
Indeterminacy in Generative Models: Characterization and Strong Identifiability

Quanhan Xi

johnny.xi@stat.ubc.ca
University of British Columbia

Benjamin Bloem-Reddy

benbr@stat.ubc.ca
University of British Columbia

Abstract

Most modern probabilistic generative models, such as the variational autoencoder (VAE), have certain indeterminacies that are unresolvable even with an infinite amount of data. Different tasks tolerate different indeterminacies, however recent applications have indicated the need for *strongly* identifiable models, in which an observation corresponds to a unique latent code. Progress has been made towards reducing model indeterminacies while maintaining flexibility, and recent work excludes many—but not all—indeterminacies. In this work, we motivate model-identifiability in terms of task-identifiability, then construct a theoretical framework for analyzing the indeterminacies of latent variable models, which enables their precise characterization in terms of the generator function and prior distribution spaces. We reveal that strong identifiability is possible even with highly flexible nonlinear generators, and give two such examples. One is a straightforward modification of iVAE (Khemakhem et al., 2020a); the other uses triangular monotonic maps, leading to novel connections between optimal transport and identifiability.

1 INTRODUCTION

In generative models, indeterminacy refers to the situation where the latent values underlying observations cannot be uniquely inferred from any amount of empirical evidence. It is a structural issue occurring in generative models such as the variational auto-encoder (VAE) (Kingma and Welling, 2013) or independent component analysis (ICA) (Comon, 1994; Hyvärinen and Pajunen, 1999). For example, it is well

known that even linear Gaussian models are plagued by a rotational indeterminacy, where arbitrary rotations of the latent space result in equivalent observation distributions. Modern nonlinear (deep) generative models inherit such indeterminacies and many more.

Characterizing and reducing the indeterminacies of generative models has been referred to as *identifiability*, and has been motivated as a way to address a variety of problems in representation learning such as disentanglement (Locatello et al., 2019; Hälvä et al., 2021; Yang et al., 2021; Klindt et al., 2021), posterior collapse (Wang et al., 2021), and causal representation learning (Wang and Jordan, 2021; Schölkopf et al., 2021; Lu et al., 2022). Compared to linear models, which restrict indeterminacies to linear transformations of the latent space, nonlinear identifiability is more challenging, and has been the subject of much recent work.

Consider a prototypical generative model,

$$Z_i \sim P_z, \quad \epsilon_i \sim \mathcal{N}(0, \sigma^2), \quad X_i = f(Z_i) + \epsilon_i,$$

where $P_z \in \mathcal{P}_z$ is a distribution on latent variables, $f \in \mathcal{F}$ is an injective generator function, and X_i is an observation. Indeterminacies arise when more than one (f, P_z) pair give rise to the same marginal distribution on observations, due to transformations of the latent space. Despite the seemingly disparate approaches taken in recent work to reduce indeterminacy, an intuitive trade-off appears: as more structure is added to the model through \mathcal{F} and \mathcal{P}_z , indeterminacy is reduced and hence stronger identifiability results are obtained. Our first contribution is a framework for analyzing indeterminacy that abstracts away model specifics and formalises the trade-off as a general phenomenon, described by separating indeterminacy transformations into two possible sources. The first source is the set of transports that turn one distribution in \mathcal{P}_z into another; denote this set by $\mathcal{A}(\mathcal{P}_z)$. The second source, $\mathcal{A}(\mathcal{F})$, is the set of automorphisms of the latent space formed from two functions in \mathcal{F} as $f_b^{-1} \circ f_a$ (See Section 2.3 for details). The main result of the analysis is that model indeterminacies must belong to *both* sets.

Theorem (Informal Statement of Theorem 2.2). *The set of indeterminacy transformations of a generative model $(\mathcal{F}, \mathcal{P}_z)$ is precisely $\mathcal{A}(\mathcal{F}) \cap \mathcal{A}(\mathcal{P}_z)$.*

This reduces the identifiability analysis for any suitable generative model to a well-defined mathematical problem. It also categorizes various methods to reduce indeterminacies in the model class—either constrain \mathcal{F} , as in linear models; constrain \mathcal{P}_z , e.g., non-Gaussians; or constrain both. Any of these approaches will reduce the size of the intersection. This describes many recent methods for nonlinear identifiability. Generator constraints include restrictions to certain optimal transport maps (Wang et al., 2021), sparsity (Moran et al., 2022; Zheng et al., 2022) and restrictions on the Jacobian (Gresele et al., 2021; Buchholz et al., 2022). Many methods also assume a diffeomorphic or analytic generator, which further reduces $\mathcal{A}(\mathcal{F})$. Direct constraints on \mathcal{P}_z include non-Gaussianity (Stühmer et al., 2020) and mixture distributions (Kivva et al., 2022). More commonly, latent distribution constraints are formulated in terms of dependence on an auxiliary variable (Hyvärinen and Morioka, 2016, 2017; Hyvärinen et al., 2018; Khemakhem et al., 2020a,b; Hälvä et al., 2021; Klindt et al., 2021), multiple views (Locatello et al., 2020), when latent variables are perturbed via interventions (Brehmer et al., 2022) or sparse mechanism shifts (Lachapelle et al., 2022; Ahuja et al., 2022b). As an application of our framework, we show explicitly in Section 4 how multiple environments (e.g., Khemakhem et al., 2020a) and multiple views (e.g., Locatello et al., 2020) can yield stronger identifiability, simply by reducing the set of indeterminacy transformations.

Despite the progress in the references above, each of the identifiability results contained therein are weak, in the sense that non-trivial indeterminacies remain. An unanswered question is what constraints are required for strong identifiability, that is, pointwise uniqueness of the latent representation. It has long been seen as unattainable without major sacrifices to model flexibility. For example, permutation and scaling indeterminacies are considered fundamental in ICA models (Comon, 1994). Our second main contribution is to use our framework in a few different ways to specify strongly identifiable nonlinear models, without restrictive sacrifices in flexibility. The simplest is to freeze the latent distribution before training the generator, as is typically done in a VAE. Specifically, we show that freezing the priors in iVAE (Khemakhem et al., 2020a), either from the outset or after some initial training, yields strong identifiability—with no further constraints on the generator class—when data from distinct environments (auxiliary information) are used (Section 4). We also show that monotonic triangular flow generators (Huang et al., 2018; Jaini et al., 2019; Wehenkel and Louppe, 2019; Irons et al., 2022), which are universal transports between fully supported distributions, are strongly identifiable even with a single environment, and with any latent distribution (Section 5).

Even when strong identifiability is unachievable, weakly identifiable models may be useful for tasks that tolerate the remaining indeterminacies. Our third main contribution is

to formally relate model identifiability to task identifiability in Section 3, providing precise conditions for when weak identifiability is good enough for identifiability of a particular downstream task, including some examples from the recent literature. One obvious conclusion is that strongly identifiable models are acceptable for any task based on the latent variables.

Before proceeding, we note that any notion of model identifiability is an asymptotic property not achievable with finite data. However, model identifiability is an important quality for statistical inference, and in particular is necessary for the typical consistency guarantees (van der Vaart, 1998). Though future work is required to assess the finite-sample properties of identifiable models, empirical evidence shows that even weak identifiability can recover ground truths in simulation studies (Khemakhem et al., 2020a; Sorrenson et al., 2020; Lu et al., 2022).

Outline In the rest of this section, we motivate our framework with two classical linear examples, and discuss identifiability at a high level from a downstream task-specific point of view, outlining when and what degree of identifiability is required. In Section 2, we define our mathematical framework and present our main technical results on model identifiability. After revisiting the task-specific view in detail in Section 3, we then apply the framework to analyze the auxiliary information setting in Section 4, showing in detail how the iVAE fits into the framework, and how it can be easily adapted for strong identifiability. In Section 5 we describe properties of triangular flows that yield strongly identifiable models, which can then be generalized to flows based on certain optimal transports.

1.1 Factor Analysis and Linear ICA

Factor analysis (Lawley and Maxwell, 1962) is a linear generative model,

$$Z_i \sim \mathcal{N}(0, I_{d_z}), \quad \epsilon_i \sim \mathcal{N}(\mu, I_{d_x}), \quad X_i = F Z_i + \epsilon_i,$$

where $\epsilon_i \perp\!\!\!\perp Z_i$, and F is a full-rank $d_x \times d_z$ matrix of so-called factor loadings. Here, F is the only learnable parameter and $\mathcal{P}_z = \{\mathcal{N}(0, I_{d_z})\}$, a singleton. Linear ICA (Comon, 1994) is structurally identical, but relaxes the assumption that Z has a Gaussian distribution. Instead, it is parametrized by (F, P_z) , where F is again full rank, and P_z is required to have independent components and is learned along with F from data. In other words, \mathcal{P}_z is some collection of fully supported distributions with independent components, typically excluding Gaussians.

It is well known that factor analysis suffers from a rotational indeterminacy due to the Gaussian P_z . On the other hand, Hastie et al. (2009) note that ICA is identical in form to factor analysis, but avoids the rotational indeterminacy via its non-Gaussian assumption. However, if scaling indeter-

minacies are fundamental to ICA (Hyvarinen et al., 2001), why are they ignored in factor analysis?

Factor analysis does not suffer from scaling indeterminacy due to a fundamental difference in how indeterminacies arise in these two models. The factor model fixes its Gaussian P_z , and the rotational indeterminacies in factor analysis can be characterized precisely as linear measure-preserving automorphisms of the standard Gaussian. Since scaling does not preserve the standard Gaussian, it is not an indeterminacy. On the other hand, ICA does not fix P_z . Due to this, any linear transformation of $P_z \in \mathcal{P}_z$ to another $P'_z \in \mathcal{P}_z$ is an indeterminacy, and a result of Comon (1994, Thm. 10) is that only scalings and permutations are possible when \mathcal{P}_z contains all independent distributions excluding Gaussians.

Our Theorem 2.2 below generalizes these special cases. It can be used to show that indeterminacies are completely characterized as measure-preserving automorphisms for fixed latent distributions (as in factor analysis), and as measure-transporting isomorphisms within \mathcal{P}_z otherwise (as in ICA). One takeaway is that strong identifiability appears to be much easier to obtain in the factor analysis setting. We investigate this further in Sections 4 and 5.

1.2 Why and How Much Identifiability

In the recent literature on generative model identifiability, the question of why we care about identifiability typically appeals to recovering ground truth latent factors (Khemakhem et al., 2020a; Ahuja et al., 2022a; Yang et al., 2021; Lu et al., 2022). Besides requiring a philosophical position that asserts the objective reality of the latent variables *and* that the model contains the “true” data generating distribution, the use of an unidentifiable model—unable to be uniquely resolved from *any* amount of empirical evidence—makes the recovery of “true” latent factors impossible without further, untestable assumptions about the model. Another perspective is to consider when model indeterminacies preserve certain observable quantities, for example distances between observations (Arvanitidis et al., 2018). More generally, we might judge the importance of model identifiability in terms of whether or not a model can be used for its intended purpose. This opens the possibility that in some cases, weak identifiability may be sufficient; in others, it may not. To our knowledge, a characterization of the distinction has not been formalized in this setting.

To address these questions, we take a pragmatic approach based on downstream tasks that use the inferred latent variables. Informally, a task t is a function of the model parameters (f, P_z) , of data, $\mathbf{x}_m = (x_1, \dots, x_m)$, and of a finite collection of points in the latent space, $\mathbf{z}_n = (z_1, \dots, z_n)$. There must be some procedure for selecting the latent points \mathbf{z}_n represented by a selection function $\mathbf{z}_n = s(f, P_z, \mathbf{x}_m)$, e.g., $s(f, P_z, \mathbf{x}_m) \approx f^{-1}(\mathbf{x}_m)$. The entire task is given by $t(f, P_z, \mathbf{x}_m, s(f, P_z, \mathbf{x}_m))$.




Figure 1: Rotational indeterminacy is fundamental in Gaussian i.i.d. models, where each point on the above orbits represent an equivalent solution. If a task gives constant output on each orbit, the task is identifiable even in the presence of model indeterminacies.

Some examples are generating synthetic or counterfactual data by shifting inferred latent variables (Higgins et al., 2017), or causal discovery through independence tests (Monti et al., 2020; Khemakhem et al., 2020a; Lu et al., 2022). Both of these examples are studied in detail in Section 3. For a model to be useful for a task, the resulting task output should be the same for all generative model parameters (f, P_z) that yield the same marginal distribution for the observations (this is made precise in Section 3). For example, Fig. 1 illustrates a task that is constant on each set of points induced by the indeterminacy transformations (rotations). That guarantees that two fits of the same model to the same (infinite) data yield the same task values; weak model identifiability can be sufficient when the relevant task is insensitive to all model indeterminacies. Strongly identifiable models make all tasks trivially insensitive and therefore identifiable.

2 IDENTIFIABILITY IN GENERATIVE MODELS

Conceptually, our results are relatively simple, but some care is required to establish them rigorously. See Appendix C and Appendix D for technical background and proofs.

A Borel space $(\mathbf{X}, \mathcal{B}(\mathbf{X}))$ is a topological space \mathbf{X} equipped with the σ -algebra generated by its open sets, denoted $\mathcal{B}(\mathbf{X})$. All measurable spaces in this paper will be Borel spaces, and for convenience we leave the σ -algebra implicit in the notation, referring to the Borel space $(\mathbf{X}, \mathcal{B}(\mathbf{X}))$ simply as \mathbf{X} . Let \mathbf{Z} and \mathbf{X} be two Borel spaces and $f: \mathbf{Z} \rightarrow \mathbf{X}$ a measurable function. If f is bijective with measurable inverse then it is called a Borel isomorphism of \mathbf{Z} and \mathbf{X} . If $\mathbf{X} = \mathbf{Z}$ then f is a Borel automorphism. We denote the set of Borel automorphisms of \mathbf{Z} as $\text{Aut}(\mathbf{Z})$.

We will use the notation \mathbf{X} and \mathbf{Z} to represent the observable space and latent space, respectively. In practice, typically $\mathbf{X} = \mathbb{R}^{d_x}$ and $\mathbf{Z} = \mathbb{R}^{d_z}$, $d_z \leq d_x$. For clarity, we will usually refer to these specific spaces, where *almost everywhere* statements and probability densities are with respect to the Lebesgue measure; our results in Section 2 and Section 3

apply to generic Borel measure spaces. We write $f_a \stackrel{\text{a.c.}}{=} f_b$ for measurable functions as shorthand for equality almost everywhere on their domain. The main objects of study in our theory are certain elements of $\text{Aut}(\mathbf{Z})$, as follows.

