

Causal Algorithmic Recourse: Foundations and Methods

Drago Plečko

*Department of Statistics & Data Science
UCLA
Los Angeles, CA 90095, USA
drago@stat.ucla.edu*

Collin Wang

Elias Bareinboim

*Department of Computer Science
Columbia University
New York, NY 10027, USA
{c1w2180, eb3304}@columbia.edu*

Abstract

The trustworthiness of AI decision-making systems is increasingly important. A key feature of such systems is the ability to provide recommendations for how an individual may reverse a negative decision, a problem known as algorithmic recourse. Existing approaches treat recourse outcomes as counterfactuals of a fixed unit, ignoring that real-world recourse involves repeated decisions on the same individual under possibly different latent conditions. We develop a causal framework that models recourse as a process over pre- and post-intervention outcomes, allowing for partial stability and resampling of latent variables. We introduce post-recourse stability conditions that enable reasoning about recourse from observational data alone, and develop a copula-based algorithm for inferring the effects of recourse under these conditions. For settings where paired observations of the same individual before and after intervention are available (called *recourse data*), we develop methods for inferring copula parameters and performing goodness-of-fit testing. When the copula model is rejected, we provide a distribution-free algorithm for learning recourse effects directly from recourse data. We demonstrate the value of the proposed methods on real and semi-synthetic datasets.

1. Introduction

Decision-making systems based on machine learning (ML) are being increasingly deployed in real-world settings with far-reaching implications to individuals and society more broadly. This includes hiring decisions, university admissions, law enforcement, credit lending, health care access, and finance (Khandani et al., 2010; Mahoney and Mohen, 2007; Brennan et al., 2009). With an ever larger number of decisions once made by humans now delegated to automated systems, increasing attention has been given to making AI systems trustworthy – explainable, robust, and fair.

A hallmark feature expected from trustworthy AI systems is the ability to provide recommendations for individuals who wish to reverse a negative decision obtained from such a system. This challenge has been studied in the literature under the rubric of *algorithmic recourse*. For example, an individual who obtained a negative decision when applying for a bank loan may be able to successfully reapply after increasing his/her savings (Caglayan et al., 2022); similarly, an individual who was declined surgical treatment due to increased risk of complications may wish to adjust his/her behavior and dietary habits and seek treatment again in hope of a positive decision (Arbous et al., 2001); further examples of such recourse settings are numerous. For automated systems that make

decisions in ways not necessarily understandable by humans, providing recourse recommendations to individuals presents an important methodological challenge.

In this paper, we will develop a causal approach to handling questions of algorithmic recourse. To intuitively ground some of the key developments in the remainder of the paper, we begin by discussing a simple example:

Example 1 (Exam Repetition) *Students at a university are taking a mandatory exam. For preparation, students are allowed to attend office hours held by the teaching assistants (variable X , where x_0 represents that the student attended office hours, and x_1 if she/he did not). The test scores of students are denoted by T , a numeric value in the reals \mathbb{R} . The outcome Y of whether a student passed the exam is a simple threshold operation, $Y = \mathbb{1}(T \geq t)$, where t is fixed by the university every year.*




Figure 1: Causal diagram for Ex. 1.

A graphical description of how variables affect each other in this setting is shown in Fig. 1. In particular, latent (unobserved) variables that affect the student’s test score T are drawn in gray, including the student’s ability, amount of stress in the last 2 weeks, and amount of sleep in the last 2 weeks.

Consider now a student who did not attend office hours ($X = x_0$) and obtained a test score of $t' < t$, thus failing the exam. The student is told to re-take the exam after two weeks, and is encouraged to attend office hours, i.e., set $X = x_1$ instead of x_0 in the coming week.

Let u_T denote the latent noise variables that determined the test outcome T for the student, together with x_0 . In the first instance, we observed the student’s test score T , which can be written as a function of the $X = x_0$ and the latent u_T , i.e., $T \leftarrow f_T(x_0, u_T)$, and T is realized as $t' = f_T(x_0, u_T)$. After two weeks, the student followed the recommendation of attending office hours, and now $X = x_1$. However, the latent noise variables that consist of ability, stress, and sleep levels may be different for the second attempt at taking the exam. Formally, the latent u_T for the first attempt may differ from u_T^* for the second attempt. Thus, in the second attempt, the evaluation of f_T will be $f_T(x_1, u_T^*)$, and the student will attain a new score t^* possibly different from t' . Importantly, the latent u_T^* may be related to u_T , since these two latent variables relate to the same individual. \square

Three key observations ensue from the above example:

- (1) In recourse settings, we may be able to observe two (or possibly more) outcomes, typically pre- and post-intervention. These outcomes, however, are for the *same individual*, and thus the noise variables u_T and u_T^* affecting them may share information,
- (2) The pre- and post-intervention outcomes, written as $f_T(x_0, u_T)$, and $f_T(x_1, u_T^*)$, may therefore be related. For instance, one could reasonably expect that some part of the latent u_T remains the same (e.g., student’s ability),
- (3) At the same time, however, other parts of the latent u_T may change. For instance, the student’s sleep and stress levels may possibly differ between the two exam attempts.

Therefore, when observing pre- and post-recourse outcomes (realizations of $f_T(x_0, u_T)$, $f_T(x_1, u_T^*)$), we need to take into account that u_T and u_T^* may share information but are not necessarily identical. In other words, the process of repeating actions in the real world adds inherent structural uncertainty, and the latent variables u_T may be perturbed to u_T^* . The ability to jointly observe outcomes of the same variable for a single individual will be a key feature of recourse data and motivates the methodological developments described in the sequel.

1.1 Relationship to Previous Literature

We now describe how the approach we present relates to previous works on algorithmic recourse. In particular, we explain the relation to (i) non-causal approaches to recourse; (ii) existing causal approaches; and (iii) the literature on combining observational and experimental data.

Contrastive Explanations. The approach of contrastive, or counterfactual, explanations seeks to find an alternative set of covariate values, close to the true ones, that would have resulted in a different decision for the individual (Wachter et al., 2017). Various notions of distance between features and the cost of manipulating them have been explored in the literature, in an attempt to measure how difficult it may be for an individual to change their attributes (Karimi et al., 2022, Sec. 3.1). If there is no causal structure among the covariates, and each feature may be manipulated by the individual, such explanations may also lead to recommendations on how the individual may change their outcome (Wachter et al., 2017; Joshi et al., 2019; Sharma et al., 2020; Ustun et al., 2019). A commonly used assumption is the *independently manipulable features* (IMF) assumption, which basically rules out any causal effects between the covariates used for prediction. However, methods that ignore the fact that actions taken by the individual may affect other decision-related characteristics may lead to suboptimal or infeasible recommendations (Barocas et al., 2020; Venkatasubramanian and Alfano, 2020). We discuss an explicit example of this in Sec. 2.1.

Causal Algorithmic Recourse. Causal recourse methods that explicitly take into account the underlying causal model of the environment have been proposed (Mahajan et al., 2019; Karimi et al., 2021; Von Kügelgen et al., 2022) in order to improve upon the possibly rigid IMF assumption. However, some of these methods require the full specification of the generating model of reality, in terms of a structural causal model (SCM, Pearl (2000)), and their assumptions are therefore too stringent for most practical applications (Bareinboim et al., 2022). The work most related to ours is that of Karimi et al. (2020), which assumes the availability of a causal diagram and data to perform inference, corresponding to the standard setting in the causality literature, and also the setting considered in this paper. The authors offer two solutions. They either make additive noise assumptions that render the joint distribution over counterfactual outcomes identifiable, or they consider an alternative method that avoids considering the joint, at the expense of recourse recommendations not being covariate-specific. Importantly, the formulation of recourse offered in this work considers the recourse outcomes to be *exact counterfactuals*, and does not consider the implications of the unit’s pre- and post-recourse latent variables possibly being different, as described in Ex. 1. Therefore, this prior work can be seen as a special case of the framework we propose, in which the pre- and post-recourse units are identical.

Combining Observational and Experimental Data. The final line of work related to ours is that of identifying interventional (Lee et al., 2020) or counterfactual (Correa et al., 2021) distributions from combinations of experimental and observational data. In the setting of recourse, however, there is a distinctive feature of observational and experimental samples being related, as they correspond to the same individual, as discussed in Ex. 1. This makes the setting unique, and any off-the-shelf inference ideas that just combine observational and experimental data will not be sufficient. In fact, our work is the first to formalize an inference problem of this kind, and the first to consider the concept of learning from *recourse data*.

After demonstrating how our work fits into the broader context, we can list the specific contributions found in this paper:

- (1) We develop a formulation of algorithmic recourse based on structural causal models that takes into account the inherent variability in the latent variables described in Ex. 1 (Def. 10). We define the so-called post-recourse stability (PRS) conditions that guarantee that, for a fixed set of causal parents, distributions over latent variables pre- and post-recourse coincide. We prove that recourse effects may be inferred from observational data under PRS (Thm. 1),

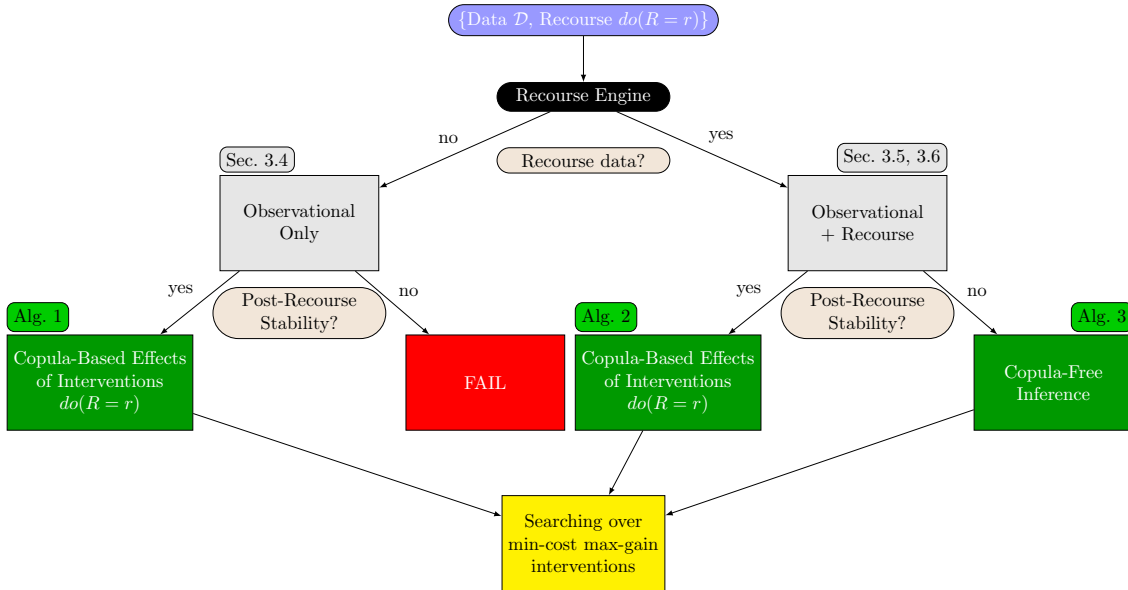


Figure 2: Overview of the framework proposed in this paper.

- (2) We develop an algorithm (Alg. 1) for non-parametrically inferring effects of recourse actions under PRS, using observational data and the assumption that the joint distribution over factual and recourse outcomes is modeled by Frank’s copula (Frank, 1979).
- (3) We introduce a novel causal notion of recourse data (Def. 13) in which we have access to an observational sample and an interventional sample for the same individual. We define a new class of learning problems based on recourse data (Def. 14). We develop an algorithm (Alg. 2) for inferring the effects of recourse actions under PRS, using observational and recourse data. Furthermore, we propose a goodness-of-fit hypothesis test that can be performed to falsify the assumption that Frank’s copula is an appropriate modeling choice.
- (4) We develop an algorithm (Alg. 3) for inferring effects of recourse actions from observational and recourse data for the case when PRS does not hold, and provide theoretical guarantees for the algorithm in the linear case (Thm. 2).

In Fig. 2, we provide a visual summary of the causal recourse framework introduced in this paper. In the following section, we cover some important preliminary notions for our discussion. Readers familiar with the language of structural causality may wish to go straight to Sec. 3.

2. Preliminaries

We use the language of structural causal models (SCMs) as our basic semantical framework (Pearl, 2000). A structural causal model (SCM) is defined as:

Definition 1 (Structural Causal Model (SCM) (Pearl, 2000)) *A structural causal model \mathcal{M} is a 4-tuple $\langle V, U, \mathcal{F}, P(u) \rangle$, where*

1. U is a set of exogenous variables, also called background variables, that are determined by factors outside the model;
2. $V = \{V_1, \dots, V_n\}$ is a set of endogenous (observed) variables, that are determined by variables in the model (i.e. by the variables in $U \cup V$);

3. $\mathcal{F} = \{f_{V_1}, \dots, f_{V_n}\}$ is the set of structural functions determining V , $v_i \leftarrow f_{V_i}(\text{pa}(v_i), u_i)$, where $\text{pa}(V_i) \subseteq V \setminus V_i$ and $U_i \subseteq U$ are the functional arguments of f_{V_i} ;
4. $P(u)$ is a distribution over the exogenous variables U .

□

The assignment mechanisms \mathcal{F} determine how each of the observed variables V_i attains its value, based on other observed variables and the latent variables U . Together with the probability distribution $P(u)$ over the exogenous variables U , it specifies the entire behavior of the underlying phenomenon. In particular, the SCM also specifies the *observational distribution* of the underlying phenomenon, defined through:

Definition 2 (Observational Distribution (Bareinboim et al., 2022)) *An SCM \mathcal{M} that is a 4-tuple $\langle V, U, \mathcal{F}, P(u) \rangle$ induces a joint probability distribution $P(V)$ such that for each $Y \subseteq V$,*

$$P^{\mathcal{M}}(y) = \sum_u \mathbb{1}(Y(u) = y) P(u), \quad (1)$$

where $Y(u)$ is the solution for Y after evaluating \mathcal{F} with $U = u$.

□

A further important notion building on the concept of the SCM is that of a submodel, which is defined next:

Definition 3 (Submodel (Pearl, 2000)) *Let \mathcal{M} be a structural causal model, X a set of variables in V , and x a particular value of X . A submodel \mathcal{M}_x (of \mathcal{M}) is a 4-tuple:*

$$\mathcal{M}_x = \langle V, U, \mathcal{F}_x, P(u) \rangle \quad (2)$$

where

$$\mathcal{F}_x = \{f_i : V_i \notin X\} \cup \{X \leftarrow x\}, \quad (3)$$

and all other components are preserved from \mathcal{M} .

□

Building on submodels, we introduce next the notion of a potential outcome:

Definition 4 (Potential Outcome / Response (Rubin, 1974; Pearl, 2000)) *Let X and Y be two sets of variables in V and $u \in \mathcal{U}$ be a unit. The potential outcome/response $Y_x(u)$ is defined as the solution for Y of the set of equations \mathcal{F}_x evaluated with $U = u$. That is, $Y_x(u)$ denotes the solution of Y in the submodel \mathcal{M}_x of \mathcal{M} .*

□

In words, $Y_x(u)$ is the value variable Y would take if (possibly contrary to observed facts) X is set to x , for a specific unit u . In Ex. 1, one could think of the potential outcome of the test score T subject to setting $X = 1$, which would be written $T_{X=1}(u)$. We further define how counterfactual distributions over various possible potential outcomes are computed:

Definition 5 (Counterfactual Distributions (Bareinboim et al., 2022)) *Consider an SCM $\mathcal{M} = \langle V, U, \mathcal{F}, P(u) \rangle$, and let $Y_1, \dots, Y_k \subset V$, and $X_1, \dots, X_k \subset V$ be subsets of the observables, and let x_1, \dots, x_k be specific values of X_i s. Denote by $(Y_i)_{x_i}$ the potential response of variables Y_i when setting $X_i = x_i$. The SCM \mathcal{M} induces a family of joint distributions over counterfactual events $(Y_1)_{x_1}, \dots, (Y_k)_{x_k}$:*

$$P^{\mathcal{M}}((y_1)_{x_1}, \dots, (y_k)_{x_k}) = \sum_u \mathbb{1}\left(\bigwedge_{i=1}^k (Y_i)_{x_i}(u) = y_i\right) P(u). \quad (4)$$

□

The l.h.s. in Eq. 4 contains variables with different subscripts, which syntactically represent different potential responses (Def. 4), or counterfactual worlds. Finally, there is one more prerequisite notion for our discussion. The mechanisms \mathcal{F} and the distribution over the exogenous variables $P(u)$ are almost never observed. However, to perform causal inference, we need a way of encoding assumptions about the underlying SCM. A common way of doing so is through an object called a causal diagram, which is defined next:

Definition 6 (Causal Diagram (Pearl, 2000; Bareinboim et al., 2022)) *Let an SCM \mathcal{M} be a 4-tuple $\langle V, U, \mathcal{F}, P(u) \rangle$. A graph \mathcal{G} is said to be a causal diagram (of \mathcal{M}) if:*

- (1) *there is a vertex for every endogenous variable $V_i \in V$,*
- (2) *there is an edge $V_i \rightarrow V_j$ if V_i appears as an argument of $f_j \in \mathcal{F}$,*
- (3) *there is a bidirected edge $V_i \leftrightarrow V_j$ if the corresponding $U_i, U_j \subset U$ are correlated or the corresponding functions f_i, f_j share some $U_{ij} \in U$ as an argument.*

□

We call $\text{pa}(V_i)$ the set of parents of V_i , while the sets of children $\text{ch}(V_i)$, ancestors $\text{an}(V_i)$, and descendants $\text{de}(V_i)$ are defined analogously.

