
 &= 0
 \end{aligned}$$

where we used the fact that $\mathbb{E}[h(\mathbf{X})|\mathbf{X}_I] = \mathbb{E}[\mathbb{E}[h(\mathbf{X})|\mathbf{X}_{-i}|\mathbf{X}_I]$ for some $i \notin I$ as $|I| \leq p-1$. Therefore $h(\mathbf{X}) = 0$, and therefore

$$g(\mathbf{x}) = \sum_{\substack{I \subset [p] \\ |I| \leq p-1}} (-1)^{p-|I|-1} \mathbb{E}[Y|\mathbf{X}_I = \mathbf{x}_I] = \sum_{i=1}^p g_i(\mathbf{x}_{-i})$$

over the support of \mathbf{X} for some choice of functions $g_i : \mathbb{R}^{p-1} \rightarrow \mathbb{R}$. ■