Let μ and ν be two measures on \mathbf{Z} , and $A \in \text{Aut}(\mathbf{Z})$. A is called a (μ, ν) -**measure isomorphism** whenever $A_{\#}\mu = \nu$ (equiv. $A_{\#}^{-1}\nu = \mu$), and a μ -**measure preserving automorphism** if $\mu = A_{\#}\mu = A_{\#}^{-1}\mu$.¹

2.1 Model Definition

We model observations denoted x_i as i.i.d. realizations of a random variable X with distribution P_x on \mathbf{X} . Define a latent random variable Z with corresponding realizations z_i , with distribution P_z on the latent space \mathbf{Z} . We further assume the presence of some noise realization ϵ_i with some fixed distribution P_ϵ applied by some noise mechanism g . Let $f: \mathbf{Z} \rightarrow \mathbf{X}$ be a measurable function, which we call the *generator*. We define a *generative model* as,

$$Z_i \sim P_z, \quad \epsilon_i \sim P_\epsilon, \quad X_i = g(f(Z_i), \epsilon_i), \quad (1)$$

with $Z_i \perp\!\!\!\perp \epsilon_i$. For example, in a VAE or factor analysis, \mathbf{X} and \mathbf{Z} are Euclidean spaces, and $g(f(z), \epsilon) = f(z) + \epsilon$ is additive noise. We will not be concerned with inferring the noise. Instead, we work with the assumption that the noise distribution and mechanism are fixed and have null effect on the probabilistic properties of the model.

Assumption 1. Assume that g and P_ϵ are such that, with $\epsilon_a \stackrel{\text{d}}{=} \epsilon_b$, $g(f(Z_a), \epsilon_a) \stackrel{\text{d}}{=} g(f(Z_b), \epsilon_b)$ if and only if $f(Z_a) \stackrel{\text{d}}{=} f(Z_b)$.

This assumption includes, for example, the noiseless case, and additive noise for a suitable noise distribution (e.g., Hälvä et al., 2021). It means that, for identifiability purposes, it is sufficient to analyze the noiseless case. We note that it rules out the possibility of discrete observations except in very limited cases; see Appendix B for a brief discussion of this point. Designing a generative model involves specifying parameter spaces for the generator and the prior. We denote these as \mathcal{F} , a set of measurable mappings $\mathbf{Z} \rightarrow \mathbf{X}$; and \mathcal{P}_z , a set of probability measures on \mathbf{Z} . For example, \mathcal{P}_z could be a singleton, as in factor analysis. In ICA, \mathcal{P}_z contains only distributions that factorize over its components.

We will also assume that the generators are bijective on their range, which is required to make the inverse problem of recovering latents well-defined, and is a standard assumption in the identifiability literature.

Assumption 2. Assume that any $f \in \mathcal{F}$ is injective, and has the same image: for any $f_a, f_b \in \mathcal{F}$, $f_a(\mathbf{Z}) = f_b(\mathbf{Z}) := \mathcal{F}(\mathbf{Z}) \subseteq \mathbf{X}$.

¹ $A_{\#}\mu$ (A pushforward μ) can also be denoted $A_*\mu$, or $\mu \circ A^{-1}$ (the image measure), and is defined by $A_{\#}\mu(B) = \mu(A^{-1}(B))$, $B \in \mathcal{B}(\mathbf{Z})$.

2.2 Model Indeterminacies

A generative model (1) induces a statistical model as

$$\mathcal{M}(\mathcal{F}, \mathcal{P}_z) = \{P_\theta \text{ on } \mathbf{X} \mid \theta = (f, P_z) \in \mathcal{F} \times \mathcal{P}_z\}. \quad (2)$$

Classical parameter identifiability can be defined via an equivalence relation \sim on parameter space, $\theta \sim \theta' \iff P_\theta = P_{\theta'}$. The equivalence classes of parameters induced by \sim are denoted by $[\theta] := \{\theta' : P_\theta = P_{\theta'}\}$, and a model is identifiable up to $[\theta]$. Some authors refer only to the previous case as partial or set identifiability (Tamer, 2010), and reserve the term identifiability for the case that $[\theta] = \{\theta\}$. We refer to this latter case as strong identifiability. Parameter identifiability does not appear to contain any information about the latent values. We take classical parameter identifiability as our starting point to formulate an alternative definition tailored to latent variable indeterminacy. Specifically, we work with transformations of the latent space that yield different generators but that leave the marginal distribution of the data unchanged.

Definition 1. For a model $\mathcal{M}(\mathcal{F}, \mathcal{P}_z)$, $A_{a,b}$ is an **indeterminacy transformation** at θ_a, θ_b if $P_{\theta_a} = P_{\theta_b}$ and $f_a \circ A_{a,b}^{-1} \stackrel{\text{a.c.}}{=} f_b$. The **indeterminacy set** of a model $\mathcal{M}(\mathcal{F}, \mathcal{P}_z)$, denoted $\mathcal{A}(\mathcal{M})$, is the collection of all indeterminacy transformations of the model.

Definition 1 describes latent variable indeterminacy; it implies that if $\mathcal{A}(\mathcal{M})$ contains non-trivial transformations, any latent z that generates $x = f(z)$ has an equivalent counterpart $A(z)$ for some $A \in \mathcal{A}(\mathcal{M})$. The identity map on \mathbf{Z} , id_z , is always (trivially) an indeterminacy transformation by taking $\theta_a = \theta_b$. A simple way to construct non-trivial candidate indeterminacy transformations at $\theta_a \neq \theta_b$ is by “pushing forward” and “pulling back” along the generators,

$$\vec{A}_{a,b}(z) := f_b^{-1}(f_a(z)), \quad z \in \mathbf{Z}. \quad (3)$$

It is possible to show that $\vec{A}_{a,b} \in \text{Aut}(\mathbf{Z})$, and thus if transforming \mathbf{Z} in this way results in $P_{\theta_a} = P_{\theta_b}$, then $\vec{A}_{a,b}$ is an indeterminacy transformation. We refer to $\vec{A}_{a,b}: \mathbf{Z} \rightarrow \mathbf{Z}$ as the **generator transform** between θ_a and θ_b . It turns out that generator transforms characterize all the possible indeterminacy transformations of a model—though not every generator transform is an indeterminacy transformation.

Lemma 2.1. *Let $\theta_a = (f_a, P_{z,a})$ and $\theta_b = (f_b, P_{z,b})$ be two parametrizations of a generative model with resulting marginal distributions P_{θ_a} and P_{θ_b} . Then, $P_{\theta_a} = P_{\theta_b}$ if and only if $\vec{A}_{a,b}$ is a $(P_{z,a}, P_{z,b})$ -measure isomorphism. Furthermore, any $A \in \text{Aut}(\mathbf{Z})$ is an indeterminacy transformation at θ_a, θ_b (Definition 1) if and only if $A \stackrel{\text{a.c.}}{=} \vec{A}_{a,b}$, and therefore all indeterminacy transformations $A_{a,b}$ must be $(P_{z,a}, P_{z,b})$ -measure isomorphisms.*

Lemma 2.1 is the technical foundation for the rest of the paper, and can be interpreted as an existence and uniqueness

result, as follows. $\vec{A}_{a,b}$ is a distinguished indeterminacy transformation at θ_a, θ_b that always exists; it is probabilistically equivalent to all other indeterminacy transformations, making it essentially unique. It implies that two parameterizations θ_a, θ_b induce the same marginal distribution on observations if and only if one can transport between the two latent measures in \mathcal{P}_z by pushing and pulling along the generators. Not only are there no other non-equivalent indeterminacy transformations, but $A_{a,b}$ cannot be an indeterminacy transformation for any other parameters θ_c, θ_d unless $\vec{A}_{a,b} = f_b^{-1} \circ f_a \stackrel{\text{a.e.}}{=} f_d^{-1} \circ f_c = \vec{A}_{c,d}$. Both of these properties follow directly from the injectivity of the generators; relaxing Assumption 2 would require a substantially different theory than the one developed here.

2.3 Characterizing Model Identifiability

Let $\tilde{\text{id}}_z$ denote the set of all Borel automorphisms that are equal a.e. to the identity mapping on \mathbf{Z} . If $\mathcal{A}(\mathcal{M}) = \tilde{\text{id}}_z$, then $f_a \stackrel{\text{a.e.}}{=} f_b$ for all θ_a, θ_b with $P_{\theta_a} = P_{\theta_b}$. We use this to define the appropriate notion of identifiability.

Definition 2. A generative model $\mathcal{M}(\mathcal{F}, \mathcal{P}_z)$ is **weakly identifiable up to** $\mathcal{A}(\mathcal{M})$, or **$\mathcal{A}(\mathcal{M})$ -identifiable**, if its indeterminacy transformations are $\mathcal{A}(\mathcal{M})$. If $\mathcal{A}(\mathcal{M}) = \tilde{\text{id}}_z$, then the model is **strongly identifiable**.

Strong identifiability means that $f_a(z) = f_b(z)$ for all $z \in \mathbf{Z}$ outside of a set of measure zero. If f_a, f_b are continuous functions, this implies that $f_a(z) = f_b(z)$ for all $z \in \mathbf{Z}$, assuming the reference measure has full support. This definition is in correspondence to the classical definition, in the sense that the equivalence classes $[\theta]$ correspond to subsets of $\mathcal{A}(\mathcal{M})$ (see Proposition A.2 for details). Note that when the generator is further parameterized (for example via deep neural networks), this identifiability does not necessarily pass to the generator parameters (network weights). Rather, we use this as a proxy to make the notion of latent variable recoverability precise.

Lemma 2.1 gives a necessary and sufficient condition for two parameterizations to correspond to the same marginal distribution on observations. It also documents the cases in which P_{θ_a} cannot be equal to P_{θ_b} , and allows us to construct the set of model indeterminacies $\mathcal{A}(\mathcal{M})$ as indeterminacies generated by \mathcal{F} and \mathcal{P}_z . To that end, define the following subsets of $\text{Aut}(\mathbf{Z})$ induced by the generative model,

$$\begin{aligned} \mathcal{A}(\mathcal{F}) &= \{A \in \text{Aut}(\mathbf{Z}) \mid \exists f_a, f_b \in \mathcal{F} \text{ s.t. } A \stackrel{\text{a.e.}}{=} f_b^{-1} \circ f_a\} \\ \mathcal{A}(\mathcal{P}_z) &= \{A \in \text{Aut}(\mathbf{Z}) \mid \exists P_a, P_b \in \mathcal{P}_z \text{ s.t. } A_{\#} P_a = P_b\}. \end{aligned}$$

$\mathcal{A}(\mathcal{F})$ consists of all possible indeterminacy transforms constructed from \mathcal{F} , which are a.e. equivalent to the generator transforms. $\mathcal{A}(\mathcal{P}_z)$ consists of all possible isomorphisms between measures in \mathcal{P}_z . Both sets always include $\tilde{\text{id}}_z$ by taking $f_a = f_b$ and $P_a = P_b$. The indeterminacies of the model are precisely their intersection.

Theorem 2.2. *The generative model $\mathcal{M}(\mathcal{F}, \mathcal{P}_z)$ is identifiable up to $\mathcal{A}(\mathcal{M}) = \mathcal{A}(\mathcal{F}) \cap \mathcal{A}(\mathcal{P}_z)$. In particular, it is strongly identifiable if and only if $\mathcal{A}(\mathcal{F}) \cap \mathcal{A}(\mathcal{P}_z) = \tilde{\text{id}}_z$.*

This expresses the identifiability of a generative model in terms of the indeterminacy transforms induced by its parameter spaces. In particular, *all model indeterminacies must be transports between distributions in \mathcal{P}_z that can be constructed by pushing and pulling along generators from \mathcal{F} as $f_b^{-1} \circ f_a$* . These are the scaling and permutation matrices in the linear ICA example from Section 1.1. When \mathcal{P}_z is a singleton $\{P_z\}$ as in factor analysis, they are P_z -measure-preserving automorphisms, such as the rotation matrices preserving the standard Gaussian.

2.4 Beyond Linear Generators

Theorem 2.2 exposes the structure of unidentifiability and indicates that there may be approaches to specifying strongly identifiable non-linear generative models. It suggests that model identifiability strengthens as we increase the number of constraints on $\mathcal{F}, \mathcal{P}_z$, or both, until the intersection $\mathcal{A}(\mathcal{F}) \cap \mathcal{A}(\mathcal{P}_z)$ contains only the identity. We demonstrate examples of leaving the generator unconstrained and only constraining $\mathcal{A}(\mathcal{P}_z)$, as well as a triangular constraint on $\mathcal{A}(\mathcal{F})$ in Sections 4 and 5, respectively.

3 TASK IDENTIFIABILITY

In the literature on model identifiability, practical implications are often ignored. How does model identifiability relate to the intended uses of a model after it is fit to data? In practice, one may have a model that is only weakly identifiable but that still allows some tasks to be identified. That is, an unidentifiable model may still be useful for some tasks. For example, a causal estimand may still be identified without identification of the full causal model. For other tasks, an unidentifiable model is not useful. We make this idea precise before giving some examples from the literature.

Let $\mathbf{x}_m = (x_1, \dots, x_m)$ be a finite collection of points in \mathbf{X} (e.g., m observations), and recall that $\theta = (f, P_z)$. Define a *task* as a pair of functions (s, t) : the function $s(\theta, \mathbf{x}_m) \in \mathbf{Z}^n$ selects a set of n points in the latent space to use for the task, e.g., estimates of $f^{-1}(x_i)$; and $t: \mathcal{F} \times \mathcal{P}_z \times \mathbf{X}^m \times \mathbf{Z}^n \rightarrow \mathbf{T}$ generates the task output as $t(\theta, \mathbf{x}_m, s(\theta, \mathbf{x}_m))$. Here, \mathbf{T} is the task output space, such as $\mathbf{T} = \mathbf{X}$ for sample generation. We consider a task to be identifiable if two equivalent model fits produce the same task output.

Definition 3. A task (s, t) is **identifiable at** $[\theta]$ if, for all $\theta \sim \theta'$ and $\mathbf{x}_m \in \mathbf{X}^m$,

$$t(\theta, \mathbf{x}_m, s(\theta, \mathbf{x}_m)) = t(\theta', \mathbf{x}_m, s(\theta', \mathbf{x}_m)). \quad (4)$$

A task is (globally) identifiable if it is identifiable at all $[\theta]$.