Prior Recourse Definition. Recourse actions will generally be indicated by $\text{do}(R = r)$, and one may be interested in the potential responses of the outcome under the recourse action for specific units, labeled $\hat{Y}_{R=r}(u)$. When clear from context, we drop the $R = r$ notation and write only r instead. A previous formulation of algorithmic recourse, proposed in (Karimi et al., 2020), is given as follows:

Definition 7 (Recourse Problem (Karimi et al., 2020)) *Let $R \subset V$ be a subset of the observables on which recourse may be performed. Let r be a fixed value of R , and let $t(r)$ be a function of $R = r$. Then, the optimal recourse action for an individual with $V = v$ is given by:*

$$\arg \min_{R, r} \quad \text{cost}(R = r) \tag{5}$$

$$\text{subject to} \quad P(\hat{Y}_{R=r} = 1 \mid v) \geq t(r). \tag{6}$$

□

The intuition behind the above formulation is simple. The goal is to minimize the cost the individual incurs from performing recourse (Eq. 5) while ensuring that the probability of success after recourse is large enough (Eq. 6).

2.1 Independently Manipulable Features

In this section, we first discuss the independently manipulable features (IMF) assumption, frequently used in the literature on algorithmic recourse and counterfactual explanations (Wachter et al., 2017; Ustun et al., 2019; Joshi et al., 2019). In particular, the IMF assumption can be represented as a causal diagram, shown in Fig. 3a, where the dots (\dots) represent variables V_2, \dots, V_{k-1} between V_1 and V_k . Each variable V_1, \dots, V_k has an edge to \hat{Y} , but no other edges exist. This assumption, often made to simplify inference of recourse recommendations, can have serious implications even in the presence of simple causal structure among the features, and may lead to wrong conclusions as witnessed by the following example:




Figure 3: Independently manipulable feature (IMF) assumption for the general case (a) and for Ex. 2 (b); (c) show the true causal diagram for Ex. 2; (d) causal diagram for Ex. 4.

Example 2 (IMF Assumption Failure) Consider the following SCM \mathcal{M} :

$$V_1 \leftarrow N(0, 1) \tag{7}$$

$$V_2 \leftarrow \alpha V_1 + N(0, 1) \tag{8}$$

$$\hat{Y} \leftarrow \mathbb{1}(V_1 - V_2 \geq 1) \tag{9}$$

where V_1 represents the income level, V_2 total expenditure, and \hat{Y} whether a person is granted a cash loan. The causal diagram induced by \mathcal{M} is shown in Fig. 3c. In words, the total expenditure V_2 is a linear function of income V_1 , and if $\alpha < 1$ the individual (on average) spends less than he/she earns, and spends more otherwise. The loan decision \hat{Y} is strictly based on whether the individual's income is larger than their expenditure by a fixed amount of 1. Now, suppose we have an individual with features $(v_1, v_2) = (1, 0.5)$ who has his/her loan declined.

Suppose that V_1 is the only manipulable feature in this scenario, and suppose that a data analyst invokes the IMF assumption, which in this case does not hold due to a causal influence $V_1 \rightarrow V_2$. By knowing the mechanism of \hat{Y} (usually known to the system designer), they would reason as follows. To increase their creditworthiness, the person needs to increase their difference in income vs. expenditure, from the current $v_1 - v_2 = 0.5$ to 1. Under IMF, a sufficient intervention would be to add 0.5 to V_1 , i.e., do the recourse action $v_1 = 0.5 \rightarrow v_1 = 1$.

The IMF assumption would entirely disregard the fact that increased income may also be linked with increased expenditure, as described by the α coefficient; the IMF assumption implicitly assumes that $\alpha = 0$. However, in the SCM \mathcal{M} , the minimal necessary intervention when taking into account the causal structure $V_1 \rightarrow V_2$ would be equal to

$$\arg \min_{\delta} v_1 + \delta - (v_2 - \alpha\delta) \geq 1, \text{ for } v_1 = 1, v_2 = 0.5, \tag{10}$$

with the solution $\delta_{\min} = \frac{1}{2(1-\alpha)}$ for $\alpha < 1$ and no solution for $\alpha \geq 1$, yielding an entirely different insight from what was obtained based on the IMF assumption. \square

This example highlights that, even in the basic linear setting, it can be rather important to consider the causal structure among variables. Therefore, the IMF assumption, commonly invoked in the literature, may have significant drawbacks.

2.2 Hardness of Recourse – Joint Counterfactual Distributions

For the discussion in this section, suppose for simplicity that the set of covariates under recourse R is the set of all root nodes in the causal diagram \mathcal{G} , and that the decision mechanism $f_{\hat{y}}$ of \hat{Y} is available to the data analyst. Let $\text{de}(R)$ denote the descendants of R , and let $R = r_0, \text{de}(R) = \text{de}$

denote the covariate values we observed naturally for an individual, while $R = r$ denotes the recourse values. Then, the constraint term in Eq. 6 can be expanded as follows:

$$P(\widehat{Y}_{R=r} = 1 \mid V = v) = \sum_{de^*} P(\widehat{Y}_{R=r, de(R)=de^*} = 1, de(R)_{R=r} = de^* \mid R = r_0, de(R) = de) \quad (11)$$

$$= \sum_{de^*} \mathbb{1}(f_{\widehat{y}}(r, de^*) = 1) P(de(R)_{R=r} = de^* \mid R = r_0, de(R) = de), \quad (12)$$

where de^* ranges over all possible values of $de(R)$ that may be obtained after implementing the recourse $R = r$. Note that the expression in Eq. 12 depends on the joint counterfactual event $\{de(R) = de, de(R)_{R=r} = de^*\}$. The induced distribution is known to be challenging to evaluate without strong additional assumptions, (Tian and Pearl, 2000), making the problem in Def. 7 difficult to solve. This is illustrated by the following example:

Example 3 (Non-Identifiability of Recourse) Consider SCMs $\mathcal{M}^{(0)}, \mathcal{M}^{(1)}$, defined as follows:

$$X \leftarrow U_x \quad (13)$$

$$W \leftarrow X + (-1)^{X \cdot i} U_w \quad (14)$$

$$\widehat{Y} \leftarrow \mathbb{1}(W \geq 1) \quad (15)$$

$$U_x \sim \text{Bern}(0.5), U_w \sim N(0, 1). \quad (16)$$

The two SCMs are equivalent apart from the W mechanism f_w . Furthermore, the SCMs generate identical observational and interventional distributions. The causal diagram induced by both SCMs is the chain graph $X \rightarrow W \rightarrow \widehat{Y}$.

In the context of recourse, consider an individual with $X(u) = 0, W(u) = 0.5$ corresponding to $(U_x = 0, U_w = 0.5)$. In $\mathcal{M}^{(0)}$ the recourse action $do(X = 1)$ would yield $W_{X=1} = 1.5$ and $\widehat{Y}_{X=1} = 1$ (i.e., successful recourse), whereas in $\mathcal{M}^{(1)}$ it would yield $W_{X=1} = 0.5$ and $\widehat{Y}_{X=1} = 0$ (unsuccessful recourse). However, no amount of observational or interventional data would allow these two cases to be distinguished. \square

The above example provides a negative result for causal recourse. Interestingly, this result is obtained despite the fact that the assignment mechanism of the predictor \widehat{Y} , written $f_{\widehat{y}}$ is known to the data analyst¹. The remainder of this paper deals with inferring recourse actions while acknowledging this issue. We begin by showing that counterfactual outcomes are not necessarily the appropriate notion for recourse in practice.

2.3 Counterfactual Reasoning Is Necessary

The formulation of recourse in Def. 7 requires reasoning about counterfactual quantities, namely queries of the form

$$P(\widehat{y}_{R=r} \mid V = v). \quad (17)$$

Generally, counterfactual queries that belong to the so-called third layer of Pearl’s Causal Hierarchy² (PCH, for short) are more difficult to infer compared to interventional queries from the second layer of the PCH (Bareinboim et al., 2022). Therefore, it may be tempting to search for a different formulation of causal recourse, based purely on interventional reasoning, such as computing:

$$P(\widehat{y} \mid do(R = r), V \setminus R = v \setminus r). \quad (18)$$

1. In this sense, we can see how the setting of algorithmic recourse may differ from the standard literature in causal effect identification, in which none of the mechanisms \mathcal{F} of an SCM are known.
2. We remind the reader that the Pearl’s Causal Hierarchy consists of the (1) *observational*, (2) *interventional*, and (3) *counterfactual* layers (Bareinboim et al., 2022).

However, the situation is more involved, and some type of counterfactual reasoning is needed. As the following example illustrates, the two quantities in Eq. 17 and Eq. 18 can differ in practice:

Example 4 (Counterfactual vs. Interventional Reasoning for Recourse) Consider the following SCM \mathcal{M} :

$$V_1 \leftarrow U_1 \tag{19}$$

$$V_2 \leftarrow U_2 \tag{20}$$

$$V_3 \leftarrow V_1 \wedge U_3 \tag{21}$$

$$V_4 \leftarrow V_2 \vee U_4 \tag{22}$$

$$\widehat{Y} \leftarrow \mathbb{1}(V_1 + V_2 + V_3 + V_4 \geq 2.5), \tag{23}$$

$$U_1, U_2, U_3, U_4 \sim \text{Bernoulli}(0.5). \tag{24}$$

Now, consider a person with $(v_1, v_2, v_3, v_4) = (0, 0, 0, 0)$ implementing the recourse recommendation $do(V_1 = 1, V_4 = 1)$. The quantities of interest may be evaluated as:

$$P(\widehat{y} \mid do(v_1 = 1, v_4 = 1), v_2 = 0, v_3 = 0) = 0 \tag{25}$$

$$P(\widehat{y}_{v_1=1, v_4=1} \mid v_1 = 0, v_2 = 0, v_3 = 0, v_4 = 0) = \frac{1}{2}. \tag{26}$$

For obtaining the quantity in Eq. 25 we proceed as follows. First, we replace the mechanisms of V_1, V_4 with fixed values $V_1 \leftarrow 1, V_4 \leftarrow 1$. Then, in the submodel $\mathcal{M}_{V_1=1, V_4=1}$, we look at the implication of conditioning on $V_2 = 0, V_3 = 0$. First, in $\mathcal{M}_{V_1=1, V_4=1}$ observing that $V_3 = 0$ implies $U_3 = 0$ (since $V_1 = 1$ in this submodel and using Eq. 21). Observing that $V_2 = 0$ implies that $U_2 = 0$, based on Eq. 20. Thus, we know that $U_3 = U_2 = 0$ and this is sufficient to compute the \widehat{Y} in the submodel $\mathcal{M}_{V_1=1, V_4=1}$, yielding $\widehat{Y} = 0$ with probability 1.

For the quantity in Eq. 26 we follow a different procedure. We first update the latent variables according to available evidence. Since we observe $V_1 = 0, V_2 = 0$, it follows that $U_1 = 0, U_2 = 0$. Then, since $V_1 = 0, V_3 = 0$, we cannot infer U_3 based on Eq. 21, but conclude that U_3 equals 1 with probability 1/2. Finally, using mechanism in Eq. 22 and evidence $V_2 = 0, V_4 = 0$, we can conclude that $U_4 = 0$. This gives us the probability distribution of U_1, U_2, U_3, U_4 given the evidence $V_1 = V_2 = V_3 = V_4 = 0$. To compute the quantity in Eq. 26, we replace the mechanisms of V_1, V_4 with $V_1 \leftarrow 1, V_4 \leftarrow 1$ and compute \widehat{Y} based on the updated distribution of the U variables. This yields the $\frac{1}{2}$ result in Eq. 26, due to the fact that $P(U_3 \mid v_1 = 0, v_2 = 0, v_3 = 0) = \frac{1}{2}$. \square

The above example illustrates how counterfactual reasoning differs from interventional reasoning, and may yield very different results in practice. In fact, queries of the form in Eq. 25 will generally not allow one to reason about effects of recourse actions. Queries as in Eq. 26 are closer to what is needed for recourse inference, but still do not capture the problem complexity fully (as discussed in Ex. 1). This motivates the developments in the next section.

3. New Model of Causal Recourse

In this section, we discuss why the potential outcome $\widehat{Y}_{R=r}(u)$ may not be the target quantity of interest for solving recourse problems. We start by examining the implicit causal assumptions needed for recourse inference, first illustrated by an example:

Example 5 (Exam Repetition continued – Unsuccessful Recourse) Consider the exam repetition setting from Ex. 1, and consider a student who did not attend office hours ($X = x_0$), obtained a test score $T = t$, and failed the exam ($Y = 0$). Subsequently, the student was given a recourse

recommendation to attend office hours (action $do(X = x_1)$). However, after two weeks, the student comes for the exam repetition without having attended office hours, i.e., $X = x_0$ still (recourse was unsuccessful). However, could it happen that the student obtains a different test score $T = t'$ with $t' \neq t$ on the second attempt, and obtains a passing grade? \square

The example illustrates that repeated instantiations of a variable over time, even in the absence of recourse, may be subject to uncertainty. While the assignment mechanisms \mathcal{F} of the SCM (in this case f_T) are commonly assumed to be stable across time, the latent $U_t = u_t$ is unlikely to remain identical. In other words, there may be other sources of variation within the latent u_t that are different pre- and post-recourse. In Ex. 1, we discussed the obtained amount of sleep and the stress level in the past two weeks as possible latent variables that may change between different exam attempts. Denote the post-recourse test score by $T_{X=x_0}(u_w^*)$ (given that recourse was unsuccessful, and $X = x_0$ after recourse). In fact, since the student's ability is also affecting the test score, there may be some relation between the latent $U_t = u_t$ pre-recourse and the latent $U_t^* = u_t^*$ post-recourse. That is, $U_t^* \not\perp U_t$ since some parts of the latent variable are invariant across time. Therefore, based on this observation, a slightly different approach may be necessary to take into account this variability between the latent variables pre- and post-recourse. Instead of assuming that the unit $U = u$ is fixed, we may assume that

$$u = (u^{(f)}, u^{(v)}), \quad (27)$$

where the first part $u^{(f)}$ represents the fixed, immutable (or intrinsic) latent information on the unit, such as the individual's biology, genetics, or ability. The second part $u^{(v)}$ represents circumstantial information, such as the amount of sleep or stress in the past week. Crucially, in modeling, one needs to account for the fact that this second, circumstantial part $u_t^{(v)}$ may be resampled after the implementation of the recourse action.