This captures the idea that two different research groups may obtain different but equivalent model fits (θ and θ') from the same data, and still reach the same conclusion for a task. Before proceeding, we give some examples.

Tasks Without Reference to Observations Consider two labs with different but equivalent models, $\theta_a \sim \theta_b$. Now both labs set $s(\theta, \mathbf{x}_m) = c$, without reference to observations. This could, for example, correspond to a ‘do’ intervention when \mathbf{Z} is endowed with a causal model (Yang et al., 2021; Shen et al., 2022; Lu et al., 2022). As the task output, each lab generates a synthetic observation, $f_a(c)$ and $f_b(c)$. Without strong model identifiability, in general $f_a(c) \neq f_b(c)$ and the task is unidentifiable due to the fact that the value c has no inherent meaning in either model. Tasks $t(\theta, \mathbf{x}_m, c)$ will not be identifiable in general: the selection of latent points must incorporate the transformation of \mathbf{Z} that occurs when $\theta_a \mapsto \theta_b$; only if the task output is constant across all equivalent θ will it be identifiable.

Disentanglement and Latent Shifts A common demonstration of disentanglement is to select a latent variable $s(\theta, x_i) = f^{-1}(x_i)$ corresponding to an observation x_i , apply a shift δ along a latent dimension k (with unit vector e_k), and then generate synthetic data $t_{\delta e_k}(\theta, x_i, s(\theta, x_i)) = f(\delta e_k + f^{-1}(x_i))$ for a sequence of different values of δ (e.g., Higgins et al., 2017; Lu et al., 2022). For most models, this task is not identifiable. To see why, observe that task identifiability requires

$$\begin{aligned} f_a(\delta e_k + f_a^{-1}(x)) &= f_b(\delta e_k + f_b^{-1}(x)) \\ &= f_a(\vec{A}_{a,b}^{-1}(\delta e_k + \vec{A}_{a,b}(f_a^{-1}(x)))) . \end{aligned}$$

In general, this will not be the case unless

$$\vec{A}_{a,b}^{-1}(\delta e_k + \vec{A}_{a,b}(f_a^{-1}(x))) = \delta e_k + f_a^{-1}(x) .$$

That is, $\vec{A}_{a,b}$ must commute with δe_k . If that property is to hold for all δ , e_k , and x then $\vec{A}_{a,b}$ must be itself a translation.

Independence Testing and Causal Discovery In applications of (nonlinear) ICA to causal discovery, tests for independence between observed variables and components of latent variables are conducted (Monti et al., 2020; Khemakhem et al., 2020a). In the simplest case, the goal is to determine the causal direction between two observed variables, X_1, X_2 . After fitting a model with independent latent components, pair-wise independence tests are conducted; if X_1 is a cause of X_2 and not vice versa, then X_1 will be independent of the second component of the corresponding latent variable, $f^{-1}(X_1, X_2)_2$. The task, based on observations $(x_i)_{i=1}^N = (x_{i,1}, x_{i,2})_{i=1}^N$ is then $s(\theta_a, (x_i)_{i=1}^N) = (f_a^{-1}(x_i))_{i=1}^N$ and

$$\begin{aligned} t(\theta_a, (x_i)_{i=1}^N, s(\theta_a, (x_i)_{i=1}^N)) \\ = \text{IndTest}((x_{i,1})_{i=1}^N, (f_a^{-1}(x_{i,1}, x_{i,2}))_{i=1}^N) . \end{aligned}$$

If the model is identifiable up to component-wise transformations then $f_a^{-1}(x_{i,1}, x_{i,2})_2 = h(f_b^{-1}(x_{i,1}, x_{i,2})_2)$, for some function $h: \mathbb{R} \rightarrow \mathbb{R}$. Since two real-valued random variables U, V are independent if and only if all transformations $k(U), k'(V)$ are independent, this task remains identifiable under component-wise indeterminacies.²

Define a local subset of indeterminacy transformations as $\mathcal{A}(\mathcal{M})|_\theta = \{A \in \mathcal{A}(\mathcal{M}): P_\theta = P_{A\theta}\}$. Our main result on task identifiability is an application of Proposition A.2.

Proposition 3.1. *A task (s, t) is identifiable at $[\theta]$ if and only if, for each $\mathbf{x}_m \in \mathbf{X}^m$,*

$$t(\theta, \mathbf{x}_m, s(\theta, \mathbf{x}_m)) = t(A\theta, \mathbf{x}_m, s(A\theta, \mathbf{x}_m)) , \quad (5)$$

for each $A \in \mathcal{A}(\mathcal{M})|_\theta$. A sufficient condition for (s, t) to be identifiable at $[\theta]$ is hence if the following holds for each $A \in \mathcal{A}(\mathcal{M})|_\theta$ and $\mathbf{x}_m \in \mathbf{X}^m$:

$$\begin{aligned} t(\theta, \mathbf{x}_m, \mathbf{z}_n) &= t(A\theta, \mathbf{x}_m, A(\mathbf{z}_n)) \quad (6) \\ \text{and } s(A\theta, \mathbf{x}_m) &= A(s(\theta, \mathbf{x}_m)) . \end{aligned}$$

Clearly, any task is identifiable if \mathcal{M} is strongly identifiable: $\mathcal{A}(\mathcal{M})|_\theta = \mathcal{A}(\mathcal{M}) = \text{id}_{\mathbf{z}}$. Weak identifiability may be enough depending on the task, Proposition 3.1 gives sufficient conditions based on the symmetries implied by $\mathcal{A}(\mathcal{M})$. In the next two sections, we apply the theory of Section 2 to obtain strongly identifiable models.

4 GENERATIVE MODELS IN MULTIPLE ENVIRONMENTS

Suppose data arise from environments indexed by $e \in E$, where E is an arbitrary set, and the environment label is assumed to be deterministic (i.e., known, or observed without noise). Each environment corresponds to a different observation random variable $X^e \sim P_x^e$ on a shared observation space \mathbf{X} . This is reflected in the generative model as $|E|$ distinct distributions on latent variables, $Z^e \sim P_z^e$ on a shared latent space \mathbf{Z} . Crucially, each environment shares the same generator f . We denote the generative model $\mathcal{M}(\mathcal{F}, \{P_z^e\}_{e \in E})$ specified as, for each $e \in E$,

$$Z_i \sim P_z^e , \epsilon_i^e \sim P_\epsilon , X_i^e = g(f(Z_i^e), \epsilon_i^e) , \quad (7)$$

with $Z_i \perp\!\!\!\perp \epsilon_i^e$ in each environment. In general, we do not assume that observations from different environments are paired in any way besides sharing a generator, e.g., they may be from separate datasets.

We use the multiple environments set-up in the same way as auxiliary information in iVAE. The definition of identifiability is subtly different here, but its implications for latent variable indeterminacy remain the same (see Appendix D.4 for details).

²ICA-type models also have permutation indeterminacy; however, there is a ‘correct’ permutation that can be distinguished within causal discovery (Shimizu et al., 2006).

Corollary 4.1. *The generative model $\mathcal{M}(\mathcal{F}, \{\mathcal{P}_z^e\}_{e \in E})$ is identifiable up to*

$$\mathcal{A}(\mathcal{F}) \cap (\cap_{e \in E} \mathcal{A}(\mathcal{P}_z^e)). \quad (8)$$

We note that permutations of the environments are ruled out because the environments are assumed to be known. Clearly, the indeterminacy set shrinks with each environment added, formalizing why auxiliary information has such significant benefits for identifiability. As a demonstration, we recast iVAE under this framework and present the necessary modification for strong identifiability in Sections 4.2 and 4.3 (see Appendix D.10 for a similar exercise for the results of Ahuja et al. 2022a).

4.1 Multiple Views

A result similar to Corollary 4.1 can be obtained under different modelling assumptions similar to those made by Gresele et al. (2019); Locatello et al. (2020). Specifically, we assume here finitely many observed views of each latent variable, indexed by the finite set E . We study the following model, denoted by $\mathcal{M}(\{\mathcal{F}^e\}_{e \in E}, \mathcal{P}_z)$: for $e \in E$,

$$Z_i \sim P_z, \quad \epsilon_i^e \sim P_\epsilon^e, \quad X_i^e = g^e(f^e(Z_i), \epsilon_i^e), \quad (9)$$

with $Z_i \perp \epsilon_i^e$. Here, for fixed i , each observed view X_i^e , $e \in E$, is generated by a shared latent variable Z_i , and hence the observations are necessarily grouped across environments (unlike the multiple environments setting). Each view X_i^e can have a different observation space \mathbf{X}^e with its own generator $f^e \in \mathcal{F}^e$, e.g., an image and its caption.

Again, the definition of identifiability here changes, but the implications for latent variable indeterminacy are identical. Intuitively, identification in any of the views yields identification of the generating latent (Appendix D.5).

Corollary 4.2. *The generative model $\mathcal{M}(\{\mathcal{F}^e\}_{e \in E}, \mathcal{P}_z)$ is identifiable up to*

$$(\cap_{e \in E} \mathcal{A}(\mathcal{F}^e)) \cap \mathcal{A}(\mathcal{P}_z). \quad (10)$$

This can be useful when, for example, when we have access to a lower-dimensional view (e.g., clinical data) which we might fit using an identifiable triangular flow (as developed in Section 5), in addition to a higher-dimensional view (e.g., genomics data), which can be reconstructed using a VAE.

4.2 (Weakly) Identifiable VAE

In the iVAE model (Khemakhem et al., 2020a), the latent variable distribution varies via an auxiliary variable u , for example a time index (Hyvärinen et al., 2018). The prior is parameterized as an exponential family density on $\mathbf{Z} = \mathbb{R}^{d_z}$,

$$p(z; \eta(u)) = m(z) \exp(\eta(u)^\top T(z) - a(\eta(u))), \quad (11)$$

with functional parameters η, T taking values in \mathbb{R}^K . We denote this distribution as $\mathcal{E}_m(\eta(u), T)$. Note that this means the prior is explicitly inferred from the data, contrary to a standard VAE. The remainder of the model design follows (7) with additive noise. We note that Khemakhem et al. (2020a) assumed the distribution to factorize over dimensions of Z (as in ICA), but that is not necessary, an observation also made recently in Lu et al. (2022).

The main identifiability result of Khemakhem et al. (2020a) (their Thm. 1) says that for two parametrizations (f_a, T_a, η_a) and (f_b, T_b, η_b) , if there exist points u_0, u_1, \dots, u_K such that $\{\eta_a(u_i) - \eta_a(u_0)\}_{i=1}^K$ are linearly independent and span the latent space (and likewise for η_b), then there are an invertible matrix L and offset vector \mathbf{c} such that for all x ,

$$T_a(f_a^{-1}(x)) = L^\top T_b(f_b^{-1}(x)) + \mathbf{c}. \quad (12)$$

Our framework allows us to show that this is purely a result of the following property of exponential families.

Proposition 4.3. *Suppose that $A \in \text{Aut}(\mathbf{Z})$ is a $(\mathcal{E}_m(\eta_a(u_i), T_a), \mathcal{E}_m(\eta_b(u_i), T_b))$ -measure isomorphism for each $i = 0, 1, \dots, K$. Suppose that both $\{\eta_a(u_i)\}_{i=0}^K$ and $\{\eta_b(u_i)\}_{i=0}^K$ are linearly independent. Then,*

$$T_b(A(z)) = L^\top T_a(z) + \mathbf{d}, \quad (13)$$

almost everywhere, where L is a $K \times K$ invertible matrix and \mathbf{d} is a K -dimensional vector not depending on x .

Although the arguments made in the proof of the proposition are similar to those originally presented in Khemakhem et al. (2020a), our framework indicates that identifiability is a result purely of the exponential family form of the prior—neither diffeomorphic decoders nor independent components are needed. Further assuming that f is diffeomorphic reduces $\mathcal{A}(\mathcal{F})$, leading to the stronger identifiability results of Theorems 2 and beyond in (Khemakhem et al., 2020a).

4.3 (Strongly) Identifiable VAE

In this section, we show how fixed priors, i.e., a nonlinear factor analysis take on VAEs, can lead to strong identifiability. By definition, with fixed latent distribution P_z , $\mathcal{A}(\{P_z\})$ is the set of P_z -preserving automorphisms. However, constructing distributions such that the only automorphism is the identity is in general very difficult; non-trivial measure automorphisms, such as via the Darrois construction, almost always exist (see Gresele et al., 2021, for an example). On the other hand, we show that multiple environments provide a simple solution to this problem.

We denote a set of full-rank exponential family densities indexed by $\eta \in \mathbb{R}^K$ as

$$\mathcal{E}_{m,T} = \{p_\eta(z) = m(z) \exp(\eta^\top T(z) - a(\eta)) \mid \eta \in \mathbb{R}^K\}. \quad (14)$$




Figure 2: The indeterminacy set (under the sufficient statistic) in red for parameter vectors η_1, η_2 .

We refer to a particular distribution in such a set as $\mathcal{E}_{m,T}(\eta)$. Compared to the isomorphisms in Proposition 4.3, there is a stronger result characterizing the automorphisms of collections of exponential families.

Proposition 4.4. *Let $\mathcal{E}_{m,T}$ be as in (14), with m strictly positive. Let $\eta_i \in \mathbb{R}^K$ for $i = 0, 1, \dots, K$ and suppose $A \in \text{Aut}(\mathbb{R}^{d_z})$ is a $\mathcal{E}_{m,T}(\eta_i)$ -measure preserving automorphism for each η_i . Then, arranging η_i into the rows of a matrix M , we have*

$$(T(z) - T(A(z))) \in \ker M, \text{ a.e.} \quad (15)$$

Figure 2 illustrates (15). The proposition provides a characterization that can be specialized in several ways. Firstly, the only *shared* automorphism for each distribution from a suitably fixed exponential family $\mathcal{E}_{m,T}$ whose parameters form a basis of \mathbb{R}^K is the identity (under T). Strong identifiability follows when these are fixed as the priors.

Theorem 4.5. *Let $\mathcal{M}(\mathcal{F}, \{\mathcal{P}_z^e\}_{e \in E})$ be the multiple environments model described in (7), with $\mathbf{Z} = \mathbb{R}^{d_z}$. For a subset of environments $E^* \subset E$, with $|E^*| = K + 1$, fix $P_z^e = \mathcal{E}_{m,T}(\eta_e)$ with m strictly positive and T injective in at least one dimension, and such that the corresponding parameters $\{\eta_e\}_{e \in E^*}$ span \mathbb{R}^K . Then $\mathcal{M}(\mathcal{F}, \{\mathcal{P}_z^e\}_{e \in E})$ is strongly identifiable.*

In fact, we do not require a spanning set of environments for a useful characterization of the indeterminacy. A second specialization of Proposition 4.4 is in the case of Gaussian priors. In such a setting, assuming common covariances across environments, the model is identifiable up to arbitrary transformations on the dimensions *not spanned* by the environment means μ_e ; see Appendix D.9 for details.

4.4 Fixed versus Learned Distributions

In ICA compared to factor analysis (Section 1.1), and now iVAE compared to a strongly identifiable version, learning the latent distributions adds indeterminacies. It is not apparent to us that that source of indeterminacies can be addressed without either: (i) fixing the latent distributions before fitting the generator; or (ii) cleverly constructing $\mathcal{F}, \mathcal{P}_z$ to

ensure that any generator transform $f_b^{-1} \circ f_a$ applied to any $P_z \in \mathcal{P}_z$ transports P_z to a different distribution not contained in \mathcal{P}_z . The latter seems difficult to do in a flexible and general way, though it poses an interesting question for future research. Fixing the latent distributions, on the other hand, is easy, but it comes with questions about flexibility and interpretation. We briefly discuss these here, and note that more work is required.

Fixing the latent distributions before fitting the generator is one way to achieve strong identifiability, though *how* the distributions ought to be specified depends on the desired use of the model, and the framework for doing so remains unresolved. A task-driven approach, as in Section 3, is one avenue for exploration. The use of auxiliary data, for example u in iVAE, also appears to be a useful way to structure the latent space. In all of the “fixed distribution” identifiability results in this paper, a mapping $u \rightarrow \mathcal{P}_z$ can be learned from data (as in iVAE), as long as it is frozen at some point during training. These frozen parameters should be published alongside any results for reproducibility. We note this does not hinder efforts to generalize to new environments—the fixed environments E^* (as in Theorem 4.5) may act as “anchors” that shape the latent space for the remaining distributions in $E \setminus E^*$.

5 IDENTIFIABILITY VIA TRANSPORTS

In Section 4.3, \mathcal{F} was left unconstrained; strong identifiability was achieved through multiple environments and restricting the environment latent variable distributions to be fixed members of an exponential family. In this section we construct strongly identifiable models in which the latent distributions can be *any* fixed distributions with strictly positive density, and without requiring observations from multiple environments (i.e., auxiliary information). This is achieved by restricting the class of generators, with the additional condition that $\mathbf{Z} = \mathbf{X} = \mathbb{R}^d$. For simplicity of presentation, results are stated for a single environment, but the results also apply in multiple environments or views.

We approach the problem indirectly, by considering what properties of $\mathcal{A}(\mathcal{F})$ would yield strong identifiability. We aim to specify \mathcal{F} such that when $\bar{A}_{a,b} = f_b^{-1} \circ f_a$ transports one latent distribution to another, it is unique in a suitable sense so that, in particular, when it transports a distribution to itself, it must be (equivalent to) the identity map.

5.1 Knöthe–Rosenblatt Transports

Triangular monotone increasing (TMI) maps are of growing interest in generative modelling (Kingma et al., 2016; Papamakarios et al., 2017; Jaini et al., 2019; Irons et al., 2022). In particular, they can be used to approximate any fully supported distribution, and can be parametrized via deep neural networks (Huang et al., 2018; Wehenkel and

Loupe, 2019). Given two probability measures P_a, P_b on \mathbb{R}^d , there always exists a (P_a, P_b) -measure isomorphism with an explicit construction as a TMI map in terms of one dimensional conditional CDF transforms. This is known as the Knöthe–Rosenblatt (KR) transport.

The KR transport has several appealing properties. 1) It is the unique (up to a.e. equivalence) TMI map that transports between P_a and P_b . 2) The class of TMI maps are closed under composition and inversion. 3) Due to its construction, if $P_a = P_b$, then it is almost everywhere equal to the identity map. These properties are ideal for specifying identifiable generative models: 1) gives us a criterion (TMI maps); 2) ensures that the resulting generator transforms are also TMI maps; and 3) says that the resulting measure preserving automorphisms are trivial (akin to triangular matrices in factor analysis). For more details, see Appendix E.

Theorem 5.1. *Let $\mathbf{Z} = \mathbf{X} = \mathbb{R}^d$ with fixed latent distribution P_z with strictly positive density. If \mathcal{F} is the set of all TMI maps, then $\mathcal{M}(\mathcal{F}, \{P_z\})$ is strongly identifiable.*

This is a generalization of linear factor models with triangular F (Geweke and Zhou, 1996; Aguilar and West, 2000). Such a generator class also provides identifiability in the ICA sense—we thus believe the following result may also be of independent interest to the ICA community.

Proposition 5.2. *Let $\mathbf{Z} = \mathbf{X} = \mathbb{R}^d$. The nonlinear ICA model where \mathcal{F} are TMI maps and \mathcal{P}_z are fully supported distributions with independent components is identifiable up to invertible, component-wise transformations.*

5.2 Optimal Transport Indeterminacies

Though the KR transport is not an optimal transport map itself, it is the limit of optimal transport maps for a sequence of appropriately weighted quadratic costs (Santambrogio, 2015, Ch. 2.4), and the properties that yield identifiability may apply to some optimal transport maps more generally.

Given two probability measures P_a and P_b on \mathbb{R}^d , the Monge formulation of the optimal transport problem with respect to the cost function $c: \mathbf{Z} \times \mathbf{Z} \rightarrow \mathbb{R}_+$ is to find a map $T: \mathbf{Z} \rightarrow \mathbf{Z}$ such that $T_{\#}P_a = P_b$ and that minimizes the total cost (Santambrogio, 2015),

$$\int_{\mathbb{R}^d} c(z, T(z)) dP_a(z). \quad (16)$$

We call T an optimal transport (OT) map with respect to c if it minimizes (16) for transporting between *some* pair of probability distributions on \mathbb{R}^d . Let \mathcal{T}_c be the set of OT maps with respect to a cost c . Most OT cost functions are derived from metrics, and thus typically $c(z_1, z_2) = 0 \iff z_1 = z_2$, in which case we say that c is separating. This has the implication that the unique OT map from a distribution P_z to itself must be equal to the identity map P_z -almost everywhere, as it incurs the minimal cost 0. We use this to formulate a sufficient condition for strong identifiability.

Theorem 5.3. *Let $\mathbf{Z} = \mathbf{X} = \mathbb{R}^d$, with fixed latent distribution P_z with full support on \mathbb{R}^d . If $\mathcal{A}(\mathcal{F}) \subseteq \mathcal{T}_c$ for a separating cost, then $\mathcal{M}(\mathcal{F}, \{P_z\})$ is strongly identifiable.*

We are unaware of any general and flexible function classes that satisfy this property. In particular, properties 1) and 3) that yield strong identifiability in TMI maps are generally true for optimal transport maps, but property 2), the closure property, typically does not hold. Detailed specification of such models is left for future work.

6 CONCLUSION

We have developed a general formal framework for analyzing the sources of indeterminacy in a broad class of generative models. The framework brings seemingly disparate approaches to the problem together: if a model satisfies our assumptions then any identifiability result *must* be a special case of Theorem 2.2. The theory is descriptive rather than constructive, in the sense that proving identifiability results for specific models still requires the non-trivial work of characterizing $\mathcal{A}(\mathcal{F})$ and $\mathcal{A}(P_z)$. To that end, Theorem 2.2 also makes the sources of indeterminacy visible, enabling more straightforward reasoning during model design. That visibility was crucial to our strong identifiability results, particularly for the transport-based models in Section 5. Those models are far from exhaustive, and we believe that our framework can be useful in developing novel strongly identifiable generative models.

6.1 Limitations and Societal Implications

As with much of the identifiability literature, our results require that the generator is injective with ignorable noise; this is a technical limitation that must be overcome in order to have a theory that is applicable to the full array of models in use. We also note that incorporating inference from finite data into the framework is an important unsolved problem; our framework is limited in that respect. Finally, generative models with stronger identifiability properties can be used to the benefit of society (e.g., latent causal models for drug development), or to its detriment (e.g., better control over harmful synthetic data generation). The present work does not specifically address either, as it is purely theoretical; the potential impact depends on the application.

Acknowledgements

The authors acknowledge helpful comments from anonymous reviewers at AISTATS 2023, as well as the organizers and participants of the *Generative Models and Uncertainty Quantification Workshop* (Copenhagen, September, 2022) for challenging and helpful discussions. BBR gratefully acknowledges the support of the Natural Sciences and Engineering Research Council of Canada (NSERC): RGPIN-2020-04995, RGPAS-2020-00095, DGEGR-2020-00343.

References

- Aguilar, O. and West, M. (2000). Bayesian dynamic factor models and portfolio allocation. *Journal of Business & Economic Statistics*, 18(3).
- Ahuja, K., Hartford, J., and Bengio, Y. (2022a). Properties from mechanisms: An equivariance perspective on identifiable representation learning. In *ICLR 2022*.
- Ahuja, K., Hartford, J., and Bengio, Y. (2022b). Weakly supervised representation learning with sparse perturbations. In *NeurIPS 2022*.
- Arvanitidis, G., Hansen, L., and Hauberg, S. (2018). Latent space oddity: On the curvature of deep generative models. In *ICLR 2018*.
- Bogachev, V. I. (2007). *Isomorphisms of measure spaces*. Springer.
- Brehmer, J., de Haan, P., Lippe, P., and Cohen, T. (2022). Weakly supervised causal representation learning. In *NeurIPS 2022*.
- Buchholz, S., Besserve, M., and Schölkopf, B. (2022). Function classes for identifiable nonlinear independent component analysis. In *NeurIPS 2022*.
- Çınlar, E. (2011). *Probability and stochastics*, volume 261. Springer.
- Comon, P. (1994). Independent Component Analysis, a new concept? *Signal Processing*, 36(3):287–314.
- Geweke, J. and Zhou, G. (1996). Measuring the pricing error of the arbitrage pricing theory. *The review of financial studies*, 9(2).
- Gresele, L., Kügelgen, J. V., Stimper, V., Schölkopf, B., and Besserve, M. (2021). Independent mechanism analysis, a new concept? In *NeurIPS 2021*.
- Gresele, L., Rubenstein, P. K., Mehrjou, A., Locatello, F., and Schölkopf, B. (2019). The incomplete Rosetta Stone problem: Identifiability results for multi-view nonlinear ICA. In *UAI 2019*.
- Hälvä, H., Corff, S. L., Lehericy, L., So, J., Zhu, Y., Gassiat, E., and Hyvärinen, A. (2021). Disentangling identifiable features from noisy data with structured nonlinear ICA. In *NeurIPS 2021*.
- Hastie, T., Tibshirani, R., and Friedman, J. H. (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, volume 2. Springer.
- Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., and Lerchner, A. (2017). beta-VAE: Learning basic visual concepts with a constrained variational framework. In *ICLR 2017*.
- Huang, C.-W., Krueger, D., Lacoste, A., and Courville, A. (2018). Neural autoregressive flows. In *ICML 2018*.
- Hyvarinen, A., Karhunen, J., and Oja, E. (2001). *Independent Component Analysis and Blind Source Separation*, volume 1. John Wiley & Sons Hoboken, NJ.
- Hyvärinen, A. and Morioka, H. (2016). Unsupervised feature extraction by time-contrastive learning and nonlinear ICA. In *NeurIPS 2016*.
- Hyvärinen, A. and Morioka, H. (2017). Nonlinear ICA of temporally dependent stationary sources. In *AISTATS 2017*.
- Hyvärinen, A. and Pajunen, P. (1999). Nonlinear Independent Component Analysis: Existence and uniqueness results. *Neural Networks*, 12(3):429–439.
- Hyvärinen, A., Sasaki, H., and Turner, R. E. (2018). Nonlinear ICA using auxiliary variables and generalized contrastive learning. In *AISTATS 2019*, pages 859–868.
- Irons, N. J., Scetbon, M., Pal, S., and Harchaoui, Z. (2022). Triangular flows for generative modeling: Statistical consistency, smoothness classes, and fast rates. In *AISTATS 2022*.
- Jaini, P., Selby, K. A., and Yu, Y. (2019). Sum-of-squares polynomial flow. In *ICML 2019*.
- Kechris, A. S. (1995). *Classical Descriptive Set Theory*. Springer.
- Khemakhem, I., Kingma, D. P., Monti, R. P., and Hyvärinen, A. (2020a). Variational autoencoders and nonlinear ICA: A unifying framework. In *AISTATS 2020*.
- Khemakhem, I., Monti, R. P., Kingma, D. P., and Hyvärinen, A. (2020b). ICE-BeeM: Identifiable conditional energy-based deep models based on nonlinear ICA. In *NeurIPS 2020*.
- Kingma, D. and Welling, M. (2013). Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114*.
- Kingma, D. P., Salimans, T., Jozefowicz, R., Chen, X., Sutskever, I., and Welling, M. (2016). Improved variational inference with inverse autoregressive flow. In *NeurIPS 2016*.
- Kivva, B., Rajendran, G., Ravikumar, P. K., and Aragam, B. (2022). Identifiability of deep generative models under mixture priors without auxiliary information. In *NeurIPS 2022*.
- Klindt, D., Schott, L., Sharma, Y., Ustyuzhaninov, I., Brendel, W., Bethge, M., and Palton, D. M. (2021). Towards nonlinear disentanglement in natural data with temporal sparse decoding. In *ICLR 2021*.
- Lachapelle, S., Rodriguez, P., Sharma, Y., Everett, K. E., Le Priol, R., Lacoste, A., and Lacoste-Julien, S. (2022). Disentanglement via mechanism sparsity regularization: A new principle for nonlinear ICA. In *Conference on Causal Learning and Reasoning*.
- Lawley, D. N. and Maxwell, A. E. (1962). Factor Analysis as a statistical method. *Journal of the Royal Statistical Society, Series D (The Statistician)*, 12(3):209–229.
- Locatello, F., Bauer, S., Lucic, M., Rätsch, G., Gelly, S., Schölkopf, B., and Bachem, O. (2019). Challenging