3.1 Causal Recourse Building Blocks

To represent these dynamics in a more fine-grained way, we introduce a new building block called recourse SCM:

Definition 8 (Recourse SCM) Let $\mathcal{M} = \langle \mathcal{F}, P(u) \rangle$ be an SCM over variables V, U . For $R \subset V$ let $R = r$ denote the value of a recourse intervention. A recourse SCM $\mathcal{M}_{R=r}^{P(u^*)}$ is a tuple

$$\langle \mathcal{F}, P(u), R = r, P(u^* | u) \rangle, \quad (28)$$

where $P(u^* | u)$ is the distribution over the units after the recourse action $do(R = r)$. A recourse SCM is said to be Markovian if $\langle \mathcal{F}, P(u) \rangle$ is Markovian and

$$P(u^* | u) = \prod_{i=1}^n P(u_i^* | u_i). \quad (29)$$

\square

Based on the definition of a recourse SCM, we can also define notions of a recourse distribution and recourse sampling, analogous to the definitions of observational and counterfactual distributions (Defs. 2 and 5):

Definition 9 (Recourse Distribution) Let $\mathcal{M}_{R=r}^{P(u^*)}$ be a recourse SCM, and let V, V^* denote variables pre- and post-recourse, respectively. The recourse distribution is then given by

$$P_{\mathcal{M}_{R=r}^{P(u^*)}}(v, v^*) = \sum_u \mathbb{1}(V(u) = v, V^*(u^*) = v^*) P(u) P(u^* | u). \quad (30)$$

\square

The definition is now illustrated through an example:

Example 6 (Exam Repetition continued) Consider the exam repetition setting from Ex. 1, and suppose it is described by the following SCM \mathcal{M} :

$$X \leftarrow U_X \tag{31}$$

$$T \leftarrow X + 2U_1 - U_2 + U_3 \tag{32}$$

$$Y \leftarrow \mathbb{1}(T \geq 1) \tag{33}$$

$$U_X \sim \text{Bernoulli}(0.5), U_1, U_2, U_3 \sim N(0, 1), \tag{34}$$

corresponding to the causal diagram in Fig. 1, with U_1, U_2, U_3 corresponding to latent ability, stress, and sleep levels. Office hours attendance (X) is a Bernoulli(0.5) random variable, meaning that half of the students attend office hours. The f_T mechanism says that students who attend office hours perform better, with the performance influenced positively by ability (U_1) and sleep (U_3), and influenced negatively by stress levels (U_2). For an individual who did not attend office hours, $X(u) = x_0$, obtained a test score $T(u) = \frac{1}{2}$, and failed the exam $Y(u) = 0$, we consider an intervention $\text{do}(X = 1)$ (attending office hours), and two alternative distributions over the post-recourse noise variables U^*

$$P^{(a)}(u^* | u) = \begin{cases} U_1^* = U_1 & (35) \\ U_2^* \sim N(0, 1) & (36) \\ U_3^* \sim N(0, 1), & (37) \end{cases} \quad \text{vs.} \quad P^{(b)}(u^* | u) = \begin{cases} U_1^* = U_1 & (38) \\ U_2^* \sim N(1, 1) & (39) \\ U_3^* \sim N(-1, 1). & (40) \end{cases}$$

The distribution $P(u^* | u)$ determines how the noise variables are sampled post-recourse, possibly conditional on the initial noise variables $U = u$. In particular, for both $P^{(a)}, P^{(b)}$ assume that the latent ability of the student (U_1) remains invariant post-recourse, written $U_1^* = U_1$. However, the two distributions differ according to the distribution over the U_2^*, U_3^* variables. In $P^{(a)}$, post-recourse stress and sleep levels U_2^*, U_3^* follow a $N(0, 1)$ distribution, the same distribution as U_2, U_3 variables pre-recourse. In $P^{(b)}$, however, post-recourse stress and sleep levels follow different distributions, $N(1, 1)$ and $N(-1, 1)$, meaning that post-recourse students have higher degrees of stress and lower levels of sleep. \square

We now discuss different possible ways to think about recourse inference. In the oracle case, if the true recourse SCM $\mathcal{M}_{R=r}^{P(u^*)}$ were available to us, we would be able to compute the recourse outcome $\hat{Y}_{R=r}(u^*)$ exactly. For each variable under recourse, labeled R_i^* , we have $R_i^* \leftarrow r_i$ (due to recourse), whereas for each variable V_i^* not under recourse, one evaluates

$$V_i^*(u_i^*) \leftarrow f_{V_i}(\text{Pa}_i^*, u_i^*). \tag{41}$$

Once all the post-recourse variables V^* are computed (made possible by the exclusive knowledge of the latent recourse variables U_i^*), we can simply evaluate the recourse outcome using the classifier \hat{Y} . This approach describes how the ground truth outcome $\hat{Y}_{R=r}(u^*)$ (which we are interested in) may be recovered from the recourse SCM. We place this outcome in the first column of Fig. 4, which provides a mental map of possible recourse inference targets. Inferring outcomes deterministically is rarely possible, however, and we thus need to resort to probabilistic reasoning. A slightly weaker inference target than obtaining $\hat{Y}_{R=r}(u^*)$ would be to condition on the unit $U = u$, that is, the pre-recourse latent variables. Since $U = u$ (pre-recourse), and $U^* = u^*$ (post-recourse) share information, conditioning on $U = u$ would provide useful information about the random outcome $\hat{Y}_{R=r}(U^*)$. This approach (corresponding to the second column in Fig. 4) would still require the knowledge of the SCM (since we need to know the values of the latent $U = u$), which is almost never available to us. Therefore, we need to use a slightly weaker probabilistic target, and instead

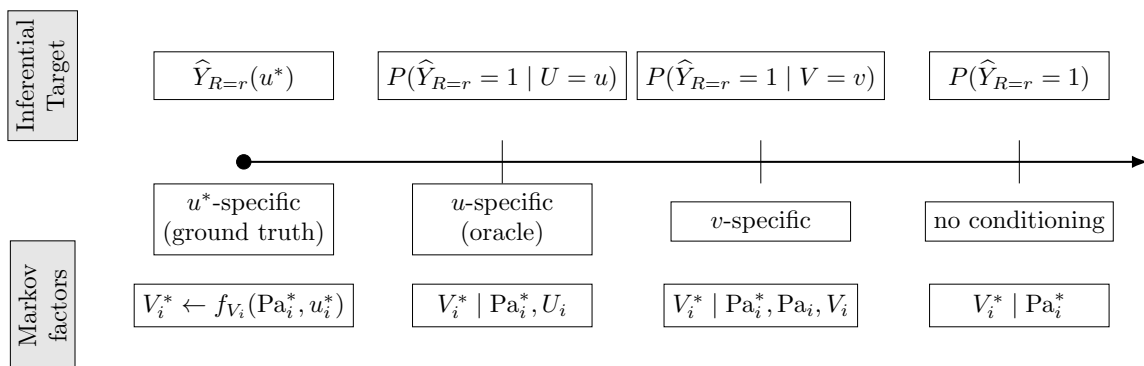


Figure 4: Visualization of different possible inference goals for recourse.

of conditioning on $U = u$, we make use of the available information on each individual (given in terms of the observed vector of covariates $V = v$) and conditioning on it instead (column three in Fig. 4). Finally, a baseline approach we mention is to drop the information on the individual pre-recourse, and simply estimate $P(\hat{Y}_{R=r} = 1)$ without conditioning on any information (Fig. 4, column four). In this way, however, we may drop useful information. In order perform recourse inference, in this manuscript we opt for conditioning on the observed covariate pre-recourse $V = v$, yielding the following problem definition:

Definition 10 (Causal Recourse Under Resampling) Let $\mathcal{M}_{R=r}^{P(u^*)}$ be a recourse SCM, and let $\text{cost}(R = r, V = v)$ be the cost of implementing a recourse action $R = r$ for an individual with attributes $V = v$. The problem of finding the optimal recourse action is given by:

$$\arg \min_{R,r} \text{cost}(R = r, V = v) \quad (42)$$

$$\text{subject to } P_{\mathcal{M}_{R=r}^{P(u^*)}}(\hat{Y}_{R=r} = 1 \mid V = v) \geq t(r), \quad (43)$$

where $t(r)$ is a threshold that may vary according to the recourse action $R = r$, and the probability $P_{\mathcal{M}_{R=r}^{P(u^*)}}$ is computed based on the recourse SCM $\mathcal{M}_{R=r}^{P(u^*)}$, or data generated from $\mathcal{M}_{R=r}^{P(u^*)}$. \square

The emphasis in the above definition is that (i) there may be a type of *distribution change* in the sampling of units $P(u^* \mid u)$ compared to the initial distribution $P(u)$; and (ii) the evaluation of the probability of success in Eq. 43 should be considered with respect to the joint distribution over $\mathcal{M}_{R=r}^{P(u^*)}$ and not just the initial \mathcal{M} . In particular, the observed $V = v$ are sampled from \mathcal{M} but evaluating $\hat{Y}_{R=r}(u^*)$ post-recourse requires information beyond just \mathcal{M} since it depends on the post-recourse distribution over the units $P(u^* \mid u)$; and finally (iii) the probability of success in Eq. 43 is conditioned on $V = v$, that is, all the observed information. Ideally, in the oracle case, we would want to compute $\hat{Y}_{R=r}(u^*)$ (or the slightly weaker counterpart $P_{\mathcal{M}_{R=r}^{P(u^*)}}(\hat{Y}_{R=r} = 1 \mid U = u)$). The latter approach, even though arguably more informative, would amount to conditioning on information that is never observed, and thus does not provide us with a practicable formulation of recourse. Therefore, conditioning on $V = v$ is a relaxation that still leverages all the available information on the specific individual while defining a possibly feasible problem. We next illustrate the challenges in evaluating the optimization problem in Def. 10 through our running example:

Example 7 (Exam Repetition continued) Consider the recourse SCM $\mathcal{M}_{R=r}^{P(u^*)}$ in Ex. 6 with a post-recourse distribution $P^{(a)}$ in Eq. 35-37. Consider an individual with the latent variables values $(u_X, u_1, u_2, u_3) = (0, 1, 1, -1)$. The individual attains the value $T(u) = 0$ and thus fails the exam

on the first attempt. We are now interested in computing the probability of success after a recourse action $\text{do}(X = 1)$. Based on $\mathcal{M}_{R=r}^{P(u^*)}$, we can see that

$$T_{X=1}(U^*) = 1 + 2U_1^* - U_2^* + U_3^*. \quad (44)$$

Based on $P^{(a)}$, conditioning on $u_1 = 1$ would imply that $u_1^* = 1$, and $U_2^*, U_3^* \sim N(0, 1)$, i.e.,

$$T_{X=1}(U^*) \mid U = u \stackrel{d}{=} 1 + 1 - U_2^* + U_3^* \sim N(2, 2), \quad (45)$$

which implies $P_{\mathcal{M}_{X=1}^{P(u^*)}}(\widehat{Y}_{R=r} = 1 \mid U = u) = P(T_{X=1}(U^*) > 1 \mid U = u) = 1 - \Phi(-\frac{\sqrt{2}}{2}) \approx 0.76$, where Φ is the cumulative distribution of $N(0, 1)$.

However, as discussed above, conditioning on $U = u$ is not feasible since this information is almost never available to the data analyst. Thus, we condition on the observed values $X = 0, T = 0$. Based on the SCM, we can infer that

$$X(u) = 0, T(u) = 0 \implies 2U_1 - U_2 + U_3 = 0. \quad (46)$$

Let $U_T = (U_1, U_2, U_3)$ and $U_T^* = (U_1^*, U_2^*, U_3^*)$. Let the vector $a = (2, -1, 1)$ correspond to the coefficients of U_1, U_2, U_3 in the f_T mechanism in Eq. 32. Based on $P(u), P(u^* \mid u)$, we know that

$$\begin{pmatrix} a^T U_T \\ a^T U_T^* \end{pmatrix} \sim N \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 6 & 4 \\ 4 & 6 \end{pmatrix} \right), \quad (47)$$

since $2U_1 - U_2 + U_3 \sim N(0, 6)$ and $\text{Cov}(2U_1 - U_2 + U_3, 2U_1^* - U_2^* + U_3^*) = 4$. Therefore, by properties of bivariate Gaussian variables, Eq. 46 implies:

$$2U_1^* - U_2^* + U_3^* \mid 2U_1 - U_2 + U_3 = 0 \sim N(0, \frac{10}{3}). \quad (48)$$

Thus, $P_{\mathcal{M}_{X=1}^{P(u^*)}}(\widehat{Y}_{R=r} = 1 \mid V = v) = 1 - \Phi(0) = \frac{1}{2}$ based on Eq. 48. The example highlights the distinction between a unit-level recourse probability (conditional on the unobserved $U = u$, Eq. 45) vs. a covariate-level recourse probability (conditional on the observed $V = v$, Eq. 48). \square

In the sequel, we discuss different approaches for inference of recourse probabilities in practice.

3.2 Markovian Factorization of the Recourse Distribution

In this section, we discuss how the recourse problem in Eqs. 42-43 may be broken down. In this paper, we work with Markovian models, meaning that the latent noise variables U_i are independent (i.e., there is no hidden confounding). Clearly, solving the recourse optimization problem requires the evaluation of

$$P_{\mathcal{M}_{R=r}^{P(u^*)}}(\widehat{Y}_{R=r} = 1 \mid V = v) \quad (49)$$

appearing in the condition in Eq. 43. Here, \widehat{Y} is a classifier, which takes the post-recourse variables V^* as an input. Therefore, evaluating the expression in Eq. 49 requires some way of inferring the joint distribution over $P(v, v^*)$. In Fig. 5, we provide a graphical representation based on which the inference approach described in this manuscript can be understood. In Fig. 5, the observed variables V_1, \dots, V_n appear, together with the recourse variables V_1^*, \dots, V_n^* . For each V_i , there is a variable part of the latent U_i , labeled $U_i^{(v)}$, and a fixed part of U_i , labeled $U_i^{(f)}$. The key idea is that the fixed part $U_i^{(f)}$ is assumed to be the same pre- and post-recourse, while the variable part $U_i^{(v)}$ may




Figure 5: Informal causal diagram over observational and recourse variables. Vertical dots indicate arbitrary arrows and variables between the observed (V) and recourse variables (V^*).

change. Therefore, the corresponding recourse variable V_i^* is influenced by $U_1^{(f)}$, and also $U_1^{(v)*}$, with the latter being possibly different than $U_1^{(v)}$. Importantly, based on Fig. 5, we can also see that

$$V_i^* \perp\!\!\!\perp V \setminus \{V_i, \text{Pa}_i\} \mid V_i, \text{Pa}_i, \text{Pa}_i^*. \quad (50)$$

In words, once we condition on V_i and its set of pre- and post-recourse parents $\text{Pa}_i, \text{Pa}_i^*$, the post-recourse V_i^* is independent of all other pre-recourse variables $V \setminus \{V_i, \text{Pa}_i\}$. This means that the conditional distribution $P(v^* \mid v)$ can be factorized as:

$$P(v^* \mid v) = \prod_{V_i \notin R} P(v_i^* \mid \text{pa}_i^*, v_i, \text{pa}_i) \times \prod_{V_i \in R} \mathbb{1}(V_i = r_{V_i}), \quad (51)$$

where r_{V_i} is the value of V_i in the vector of recourse values r . We refer to the terms $P(v_i^* \mid \text{pa}_i^*, v_i, \text{pa}_i)$ as Markov factors, corresponding to the distribution of V_i^* in the joint distribution $P(v^* \mid v)$ (see bottom row of Fig. 4). We remark that similar reasoning as above could be used if we considered conditioning on $U_i = u_i$ instead of the pair (V_i, Pa_i) . Note that

$$V_i^* \perp\!\!\!\perp U \setminus U_i \mid U_i, \text{Pa}_i^*, \quad (52)$$

which then implies the factorization

$$P(v^* \mid u) = \prod_{V_i \notin R} P(v_i^* \mid \text{pa}_i^*, u_i) \times \prod_{V_i \in R} \mathbb{1}(V_i = r_{V_i}). \quad (53)$$

However, in this paper, we are generally interested in conditioning on $V = v$, although we discuss some alternative approaches later on.

Various inferential challenges arise when considering the factorization in Eq. 51. Firstly, one may ask whether the conditional distribution $P(v_i^* \mid \text{pa}_i^*, v_i, \text{pa}_i)$ can be recovered in the case when there is no post-recourse data, i.e., V^* is not observed. Secondly, one may ask how to infer this conditional distribution if some (limited) amount of recourse data is available. The remainder of the section deals with these inferential challenges, and is organized as follows (see Fig. 2 for a schematic overview). First, in Sec. 3.3 we formally introduce post-recourse stability (PRS), which tells us whether the distribution over the latent variables pre- and post-recourse is the same (intuitively, in Fig. 5, whether $U_i^{(v)}$ and $U_i^{(v)*}$ are distributionally the same). Assuming we have access to observational data (pre-recourse), the ability to infer the effects of recourse actions depends on PRS. In Sec. 3.4, we show

how to perform recourse inference from observational data under PRS, based on a copula model. In Sec. 3.5, we introduce the concept of *recourse data* – assuming that data is available on individuals after they have performed recourse actions (that is, we have both pre-recourse samples V and post-recourse samples V^* for the same set of individuals). Then, we discuss how to perform recourse inference from combinations of observational and recourse data under PRS. Finally, in Sec. 3.6, we demonstrate that recourse inference is possible from combinations of observational and recourse data even if the PRS does not hold.

3.3 Post-Recourse Stability (PRS)

We begin the discussion by stating a set of conditions under which reasoning about recourse will be possible from observational data alone:

Definition 11 (Post-Recourse Stability) *Let $\mathcal{M} = \langle \mathcal{F}, P(u) \rangle$ be the SCM pre-recourse, and $\mathcal{M}^* = \langle \mathcal{F}_{R=r}, P^*(u) \rangle$ post-recourse. Suppose the unit $u = (u^{(f)}, u^{(v)})$ is partitioned into fixed and variable exogenous features, that is, only $u^{(v)}$ is resampled after recourse. Then, post-recourse stability (PRS) is defined as:*

- (a) *the exogenous variables of non-descendant variables ($\text{nd}(R)$) of the recourse action $R = r$ remain unchanged*

$$U_{\text{nd}(R)}^* = U_{\text{nd}(R)}, \quad (54)$$

- (b) *for all V_i descendants of R , we have that*

$$U_i^{*(f)} = U_i^{(f)}, \text{ and} \quad (55)$$

$$P(U_i^{*(v)} = u_i^{(v)} \mid U_i^{*(f)} = u_i^{(f)}) = P(U_i^{(v)} = u_i^{(v)} \mid U_i^{(f)} = u_i^{(f)}). \quad (56)$$

□

Despite the non-trivial notation, the two conditions are quite intuitive. Firstly, Eq. 54 states that for all non-descendants of the recourse action $do(R = r)$, the latent variables remain unchanged. This is natural because these variables are not intervened upon, and neither are any of their ancestors (hence, we expect them to remain invariant). Secondly, Eq. 55 states that the fixed part of the latent u_i , labeled $u_i^{(f)}$, remains invariant, while Eq. 56 states that the sampling of the variable (coincidental) part $u^{(v)}$ of the unit u is the same as in the original SCM \mathcal{M} , given the fixed part of the unit $u^{(f)}$. This condition precludes certain scenarios, for example, settings in which applicants perform better on repeated tries of a standardized test; in this case, the $u^{(v)}$ distribution post-recourse may change, even though the $u^{(f)}$ (intrinsic ability) would not.

Example 8 (Exam Repetition – Post-Recourse Stability) *Going back to the recourse SCM from Ex. 6, consider the two distributions over post-recourse latent variables, $P^{(a)}$ and $P^{(b)}$. For both distributions, the latent U_1 remains unchanged pre- and post-recourse. For $P^{(a)}$, the post-recourse latent variables U_2^*, U_3^* are distributed as $N(0, 1)$, same as pre-recourse. However, this is not true for $P^{(b)}$. The distribution of U_2^* is $N(1, 1)$ post-recourse, whereas it was $N(0, 1)$ pre-recourse. □*

As shown next, settings where PRS holds are interesting since reasoning about recourse from observational data is feasible (the theorem’s proof is given in Appendix C):

Theorem 1 (Post-Recourse Stability \implies Margin Stability) *Let $U \sim P(u)$ with $P(u)$ the distribution over units before recourse, and $U^* \sim P^*(u)$ with $P^*(u)$ the distribution after recourse. Post-recourse stability (Def. 11) implies margin stability, written as*

$$f_{V_i}(\text{pa}_i, U_i) \stackrel{d}{=} f_{V_i}(\text{pa}_i, U_i^*) \quad \forall i, \text{pa}_i \text{ fixed}. \quad (57)$$

□

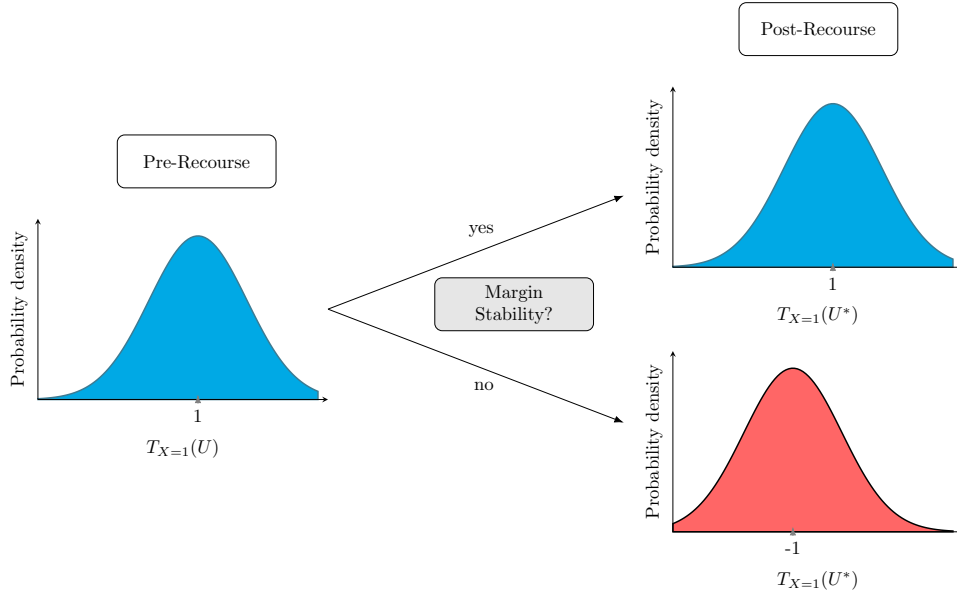


Figure 6: Schematic representation of margin stability conditions (Thm. 1) related to Ex. 9. On the left we have the marginal distribution pre-recourse, whereas on the right we have the possible post-recourse distributions. If MSC hold, then the post-recourse distribution will be equal to the pre-recourse one (in blue). However, if MSC do not hold, then the post-recourse distribution may differ (in red).

The importance of this result stems from the following. Suppose there is a unit u with $V(u) = v$ implementing a recourse action $R = r$. Let V_i^* , Pa_i^* denote a variable and its parents post-recourse, respectively. Based on Eq. 57 we know that, given the parents under the action $do(R = r)$, denoted by $\text{Pa}_i^*(u^*) = \text{pa}_i^*$, the distribution of V_i^* after recourse is the same as the observational distribution of V_i given $\text{Pa}_i = \text{pa}_i^*$, i.e.,

$$V_i^* \mid \text{Pa}_i^* = \text{pa}_i^* \stackrel{d}{=} V_i \mid \text{Pa}_i = \text{pa}_i^*. \quad (58)$$

The l.h.s. of Eq. 58 is the marginal target distribution after recourse, and the r.h.s. is available from observational data. Thus, we know the marginal distribution of V_i^* conditional on parents Pa_i^* . Nonetheless, note that we have still not leveraged the fact that the values of V_i , Pa_i before recourse were observed, which may further improve our inference.

Example 9 (Exam Repetition – MSC Implication) Consider the recourse SCM from Ex. 6 and the two post-recourse distributions $P^{(a)}, P^{(b)}$. Suppose now we are interested in the distribution of test scores after recourse, assuming all units of the population implemented the recourse action $do(X = 1)$. Then, using the f_T mechanism in Eq. 32, we have that for $P^{(a)}$

$$T_{X=1}(U_T^*) \sim N(1, 6) \quad (59)$$

since $X = 1$ and $2U_1^* - U_2^* + U_3^* \sim N(0, 6)$ according to $P^{(a)}$. Crucially, note that the pre-recourse distribution of test scores equals

$$T_{X=1}(U_T) \sim N(1, 6), \quad (60)$$

that is, the marginal distributions pre- and post-recourse are equal. This is a key implication of post-recourse stability.

On the other hand, for $P^{(b)}$, we have that $2U_1^* - U_2^* + U_3^* \sim N(-2, 6)$, and therefore

$$T_{X=1}(U_T^*) \sim N(-1, 6) \quad (61)$$

In this case, the post-recourse marginal differs from the pre-recourse marginal. A schematic representation is shown in Fig. 6. \square

The key observation is that, in absence of data on individuals implementing recourse actions, and without PRS, the marginal distribution after recourse is not recoverable. Therefore, in absence of PRS and recourse data, inference on recourse actions would generally not be feasible. Under PRS, at least the marginal distribution following from recourse action can be identified. We will leverage this insight in the sequel.

Connection to Unobservability of Potential Outcomes. Before continuing, we draw an important connection to the inherent unobservability of joint potential outcomes $\widehat{Y}(u), \widehat{Y}_{R=r}(u)$. Previously, we discussed how a unit u can be decomposed into its fixed part $u^{(f)}$ and variable part $u^{(v)}$ that is resampled. For cases in which the proportion of the variance explained by $u^{(v)}$ tends to 0, that is, the unit’s value is almost entirely fixed, we expect that the counterfactual outcome $\widehat{Y}_{R=r}(u)$ is equal to the post-recourse outcome $\widehat{Y}_{R=r}(u^*)$ obtained by the unit u after implementing recourse. Therefore, in such a setting, if the post-recourse sample $\widehat{Y}_{R=r}(u^*)$ is available, it would allow us to effectively observe the counterfactual $\widehat{Y}_{R=r}(u)$. The opposite extreme would be when the proportion of variance explained by $u^{(f)}$ tends to 0, in which case the unit is considered as entirely resampled, i.e., $U \perp\!\!\!\perp U^*$ in the previous notation.

3.4 Non-Parametric Inference Under Post-Recourse Stability – Observational Data

In Sec. 3.2 we discussed several different ways of incorporating the pre-recourse information of the individual when computing recourse probabilities (recall Fig. 4). As argued, when inferring the distribution of the post-recourse variable V_i^* , in the oracle case, we would consider the post-recourse latent variables U_i^* , or their pre-recourse counterparts U_i . In practice, however, we would have to resort to conditioning on the pair (V_i, Pa_i) according to the factorization in Eq. 51. However, there is another route to inferring recourse probabilities. Importantly, we know that the latent U_i shares information with the U_i^* , written $U_i \not\perp\!\!\!\perp U_i^*$. In this section, we discuss a relaxed approach for doing so, in which instead of conditioning on U_i itself, which is unobserved, we condition on an important proxy of U_i , namely the quantile of $V_i = v_i$ in the distribution $P(V_i | \text{Pa}_i = \text{pa}_i)$:

Definition 12 (Distribution Quantile) Let f_{V_i} be a mechanism of the SCM, and let U_i be the latent variable associated with V_i . We define the quantile of $V_i = v_i$ given parents pa_i as the

$$Q(v_i | \text{pa}_i) = P(V_i \leq v_i | \text{pa}_i). \quad (62)$$

\square

The quantile q_i results from a many-to-one mapping of latents $u_i \mapsto q_i$. Nonetheless, the quantile q_i is a function of and contains information on the latent u_i , i.e., $q_i = q_i(u_i)$. Similarly for the post-recourse quantile, $q_i^* = q_i^*(u_i^*)$. Since in general $U_i \not\perp\!\!\!\perp U_i^*$, it also implies that $Q_i \not\perp\!\!\!\perp Q_i^*$, meaning that the quantiles pre- and post-recourse share information. To leverage this fact for inference, we introduce a model for the coupling of pre-recourse quantiles q_i and post-recourse quantiles q_i^* .

Applying Sklar’s Theorem. To model the coupling of quantiles we apply Sklar’s theorem (Sklar, 1959). The theorem states that a bivariate cumulative distribution function $H(x_1, x_2)$ can be expressed in terms of its marginals F_1, F_2 and a copula C , i.e.,

$$H(x_1, x_2) = C(F_1(x_1), F_2(x_2)). \quad (63)$$




Figure 7: Frank’s copula coupling of quantiles for different values of $\tau \in \{-0.5, 0, 0.5, 1\}$.

The copula C represents the coupling of the quantiles of the marginal distributions F_1, F_2 that defines the joint distribution and is a commonly used modeling tool in statistics, finance, econometrics, and a number of other domains (Jaworski et al., 2010). By an application of Sklar’s theorem, for any pair of distributions

$$V_i \mid \text{Pa}_i = \text{pa}_i \text{ and } V_i^* \mid \text{Pa}_i^* = \text{pa}_i^*, \quad (64)$$

there exists a copula that relates their quantiles.

Why Quantiles? In the above construction, quantiles play a key role as a connection point with the latent variables U . Quantiles can be thought of as the best possible non-parametric proxy for the unobserved unit u , since the quantiles are (i) independent of the value of the parents $\text{Pa}_i = \text{pa}_i$; (ii) independent of the mechanism f_{V_i} of V_i . However, the many-to-one mapping of a unit $Q_i : u_i \mapsto q_i$ can be understood as a type of quotient operation, with all units $\{u_i : Q_i(u_i) = q_i\}$ considered as congruent. In other words, no amount of data will allow an analyst to distinguish among the units in the set $\{u_i : Q_i(u_i) = q_i\}$ regardless of the level of sophistication of the inference method. However, units corresponding to different values of q_i, q_i' may be distinguished, and this fact may allow us to better infer recourse probabilities.

Example 10 (Exam Repetition – Congruence of Units) Consider the recourse SCM in Ex. 6 with a post-recourse distribution $P^{(a)}$ and an individual with the latent variables values $(u_X, u_1, u_2, u_3) = (0, 1, 1, -1)$. For this individual, we have that

$$X(u) = 0, T(u) = 0. \quad (65)$$

The marginal distribution of $T \mid X = 0$ is $N(0, 6)$. Therefore, the unit $(0, 1, 1, -1)$ corresponds to the 50% quantile (i.e., the median). Importantly, any other unit satisfying

$$2u_1 - u_2 + u_3 = 0 \quad (66)$$

also corresponds to the 50% quantile. Thus, the 50% quantile corresponds to the subspace given by $\{(u_1, u_2, u_3) \in \mathbb{R}^3 : 2u_1 - u_2 + u_3 = 0\}$. \square

The above example illustrates that quantiles are mixing different kinds of units. However, these units may share some similarity. In the sequel, the Kendall’s τ parameter of the copula will represent a measure of how similar the units within a fixed quantile are.

Inference Based on Frank’s Copula. In this paper, we use Frank’s copula (Frank, 1979) parameterized by the Kendall’s τ correlation coefficient (Kendall, 1948) for the coupling of quantiles of $V_i \mid \text{Pa}_i = \text{pa}_i$ and $V_i^* \mid \text{Pa}_i^* = \text{pa}_i^*$. The joint distributions over the quantiles for values of $\tau \in \{-0.5, 0, 0.5, 1\}$ are shown in Fig. 7. As can be seen from the figure, the larger the τ , the more closely related the quantiles are. For $\tau = 1$, the quantiles pre- and post-recourse are identical. Values of $\tau \in (0, 1)$ still have the property that the expectation of $\mathbb{E}[Q_i^* \mid Q_i = q_i] = q_i$, but the spread

Algorithm 1 Copula Inference from Observational Data

- **Inputs:** Causal Diagram \mathcal{G} , Observational Data \mathcal{D} , Kendall's τ Parameter of the Frank copula, Recourse Action $R = r$, Individual $V = v$, Number of Monte Carlo samples N .
- 1: **for** $V_i \in V$ in topological order **do**
- 2: perform quantile regression $V_i \sim \text{Pa}_i$ to learn the quantile function $Q(v_i | \text{pa}_i)$ using \mathcal{D}
- 3: **end for**
- 4: **for** $\text{de}(R) \setminus R \in V$ in topological order **do**
- 5: infer the quantile q_i of $v_i | \text{pa}_i$ from $Q(v_i | \text{pa}_i)$, labeled q_i
- 6: draw N samples from the conditional Frank's copula $C_\tau(\cdot | Q_1 = q_i)$ labeled $\{q_i^{(k)}\}_{k=1:n}$
- 7: infer the N values of $v_i^{(k)*}$ as

$$v_i^{(k)*} \leftarrow Q^{-1}(q_i^{(k)} | \text{pa}_i^{(k)*}) \quad (67)$$

where parents $\text{pa}^{(k)*}$ are either known or possibly obtained in the previous steps.

8: **end for**

- **Output:** N Monte Carlo values $\{v^{(k)*}\}_{k=1:n}$ under recourse $R = r$ for individual $V = v$.
-

around the expected value increases as τ approaches 0. For $\tau = 0$, the quantiles are independent, representing the so-called independence copula. In this case, given the value of $\text{Pa}_i = \text{pa}_i$ pre-recourse, knowing the value of $V_i = v_i$ does not tell us anything additional about the post-recourse value $V_i^* = v_i^*$ (in other words, only the pre-recourse parent set Pa_i possibly contains information about V_i^*). For $\tau < 0$, the quantile coupling is reversed, so that $\mathbb{E}[Q_i^* | Q_i = q_i] = 1 - q_i$.

In Alg. 1 we describe the procedure for estimating the effect of a recourse action under post-recourse stability based on observational data. The τ parameter of the copula is an input to the algorithm. We remark that at each node V_i , the τ parameter may depend on the parents $\text{Pa}_i = \text{pa}_i$, i.e., $\tau = \tau(\text{pa}_i)$. However, for simplicity, we assume τ is constant for each node V_i and each value of the parents $\text{Pa}_i = \text{pa}_i$. The τ parameter is an input since it cannot be inferred purely from observational data. Therefore, Alg. 1 can be run with varying levels of τ , and we can inspect how this affects the post-recourse distribution through a sensitivity-type analysis with τ acting as the sensitivity parameter (demonstrated empirically in Sec. 4).

The algorithm first iterates through all nodes in the graph and learns the conditional distribution of $V_i | \text{Pa}_i$. Then, for each descendant node of the recourse action $R = r$ that is not subject to intervention itself, we first infer the quantile q_i of the observation $V_i = v_i$ for the individual. Then, based on the Frank's copula C_τ , we sample the N Monte Carlo quantiles conditional on q_i and under the recourse action, labeled $\{q_i^{(k)}\}_{k=1:n}$. We then obtain the Monte Carlo values for $v_i^{(k)*}$ based on the quantiles $q_i^{(k)}$ and the quantile function $Q(\cdot | \text{pa}^{(k)*})$ of the distribution $V_i | \text{pa}_i^{(k)*}$. The main drawback of the procedure is the inherent lack of identifiability of τ from observational data. However, if we have the resulting data from implementing recourse decisions in practice, τ may in fact be inferred. Moreover, we can also test whether the Frank copula model holds for our data, as discussed in the sequel.

3.5 Non-Parametric Inference Under Post-Recourse Stability – Recourse Data

To describe the problem of learning from recourse data, we first provide a formal definition of recourse data:

Definition 13 (Recourse Data) Let $\mathcal{M}_{R=r}^{P(u^*)} = \langle \mathcal{F}, P(u), R = r, P(u^* | u) \rangle$ be a recourse SCM. Samples are drawn from $\mathcal{M}_{R=r}^{P(u^*)}$ as follows:

(S1) Draw $U = u$ according to $P(u)$,

(S2) Evaluate $V(u) = v$ according to \mathcal{F} ,

(S3) If $\widehat{Y}(u) = 0$, then

(S3a) Draw $U^* = u^*$ according to $P(u^* | u)$,

(S3b) Evaluate $V^*(u^*)$ based on mechanisms $\mathcal{F}_{R \leftarrow r}$.