- common assumptions in the unsupervised learning of disentangled representations. In *ICML 2019*.
- Locatello, F., Poole, B., Rätsch, G., Schölkopf, B., Bachem, O., and Tshannen, M. (2020). Weakly-supervised disentanglement without compromises. In *ICML 2020*.
- Lu, C., Wu, Y., Hernández-Lobato, J. M., and Schölkopf, B. (2022). Invariant causal representation learning for out-of-distribution generalization. In *ICLR 2022*.
- Monti, R. P., Zhang, K., and Hyvärinen, A. (2020). Causal discovery with general non-linear relationships using non-linear ICA. In *UAI 2020*.
- Moran, G. E., Sridhar, D., Wang, Y., and Blei, D. M. (2022). Identifiable variational autoencoders via sparse decoding. *Transactions on Machine Learning Research*.
- Papamakarios, G., Pavlakou, T., and Murray, I. (2017). Masked autoregressive flow for density estimation. In *NeurIPS 2017*.
- Santambrogio, F. (2015). *Optimal Transport for Applied Mathematicians*. Birkhäuser Cham.
- Schilling, R. L. (2005). *Measures, Integrals and Martingales*. Cambridge University Press.
- Schölkopf, B., Locatello, F., Bauer, S., Ke, N. R., Kalchbrenner, N., Goyal, A., and Bengio, Y. (2021). Toward causal representation learning. *Proceedings of the IEEE*, 109(5).
- Shen, X., Liu, F., Dong, H., Lian, Q., Chen, Z., and Zhang, T. (2022). Weakly supervised disentangled generative causal representation learning. *Journal of Machine Learning Research*, 23.
- Shimizu, S., Hoyer, P. O., Hyvärinen, A., Kerminen, A., and Jordan, M. (2006). A linear non-Gaussian acyclic model for causal discovery. *Journal of Machine Learning Research*, 7(10).
- Sorenson, P., Rother, C., and Köthe, U. (2020). Disentanglement by non-linear ICA with general incompressible-flow networks (GIN). In *ICLR 2020*.
- Stühmer, J., Turner, R., and Nowozin, S. (2020). Independent subspace analysis for unsupervised learning of disentangled representations. In *AISTATS 2020*.
- Tamer, E. (2010). Partial identification in econometrics. *Annu. Rev. Econ.*
- van der Vaart, A. (1998). *Asymptotic Statistics*. Cambridge University Press.
- Wang, Y., Blei, D., and Cunningham, J. P. (2021). Posterior collapse and latent variable non-identifiability. In *NeurIPS 2021*.
- Wang, Y. and Jordan, M. I. (2021). Desiderata for representation learning: A causal perspective. *arXiv preprint arXiv:2109.03795*.
- Wehenkel, A. and Louppe, G. (2019). Unconstrained monotonic neural networks. In *NeurIPS 2019*.
- Yang, M., Liu, F., Chen, Z., Shen, X., Hao, J., and Wang, J. (2021). CausalVAE: Disentangled representation learning via neural structural causal models. In *CVPR 2021*.
- Zheng, Y., Ng, I., and Zhang, K. (2022). On the identifiability of nonlinear ICA: Sparsity and beyond. In *NeurIPS 2022*.
- Zhou, D. and Wei, X.-X. (2020). Learning identifiable and interpretable latent models of high-dimensional neural activity using pi-VAE. In *NeurIPS 2020*.

A INDETERMINACIES AND PARAMETER IDENTIFIABILITY

In this section, we justify the study of indeterminacy transformations, which are mappings on \mathbf{Z} , as a proxy for identifiability in parameter space. For any $A \in \text{Aut}(\mathbf{Z})$ and $\theta = (f, P_z) \in \mathcal{F} \times \mathcal{P}_z$, define $A\theta = (f \circ A^{-1}, A_{\#}P_z)$. We say that $A\theta_a = \theta_b$ if $A_{\#}P_{z,a} = P_{z,b}$ and $f_a \circ A^{-1} \stackrel{\text{a.c.}}{=} f_b$. Finally, let $\mathcal{A}(\mathcal{M})|_{\theta} = \{A \in \mathcal{A}(\mathcal{M}) : P_{\theta} = P_{A\theta}\}$ —we can show that this is a local set of indeterminacy transformations “originating” at θ (Lemma A.1).

Lemma A.1. $\mathcal{A}(\mathcal{M})|_{\theta_a}$ consists exactly the collection of indeterminacy transformations at θ_a and all $\theta_{\bullet} \sim \theta_a$.

Proof. Let \mathcal{A}' denote the set of indeterminacy transformations at θ_a and some arbitrary $\theta_{\bullet} \sim \theta_a$. We show that $\mathcal{A}' = \mathcal{A}(\mathcal{M})|_{\theta_a}$.

We first show that $\mathcal{A}' \subset \mathcal{A}(\mathcal{M})|_{\theta_a}$. Let $\theta_{\bullet} \sim \theta_a$, i.e., $P_{\theta_a} = P_{\theta_{\bullet}}$. Then $A_{a,\bullet} \in \mathcal{A}'$ is an indeterminacy transformation at $\theta_a, \theta_{\bullet}$. By Lemma 2.1, $A_{a,\bullet}P_{z,a} = P_{z,\bullet}$, and by definition of an indeterminacy transformation, $f_a \circ A_{a,\bullet}^{-1} \stackrel{\text{a.c.}}{=} f_{\bullet}$. Hence $A_{a,\bullet}\theta_a = \theta_{\bullet}$, which shows that $A_{a,\bullet} \in \mathcal{A}(\mathcal{M})|_{\theta_a}$.

We now show that $\mathcal{A}(\mathcal{M})|_{\theta_a} \subset \mathcal{A}'$. Let $A \in \mathcal{A}(\mathcal{M})|_{\theta_a}$ and denote $A\theta = \theta_{\bullet}$, so that $P_{\theta} = P_{A\theta} = P_{\theta_{\bullet}}$. By definition of $A\theta$, we have $f_a \circ A^{-1} \stackrel{\text{a.c.}}{=} f_{\bullet}$. Hence, A is an indeterminacy map at θ, θ_{\bullet} , i.e., $A \in \mathcal{A}'$. \square

We are now ready to prove Proposition A.2.

Proposition A.2. The equivalence class of each $\theta \in \mathcal{F} \times \mathcal{P}_z$ is generated by $\mathcal{A}(\mathcal{M})|_{\theta}$. That is, $[\theta] = \{A\theta : A \in \mathcal{A}(\mathcal{M})|_{\theta}\}$.

Proof of Proposition A.2. For notational purposes, we denote an arbitrary θ by θ_a .

We wish to show that $[\theta_a] = \{A\theta_a : A \in \mathcal{A}(\mathcal{M})|_{\theta_a}\}$. Note that by Lemma A.1, $\mathcal{A}(\mathcal{M})|_{\theta_a}$ consists of indeterminacy transformations at θ and some arbitrary $\theta_{\bullet} \sim \theta_a$.

We first show that $[\theta_a] \subset \{A\theta_a : A \in \mathcal{A}(\mathcal{M})|_{\theta_a}\}$. Let $\theta_{\bullet} \in [\theta_a]$. Then, $P_{\theta_a} = P_{\theta_{\bullet}}$ which by Lemma 2.1 implies that $\vec{A}_{a,\bullet}\theta_a = \theta_{\bullet}$. Clearly, $\vec{A}_{a,\bullet}$ is an indeterminacy transformation at θ, θ_{\bullet} , and so $\vec{A} \in \mathcal{A}(\mathcal{M})|_{\theta_a}$, implying that $\theta_{\bullet} \in \{A\theta_a : A \in \mathcal{A}(\mathcal{M})|_{\theta_a}\}$.

We now show that $\{A\theta_a : A \in \mathcal{A}(\mathcal{M})|_{\theta_a}\} \subset [\theta_a]$. Let $\theta_{\bullet} \in \{A\theta_a : A \in \mathcal{A}(\mathcal{M})|_{\theta_a}\}$. Then, $\theta_{\bullet} = A_{a,\bullet}\theta_a$ for some indeterminacy transformation $A_{a,\bullet}$. Since $A_{a,\bullet}$ is an indeterminacy transformation, $P_{\theta_a} = P_{\theta_{\bullet}}$ by definition and hence $\theta_{\bullet} \in [\theta_a]$. \square

The above result states that the indeterminacy transforms (mappings on \mathbf{Z}) can be used to generate the equivalence classes of parameters. This justifies the study of indeterminacy transformations of \mathbf{Z} (and their extension to the parameter space), rather than studying the parameter space directly, and puts identifiability up to indeterminacy transformations (Definition 2) in correspondence with parameter identifiability.

B DISCRETE OBSERVATIONS

In this section, we discuss models with discrete observations, e.g., Bernoulli with probability parameter given by $f(z)$, or Poisson with mean parameter $f(z)$ (such models were briefly discussed in iVAE (Khemakhem et al., 2020a), as well as in follow-up work such as the pi-VAE (Zhou and Wei, 2020)). In short, the framework developed in this paper rests on bijective generators which enable the recovery of unique latent codes for each observed value. As noted in a correction in Khemakhem et al. (2020a), this task seems fundamentally impossible for example when the latent space is uncountable and the outcome is discrete, due to the lack of an bijective map between spaces of different cardinality. However, generative models do not typically send a latent variable to the outcome, but rather to a parameter value of a conditional distribution. This allows us to reformulate the assumptions required for our theory, although as we will see shortly, most discrete outcome models do not satisfy these assumptions. Formally, suppose \mathbf{X} is either finite or countable. Let X be a random variable on \mathbf{X} and denote by $P_x := P(X = x) : \mathbf{X} \rightarrow [0, 1]$ the probability mass function, which satisfies $\sum_{x \in \mathbf{X}} P_x(x) = 1$. Two random variables X_a, X_b are said to be equal in distribution if and only if their respective PMFs satisfy $P_{x,a}(x) = P_{x,b}(x)$ for all $x \in \mathbf{X}$.

The observation model is then described by a conditional PMF $P(X = x|z)$. We assume this model has a topological parameter space Θ and also pair it with the Borel σ -algebra, e.g., $\Theta = [0, 1]^n$ for an n -dimensional Bernoulli. Let $f : \mathbf{Z} \rightarrow \Theta$

be an injective generator (note this implies Θ has cardinality at least that of \mathbf{Z}). Then the generative model is as follows:

$$Z \sim P_z, \quad P(X = x|z) = g_x(f(z)), \quad (17)$$

where $g_x(\theta)$ is the PMF of the observation model with parameter θ at x .

Recall Assumption 1 of the main text. We now introduce its discrete analogue, which would be required for the theory developed in this paper to apply in the discrete formulation. First, note that the marginal PMF on \mathbf{X} is given as follows:

$$P_x(x) = \int_{\mathbf{Z}} g_x(f(z))P_z(dz) = \int_{\Theta} g_x(\theta)P_z(f^{-1}(d\theta)) = \mathbb{E}_{\theta} [g_x(\theta)], \quad (18)$$

where as a random variable, $\theta = f(Z)$. The assumption is then as follows:

Assumption 3 (Discrete analogue to main text: Assumption 1). Assume that $\mathbb{E}_{\theta_a} [g_x(\theta_a)] = \mathbb{E}_{\theta_b} [g_x(\theta_b)]$ for each $x \in \mathbf{X}$ if and only if $\theta_a \stackrel{\Delta}{=} \theta_b$.

In other words, the distribution of $\theta = f(Z)$ must be characterized by the moments $\mathbb{E}[g_x(\theta_a)]$, for each $x \in \mathbf{X}$. Indeed, for observational equivalence to imply anything about the latent spaces, such an assumption would be needed. However, it appears that this assumption is rarely satisfied for any reasonable models. For example, the Bernoulli observation model with $P(X = 1|z) = g_1(\theta) = f(z)$, $\Theta = [0, 1]$, requires that the distribution of θ be characterized by just its first moment, $\mathbb{E}[\theta]$. Of course, this is highly unlikely unless very strict restrictions are placed on f and P_z , such as $f_{\#}P_z$ being Gaussian, indicating that identifiability in this case may be restricted to linear generators.

More generally, the core of the issue remains the cardinality mismatch between \mathbf{Z} and \mathbf{X} . A necessary condition for $\theta_a \stackrel{\Delta}{=} \theta_b$ is that $\mathbb{E}[g(\theta_a)] = \mathbb{E}[g(\theta_b)]$ for all bounded continuous $g : \Theta \rightarrow \mathbb{R}$ (test functions). For Θ uncountable, there are clearly uncountably many test functions, while in our discrete assumption above, there are countably many test functions at best. Though we do not make this notion precise here, we believe this makes the discrete assumption above unlikely to be satisfied, and hence any notion of identifiability, at least under our framework (which we believe to be reasonably general), is highly unlikely for discrete outcomes with uncountable latent spaces.

C DEFINITIONS AND PRELIMINARIES

The rest of this Appendix takes place in the mathematical setting of measure-theoretic probability. This section will review some relevant definitions. Our standard reference here will be Çinlar (2011), though we will sometimes refer to other texts depending on the specific topic (Kechris, 1995; Schilling, 2005; Bogachev, 2007).

C.1 Basic Definitions

Let (E, τ) be a topological space with E a set and τ its collection of open sets.

Definition 4 (σ -algebras, Çinlar, 2011, Eq. 1.3). A collection \mathcal{E} of subsets of E is called a σ -algebra on E if it is closed under complements and countable unions:

$$B \in \mathcal{E} \implies E \setminus B \in \mathcal{E}, \quad B_1, B_2, \dots \in \mathcal{E} \implies \cup_n A_n \in \mathcal{E}. \quad (19)$$

Note that a σ -algebra always contains the empty set and E itself. The pair (E, \mathcal{E}) defines a measurable space. The elements of \mathcal{E} in this context are called measurable sets. When the σ -algebra is insignificant, or obvious by context, we will simply refer to the space by E .

Definition 5 (Generated σ -algebras, Çinlar, 2011, Sec. 1). The σ -algebra generated by a collection of subsets \mathcal{E}' , denoted $\sigma(\mathcal{E}')$ is the smallest σ -algebra that contains \mathcal{E}' .

In this work, we will always work with what is known as the Borel σ -algebra.

Definition 6 (Borel σ -algebras, Çinlar, 2011, Sec. 1). Let (E, τ) be a topological space. The Borel σ -algebra of E is generated by the collection of open sets, $\sigma(\tau)$. We denote it by $\mathcal{B}(E)$.

An element $B \in \mathcal{B}(E)$ is then said to be a Borel set.

Let $(E, \mathcal{E}), (F, \mathcal{F})$ be two measurable spaces, and $f : E \rightarrow F$ a mapping between them. The image of $A \subset E$ is defined as

$$f(A) = \{f(a) \mid a \in A\} \subset F. \quad (20)$$

Similarly, the preimage of $B \subset F$ is defined as

$$f^{-1}(B) = \{x \in E \mid f(x) \in B\} \subset E. \quad (21)$$

Most mappings we will be concerned with will be assumed to be *measurable*.

Definition 7 (Measurable Mappings, [Çinlar, 2011](#), Sec. 2). A mapping $f : E \rightarrow F$ is said to be $(\mathcal{E}, \mathcal{F})$ -measurable if $f^{-1}(B) \in \mathcal{E}$ for each $B \in \mathcal{F}$.

If $(E, \mathcal{E}) = (F, \mathcal{F})$, we will refer to f as simply \mathcal{E} -measurable.

Definition 8 (Measures, [Çinlar, 2011](#), Sec. 3). Given a measurable space (E, \mathcal{E}) , a mapping $\mu : \mathcal{E} \rightarrow [0, \infty]$ is a measure if it satisfies

$$\mu(\emptyset) = 0, \quad \mu(\cup_n A_n) = \sum_n \mu(A_n), \quad (A_n) \text{ a disjoint sequence in } \mathcal{E}. \quad (22)$$

In particular, if $\mu(E) = 1$, then μ is said to be a probability measure.

The triplet (E, \mathcal{E}, μ) is known as a measure space, and, if μ is a probability measure, as a probability space. When there is no risk for confusion, we simply identify a measure space by its measure μ . Two measures μ and ν on the same measurable space are equal whenever

$$\mu(B) = \nu(B), \text{ for all } B \in \mathcal{E}. \quad (23)$$

C.2 Random Variables, Pushforward measures

([Çinlar, 2011](#), Ch. 2) covers probability spaces in depth. Here, we only review the relevant notion pertaining to random variables and their distributions. A random variable X on (E, \mathcal{E}) is associated to a probability measure μ , called its distribution, defined as³

$$\mu(B) = P(X \in B), \text{ for all } B \in \mathcal{E}. \quad (24)$$

Random variables X, Y defined on the same measurable space with distributions μ, ν are said to be equal in distribution if $\mu = \nu$ as probability measures, denoted $X \stackrel{d}{=} Y$.

Any $(\mathcal{E}, \mathcal{F})$ -measurable function $f : E \rightarrow F$ applied to X , denoted $f(X)$, defines a random variable on (F, \mathcal{F}) with distribution $\mu \circ f^{-1}$, where f^{-1} denotes the preimage of f as a set function $\mathcal{F} \rightarrow \mathcal{E}$. Whenever it is convenient, we will use the more streamlined *pushforward* notation for the distribution of $f(X)$, as follows:

$$f_{\#}\mu = \mu \circ f^{-1}. \quad (25)$$

Finally, we define the notion of a pushforward σ -algebra.

Definition 9 (Pushforward σ -algebras, [Çinlar, 2011](#), Sec. 2). Let (E, \mathcal{E}) be a measurable space, F be a set, and f be a mapping $f : E \rightarrow F$. The pushforward σ -algebra of f is defined as

$$\sigma(f) = \{B \subset F; f^{-1}(B) \in \mathcal{E}\} \quad (26)$$

It is easily shown that $\sigma(f)$ is a σ -algebra on F , and that f is measurable with respect to $\sigma(f)$ ([Çinlar, 2011](#), Exercise 2.20). In fact, it is the smallest σ -algebra that makes f measurable.

C.3 Null sets and absolute continuity

Definition 10 (Null Sets, [Çinlar, 2011](#), Sec. 3). Given a measurable space (E, \mathcal{E}) , a measurable set $A \in \mathcal{E}$ is said to be **null** with respect to a measure μ , or μ -null, if $\mu(A) = 0$.

³Technically, we require a background measure space (Ω, \mathcal{F}, P) , and a random variable is defined as a measurable function $X : \Omega \rightarrow E$.

The empty set is always a null set by the definition of a measure, but there can be many more null sets, depending on the measure. Any countable union of null sets is again null. For example, the Lebesgue measure λ on \mathbb{R} assigns null measure to a singleton $\{x\}$ for $x \in \mathbb{R}$, and so both the set of natural numbers and rational numbers are λ -null sets.

For a measure space (E, \mathcal{E}, μ) , most properties that are consequences of measure-theoretic manipulations can only hold μ -almost everywhere. This is a weaker notion than a property holding pointwise, and the strength of a result can depend on the measure μ .

Definition 11 (Almost Everywhere, [Çinlar, 2011](#), Sec. 3). Given a measure space, a property that is stated for $x \in E$ is said to hold μ -almost everywhere if there exists a measurable set N with $\mu(N) = 0$ such that P holds for all $x \in E \setminus N$.

Furthermore, two measures are often compared with respect to their null sets.

Definition 12 (Absolute Continuity, [Çinlar, 2011](#), p. 31). A measure μ is said to be absolutely continuous with respect to ν defined on the same measurable space, denoted $\mu \ll \nu$, if for any $A \in \mathcal{E}$ such that $\nu(A) = 0$, we also have $\mu(A) = 0$.

An equivalent property is given by the Radon-Nikodym theorem.

Theorem C.1 (Radon-Nikodym, [Çinlar, 2011](#), Thm. 5.11). $\mu \ll \nu$ on (E, \mathcal{E}) if and only if there exists a measurable function $p : E \rightarrow [0, \infty)$, uniquely defined ν -almost everywhere, such that for any measurable set $A \in \mathcal{E}$ and measurable function f , we have

$$\int_{x \in A} f(x) \mu(dx) = \int_{x \in A} p(x) f(x) \nu(dx). \quad (27)$$

We call p the density of μ with respect to ν .

Typically, when discussing a probability measure P on \mathbb{R}^d , p is the density of P with respect to the Lebesgue measure λ , implicitly assuming that $P \ll \lambda$. In this work, we will also be working in this context, referring to p as a probability density, unless stated otherwise.

Definition 13 (Equivalence, [Schilling, 2005](#), Problem 19.5). Two measures μ, ν are said to be **equivalent** if $\mu \ll \nu$ and $\nu \ll \mu$.

Clearly, equivalent measures assign the exact same null sets, and imply the same “almost everywhere” statements. There is again an analogous definition in terms of densities.

Lemma C.2. [[Schilling, 2005](#), Exercise 19.5] If μ and ν are two measures defined on the same measurable space, then any density p of one with respect to the other satisfies $p(x) > 0$ μ -almost everywhere (equiv. ν -almost everywhere) if and only if they are equivalent.

When working on Euclidean spaces, we will simply write almost everywhere (or a.e.) to mean λ -almost everywhere—if a probability measure P_z is equivalent to λ , then P_z -almost everywhere is equivalent.

Finally, it should be intuitively obvious that if f and g are μ -almost-everywhere equal, then pushing forward μ by either f or g results in the same measure.

Proposition C.3. Let (E, \mathcal{E}, μ) be a measure space and (F, \mathcal{F}) a measurable space. Let $f, g : E \rightarrow F$ be measurable. Suppose $f = g$ μ -almost everywhere. Then, we have $f_{\#}\mu = g_{\#}\mu$ on \mathcal{F} .

Proof. Let $B \in \mathcal{F}$, and $\mathbf{1}_B$ be the indicator function for B . Clearly, $\mathbf{1}_B \circ f = \mathbf{1}_B \circ g$ μ -almost everywhere. Then,

$$f_{\#}\mu(B) = \int \mathbf{1}_B(x) \mu(f^{-1}(dx)) = \int \mathbf{1}_B \circ f(x) \mu(dx) \quad (28)$$

$$= \int \mathbf{1}_B \circ g(x) \mu(dx) = \int \mathbf{1}_B(x) \mu(g^{-1}(dx)) = g_{\#}\mu(B), \quad (29)$$

since $\int f d\mu = \int g d\mu$ for $f = g$ μ -almost everywhere. □

C.4 Bijective Mappings, Borel Isomorphisms

This section reviews relevant facts about Borel isomorphisms, the main reference is ([Kechris, 1995](#), Ch. 15). In this section, we refer to the Borel spaces measurable spaces $(E, \mathcal{B}(E))$ and $(F, \mathcal{B}(F))$ simply by E and F .

Bijjective mappings $f : E \rightarrow F$ can enjoy some additionally nice properties. First, it can be easily shown that

$$f(f^{-1}(A)) = A, \text{ for all } A \subset E. \quad (30)$$

Furthermore, the image $f(A)$ defines the pre-image of the inverse mapping f^{-1} . This is not necessarily true for non-bijjective mappings.

We now define the notion of isomorphism between Borel measurable spaces.

Definition 14 (Borel Isomorphism, [Kechris, 1995](#), Ch. 10.B). Let $f : E \rightarrow F$ be a bijjective mapping. If f and f^{-1} are both measurable, it is known as an **Borel isomorphism**, and E, F are said to be isomorphic. If $E = F$, f is called an **Borel automorphism**.

We denote the space of Borel automorphisms of E as $\text{Aut}(E)$.

Lemma C.4 ([Kechris, 1995](#), Corollary 15.2). For $f : E \rightarrow F$ Borel and injective, $f(B) \in \mathcal{B}(F)$ for any $B \in \mathcal{B}(E)$.

Since compositions of measurable functions remain measurable, the composition of Borel isomorphisms is again a Borel isomorphism of the appropriate spaces. Borel-measurable bijections are particularly nice to work with—they are automatically Borel isomorphisms.

Lemma C.5. Let E, F be Borel spaces and $f : E \rightarrow F$ be bijective. Then, f is measurable if and only if it is a Borel isomorphism.

Proof. We only prove the forward direction, i.e., showing f^{-1} is measurable if f is measurable. The reverse direction is identically proved. By ([Kechris, 1995](#), Theorem 15.1), for f Borel-measurable and injective, $f(B)$ is a Borel set of F for any Borel set B of E . Since $f(B)$ defines the pre-image of f^{-1} for any Borel set B of E , it immediately follows that f^{-1} is also Borel-measurable. \square

Finally, we define the notion of a measure auto/isomorphism.

Definition 15 (Measure Isomorphism, [Bogachev, 2007](#), Sec 9.2). Let (E, \mathcal{E}, μ) and (F, \mathcal{F}, ν) be two Borel measure spaces. A Borel isomorphism $f : E \rightarrow F$ is called a (μ, ν) -measure isomorphism if

$$f_{\#}\mu = \nu, \quad f^{-1}_{\#}\nu = \mu. \quad (31)$$

It is called a μ -measure-preserving automorphism if $(E, \mathcal{E}, \mu) = (F, \mathcal{F}, \nu)$.

Furthermore, by Proposition [C.3](#), if f is a measure auto/isomorphism and $f = g$ μ -almost everywhere, then g is also a measure auto/isomorphism.

D PROOFS

Before proceeding into the proofs of any specific result, we first state a useful technical fact.

Lemma D.1. Let $f_a, f_b \in \mathcal{F}$. Then, $\vec{A}_{a,b} = f_b^{-1} \circ f_a \in \text{Aut}(\mathbf{Z})$.

Proof. The pre-image of $\vec{A}_{a,b}$ is $f_a^{-1} \circ f_b$, since \mathcal{F} is a family of injective functions. By definition of measurability, it suffices to show that $f_a^{-1}(f_b(B)) \subset \mathbf{Z}$ is still Borel for any Borel set $B \in \mathcal{B}(\mathbf{Z})$. By Lemma [C.4](#), $f_b(B) \subset \mathbf{X}$ is Borel. Then, by measurability of f_a , it follows that $f_a^{-1}(f_b(B))$ is Borel and hence $\vec{A}_{a,b}$ is measurable. \square

D.1 Proof of Lemma [2.1](#)

The proof of Lemma [2.1](#) is conceptually simple, but requires some careful book-keeping to make the measure theoretic arguments precise. Specifically, we are required to analyze $\vec{A}_{a,b}$, involving f_b^{-1} , which is only well-defined on $\mathcal{F}(\mathbf{Z})$. As a result, we need to first establish a measurable space on $\mathcal{F}(\mathbf{Z})$.

Recall the definition of a pushforward σ -algebra of f (Definition [9](#)). This defines a σ -algebra on $\mathcal{F}(\mathbf{Z})$, and by construction, f is $(\mathcal{B}(\mathbf{Z}), \mathcal{F}(\mathbf{Z}))$ -measurable. Further, all generators in \mathcal{F} induce the same pushforward σ -algebra.

Lemma D.2. For f_a, f_b in \mathcal{F} , $\sigma(f_a) = \sigma(f_b)$.

Proof. To see that $\sigma(f_a) \subset \sigma(f_b)$, suppose $C \in \sigma(f_a)$. By Lemma D.3, C is Borel, which means that $f_b^{-1}(C)$ is Borel by measurability. Hence, $C \in \sigma(f_b)$. We have $\sigma(f_b) \subset \sigma(f_a)$ by the exact same argument, which implies that $\sigma(f_a) = \sigma(f_b)$. \square

We denote this shared σ -algebra by $\sigma(\mathcal{F})$. Importantly, $\sigma(\mathcal{F})$ is a subset of the Borel sets of \mathbf{X} .

Lemma D.3. $\sigma(\mathcal{F})$ contains only Borel sets. In other words, $\sigma(\mathcal{F}) \subset \mathcal{B}(\mathbf{X})$.

Proof. Let $C \in \sigma(\mathcal{F})$, and let $f \in \mathcal{F}$ be any generator. By definition, $f^{-1}(C)$ is Borel. Since f is injective, we have $f(f^{-1}(C)) = C$. By Lemma C.4, C must be Borel. \square

We are now ready to construct the measurable space $(\mathcal{F}(\mathbf{Z}), \sigma(\mathcal{F}))$. Note the following facts about $(\mathcal{F}(\mathbf{Z}), \sigma(\mathcal{F}))$:

- For any $f \in \mathcal{F}$, $f : \mathbf{Z} \rightarrow \mathcal{F}(\mathbf{Z})$ is bijective, and $f^{-1} : \mathcal{F}(\mathbf{Z}) \rightarrow \mathbf{Z}$ is well defined.
- For any $f \in \mathcal{F}$ and a Borel set $B \in \mathcal{B}(\mathbf{Z})$, its image $f(B) \subset \mathcal{F}(\mathbf{Z})$ is also the pre-image of f^{-1} —that is, $(f^{-1})^{-1}(f(B)) = B$.
- Since $\sigma(\mathcal{F}) \subset \mathcal{B}(\mathbf{X})$, if any measures are equal on $\mathcal{B}(\mathbf{X})$, then they are also equal on $\sigma(\mathcal{F})$.