Let $P_{R=r}^*(V)$ be the described post-recourse distribution, and let $V_{R=r}^*$ be the random variable describing the post-recourse covariates. \square

The definition of recourse data captures several key characteristics of the setting. Firstly, in (S1)-(S2), the pre-recourse observational sample is obtained. After this, in (S3), a post-recourse sample is drawn based on $P(u^* | u)$ and $\mathcal{F}_{R \leftarrow r}$, but this happens only for *individuals with $\widehat{Y}(u) = 0$* . In Def. 13, we suppose that a recourse sample is drawn for all individuals with $\widehat{Y}(u) = 0$, while in practice only a subset of these individuals may implement recourse actions (suppose $S = 1$ is an indicator of whether recourse is implemented). Cases where the probability of implementing recourse, written $P(S = 1 | U = u)$ depends only on the observed variables,

$$P(S = 1 | U = u) = P(S = 1 | V(u)) \quad (68)$$

can also be handled using the methods described in this paper, assuming that we have

$$\delta < P(S = 1 | V(u)) < 1 - \delta \quad (69)$$

for all u with $\widehat{Y}(u) = 0$ (in this case, a simple reweighing of the recourse data would be, in infinite samples, equivalent to obtaining recourse data on every individual with $\widehat{Y}(u) = 0$). More complex cases, where the probability of implementing recourse depends on U or even U^* are not considered, and are left for future work.

While the full recourse SCM $\mathcal{M}_{R=r}^{P(u^*)}$ would, in principle, allow us to generate post-recourse samples for any unit u , in practice, no post-recourse samples will be available for units with $\widehat{Y}(u) = 1$ since those with a positive outcome will not implement recourse actions. Therefore, any recourse data collected in practice will always be conditional on $\widehat{Y} = 0$, representing a kind of selection bias (Hernán et al., 2004; Bareinboim and Pearl, 2012; Bareinboim et al., 2014).

Example 11 (Exam Repetitions – Recourse Data) Consider the recourse SCM from Ex. 6 and the distribution over post-recourse latent variables $P^{(a)}$. Then, consider all units with $X(u) = 0$. Based on the f_T mechanism in Eq. 32, only the units with $2u_1 - u_2 + u_3 < 1$ will have $T(u) < 1$ and would, therefore, possibly be interested in implementing recourse. For those with $X(u) = 1$, only units with $2u_1 - u_2 + u_3 < 0$ would have $T(u) < 1$ and thus may implement recourse. In conclusion, recourse data would be available for

$$u \in \mathcal{U} \text{ s.t. } \{u_X = 0, 2u_1 - u_2 + u_3 < 1\} \vee \{u_X = 1, 2u_1 - u_2 + u_3 < 0\} \quad (70)$$

\square

Based on the definition of recourse data, we can now define the concept of recourse learning:

Definition 14 (Recourse Learning) Consider an SCM \mathcal{M} and its observational distribution $P(V)$. Further, consider a collection of recourse distributions $\{P_{R=r}^*(V)\}_{R=r \in \mathcal{R}}$ as described in Def. 13. The task of learning from recourse data is to recover the conditional distribution

$$V_{R=r}^* | V = v \quad (72)$$

based on samples from $P(V)$ and $\{P_{R'=r'}^*(V)\}_{R'=r' \in \mathcal{R}}$. \square

Algorithm 2 Copula Inference from Observational and Recourse Data

- **Inputs:** Causal Diagram \mathcal{G} , Observational Data \mathcal{D} , Recourse $do(R = r)$, Recourse Data \mathcal{D}^* .
 - 1: **for** $V_i \in \{\text{de}(R) \setminus R\}$ in topological order **do**
 - 2: learn quantile function of $v_i \mid \text{pa}_i$, labeled $Q(v_i \mid \text{pa}_i)$, using \mathcal{D}
 - 3: infer pre-recourse quantile $\hat{q}_i^{(k)}$ of $v_i^{(k)} \mid \text{pa}_i$ from $Q(v_i \mid \text{pa}_i) \forall k$ appearing in both $\mathcal{D}, \mathcal{D}^*$
 - 4: infer post-recourse quantile $\hat{q}_i^{(k)*}$ of $v_i^{(k)*} \mid \text{pa}_i^{(k)*}$ from $Q(v_i^* \mid \text{pa}_i^*)$
 - 5: compute $\hat{\tau}_i$ as maximizer of conditional copula likelihood,

$$\hat{\tau}_i = \arg \max_{\tau} \log f_{\tau}(\{\hat{q}_i^{(k)*}\}_{k=1}^{n_{\text{rec}}} \mid \{\hat{q}_i^{(k)}\}_{k=1}^{n_{\text{rec}}}). \quad (71)$$
 - 6: compute baseline Cramer-von Mises statistic $S_{i,n}$ as in Eq. 74
 - 7: **for** $m = 1, \dots, M$ **do** ▷ Bootstrap loop
 - 8: draw $\hat{q}_i^{(k,m)*} \sim \text{Frank}(\hat{\tau}_i \mid \hat{q}_i^{(k)})$ for all k ▷ Simulate from H_0
 - 9: set $\tilde{v}_i^{(k,m)*} = Q^{-1}(\hat{q}_i^{(k,m)*} \mid \text{pa}_i^{(k)*})$ ▷ Map quantile to V_i value
 - 10: refit $Q^{(m)}$ on a bootstrapped dataset $\mathcal{D}^{(m)}$
 - 11: re-estimate $\hat{q}_i^{(k,m)}, \hat{q}_i^{(k,m)*}$ from $Q^{(m)}$ applied to $v_i^{(k)}, \tilde{v}_i^{(k,m)*}$, respectively
 - 12: compute $\hat{\tau}_i^{(m)}$ from pairs $(\hat{q}_i^{(k,m)}, \hat{q}_i^{(k,m)*})$
 - 13: compute $S_{i,n}^{(m)}$ as the Cramer-von Mises statistic for the m -th bootstrap sample (Eq. 75)
 - 14: **end for**
 - 15: compute the p-value $p_i = \frac{1}{M} \sum_{m=1}^M \mathbb{1}(S_{i,n} < S_{i,n}^{(m)})$
 - 16: **end for**
 - **Output:** τ estimates $\hat{\tau}_i$, p-values p_i
-

The task of learning from recourse data represents a natural setting in algorithmic recourse. Consider an institution interested in implementing a recourse policy. At first, they may attempt to infer what would happen to individuals under recourse based on observational data from $P(V)$ only. Then, they may start issuing recourse recommendations $\{R = r\} \in \mathcal{R}$, and recording samples of the variables after recourse was implemented, i.e., from $P_{R=r}^*(V)$. These additional samples can help a great deal in inferring the effects of recourse actions, since they allow one to verify whether post-recourse stability holds and whether Frank’s copula model holds.

Selection Bias Based on $\hat{Y} = 0$. As indicated in the step (S3) of recourse data sampling in Def. 13, a post-recourse sample will only be available for units u for whom $\hat{Y}(u) = 0$. Therefore, the recourse data we have access to will always be subject to *selection bias* (Hernán et al., 2004; Bareinboim and Pearl, 2012; Bareinboim et al., 2014) based on the outcome \hat{Y} . The key consequence of this is that pre-recourse quantiles for individuals who have post-recourse data *will not follow a* $\text{Unif}[0, 1]$ *distribution*. By definition, if for a variable $V_i = v_i$, we computed its quantile conditional on the parents $\text{Pa}_i = \text{pa}_i$, and we pooled all the obtained quantiles across the population, the resulting distribution of quantiles would be $\text{Unif}[0, 1]$. However, when we focus on the quantiles of those for whom ultimately $\hat{Y} = 0$, the resulting distribution of quantiles need not equal $\text{Unif}[0, 1]$. Care needs to be taken to account for this key feature of recourse data. Copula models are usually intended for distributional coupling where marginal distributions are $\text{Unif}[0, 1]$, which is not the case for our setting.

Inferring τ from Recourse Data. A procedure for inference from observational and recourse data is described in Alg. 2. This algorithm is intended for cases when a large amount of observational together with a small amount of recourse data is available. First, based on observational data, the quantile function $Q(v_i \mid \text{pa}_i)$ is learned. Then, for individuals with recourse data, pre and post-recourse quantiles for the sample v_i, v_i^* are inferred from the quantile functions $Q(\cdot \mid \text{pa}_i), Q(\cdot \mid \text{pa}_i^*)$.

Based on the coupling of quantiles (since pre and post-recourse samples are paired) we can perform conditional maximum likelihood estimation to infer the most likely τ for the generated data, such that the likelihood $\log f_\tau(q_i^* | q_i)$ is maximized (see Appendix A for conditional likelihood expressions).

Goodness-of-fit for Frank’s copula. After estimating τ for each node V_i , we verify whether Frank’s copula is an appropriate model for the data by performing a goodness-of-fit test. We adapt the approach of Genest et al. (2009) to account for estimation error in the quantile estimates \hat{q}_i, \hat{q}_i^* , which are obtained by fitting a quantile function $Q(v_i | \text{pa}_i)$ (see Alg. 2). Since the quantiles are estimated rather than computed from known marginals, the estimation error, if unaccounted for, may lead to a spurious rejection of the null hypothesis. We first compute the empirical cumulative distribution function (ECDF) for the estimated quantile pairs \hat{q}_i, \hat{q}_i^* pre- and post-recourse, defined as

$$C_{i,n}(u, v) = \frac{1}{n_{\text{rec}}} \sum_{k=1}^{n_{\text{rec}}} \mathbb{1}(u \leq \hat{q}_i^{(k)} \wedge v \leq \hat{q}_i^{(k)*}). \quad (73)$$

We estimate $\hat{\tau}_i$ by maximizing the conditional Frank copula likelihood, and denote by $C_{i,\hat{\tau}}$ the ECDF of the Frank($\hat{\tau}_i$) copula conditional on $\hat{q}_i^{(k)}$, approximated via Monte Carlo. The baseline Cramer-von Mises statistic is

$$S_{i,n} := \sum_{k=1}^{n_{\text{rec}}} \left(C_{i,n}(\hat{q}_i^{(k)}, \hat{q}_i^{(k)*}) - C_{i,\hat{\tau}}(\hat{q}_i^{(k)}, \hat{q}_i^{(k)*}) \right)^2. \quad (74)$$

To construct the null distribution for the Cramer-von Mises statistic, we employ a bootstrap procedure. For each bootstrap repetition m , we: (i) draw post-recourse quantiles $\hat{q}_i^{(k,m)*}$ from the fitted $C_{i,\hat{\tau}}(\cdot | \hat{q}_i^{(k)})$ copula; (ii) map the quantiles using the inverse of the learned quantile function to obtain bootstrap samples $\tilde{v}_i^{(k,m)*} = Q^{-1}(\hat{q}_i^{(k,m)*} | \text{pa}_i^{(k)*})$; (iii) refit the quantile function $Q^{(m)}$ on bootstrapped sample of the data $\mathcal{D}^{(m)}$; and (iv) re-estimate the pre- and post-recourse quantiles $\hat{q}_i^{(k,m)}, \hat{q}_i^{(k,m)*}$ using $Q^{(m)}$. The sampling in steps (i) and (iii) ensures that the bootstrap null distribution of the test statistic reflects both the estimation noise in the quantile function and the randomness of the post-recourse quantiles. Note that the bootstrap procedure conditions on the pre-recourse quantiles $\hat{q}_i^{(k)}$ throughout, which ensures that selection on $\hat{Y} = 0$ does not affect inference (i.e., we do not assume uniform margins at any point). The Cramer-von Mises statistic for the m -th bootstrap sample is given by

$$S_{i,n}^{(m)} := \sum_{k=1}^{n_{\text{rec}}} \left(C_{i,n}^{(m)}(\hat{q}_i^{(k,m)}, \hat{q}_i^{(k,m)*}) - C_{i,\hat{\tau}^{(m)}}(\hat{q}_i^{(k,m)}, \hat{q}_i^{(k,m)*}) \right)^2, \quad (75)$$

where $\hat{\tau}_i^{(m)}$ is re-estimated from the bootstrap pairs $(\hat{q}_i^{(k,m)}, \hat{q}_i^{(k,m)*})$. The p-value is the empirical quantile of $S_{i,n}$ within $\{S_{i,n}^{(m)}\}_{m=1}^M$, computed in Line 15. The key question we discuss in the sequel is how to proceed when the null hypothesis of Frank’s copula is rejected.

3.6 Non-Parametric Inference Under Post-Recourse Instability – Recourse Data

The goodness-of-fit hypothesis test described in Alg. 2 may be rejected for several reasons³:

- (a) The quantile coupling cannot be represented by Frank’s copula,
- (b) Post-recourse stability from Def. 11 does not hold.

For case (a), a different copula family may be more appropriate, or a non-parametric model could be introduced (such extensions would be variations of Algs. 1 and 2). More interestingly, we now discuss what happens for case (b) when Post-recourse stability is violated.

3. The test may be rejected if the causal diagram is misspecified and latent confounding exists (i.e., the true model is Semi-Markovian). However, we do not focus on this case, and assume a correctly specified diagram.

Algorithm 3 Copula-Free Inference from Observational and Recourse Data

- **Inputs:** Causal Diagram \mathcal{G} , Observational Data \mathcal{D} , Recourse Action $do(R = r)$, Recourse Data \mathcal{D}^* , Number of Monte Carlo samples N .
- 1: **for** $de(R) \setminus R \in V$ in topological order **do**
- 2: regress $V_i^* \sim V_i + Pa_i + Pa_i^*$ to learn quantile function $Q(v_i^* | v_i, pa_i, pa_i^*)$ using $\mathcal{D}, \mathcal{D}^*$
- 3: draw N samples from $\text{Unif}[0, 1]$, labeled $\{q^{(k)}\}_{k=1:n}$, and infer values of $v_i^{(k)*}$ as

$$v_i^{(k)*} \leftarrow Q^{-1}(q^{(k)} | v_i, pa_i, pa_i^{(k)*}) \quad (76)$$

where values $v_i, pa_i, pa_i^{(k)*}$ are either known or obtained in the previous steps.

- 4: **end for**

- **Output:** N Monte Carlo values $\{v^{(k)*}\}_{k=1}^N$ under recourse $R = r$ for individual $V = v$.

The procedure for this case is described in Alg. 3. It assumes a large amount of both observational and recourse data. Crucially, the quantile function $Q(v_i | pa_i)$ learned in the observational data is no longer applicable to recourse data (since margin stability does not hold, see discussion around Ex. 9). Nonetheless, we may still be able to learn the correct quantile function based on *recourse data*. The key question here is whether observational samples could still be useful for inference even if margin stability does not hold. Even though margin stability may not be satisfied, it is still the case that the quantile of $V_i = v_i$ in the observational distribution $V_i | Pa_i = pa_i$ shares information with the quantile of $V_i^* = v_i^*$ in the interventional, post-recourse distribution $V_i^* | Pa_i^* = pa_i^*$. To leverage this connection, we learn the quantile of the recourse random variable V_i^* conditional on (i) the observed parents pa_i ; (ii) the initial observed value v_i ; (iii) the recourse parents pa_i^* . The usage of (v_i, pa_i) in the quantile regression in Line 2 serves as conditioning on the quantile q_i of $v_i | pa_i$, since we know this quantile may be correlated with the post-recourse quantile q_i^* , and may thus improve inference. A theoretical basis for the regression in Line 2 is developed in the following theorem, in the case of linear models:

Theorem 2 (EMSPE Reduction) *Let Y denote a variable in a linear Gaussian SCM, and let X denote $pa(Y)$. Pre-recourse values are denoted by a subscript 0, while post-recourse values have no subscript. The pre-recourse and post-recourse models can be written as*

$$Y_0 \leftarrow X_0 \beta_0 + \varepsilon_0, \quad Y \leftarrow X \beta + \varepsilon, \quad (77)$$

where $(X_{0,i}, X_i)$ are i.i.d. jointly Gaussian with $\mathbb{E}[X_{0,i} X_{0,i}^\top] = \Sigma_0 \succ 0$ and $\mathbb{E}[X_i X_i^\top] = \Sigma \succ 0$. Further, assume the noise terms $\varepsilon_0, \varepsilon$ are independent of covariates, and have a correlation ρ , so that $\varepsilon = \rho \varepsilon_0 + \sqrt{1 - \rho^2} \bar{\varepsilon}$ with $\varepsilon_0 \sim N(0, \sigma^2 I)$, $\bar{\varepsilon} \sim N(0, \sigma^2 I)$, and $\varepsilon_0 \perp \bar{\varepsilon}$. Let $\hat{\beta}_c$ denote the two-stage OLS estimator obtained via:

$$\text{Stage I: } Y_0 \stackrel{\text{OLS}}{\sim} X_0 \text{ to obtain } \hat{\beta}_0 \text{ and } \hat{\varepsilon}_0 = Y_0 - X_0 \hat{\beta}_0, \quad (78)$$


$$\text{Stage II: } Y \stackrel{\text{OLS}}{\sim} \bar{X} = (X \ \hat{\varepsilon}_0) \text{ to obtain } \hat{\beta}_c, \quad (79)$$

Next, we label training data with a superscript a , and test data with b . Let the test prediction risk (conditional on training data) be defined as


$$R_{\text{aug}} := \frac{1}{n_b} \mathbb{E}_{\text{test}} [\|Y^b - \bar{X}^b \hat{\beta}_c\|^2 | \mathcal{D}^t], \quad (80)$$

where $\mathcal{D}^t = \{X_0^a, X^a, \varepsilon_0^a, \bar{\varepsilon}^a\}$ denotes the training data. For the number of training samples $n_a \gg d$, R_{aug} satisfies

$$R_{\text{aug}} = (1 - \rho^2) \sigma^2 + \mathcal{O}_P \left(\sqrt{\frac{d}{n_a}} \right). \quad (81)$$



(a) HELOC causal diagram.



(b) Spread of MC samples for $do(C_3 = 85, D_2 = 20)$.

Figure 8: Causal diagram and spread of $\hat{Y}(u^*)$ values for different values of $\tau \in \{0.7, 0.8, 0.9, 1\}$. Variables under recourse are marked in green.

Since the risk of the standard single-stage $Y \stackrel{\text{OLS}}{\sim} X$, written $R_{\text{std}} := \frac{1}{n_b} \mathbb{E}_{\text{test}} [\|Y^b - X^b \hat{\beta}\|^2 \mid \mathcal{D}^t]$, equals $\sigma^2 + \mathcal{O}_P\left(\frac{d}{n_a}\right)$, we have $R_{\text{aug}} < R_{\text{std}}$ for $\rho \neq 0$ and n_a sufficiently large. \square

The proof is given in Appendix D. The theorem compares two approaches: (i) the standard approach, in which assuming access to recourse data, we predict a recourse variable based only on its recourse parents; and (ii) the augmented two-stage approach that leverages the fact that pre- and post-recourse noise terms are correlated (e.g., $U \not\perp U^*$, as discussed in Sec. 3), and computes pre-recourse residuals $\hat{\varepsilon}$, and adds these residuals into the regression of the post-recourse outcome on its parents, in order to improve inference. In terms of EMSPE on a test set, the standard approach has an irreducible loss of σ^2 and an excess loss of $\mathcal{O}_P\left(\frac{d}{n_a}\right)$. The two stage approach is different, however: the irreducible loss is reduced by a factor depending on the correlation of noise terms, to a level $(1 - \rho^2)\sigma^2$, whereas the excess loss vanishes more slowly compared to the usual case, at the rate $\mathcal{O}_P\left(\sqrt{\frac{d}{n_a}}\right)$. Therefore, for large enough samples, the augmented approach outperforms the standard approach, and in this way Thm. 2 provides a basis for Alg. 3.

4. Experiments

In this section, we apply the procedures in Alg. 1, 2, and 3 to real and semi-synthetic data. Throughout, we use the Home Equity Line of Credit (HELOC) dataset (Fair Isaac Corporation (FICO), 2016). The dataset contains anonymized information on $n = 10459$ applicants who applied for a home equity line of credit. It includes features such as risk scores, income levels, loan-to-value ratios, and others that reflect the creditworthiness of the applicants. The goal is determine whether the applicant was deemed high (“Bad”) or low (“Good”) credit risk. The experiments are accompanied with vignettes for Algs. 1, 2, 3 available in our code repository.

4.1 Alg. 1 – A Sensitivity Approach.

We begin by applying Alg. 1 to the HELOC dataset. The outcome variable $Y \in \{0, 1\}$ indicates whether the individual’s risk for repaying a loan is considered bad or good. First, we fit a random forest (Breiman, 2001) model to predict the outcome Y . We select the top 10 variables according to

variable importance based on the Gini index. We then construct the causal diagram over these 10 variables, by grouping them as follows:

- (G1) months since oldest trade (A_1), months in file (A_2),
- (G2) number of total trades (B_1), number of satisfactory trades (B_2),
- (G3) % installment trades (C_1), % trades with balance (C_2), % trades never delinquent (C_3),
- (G4) months since last credit report inquiry (D_1), revolving balance to credit limit ratio (D_2),
- (G5) external estimate of lending risk (E_1).

In particular, we assume that each group G_i points to each downstream group G_j for $j > i$. The causal diagram is shown in Fig. 8a, where each arrow from one cluster to another corresponds to multiple arrows in the fully expanded causal diagram (full diagram shown in Appendix B). We then apply Alg. 1 and learn the quantile functions $Q(v_i \mid \text{pa}_i)$ using quantile regression forests (Meinshausen and Ridgeway, 2006). We then pick individuals with percent trades never delinquent $C_3 < 70$, and revolving balance ratio $D_2 > 50$, who obtained a negative decision from \widehat{Y} . We pick the recourse action $do(C_3 = 85, D_2 = 20)$, and we generate $N = 100$ recourse MC samples for each individual, and with different values of τ . For such individuals, we expect the recourse samples to have higher value of the predictor \widehat{Y} , with the spread increasing for τ values closer to 0. A prototypical output of such a sensitivity analysis is shown in Fig. 8b. As expected, τ values closer to 0 correspond to a larger spread, and from the MC samples we can also compute the probability of crossing the decision boundary, i.e., $P_\tau(\widehat{Y}_{R=\tau}(u^*) > \frac{1}{2})$, which in general depends on τ .

4.2 Alg. 2 – Copula Goodness-of-fit.

For applying Alg. 2, access to recourse data samples is needed. To do so, we create a semi-synthetic HELOC dataset. To simplify the setting slightly, we further reduce the number of variables, by dropping variables A_1, B_1, B_2 , and C_1 . The assumed causal diagram for the semi-synthetic (SeS) HELOC dataset is given in Fig. 9a. Variables C_2, C_3, D_2, E_1 that are percentages are scaled into the $[0, 1]$ interval. Then, we fit the maximum likelihood estimator for the following SCM functional form:

$$V_i = A_1 : \quad A_1 \leftarrow N(\mu, \sigma^2) \quad (82)$$

$$V_i \in \{C_2, C_3, D_2, E_1\} : V_i \leftarrow \text{Beta}(\alpha(\text{pa}_i), \beta(\text{pa}_i)) \quad (83)$$

$$V_i = D_1 : \quad D_1 \leftarrow \text{Geom}(p(\text{pa}_i)) \quad (84)$$

In particular, μ, σ^2 are fixed parameters, whereas functions $\alpha(\cdot), \beta(\cdot), p(\cdot)$ are linear functions of the respective arguments. For each variable, we then assume that the mechanisms \mathcal{F} pre- and post-recourse remain the same. However, for the quantile coupling, we distinguish two versions. First, in the SeS-HELOC A we set that quantiles q_i, q_i^* are coupled by Frank’s copula with $\tau = \frac{2}{3}$. For the SeS-HELOC B, we set that $Q_i^* \mid Q_i = q_i \sim \text{Unif}[q_i, 1]$, i.e., the post-recourse quantile is strictly larger than the initial quantile. For this dataset, Frank’s copula is not valid, and post-recourse stability is not satisfied. In Fig. 9b we show the empirical cumulative distribution function (ECDF) of the p-values for the hypothesis test for datasets A and B over $n_{\text{rep}} = 100$ repetitions of $5 \cdot 10^3$ observational samples and 500 recourse samples from individuals with a negative decision. For dataset A, where the null hypothesis of Frank’s copula is true, the distribution is nearly uniform, confirming that the bootstrap procedure in Alg. 2 correctly accounts for quantile estimation error. For dataset B, the distribution is skewed heavily towards 0, indicating that we can detect violations of the copula coupling from recourse data.




Figure 9: (a) causal diagram for SeS-HELOC datasets; (b) spread of p-values for H_0 for cases A, B; (c) spread of Kolmogorov-Smirnov statistics for copula-free and copula-based learning for case B.

4.3 Alg. 3 – Copula-free Learning.

The final challenge we have is inferring effects of recourse actions when post-recourse stability does not hold, such as in the SeS-HELOC B dataset. In this case, we apply Alg. 3 with 10^4 pre-recourse samples, and assume that for all individuals with $\hat{Y} = 0$ we also have post-recourse samples. For each individual in a separate test set of size $n_{\text{test}} = 50$, we compute 100 MC samples of $\hat{Y}_{R=r}^{\text{Alg.3}}(u^*)$ based on Alg. 3. As a comparison, we also compute 100 MC samples by assuming (i) Frank’s copula holds true, and estimating the τ parameter, labeled $\hat{Y}_{R=r}^F(u^*)$; (ii) by directly regressing $V_i^* \sim \text{Pa}_i^*$, labeled $\hat{Y}_{R=r}^{\text{exp}}(u^*)$. From the SCM itself, we compute 100 MC samples from the true underlying post-recourse distribution, labeled $\hat{Y}_{R=r}(u^*)$. For each individual in the test set, we compute the Kolmogorov-Smirnov statistics D of the MC samples $\hat{Y}_{R=r}^{\text{Alg.3}}(u^*)$, $\hat{Y}_{R=r}^F(u^*)$, and $\hat{Y}_{R=r}^{\text{exp}}(u^*)$ compared to $\hat{Y}_{R=r}(u^*)$. The distributions of the statistics D for the 50 individuals are shown in Fig. 9c. As the figure illustrates, the distance from the true distribution $\hat{Y}_{R=r}(u^*)$ is smaller for $\hat{Y}_{R=r}^{\text{Alg.3}}(u^*)$ from Alg. 3 than for $\hat{Y}_{R=r}^{\text{exp}}(u^*)$, giving empirical verification of Thm. 2 in a non-parametric setting. Furthermore, $\hat{Y}_{R=r}^{\text{exp}}(u^*)$ is still better than the copula-based estimator $\hat{Y}_{R=r}^F(u^*)$, since the underlying conditions for copula-based learning from Def. 11 are violated.

5. Conclusion

In this paper, we introduced a novel causal framework for combining observational and experimental data on the same individuals in the context of algorithmic recourse (Defs. 10, 13, 14). We discussed conditions that allow inference from observational data (Def. 11 and Thm. 1), and provided a copula-based sensitivity analysis for effects of recourse (Alg. 1). When recourse data is available, copula parameters can be inferred, and the copula model can be tested (Alg. 2). Finally, when the copula model does not hold, a more general learning approach requiring more experimental recourse data is still feasible (Alg. 3). These claims were empirically validated on both real and semi-synthetic HELOC data.

References

- Mendi Sesmu Arbous, DE Grobbee, JW Van Kleef, JJ De Lange, HHAJM Spoormans, P Touw, FM Werner, and Anneke Elina Elvira Meursing. Mortality associated with anaesthesia: a qualitative analysis to identify risk factors. *Anaesthesia*, 56(12):1141–1153, 2001.
- E. Bareinboim, J. Tian, and J. Pearl. Recovering from selection bias in causal and statistical inference. In C. E. Brodley and P. Stone, editors, *Proceedings of the 28th AAAI Conference on Artificial Intelligence*, pages 2410–2416, Menlo Park, CA, 2014. AAAI Press.
- Elias Bareinboim and Judea Pearl. Controlling selection bias in causal inference. In *Artificial Intelligence and Statistics*, pages 100–108. PMLR, 2012.
- Elias Bareinboim, Juan D. Correa, Duligur Ibeling, and Thomas Icard. On pearl’s hierarchy and the foundations of causal inference. In *Probabilistic and Causal Inference: The Works of Judea Pearl*, page 507–556. Association for Computing Machinery, New York, NY, USA, 1st edition, 2022.
- Solon Barocas, Andrew D Selbst, and Manish Raghavan. The hidden assumptions behind counterfactual explanations and principal reasons. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, pages 80–89, 2020.
- Leo Breiman. Random forests. *Machine learning*, 45:5–32, 2001.
- Tim Brennan, William Dieterich, and Beate Ehret. Evaluating the predictive validity of the compass risk and needs assessment system. *Criminal Justice and Behavior*, 36(1):21–40, 2009.
- Mustafa Caglayan, Oleksandr Talavera, Lin Xiong, and Jing Zhang. What does not kill us makes us stronger: the story of repetitive consumer loan applications. *The European journal of finance*, 28(1):46–65, 2022.
- Juan Correa, Sanghack Lee, and Elias Bareinboim. Nested counterfactual identification from arbitrary surrogate experiments. *Advances in Neural Information Processing Systems*, 34:6856–6867, 2021.
- Fair Isaac Corporation (FICO). Home equity line of credit, 2016. URL <https://community.fico.com/s/explainable-machine-learning-challenge>.
- Maurice J Frank. On the simultaneous associativity of $f(x, y)$ and $x + y - f(x, y)$. *Aequationes mathematicae*, 19:194–226, 1979.
- Christian Genest, Bruno Rémillard, and David Beaudoin. Goodness-of-fit tests for copulas: A review and a power study. *Insurance: Mathematics and economics*, 44(2):199–213, 2009.
- Miguel A Hernán, Sonia Hernández-Díaz, and James M Robins. A structural approach to selection bias. *Epidemiology*, pages 615–625, 2004.
- Piotr Jaworski, Fabrizio Durante, Wolfgang Karl Hardle, and Tomasz Rychlik. *Copula theory and its applications*, volume 198. Springer, 2010.
- Shalmali Joshi, Oluwasanmi Koyejo, Warut Vijitbenjaronk, Been Kim, and Joydeep Ghosh. Towards realistic individual recourse and actionable explanations in black-box decision making systems. *arXiv preprint arXiv:1907.09615*, 2019.
- Amir-Hossein Karimi, Julius Von Kügelgen, Bernhard Schölkopf, and Isabel Valera. Algorithmic recourse under imperfect causal knowledge: a probabilistic approach. *Advances in neural information processing systems*, 33:265–277, 2020.

- Amir-Hossein Karimi, Bernhard Schölkopf, and Isabel Valera. Algorithmic recourse: from counterfactual explanations to interventions. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, pages 353–362, 2021.
- Amir-Hossein Karimi, Gilles Barthe, Bernhard Schölkopf, and Isabel Valera. A survey of algorithmic recourse: contrastive explanations and consequential recommendations. *ACM Computing Surveys*, 55(5):1–29, 2022.
- Maurice George Kendall. *Rank correlation methods*. Griffin, 1948.
- Amir E Khandani, Adlar J Kim, and Andrew W Lo. Consumer credit-risk models via machine-learning algorithms. *Journal of Banking & Finance*, 34(11):2767–2787, 2010.
- Beatrice Laurent and Pascal Massart. Adaptive estimation of a quadratic functional by model selection. *Annals of statistics*, pages 1302–1338, 2000.
- Sanghack Lee, Juan D Correa, and Elias Bareinboim. General identifiability with arbitrary surrogate experiments. In *Uncertainty in artificial intelligence*, pages 389–398. PMLR, 2020.
- Divyat Mahajan, Chenhao Tan, and Amit Sharma. Preserving causal constraints in counterfactual explanations for machine learning classifiers. *arXiv preprint arXiv:1912.03277*, 2019.
- John F Mahoney and James M Mohen. Method and system for loan origination and underwriting, October 23 2007. US Patent 7,287,008.
- Nicolai Meinshausen and Greg Ridgeway. Quantile regression forests. *Journal of machine learning research*, 7(6), 2006.
- Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, New York, 2000. 2nd edition, 2009.
- Donald B Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology*, 66(5):688, 1974.
- Shubham Sharma, Jette Henderson, and Joydeep Ghosh. Certifai: A common framework to provide explanations and analyse the fairness and robustness of black-box models. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, pages 166–172, 2020.
- M Sklar. Fonctions de répartition à n dimensions et leurs marges. In *Annales de l’ISUP*, volume 8, pages 229–231, 1959.
- Jin Tian and Judea Pearl. Probabilities of causation: Bounds and identification. *Annals of Mathematics and Artificial Intelligence*, 28(1):287–313, 2000.
- Berk Ustun, Alexander Spangher, and Yang Liu. Actionable recourse in linear classification. In *Proceedings of the conference on fairness, accountability, and transparency*, pages 10–19, 2019.
- Suresh Venkatasubramanian and Mark Alfano. The philosophical basis of algorithmic recourse. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, pages 284–293, 2020.
- Roman Vershynin. High-dimensional probability. *University of California, Irvine*, 10(11):31, 2020.
- Julius Von Kügelgen, Amir-Hossein Karimi, Umang Bhatt, Isabel Valera, Adrian Weller, and Bernhard Schölkopf. On the fairness of causal algorithmic recourse. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, pages 9584–9594, 2022.
- Sandra Wachter, Brent Mittelstadt, and Chris Russell. Counterfactual explanations without opening the black box: Automated decisions and the gdpr. *Harv. JL & Tech.*, 31:841, 2017.