We are now ready to prove Lemma 2.1. Recall that to say $A_{a,b}$ is a $(P_{z,a}, P_{z,b})$ -measure isomorphism is to say that $A_{a,b\#}P_{z,a} = P_{z,b}$ —equivalently, for each $B \in \mathcal{B}(\mathbf{Z})$, $P_{z,a}(A_{a,b}^{-1}(B)) = P_{z,b}(B)$.

Proof of Lemma 2.1. We first show the claim that $P_{\theta_a} = P_{\theta_b}$ if and only if $\vec{A}_{a,b}$ is a $(P_{z,a}, P_{z,b})$ -measure isomorphism.

\Rightarrow Recall that $\vec{A}_{a,b} \in \text{Aut}(\mathbf{Z})$, by Lemma D.1. By Assumption 1 in the main text, $P_{\theta_a} = P_{\theta_b}$ implies that $f_{a\#}P_{z,a} = f_{b\#}P_{z,b}$ as measures on $\mathcal{B}(\mathbf{X})$. By Lemma D.3, this implies $f_{a\#}P_{z,a} = f_{b\#}P_{z,b}$ also as measures on $\sigma(\mathcal{F})$. Now, let $B \in \mathcal{B}(\mathbf{Z})$. Then,

$$P_{z,a}(\vec{A}_{a,b}^{-1}(B)) = P_{z,a}(f_a^{-1}(f_b(B))) = P_{z,b}(f_b^{-1}(f_b(B))) = P_{z,b}(B), \quad (32)$$

where the first equality is by definition (working on $\sigma(\mathcal{F})$), the second equality is due to $f_{a\#}P_{z,a} = f_{b\#}P_{z,b}$, and the third equality is due to injectivity (on the measurable space $(\mathcal{F}(\mathbf{Z}), \sigma(\mathcal{F}))$). Since B was arbitrary, this shows that $P_{z,a} \circ \vec{A}_{a,b}^{-1} = P_{z,b}$.

\Leftarrow We show the contrapositive statement, i.e.,

$$\vec{A}_{a,b\#}P_{z,a} \neq P_{z,b} \implies P_{\theta_a} \neq P_{\theta_b} \quad (33)$$

Note that by Assumption 1, the hypothesis is equivalent to $f_{a\#}P_{z,a} \neq f_{b\#}P_{z,b}$. This means that there exists $B \in \mathcal{B}(\mathbf{X})$ such that

$$P_{z,a}(f_a^{-1}(B)) \neq P_{z,b}(f_b^{-1}(B)). \quad (34)$$

To show that $\vec{A}_{a,b\#}P_{z,a} \neq P_{z,b}$ is to find $B^* \in \mathcal{B}(\mathbf{Z})$ such that

$$P_{z,a}(f_a^{-1}(f_b(B^*))) \neq P_{z,b}(B^*). \quad (35)$$

Now, $f_b^{-1}(B) \in \mathcal{B}(\mathbf{Z})$ by measurability, and hence taking $B^* = f_b^{-1}(B)$, we obtain

$$P_{z,a}(f_a^{-1}(f_b(f_b^{-1}(B)))) \neq P_{z,b}(f_b^{-1}(B)), \quad (36)$$

and hence $\vec{A}_{a,b\#}P_{z,a} \neq P_{z,b}$, as required.

Now, note that by definition of $\vec{A}_{a,b}$, it is the unique such map that $f_a(\vec{A}_{a,b}^{-1}(z)) = f_b(z)$ pointwise. Therefore, all possible indeterminacy transformations must be equivalent almost everywhere to $\vec{A}_{a,b}$. By Proposition C.3, almost everywhere equivalence preserves measure isomorphisms, and hence any indeterminacy transformation $A_{a,b}$ must be a $(P_{z,a}, P_{z,b})$ -measure isomorphism. \square

D.2 Proof of Theorem 2.2

As a result of Lemma 2.1, the proof of Theorem 2.2 is straight forward.

Proof of Theorem 2.2. Recall that, for the generative model to be identifiable up to a set of measurable functions $\mathcal{A}(\mathcal{M})$ is to say that, for all $(f_a, P_{z,a}), (f_b, P_{z,b}) \in \mathcal{F} \times \mathcal{P}_z$ such that $P_{\theta_a} = P_{\theta_b}$, there exists some $A \in \mathcal{A}(\mathcal{M})$ such that $A \stackrel{\text{a.e.}}{=} f_b^{-1} \circ f_a$.

We first show that for any parameter spaces \mathcal{F} and \mathcal{P}_z , we have that $\mathcal{A}(\mathcal{M}) \subseteq \mathcal{A}(\mathcal{F}) \cap \mathcal{A}(\mathcal{P}_z)$. Suppose $A \in \mathcal{A}(\mathcal{M})$. That is, $A \stackrel{\text{a.e.}}{=} f_b^{-1} \circ f_a$ for some $P_{z,a}, P_{z,b}$ such that $P_{\theta_a} = P_{\theta_b}$, where $\theta_a = (f_a, P_{z,a}), \theta_b = (f_b, P_{z,b})$. By definition of A , we have $A \in \mathcal{A}(\mathcal{F})$. By Lemma 2.1, we must have that $A \in \mathcal{A}(\mathcal{P}_z)$ also.

We now show that $\mathcal{A}(\mathcal{F}) \cap \mathcal{A}(\mathcal{P}_z) \subseteq \mathcal{A}(\mathcal{M})$. Suppose $A \in \mathcal{A}(\mathcal{F}) \cap \mathcal{A}(\mathcal{P}_z)$. Since $A \in \mathcal{A}(\mathcal{F})$, we can write $A \stackrel{\text{a.e.}}{=} f_b^{-1} \circ f_a$ for some $f_a, f_b \in \mathcal{F}$, e.g., $A \stackrel{\text{a.e.}}{=} \tilde{A}_{a,b}$. Since $A \in \mathcal{A}(\mathcal{P}_z)$, there exist $P_{z,a}$ and $P_{z,b}$ such that $A_{\#} P_{z,a} = P_{z,b}$. By Lemma 2.1, $\theta_a = (f_a, P_{z,a}), \theta_b = (f_b, P_{z,b})$ are such that $P_{\theta_a} = P_{\theta_b}$, and hence $A \in \mathcal{A}(\mathcal{M})$. \square

D.3 Proof of Proposition 3.1

Proof of Proposition 3.1. Equation (5) is just the definition of task identifiability (Definition 3) with $\theta' = A\theta$, along with $[\theta] = \{A\theta : A \in \mathcal{A}(\mathcal{M}) |_{\theta}\}$ from Proposition A.2. Now, assume that both equations in (6) hold. Then

$$\begin{aligned} t(A\theta, \mathbf{x}_m, s(A\theta, \mathbf{x}_m)) &= t(A\theta, \mathbf{x}_m, A(s(\theta, \mathbf{x}_m))) \\ &= t(\theta, \mathbf{x}_m, s(\theta, \mathbf{x}_m)). \end{aligned}$$

\square

D.4 Proof of Corollary 4.1

We first define an appropriate notion of identifiability in the multiple environments model indexed by $e \in E$, $\mathcal{M}(\mathcal{F}, \{\mathcal{P}_z^e\}_{e \in E})$. Since the environments are known deterministically, we view the overall model as a statistical model over each environment with parameter θ^e . We have $\theta^e = (f, P_z^e) \in \mathcal{F} \times \mathcal{P}_z^e$, where \mathcal{P}_z^e is not necessarily the same for each $e \in E$. We will denote $\theta = \{\theta^e\}_{e \in E}$, and denote the marginal distribution over \mathbf{X} in environment e as P_θ^e (recall that each of these are over \mathbf{X}). We adapt the definition of an indeterminacy transform as follows.

Definition 16 (Multiple Environments Analogue to Main Text: Definition 1). $A_{a,b}$ is an **indeterminacy transformation** at θ_a, θ_b if $P_{\theta_a}^e = P_{\theta_b}^e$ for each $e \in E$ and $f_a \circ A_{a,b}^{-1} \stackrel{\text{a.e.}}{=} f_b$.

That is, $A_{a,b}$ should be an indeterminacy transformation for the model in each environment. The indeterminacy set is again the collection of all indeterminacy transforms, and weak and strong identifiability remain as in Definition 1 and Definition 2 in this setting. Since there are more constraints on θ_a and θ_b compared to the single environment case, it should be intuitively clear that the indeterminacy set is smaller in this case. Since the generator and hence the generator transforms do not change across environments, Theorem 2.2 applies in each environment. This forms the basic idea of the following proof of Corollary 4.1.

Proof of Corollary 4.1. Fix an environment $e \in E$. Recall that an indeterminacy transformation at θ_a, θ_b in each environment is an automorphism $A_{a,b} \in \text{Aut}(\mathbf{Z})$ such that $P_{\theta_a}^e = P_{\theta_b}^e$ and $f_a \circ A_{a,b}^{-1} \stackrel{\text{a.e.}}{=} f_b$. By Theorem 2.2, any $A_{a,b}$ satisfying the above lies in $\mathcal{A}(\mathcal{F}) \cap \mathcal{A}(\mathcal{P}_z^e)$.

Now, by definition, the indeterminacy transformations as defined in Definition 16 are necessarily also indeterminacy transformations for each environment $e \in E$. Hence, for any indeterminacy transformation $A_{a,b}$, we have

$$A_{a,b} \in \bigcap_{e \in E} (\mathcal{A}(\mathcal{F}) \cap \mathcal{A}(\mathcal{P}_z^e)) = \mathcal{A}(\mathcal{F}) \cap (\bigcap_e \mathcal{A}(\mathcal{P}_z^e)), \quad (37)$$

and hence the indeterminacy set is at most $\mathcal{A}(\mathcal{F}) \cap (\bigcap_e \mathcal{A}(\mathcal{P}_z^e))$. \square

D.5 Proof of Corollary 4.2

In the multiple environments setting, the model could be viewed $|E|$ separate models which jointly identify the generator. In the multiple views case, $\mathcal{M}(\{\mathcal{F}^e\}_{e \in E}, \mathcal{P}_z)$, we assume $E = (e_1, \dots, e_n)$ and combine all views into one unifying model, on which we apply our previously developed theory.

Suppose we have the observation spaces \mathbf{X}^{e_i} and generator classes \mathcal{F}^{e_i} (all of which are injective), which may be different across views (indeed, this should be the case for the result to be interesting). Our model is then parametrized in each view by $\theta^{e_i} = (f^{e_i}, P_z)$, where $f^{e_i} \in \mathcal{F}^{e_i}$ and $P_z \in \mathcal{P}_z$.

Note that each observation $\{x_{e_i}\}_{i=1}^n$ is assumed to be generated by the same latent point z , i.e., the noiseless views are $\{f^{e_i}(z)\}_{i=1}^n$. We can equivalently state this assumption by stacking the views into a multivariate function:

$$f(z) = \begin{bmatrix} f^{e_1}(z) \\ \vdots \\ f^{e_n}(z) \end{bmatrix}, \quad (38)$$

where we have $f : \mathbf{Z} \rightarrow \times_{i=1}^n \mathbf{X}^{e_i}$. Note the image of f is significantly smaller than the product of the respective images. Writing \mathcal{F} to denote the implied parameter space of the stacked f , we can write the shared image as:

$$\mathcal{F}(\mathbf{Z}) = \{(x_1, \dots, x_n) \in \times_{i=1}^n \mathbf{X}^{e_i} \mid (f^{e_i})^{-1}(x_i) = (f^{e_j})^{-1}(x_j), \text{ for all } i, j = 1, \dots, n\}. \quad (39)$$

As an example, consider $f : \mathbb{R} \rightarrow \mathbb{R}^2$ defined by $f(z) = (z, z^3)$. Writing $x = z, y = z^3$, the image is defined by the graph of the function $y = x^3$ in \mathbb{R}^2 . Now, f is invertible as a mapping $\mathbf{Z} \rightarrow \mathcal{F}(\mathbf{Z})$, and we have

$$f^{-1}(x_1, \dots, x_n) = (f^{e_1})^{-1}(x_1) = \dots = (f^{e_n})^{-1}(x_n). \quad (40)$$

Clearly, identifying f in the usual sense is equivalent to identifying each underlying generator. Hence, we follow Definition 1 and Definition 2 in the main text, and do not need any additional definitions. For an $f \in \mathcal{F}$, $f_{\#} P_z$ is the joint distribution on $\times_{i=1}^n \mathbf{X}^{e_i}$ with marginals $f_{\#}^{e_i} P_z$.⁴ Now, let $\theta = (f = \{f^e\}_{e \in E}, P_z)$, we denote the resulting distribution over $\times_{i=1}^n \mathbf{X}^{e_i}$ by \vec{P}_{θ} , and the corresponding marginals by P_{θ}^e . Note that if θ_a and θ_b result in the same joint distribution, then their corresponding marginals match also. We now prove Corollary 4.2 (for the case $E = (e_1, \dots, e_n)$).

Proof of Corollary 4.2. Suppose $A_{a,b}$ is an indeterminacy transform at θ_a, θ_b . By Lemma 2.1, $A_{a,b} \stackrel{\text{a.e.}}{=} \vec{A}_{a,b}$. The generator transform $\vec{A}_{a,b} = f_b^{-1} \circ f_a$ is given by, for any $z \in \mathbf{Z}$

$$f_b^{-1}(f_a^{e_1}(z), \dots, f_a^{e_n}(z)) = (f_b^{e_1})^{-1}(f_a^{e_1}(z)) = \dots = (f_b^{e_n})^{-1}(f_a^{e_n}(z)). \quad (41)$$

In other words, letting $\vec{A}_{a,b}^e$ denote the generator transform for each view, we have

$$\vec{A}_{a,b} = \vec{A}_{a,b}^{e_1} = \dots = \vec{A}_{a,b}^{e_n}. \quad (42)$$

Since the marginals match at θ_a, θ_b also, $A_{a,b}^{e_i} \stackrel{\text{a.e.}}{=} \vec{A}_{a,b}^{e_i} = \vec{A}_{a,b} \stackrel{\text{a.e.}}{=} A_{a,b}$ for any indeterminacy transform in each view e_i , i.e., since the generator transforms characterize the indeterminacy transforms, $A_{a,b}$ must be an indeterminacy transformation for each e_i . By Theorem 2.2, we have $A_{a,b}^{e_i} \in \mathcal{A}(\mathcal{F}^{e_i}) \cap \mathcal{A}(P_z)$ for each e_i , which implies that

$$A_{a,b} \in \left(\bigcap_{e \in E} [\mathcal{A}(\mathcal{F}^e) \cap \mathcal{A}(P_z)] \right) = \left(\bigcap_{e \in E} \mathcal{A}(\mathcal{F}^e) \cap \mathcal{A}(P_z) \right). \quad (43)$$

□

D.6 Proof of Proposition 4.3

This section makes heavy use of probability densities and absolute continuity. Refer to Appendix C for precise definitions. We will reproduce the following Lemma for convenience.