Appendix A. Conditional Likelihood of Frank's copula

In this appendix, we provide the expression for the conditional likelihood of the bivariate Frank's copula. Frank's copula with parameter θ and variables u, v has the following cumulative distribution function:

$$C_\theta(u, v) = -\frac{1}{\theta} \log \left[1 + \frac{(\exp(-\theta u) - 1)(\exp(-\theta v) - 1)}{\exp(-\theta) - 1} \right]. \quad (85)$$

To compute the conditional likelihood $P(V = v | U = u)$, we need to find the joint density of $f_\theta(u, v)$ of $C_\theta(u, v)$. Note that

$$P(V = v | U = u) = \frac{P(V = v, U = u)}{P(U = u)} = P(V = v, U = u) \quad (86)$$

since $P(U = u) = 1$ is implied by the fact that the margins of a copula are $\text{Unif}[0, 1]$. The density of the copula $P(V = v, U = u)$ can be obtained as

$$f_\theta(u, v) = \frac{\partial^2}{\partial u \partial v} C_\theta(u, v) = \frac{-\theta \exp(-\theta(u + v)) \cdot (\exp(-\theta) - 1)}{((\exp(-\theta) - 1) + (\exp(-\theta u) - 1)(\exp(-\theta v) - 1))^2}. \quad (87)$$

The conditional log-likelihood given samples $\{u_i, v_i\}_{i=1:n}$, $\log L_\theta(\{v_i\}_{i=1}^n | \{u_i\}_{i=1}^n)$ can thus be computed as

$$\sum_{i=1}^n \log \left[\frac{-\theta \exp(-\theta(u_i + v_i)) \cdot (\exp(-\theta) - 1)}{((\exp(-\theta) - 1) + (\exp(-\theta u_i) - 1)(\exp(-\theta v_i) - 1))^2} \right], \quad (88)$$

and conditional maximum likelihood estimation is performed as

$$\hat{\theta} = \arg \max_{\theta} \log L_\theta(\{v_i\}_{i=1}^n | \{u_i\}_{i=1}^n). \quad (89)$$

We remark that the Kendall's τ parameter for Frank's copula is related to the θ parameter as:

$$\tau(\theta) = 1 + \frac{4}{\theta} [D_1(\theta) - 1] \quad (90)$$

where $D_1(\theta)$ is the Debye function $D_1(\theta) = \frac{1}{\theta} \int_0^\theta \frac{t}{e^t - 1} dt$.

Appendix B. HELOC Causal Diagram

The causal diagram of the HELOC dataset used in Sec. 4 and the application of Alg. 1 is given in Fig. 10.

Appendix C. Thm. 1 Proof

Proof [Thm. 1 proof] To show that

$$f_{V_i}(\text{pa}_i, U_i) \stackrel{d}{=} f_{V_i}(\text{pa}_i, U_i^*) \quad \forall i, \text{pa}_i \text{ fixed}, \quad (91)$$

we show that U_i^* and U_i follow the same distributions according to the theorem assumptions, i.e.,

$$U_i^* \stackrel{d}{=} U_i. \quad (92)$$




Figure 10: HELOC full causal diagram.

From Eq. 92, the main claim follows, since $U_i^* \stackrel{d}{=} U_i \implies f(U_i^*) \stackrel{d}{=} f(U_i)$ for any function f . Now, we expand the distribution $P(U_i^*)$ as follows:

$$P(U_i^* = u_i) = P(U_i^{*(v)} = u_i^{(v)}, U_i^{*(f)} = u_i^{(f)}) \quad (93)$$

$$= P(U_i^{*(f)} = u_i^{(f)})P(U_i^{*(v)} = u_i^{(v)} | U_i^{*(f)} = u_i^{(f)}) \quad (94)$$

$$= P(U_i^{(f)} = u_i^{(f)})P(U_i^{*(v)} = u_i^{(v)} | U_i^{*(f)} = u_i^{(f)}) \quad (\text{Eq. 55}) \quad (95)$$

$$= P(U_i^{(f)} = u_i^{(f)})P(U_i^{(v)} = u_i^{(v)} | U_i^{(f)} = u_i^{(f)}) \quad (\text{Eq. 56}) \quad (96)$$

$$= P(U_i^{(v)} = u_i^{(v)}, U_i^{(f)} = u_i^{(f)}) \quad (97)$$

$$= P(U_i = u_i). \quad (98)$$

Therefore, we have that $U_i^* \stackrel{d}{=} U_i$ and the claim follows. ■

Appendix D. Linear Recourse – Thm. 2

In this section, we prove Thm. 2. We first introduce the required notation and explain the theorem statement in more detail.

D.1 Setting and Notation

We observe data from two linear models, referred to as *pre-recourse* (subscript 0) and *post-recourse*, sharing a common set of n units. The data are split into a training set (A) of size n_a and a test set (B) of size n_b , with $n = n_a + n_b$.

Pre-recourse model. The pre-recourse response Y_0 satisfies

$$Y_0 = X_0\beta_0 + \varepsilon_0, \quad (99)$$

where $X_0 \in \mathbb{R}^{n \times d}$ are covariates, $\beta_0 \in \mathbb{R}^d$ is the coefficient vector, and ε_0 is Gaussian noise.

Post-recourse model. The post-recourse response satisfies

$$Y = X\beta + \varepsilon, \quad (100)$$

where $X \in \mathbb{R}^{n \times d}$ are covariates, $\beta \in \mathbb{R}^d$ is the coefficient vector, and ε is Gaussian noise.

Noise correlation. The post-recourse noise decomposes as

$$\varepsilon = \rho \varepsilon_0 + \sqrt{1 - \rho^2} \bar{\varepsilon}, \quad (101)$$

where $\rho \in (0, 1)$ is the correlation between ε and ε_0 , and $\bar{\varepsilon}$ is an independent noise term.

Matrix notation. We write $X_0^a \in \mathbb{R}^{n_a \times d}$, $X^a \in \mathbb{R}^{n_a \times d}$ for the training covariate matrices, and similarly with superscript b for the test set. We write $\varepsilon_0^a, \varepsilon^a, \bar{\varepsilon}^a \in \mathbb{R}^{n_a}$ for the corresponding noise vectors. Define the projection matrix

$$P_0^a := X_0^a (X_0^{a\top} X_0^a)^{-1} X_0^{a\top}. \quad (102)$$

Two-stage procedure.

Stage I. Fit OLS on the pre-recourse training data:

$$\hat{\beta}_0 = (X_0^{a\top} X_0^a)^{-1} X_0^{a\top} Y_0^a, \quad \hat{\varepsilon}_0^a = Y_0^a - X_0^a \hat{\beta}_0 = (I - P_0^a) \varepsilon_0^a. \quad (103)$$

Stage II. Form the augmented design matrix $\bar{X}^a = (X^a \ \hat{\varepsilon}_0^a) \in \mathbb{R}^{n_a \times (d+1)}$ and fit OLS:

$$\hat{\beta}_c = (\bar{X}^{a\top} \bar{X}^a)^{-1} \bar{X}^{a\top} Y^a. \quad (104)$$

True augmented coefficient. Substituting (101) into (100) and using the decomposition $\varepsilon_0 = \hat{\varepsilon}_0 + (\varepsilon_0 - \hat{\varepsilon}_0)$, the post-recourse response can be written as

$$Y = X\beta + \rho \hat{\varepsilon}_0 + N, \quad (105)$$

where the residual noise is

$$N := \rho (\varepsilon_0 - \hat{\varepsilon}_0) + \sqrt{1 - \rho^2} \bar{\varepsilon}. \quad (106)$$

In vector form, $Y^a = \bar{X}^a \beta_c + N^a$ with the true augmented coefficient

$$\beta_c := \begin{pmatrix} \beta \\ \rho \end{pmatrix} \in \mathbb{R}^{d+1}. \quad (107)$$

Prediction on test set. Let the training data $X_0^a, X^a, \varepsilon_0^a, \bar{\varepsilon}^a$ be denoted by \mathcal{D}^t . For the test set, define $\hat{\varepsilon}_0^b = Y_0^b - X_0^b \hat{\beta}_0$ and $\bar{X}^b = (X^b \ \hat{\varepsilon}_0^b) \in \mathbb{R}^{n_b \times (d+1)}$. The EMSPE of the two-stage estimator is

$$R_{\text{aug}} := \frac{1}{n_b} \mathbb{E}[\|Y^b - \bar{X}^b \hat{\beta}_c\|^2 \mid \mathcal{D}^t], \quad (108)$$

and the baseline EMSPE using standard OLS on the post-recourse model alone is

$$R_{\text{std}} := \frac{1}{n_b} \mathbb{E}[\|Y^b - X^b \hat{\beta}\|^2 \mid \mathcal{D}^t], \quad \hat{\beta} = (X^{a\top} X^a)^{-1} X^{a\top} Y^a. \quad (109)$$

All expectations are conditional on \mathcal{D}^t .

D.2 Main Result

Theorem 3 (EMSPE reduction) *Assuming $n_a \gg d$, we have*

$$R_{\text{aug}} = (1 - \rho^2)\sigma^2 + \mathcal{O}_P\left(\sqrt{\frac{d}{n_a}}\right). \quad (110)$$

For n_a sufficiently large, $R_{\text{aug}} < R_{\text{std}}$. \square

Proof [Thm. 2] Let $\mathcal{D}^t := \{X_0^a, X^a, \varepsilon_0^a, \bar{\varepsilon}^a\}$ denote the training data. We start by expanding R_{aug} . First, note that

$$Y^b = X^b\beta + \varepsilon^b = X^b\beta + \rho\hat{\varepsilon}_0^b + \rho(\varepsilon_0^b - \hat{\varepsilon}_0^b) + \sqrt{1 - \rho^2}\bar{\varepsilon}^b \quad (111)$$

$$= \bar{X}^b\beta_c + \underbrace{\rho(\varepsilon_0^b - \hat{\varepsilon}_0^b)}_{:=N^b} + \sqrt{1 - \rho^2}\bar{\varepsilon}^b = \bar{X}^b\beta_c + N^b. \quad (112)$$

Therefore, we have that

$$\frac{1}{n_b}\mathbb{E}[\|Y^b - \bar{X}^b\hat{\beta}_c\|^2 \mid \mathcal{D}^t] = \frac{1}{n_b}\mathbb{E}[\|\bar{X}^b(\beta_c - \hat{\beta}_c) + N^b\|^2 \mid \mathcal{D}^t] \quad (113)$$

$$= \underbrace{\frac{1}{n_b}\mathbb{E}[\|\bar{X}^b(\beta_c - \hat{\beta}_c)\|^2 \mid \mathcal{D}^t]}_{T_1} + \underbrace{\frac{1}{n_b}\mathbb{E}[\|N^b\|^2 \mid \mathcal{D}^t]}_{T_2} \quad (114)$$

$$+ \underbrace{\frac{2}{n_b}\mathbb{E}[N^{b\top}\bar{X}^b(\beta_c - \hat{\beta}_c) \mid \mathcal{D}^t]}_{T_3}. \quad (115)$$

We bound the terms T_1, T_2 with high probability, specifically showing that

$$T_1 = \sigma^2\mathcal{O}_P\left(\frac{d}{n_a}\right) \quad (116)$$

in Lem. 2, and showing that

$$T_2 = (1 - \rho^2)\sigma^2 + \mathcal{O}_P\left(\frac{d}{n_a}\right) \quad (117)$$

in Lem. 1. After bounding T_1, T_2 , for T_3 we have that

$$T_3 \stackrel{CS}{\leq} \frac{2}{n_b}\sqrt{\mathbb{E}[\|N^b\|^2 \mid \mathcal{D}^t] \cdot \mathbb{E}[\|\bar{X}^b(\hat{\beta}_c - \beta_c)\|^2 \mid \mathcal{D}^t]} \quad (118)$$

$$= 2\sqrt{\frac{1}{n_b}\mathbb{E}[\|N^b\|^2 \mid \mathcal{D}^t] \cdot \frac{1}{n_b}\mathbb{E}[\|\bar{X}^b(\hat{\beta}_c - \beta_c)\|^2 \mid \mathcal{D}^t]} \quad (119)$$

$$= 2\sqrt{T_1 T_2} = \mathcal{O}_P\left(\sqrt{\frac{d}{n_a}}\right), \quad (120)$$

and the claim follows. \blacksquare

Lemma 1 (Term T_2)

$$T_2 = (1 - \rho^2)\sigma^2 + \mathcal{O}_P\left(\frac{d}{n_a}\right). \quad (121)$$

\square

Proof [Lem. 1] Since $N^b = \rho(\varepsilon_0^b - \hat{\varepsilon}_0^b) + \sqrt{1 - \rho^2}\varepsilon^b$, we can write $\frac{1}{n_b}\mathbb{E}[\|N^b\|^2 \mid \mathcal{D}^t]$ as

$$\underbrace{\frac{\rho^2}{n_b}\mathbb{E}[\|\varepsilon_0^b - \hat{\varepsilon}_0^b\|^2 \mid \mathcal{D}^t]}_{T_{2(i)}} + \underbrace{\frac{(1 - \rho^2)}{n_b}\mathbb{E}[\|\varepsilon^b\|^2 \mid \mathcal{D}^t]}_{T_{2(ii)}} + \underbrace{\frac{2\rho\sqrt{1 - \rho^2}}{n_b}\mathbb{E}[(\varepsilon_0^b - \hat{\varepsilon}_0^b)^\top \varepsilon^b \mid \mathcal{D}^t]}_{T_{2(iii)}}. \quad (122)$$

Note that ε^b is independent of $(\varepsilon_0^b - \hat{\varepsilon}_0^b)$, so $T_{2(iii)}$ vanishes. $T_{2(ii)}$ is the expectation of a $\frac{(1 - \rho^2)\sigma^2}{n_b}\chi_{n_b}^2$ distribution, equal to $(1 - \rho^2)\sigma^2$. $T_{2(i)}$ satisfies

$$\frac{\rho^2}{n_b}\mathbb{E}[\|\varepsilon_0^b - \hat{\varepsilon}_0^b\|^2 \mid \mathcal{D}^t] = \frac{\rho^2}{n_b}\mathbb{E}[\|X_0^b(\beta_0 - \hat{\beta}_0)\|^2 \mid \mathcal{D}^t] \quad (123)$$

$$= \rho^2(\beta_0 - \hat{\beta}_0)^\top \Sigma_0(\beta_0 - \hat{\beta}_0) \quad (124)$$

$$= \rho^2 \varepsilon_0^{a\top} X_0^a (X_0^{a\top} X_0^a)^{-1} \Sigma_0^{1/2} \Sigma_0^{1/2} (X_0^{a\top} X_0^a)^{-1} X_0^{a\top} \varepsilon_0^a \quad (125)$$

$$= \rho^2 ((\tilde{X}_0^{a\top} \tilde{X}_0^a)^{-1/2} \tilde{X}_0^{a\top} \varepsilon_0^a)^\top (\tilde{X}_0^{a\top} \tilde{X}_0^a)^{-1} (\tilde{X}_0^{a\top} \tilde{X}_0^a)^{-1/2} \tilde{X}_0^{a\top} \varepsilon_0^a, \quad (126)$$

where $\tilde{X}_0^a := X_0^a \Sigma_0^{-1/2}$. Now, note that $(\tilde{X}_0^{a\top} \tilde{X}_0^a)^{-1}$ follows an inverse-Wishart $IW(I_d, n_a)$ distribution, while $(\tilde{X}_0^{a\top} \tilde{X}_0^a)^{-1/2} \tilde{X}_0^{a\top} \varepsilon_0^a$ is a Gaussian with covariance $\sigma^2 I_d$, so by recognizing Hotelling's T^2 distribution we have

$$T_{2(i)} \sim \rho^2 \sigma^2 \frac{1}{n_a} T^2(d, n_a) \sim \rho^2 \sigma^2 \frac{d}{n_a - d + 1} F_{d, n_a - d + 1} \quad (127)$$

which is $\rho^2 \sigma^2 \mathcal{O}_P\left(\frac{d}{n_a}\right)$ as $n_a \rightarrow \infty$. The claim follows by putting together terms $T_{2(i)}$ and $T_{2(ii)}$ together. \blacksquare

Lemma 2 (Term T_1) *Term T_1 satisfies*

$$T_1 = \mathcal{O}_P\left(\frac{d}{n_a}\right). \quad (128)$$

\square

Proof Note that

$$T_1 = (\hat{\beta}_c - \beta_c)^\top \mathbb{E}\left[\frac{1}{n_b} \bar{X}^{b\top} \bar{X}^b \mid \mathcal{D}^t\right] (\hat{\beta}_c - \beta_c) \quad (129)$$

$$= \|\hat{\beta}_c - \beta_c\|_{\tilde{\Sigma}}^2 \quad (130)$$

where $\|x\|_{\tilde{\Sigma}} := \sqrt{x^\top \tilde{\Sigma} x}$, and the matrix $\tilde{\Sigma}$ is the covariance matrix of $\bar{X}^b = (X^b \quad \varepsilon_0^b)$, given by

$$\begin{pmatrix} \Sigma & \Sigma_{X, X_0}(\beta_0 - \hat{\beta}_0) \\ (\beta_0 - \hat{\beta}_0)^\top \Sigma_{X_0, X} & \sigma^2 + \|\beta_0 - \hat{\beta}_0\|_{\Sigma_0}^2 \end{pmatrix}. \quad (131)$$

Since $\hat{\beta}_c - \beta_c = \frac{1}{n_a} \hat{\Sigma}_{\bar{X}^a}^{-1} \bar{X}^{a\top} N^a$, we have that

$$\|\hat{\beta}_c - \beta_c\|_{\tilde{\Sigma}}^2 = \frac{1}{n_a^2} N^{a\top} \bar{X}^a \hat{\Sigma}_{\bar{X}^a}^{-1} \tilde{\Sigma} \hat{\Sigma}_{\bar{X}^a}^{-1} \bar{X}^{a\top} N^a \quad (132)$$

$$= \left\| \frac{1}{n_a} \tilde{\Sigma}^{1/2} \hat{\Sigma}_{\bar{X}^a}^{-1/2} \hat{\Sigma}_{\bar{X}^a}^{-1/2} \bar{X}^{a\top} N^a \right\|^2 \quad (133)$$

$$\leq \|\tilde{\Sigma}^{1/2} \hat{\Sigma}_{\bar{X}^a}^{-1/2}\|^2 \left\| \frac{1}{n_a} \hat{\Sigma}_{\bar{X}^a}^{-1/2} \bar{X}^{a\top} N^a \right\|^2 \quad (134)$$

$$= \underbrace{\|\tilde{\Sigma}^{1/2} \hat{\Sigma}_{\bar{X}^a}^{-1/2}\|}_{\mathcal{K}} \underbrace{\left\| \frac{1}{n_a} \hat{\Sigma}_{\bar{X}^a}^{-1/2} \bar{X}^{a\top} N^a \right\|^2}_{\mathcal{F}}, \quad (135)$$

where the last line uses the fact that $\|A^\top\|^2 = \|A^\top A\|$, applied to $A = \hat{\Sigma}_{\bar{X}^a}^{-1/2} \tilde{\Sigma}^{1/2}$. In Lem. 3 we show that $\mathcal{K} = 1 + \mathcal{O}_P\left(\sqrt{\frac{d}{n_a}}\right)$, while in Lem. 4 we show that $\mathcal{F} = \mathcal{O}_P\left(\frac{d}{n_a}\right)$, which together imply that $T_1 = \mathcal{O}_P\left(\frac{d}{n_a}\right)$. \blacksquare

Lemma 3 (Covariance Mismatch)

$$\mathcal{K} = 1 + \mathcal{O}_P\left(\sqrt{\frac{d}{n_a}}\right). \quad (136)$$

□

Proof We define the matrix $\bar{\Sigma}$ as

$$\bar{\Sigma} = \begin{pmatrix} \Sigma & 0 \\ 0 & \sigma^2 + \|\beta_0 - \hat{\beta}_0\|_{\Sigma_0}^2 \end{pmatrix}, \quad (137)$$

and denote $\tau^2 := \sigma^2 + \|\beta_0 - \hat{\beta}_0\|_{\Sigma_0}^2$, and $E := \tilde{\Sigma} - \bar{\Sigma}$. Note that we have

$$\mathcal{K} = \|\tilde{\Sigma}^{1/2} \hat{\Sigma}_{\bar{X}^a}^{-1} \tilde{\Sigma}^{1/2}\| = \|(\tilde{\Sigma}^{1/2} \bar{\Sigma}^{-1/2}) (\bar{\Sigma}^{1/2} \hat{\Sigma}_{\bar{X}^a}^{-1} \bar{\Sigma}^{1/2}) (\bar{\Sigma}^{-1/2} \tilde{\Sigma}^{1/2})\| \quad (138)$$

$$\leq \|\tilde{\Sigma}^{1/2} \bar{\Sigma}^{-1/2}\|^2 \underbrace{\|\bar{\Sigma}^{1/2} \hat{\Sigma}_{\bar{X}^a}^{-1} \bar{\Sigma}^{1/2}\|}_{\mathcal{K}_0}. \quad (139)$$

Further, we have that

$$\|\tilde{\Sigma}^{1/2} \bar{\Sigma}^{-1/2}\|^2 = \|\bar{\Sigma}^{-1/2} \tilde{\Sigma} \bar{\Sigma}^{-1/2}\| \quad (140)$$

$$= \|I + \bar{\Sigma}^{-1/2} E \bar{\Sigma}^{-1/2}\| \quad (141)$$

$$\leq 1 + \|\bar{\Sigma}^{-1/2} E \bar{\Sigma}^{-1/2}\| \quad (142)$$

and the matrix $\bar{\Sigma}^{-1/2} E \bar{\Sigma}^{-1/2}$ has only off-diagonal non-zero entries $\frac{1}{\tau} \Sigma^{-1/2} \Sigma_{X, X_0} (\beta_0 - \hat{\beta}_0)$, for which we can write

$$\frac{\|\Sigma^{-1/2} \Sigma_{X, X_0} (\beta_0 - \hat{\beta}_0)\|}{\tau} = \frac{\|\Sigma^{-1/2} \Sigma_{X, X_0} \Sigma_0^{-1/2} \Sigma_0^{1/2} (\beta_0 - \hat{\beta}_0)\|}{\tau} \quad (143)$$

$$\leq \frac{\|\Sigma^{-1/2} \Sigma_{X, X_0} \Sigma_0^{-1/2}\|}{\tau} \cdot \|\beta_0 - \hat{\beta}_0\|_{\Sigma_0} \quad (144)$$

$$\leq \frac{\|\Sigma^{-1/2} \Sigma_{X, X_0} \Sigma_0^{-1/2}\|}{\sigma} \cdot \|\beta_0 - \hat{\beta}_0\|_{\Sigma_0} = \mathcal{O}_P\left(\sqrt{\frac{d}{n_a}}\right) \quad (145)$$

since $\frac{\|\Sigma^{-1/2} \Sigma_{X, X_0} \Sigma_0^{-1/2}\|}{\sigma}$ is a constant and $\|\beta_0 - \hat{\beta}_0\|_{\Sigma_0}$ is $\mathcal{O}_P\left(\sqrt{\frac{d}{n_a}}\right)$ as shown in the proof of Lem. 1 (term $T_{2(i)}$). Therefore, it only remains to show that $\mathcal{K}_0 = 1 + \mathcal{O}_P\left(\sqrt{\frac{d}{n_a}}\right)$, from which follows that that $\mathcal{K} = 1 + \mathcal{O}_P\left(\sqrt{\frac{d}{n_a}}\right)$.

For bounding \mathcal{K}_0 , let $\hat{S} := \frac{1}{n_a} X^{a\top} \hat{\varepsilon}_0^a$, $\hat{\tau}^2 = \frac{1}{n_a} \|\hat{\varepsilon}_0^a\|^2$. Using this notation, we can write

$$\bar{\Sigma}^{1/2} \hat{\Sigma}_{\bar{X}^a}^{-1} \bar{\Sigma}^{1/2} = (I + \Delta)^{-1}, \quad \Delta := \begin{pmatrix} \Sigma^{-1/2} (\hat{\Sigma}_{X^a} - \Sigma) \Sigma^{-1/2} & \Sigma^{-1/2} \hat{S} / \tau \\ \hat{S}^\top \Sigma^{-1/2} / \tau & (\hat{\tau}^2 - \tau^2) / \tau^2 \end{pmatrix}. \quad (146)$$

If $\|\Delta\| < 1$, then

$$\mathcal{K}_0 = \|(I + \Delta)^{-1}\| \leq \frac{1}{\lambda_{\min}(I + \Delta)} \leq (1 - \|\Delta\|)^{-1}, \quad (147)$$

so we want to show that $\|\Delta\|$ is $\mathcal{O}_P\left(\sqrt{\frac{d}{n_a}}\right)$, which implies $\mathcal{K}_0 = \mathcal{O}_P\left(\sqrt{\frac{d}{n_a}}\right)$. Write $\Delta = \begin{pmatrix} A & B \\ B^\top & D \end{pmatrix}$, and by using the fact that $\|\Delta\| \leq \max(\|A\| + \|B\|, \|B\| + |D|)$, it is enough to bound the norm of each of the submatrices.

Block (1, 1) – matrix A. For the matrix $A = \Sigma^{-1/2}(\hat{\Sigma}_{X^a} - \Sigma)\Sigma^{-1/2}$, we have that

$$\|\Sigma^{-1/2}(\hat{\Sigma}_{X^a} - \Sigma)\Sigma^{-1/2}\| \leq \|\Sigma^{-1}\| \cdot \|\hat{\Sigma}_{X^a} - \Sigma\| = \mathcal{O}_P\left(\sqrt{\frac{d}{n_a}}\right) \quad (148)$$

since $\|\hat{\Sigma}_{X^a} - \Sigma\| = \mathcal{O}_P\left(\sqrt{\frac{d}{n_a}}\right)$ using a Wishart concentration inequality (Vershynin, 2020).

Block (2, 2) – scalar D. Note that $\hat{\tau}^2 = \frac{1}{n_a}\varepsilon_0^{a\top}(I - P_0^a)\varepsilon_0^a$, which conditional on X_0^a follows a $\frac{\sigma^2}{n_a}\chi_{n_a-d}^2$ distribution, and this equals $\sigma^2 + \mathcal{O}_P\left(\frac{1}{\sqrt{n_a}}\right)$ since $\frac{\chi_k^2}{k} = 1 + \mathcal{O}_P\left(\frac{1}{\sqrt{k}}\right)$ (Laurent and Massart, 2000). From before, we obtained that $\tau^2 = \sigma^2 + \|\hat{\beta}_0 - \beta_0\|_{\Sigma_0}^2 = \sigma^2 + \mathcal{O}_P\left(\frac{d}{n_a}\right)$. Putting everything together

$$\frac{|\hat{\tau}^2 - \tau^2|}{\tau^2} \leq \frac{1}{\sigma^2}|\hat{\tau}^2 - \tau^2| = \mathcal{O}_P\left(\sqrt{\frac{1}{n_a}}\right). \quad (149)$$

Off-diagonal block (1, 2) – vector B. Finally, we need to bound $\|\Sigma^{-1/2}\hat{S}/\tau\|$. Let $\Sigma_{X, X_0} := \text{Cov}(X_i^a, X_{0,i}^a)$. Note that

$$X_i^a \mid X_{0,i}^a \sim N(\Sigma_{X, X_0}\Sigma_0^{-1}X_{0,i}^a, \Sigma_{X|X_0}), \quad (150)$$

where $\Sigma_{X|X_0} = \Sigma - \Sigma_{X, X_0}\Sigma_0^{-1}\Sigma_{X_0, X}$. Therefore, we have that

$$\frac{1}{n_a}X^{a\top}\varepsilon_0^a = \frac{1}{n_a}\sum_{i=1}^{n_a}\varepsilon_{0,i}^aX_i^a, \quad (151)$$

showing that conditional of $X_0^a, \varepsilon_0^a, \hat{S}$ is a Gaussian, with a mean

$$\mathbb{E}[\hat{S} \mid X_0^a, \varepsilon_0^a] = \frac{1}{n_a}\Sigma_{X, X_0}\Sigma_0^{-1}X_0^{a\top}(I - P_0)\varepsilon_0^a = 0, \quad (152)$$

since $I - P_0$ projects onto the orthogonal complement of the column space of X_0^a . The covariance of \hat{S} (conditional on X_0^a, ε_0^a) is

$$\text{Var}\left(\frac{1}{n_a}\sum_{i=1}^{n_a}\varepsilon_{0,i}^aX_i^a \mid X_0^a, \varepsilon_0^a\right) = \frac{1}{n_a^2}\sum_{i=1}^{n_a}(\varepsilon_{0,i}^a)^2\text{Var}(X_i^a \mid X_{0,i}^a) \quad (153)$$

$$= \Sigma_{X|X_0}\frac{1}{n_a^2}\|\varepsilon_0^a\|^2 = \frac{\hat{\tau}^2}{n_a}\Sigma_{X|X_0}. \quad (154)$$

Putting together, $\tilde{S} := \frac{1}{\tau}\Sigma^{-1/2}\hat{S} \mid X_0 \sim N(0, \frac{\hat{\tau}^2}{\tau^2 n_a}\Sigma^{-1/2}\Sigma_{X|X_0}\Sigma^{-1/2})$. Therefore,

$$\|\tilde{S}\|^2 \stackrel{d}{=} \frac{\hat{\tau}^2}{\tau^2 n_a}\sum_{i=1}^d \lambda_i Z_i^2, \quad (155)$$

where λ_i are the eigenvalues of $\Sigma^{-1/2}\Sigma_{X|X_0}\Sigma^{-1/2}$ and Z_i are i.i.d. Gaussians. Note that $\Sigma_{X|X_0} = \Sigma - \Sigma_{X,X_0}\Sigma_0^{-1}\Sigma_{X_0,X} \preceq \Sigma$, which implies $\Sigma^{-1/2}\Sigma_{X|X_0}\Sigma^{-1/2} \preceq I_d$, meaning that all eigenvalues λ_i in Eq. 155 are smaller than 1, so that $\|\tilde{S}\|^2$ is stochastically smaller than $\frac{\hat{\tau}^2}{\tau^2 n_a} \chi_d^2$. Since $\frac{\hat{\tau}^2}{\tau^2} = 1 + \mathcal{O}_P\left(\sqrt{\frac{1}{n_a}}\right)$ and $\chi_d^2 = d + \mathcal{O}_P\left(\sqrt{d}\right)$, we have that $\|\tilde{S}\|^2 = \mathcal{O}_P\left(\frac{d}{n_a}\right)$, meaning that $\|S\| = \mathcal{O}_P\left(\sqrt{\frac{d}{n_a}}\right)$, completing the proof. \blacksquare

Lemma 4 (Fixed Design Loss) *The fixed design loss, denoted by \mathcal{F} , satisfies*

$$\left\| \frac{1}{n_a} \hat{\Sigma}_{\bar{X}^a}^{-1/2} \bar{X}^{a\top} N^a \right\|^2 = \mathcal{O}_P(?) \quad (156)$$

\square

Proof Note that

$$\left\| \frac{1}{n_a} \hat{\Sigma}_{\bar{X}^a}^{-1/2} \bar{X}^{a\top} N^a \right\|^2 = \frac{1}{n_a} N^{a\top} \bar{X}^a (\bar{X}^{a\top} \bar{X}^a)^{-1} \bar{X}^{a\top} N^a = \frac{1}{n_a} N^{a\top} H_a N^a, \quad (157)$$

where H_a is the projection onto the column space of \bar{X}^a . By definition $N^a = \rho P_0^a \varepsilon_0^a + \sqrt{1 - \rho^2} \bar{\varepsilon}^a$, meaning that

$$\mathcal{F} = \underbrace{\frac{\rho^2}{n_a} \varepsilon_0^{a\top} P_0^a H_a P_0^a \varepsilon_0^a}_{\mathcal{F}_1} + \underbrace{\frac{1 - \rho^2}{n_a} \bar{\varepsilon}^{a\top} H_a \bar{\varepsilon}^a}_{\mathcal{F}_2} + \underbrace{\frac{\rho \sqrt{1 - \rho^2}}{n_a} \varepsilon_0^{a\top} P_0^a H_a \bar{\varepsilon}^a}_{\mathcal{F}_3}. \quad (158)$$

Conditional on $X^a, X_0^a, \varepsilon_0^a$ (meaning \bar{X}^a fixed), note that \mathcal{F}_2 follows a $\frac{1 - \rho^2}{n_a} \chi_{d+1}^2$ distribution, meaning it is $\mathcal{O}_P\left(\frac{d}{n_a}\right)$. For term \mathcal{F}_1 , we note that conditional on $X^a, X_0^a, \hat{\varepsilon}_0^a$, both H^a, P_0^a are fixed. Since $\hat{\varepsilon}_0^a = (I - P_0^a) \varepsilon_0^a$ is independent of $P_0^a \varepsilon_0^a$, the conditioning on $\hat{\varepsilon}_0^a$ does not affect $P_0^a \varepsilon_0^a$, which follows a distribution $N(0, \sigma^2 P_0^a)$, meaning that $H^a P_0^a \varepsilon_0^a$ follows a $N(0, \sigma^2 P_0^a H^a)$ distribution. Therefore, we conclude that

$$\mathcal{F}_1 \stackrel{d}{=} \frac{\rho^2 \sigma^2}{n_a} \sum_{i=1}^{d+1} \lambda_i Z_i^2, \quad (159)$$

where $\lambda_i \in [0, 1]$ are the top $d + 1$ eigenvalues of $P_0^a H^a$, and Z_i are independent. Therefore, \mathcal{F}_1 is $\mathcal{O}_P\left(\frac{d}{n_a}\right)$. Finally, for the term \mathcal{F}_3 , we condition on $X^a, X_0^a, \hat{\varepsilon}_0^a$, meaning that H^a, P_0^a are fixed. Denote by $w := \frac{1}{n_a} H^a P_0^a \varepsilon_0^a$. Conditional on w , we have that $w^\top \bar{\varepsilon}^a \sim N(0, \sigma^2 \|w\|^2)$, so

$$P(|w^\top \bar{\varepsilon}^a| > \sigma \|w\| \sqrt{2 \log \frac{1}{\delta}}) \leq 2\delta. \quad (160)$$

As we argued before, $\|H^a P_0^a \varepsilon_0^a\|^2$ is stochastically dominated by a χ_{d+1}^2 , so that $\|w\|$ is $\mathcal{O}_P\left(\frac{\sqrt{d}}{n_a}\right)$, and by using this fact with Eq. 160 and applying a union bound, we have that \mathcal{F}_3 is $\mathcal{O}_P\left(\frac{\sqrt{d}}{n_a}\right)$. Putting everything together, we obtain that \mathcal{F} is $\mathcal{O}_P\left(\frac{d}{n_a}\right)$. \blacksquare