Lemma C.2. [Schilling, 2005, Exercise 19.5] *If μ and ν are two measures defined on the same measurable space, then any density p of one with respect to the other satisfies $p(x) > 0$ μ -almost everywhere (equiv. ν -almost everywhere) if and only if they are equivalent.*

Now, we may state an intermediate result.

⁴It is sufficient for this section to characterize $f_{\#} P_z$ up to a coupling of its marginals.

Lemma D.4. Suppose probability measures $P_{z,a}, P_{z,b}$ admits strictly positive densities p_a, p_b . Suppose A is a $(P_{z,a}, P_{z,b})$ -measure isomorphism. Then,

$$p_b(A(x))k_A(x) = p_a(x) \quad \text{a.e.}, \quad (44)$$

where k_A depends only on A and is strictly positive a.e..

Proof. Since A is a $(P_{z,a}, P_{z,b})$ -measure isomorphism and $P_{z,a}, P_{z,b}$ are equivalent to λ_z , we have that for a Borel set B ,

$$\lambda_z(B) = 0 \iff P_{z,a}(B) = 0 \iff P_{z,b}(A(B)) = 0 \iff \lambda_z(A(B)) = 0, \quad (45)$$

where the first and third equivalences are because $P_{z,a}$ and $P_{z,b}$ are equivalent to λ_z . This shows that $\lambda_z \circ A$ is equivalent to λ_z , and hence it has an a.e.-strictly positive density k_A . Then, by the definition of the density (Theorem C.1), we have for a Borel set B ,

$$P_{z,a}(B) = P_{z,b}(A(B)) \quad (46)$$

$$\iff \int_B p_a(x)\lambda_z(dx) = \int_{A(B)} p_b(x)\lambda_z(dx) = \int_B p_b(A(x))\lambda_z(A(dx)), \quad (47)$$

where the last equality is by the standard change of variables formula, noting that $B = A^{-1}(A(B))$ since A is invertible. Now, we have that

$$\int_B p_a(x)\lambda_z(dx) = \int_B p_b(A(x))k_A(x)\lambda_z(dx), \quad (48)$$

by invoking the definition of the density again. Since the above holds for any B , we have

$$p_b(A(x))k_A(x) = p_a(x) \quad \text{a.e.}, \quad (49)$$

where $k_A(x)$ is strictly positive a.e.. \square

Corollary D.5. Suppose four probability measures $P_{1,a}, P_{2,a}, P_{1,b}, P_{2,b}$ have strictly positive densities $p_{1,a}, p_{2,a}, p_{1,b}, p_{2,b}$. For A both a $(P_{1,a}, P_{1,b})$ -measure isomorphism and a $(P_{2,a}, P_{2,b})$ -measure isomorphism, we have

$$\frac{p_{1,a}}{p_{2,a}}(x) = \frac{p_{1,b}}{p_{2,b}}(A(x)) \quad \text{a.e.} \quad (50)$$

Proof. This follows immediately from Lemma D.4 from the fact that k_A is strictly positive a.e. and depends only on A . \square

We are now ready to prove Proposition 4.3.

Proof of Proposition 4.3. Let $A \in \text{Aut}(\mathbf{Z})$ be $(\mathcal{E}_m(\eta_a(u_i), T_a), \mathcal{E}_m(\eta_b(u_i), T_b))$ -measure isomorphisms for all $i = 0, 1, \dots, K$. Suppose that both $\{\eta_a(u_i)\}_{i=0}^K$ and $\{\eta_b(u_i)\}_{i=0}^K$ are linearly independent. Fix j arbitrarily and note that $\{\eta_a(u_i) - \eta_a(u_j)\}_{i \neq j}$ and $\{\eta_b(u_i) - \eta_b(u_j)\}_{i \neq j}$ both still span \mathbb{R}^K . From Corollary D.5 and by taking logarithms, we have for each $i \neq j$,

$$\eta_a(u_i)^\top T_a(z) - a(\eta_a(u_i)) - (\eta_a(u_j)^\top T_a(z) - a(\eta_a(u_j))) \quad (51)$$

$$= \eta_b(u_i)^\top T_b(A(z)) - a(\eta_b(u_i)) - (\eta_b(u_j)^\top T_b(A(z)) - a(\eta_b(u_j))), \quad (52)$$

almost everywhere, which simplifies to

$$(\eta_a(u_i) - \eta_a(u_j))^\top T_a(z) - c_a(u_i) = (\eta_b(u_i) - \eta_b(u_j))^\top T_b(A(z)) - c_b(u_i), \quad (53)$$

almost everywhere. c_a, c_b are differences in the normalizing constants $a(\eta_a)$, and do not depend on z —we suppress the dependency on u for convenience. Written in matrix form, we have

$$\begin{bmatrix} \eta_a(u_0) - \eta_a(u_j) \\ \vdots \\ \eta_a(u_K) - \eta_a(u_j) \end{bmatrix}^\top T_a(z) = \begin{bmatrix} \eta_b(u_0) - \eta_b(u_j) \\ \vdots \\ \eta_b(u_K) - \eta_b(u_j) \end{bmatrix}^\top T_b(A(z)) + \mathbf{c}, \quad (54)$$

almost everywhere, where \mathbf{c} is the vector of differences $c_a - c_b$. Following (Khemakhem et al., 2020a), we will call these two matrices L_a and L_b , noting that they are invertible since their rows are linearly independent by assumption. Then, we obtain

$$L_a^\top T_a(z) = L_b^\top T_b(A(z)) + \mathbf{c} \quad (55)$$

$$\implies T_b(A(z)) = (L_b^{-1} L_a)^\top T_a(z) - (L_b^{-1} L_a)^\top \mathbf{c} \quad (56)$$

$$\implies T_b(A(z)) = L^\top T_a(z) + \mathbf{d}, \quad (57)$$

almost everywhere, where $L = L_b^{-1} L_a$ is invertible and $\mathbf{d} = -L^\top \mathbf{c}$. \square

For completeness, we will state the iVAE identifiability result within these terms.

Proposition D.6. *Suppose a generative model is described by (7) with latent distributions described by (11), and that m is strictly positive. Suppose we observe at least $K + 1$ distinct values of u_i such that the corresponding natural parameters $\{\eta(u_i)\}_{i=0}^K$ are linearly independent. Then, any indeterminacy transformation satisfies*

$$T_b(A_{a,b}(x)) = L^\top T_a(x) + \mathbf{d}, \quad (58)$$

almost everywhere, where L is an invertible $K \times K$ matrix and \mathbf{d} is a K -dimensional vector.

Proof. By Corollary 4.1 and since we do not constrain \mathcal{F} , the generator is identifiable up to the transformations described in Proposition 4.3. \square

D.7 Proof of Proposition 4.4

Here, we present proofs for the automorphism version of Proposition 4.3. These results are simply special cases of the theory developed above by setting $P = P_{z,a} = P_{z,b}$ and where $(P_{z,a}, P_{z,b})$ -measure isomorphisms are replaced with P -measure automorphisms. We begin with an intermediate result.

Lemma D.7. *Let $\mathcal{E}_{m,T}$ be as in (14), with m strictly positive. Suppose A is simultaneously a $\mathcal{E}_{m,T}(\eta_1)$ and $\mathcal{E}_{m,T}(\eta_2)$ -measure preserving automorphism. Then, we have*

$$(\eta_1 - \eta_2)^\top T(z) = (\eta_1 - \eta_2)^\top T(A(z)) \quad a.e.. \quad (59)$$

Proof. Denote the densities for $\mathcal{E}_{m,T}(\eta_1)$ and $\mathcal{E}_{m,T}(\eta_2)$ as p_1, p_2 . The expression is a direct consequence of Corollary D.5 by plugging in the exponential family densities $p_{1,a} = p_{1,b} = p_1$ and likewise for p_2 . Taking logarithms on both sides, we have

$$\eta_1^\top T(z) - \eta_2^\top T(z) - a(\eta_1) + a(\eta_2) = \eta_1^\top T(A(z)) - \eta_2^\top T(A(z)) - a(\eta_1) + a(\eta_2) \quad (60)$$

$$\implies (\eta_1 - \eta_2)^\top T(z) = (\eta_1 - \eta_2)^\top T(A(z)) \quad a.e.. \quad (61)$$

\square

This result implies Proposition 4.4.

Proof of Proposition 4.4. Lemma D.7 applies to the K contrast vectors $(\eta_i - \eta_0)$, so we have for each $i = 1, \dots, K$,

$$(\eta_i - \eta_0)^\top (T(z) - T(A(z))) = 0 \quad a.e.. \quad (62)$$

Any vector $v \in \text{span}\{(\eta_i - \eta_0)\}_{i=1}^K$ is of the form $v = \sum_{i=1}^K a_i (\eta_i - \eta_0)$. Clearly,

$$v^\top (T(z) - T(A(z))) = 0 \quad a.e.. \quad (63)$$

Since $\text{span}\{(\eta_i - \eta_0)\}_{i=1}^K$ is the row space of M , we have $(T(z) - T(A(z))) \in \ker M$, almost everywhere. \square

D.8 Proof of Theorem 4.5

A more useful adaptation of Proposition 4.4 for strong identifiability is the following result.

Proposition D.8. *Let $\mathcal{E}_{m,T}$ be a fixed exponential family, with dimension K , such that m is strictly positive and T is injective. Suppose that $\eta_i \in \mathbb{R}^K$, $i = 0, 1, \dots, K$ span \mathbb{R}^K . Suppose $A : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a simultaneously a $\mathcal{E}_{m,T}(\eta_i)$ -measure automorphism for each η_i . Then, $A(z) = z$, almost everywhere.*

Proof. Without loss of generality, assume that η_0 is such that $\{\eta_i - \eta_0\}$ forms a basis of \mathbb{R}^K . By Proposition 4.4, $(T(z) - T(A(z))) \in \ker M = \{0\}$, since M is full rank. This shows that $T(A(z)) = T(z)$ almost everywhere. If T is injective, then we have

$$A(z) = z \quad a.e.. \quad (64)$$

□

Strong identifiability of our proposed method then follows.

Proof of Theorem 4.5. In this model, we have $\mathcal{P}_z^e = \{\mathcal{E}_{m,T}(\eta_e)\}$. By Corollary 4.1, the generator is identifiable up to $\mathcal{A}(\mathcal{F}) \cap (\cap_e \mathcal{A}(\{\mathcal{E}_{m,T}(\eta_e)\}))$. By Proposition D.8, $\cap_e \mathcal{A}(\{\mathcal{E}_{m,T}(\eta_e)\})$ contains only functions that are equal to the identity almost everywhere, and hence the model is strongly identifiable. □

D.9 Geometric characterization of iVAE indeterminacies

We can further characterize the indeterminacy set geometrically in terms of the orthogonal complements of the subspace spanned by the natural parameter vectors; this is the Gaussian specialization of Proposition 4.4. See Fig. 3 for a visualization.

Proposition D.9. *In the multiple environments model described in Theorem 4.5 (with $\mathbf{Z} = \mathbb{R}^d$), fix a base environment with distribution $\mathcal{N}(0, \Sigma)$ and a subset E^* of environments where $|E^*| = d' \leq d$, with distributions $\mathcal{N}(\mu_e, \Sigma)$. Suppose $\{\mu_e\}_{e \in E^*}$ are linearly independent, and $(\mu_e)_i = 0$ for each e and $i \notin d^*$ for some collection of dimensions d^* . Then, for any indeterminacy transformation $A_{a,b}$ it holds that*

$$(A_{a,b}(x))_{i \in d^*} = (x)_{i \in d^*} \quad a.e. \quad (65)$$

Proof. Similarly to the previous identifiability proofs, we appeal to Corollary 4.1 and analyze the shared automorphisms. For Gaussian distributions with a fixed covariance matrix varying by its mean, we have $T(x) = x$, and $\eta_i = \mu_e$. Using the base environment we have $\mu_0 = 0$. Now, arrange $\{\mu_i\}_{i=1}^K$ into the rows of a matrix M . By Proposition 4.4, we have $(z - A(z)) \in \ker M$.

M has columns of 0 corresponding to d^* and linearly independent rows. Together, standard Gaussian elimination reveals




Figure 3: The Gaussian specialization (Proposition D.9) of Proposition 4.4. The indeterminacy set (in red) concentrates entirely on the z -axis, with means $\mu_1 = (1, 0, 0)$, $\mu_2 = (0, 1, 0)$. The kernel of a plane in \mathbb{R}^3 is the perpendicular line through the origin.

that the reduced row echelon form of M has the following form:

$$\begin{bmatrix} j \in d^* & j \notin d^* & j \in d^* & \cdots \\ 1 & 0 & 0 & \cdots \\ 0 & 0 & 1 & \cdots \\ \vdots & \vdots & \vdots & \cdots \\ 0 & 0 & 0 & \cdots \end{bmatrix} \begin{matrix} \mu_1 \\ \mu_2 \\ \vdots \\ \mu_d \end{matrix} \quad (66)$$

The corresponding kernel is the space of vectors with 0 entries for the indices d^* . Hence, $(z - A(z)) \in \ker M$ implies $(A(z))_{i \in d^*} = (z)_{i \in d^*}$. □

D.10 Identifiability of Equivariant Stochastic Mechanisms

In this section, we repeat the exercise of iVAE for the equivariant stochastic mechanisms model (Ahuja et al., 2022a). In (Ahuja et al., 2022a), weak identifiability of a temporal generative model is established. Adapted to our notation (note that the time indices represent environments), the model can be described:

$$X_t = f(Z_t), \quad Z_{t+1} = m_t(Z_t, U_t), \quad Z_t \perp\!\!\!\perp U_t, \quad t = 1, 2, \dots, \quad (67)$$

where $m_t \in \mathbf{M} : \mathbf{Z} \times [0, 1] \rightarrow \mathbf{Z}$ are unknown mechanisms and U_t are auxiliary noise variables. Note that this is in fact a noiseless generative model at the level of the observations, but where the underlying latent variable evolves according to a deterministic, but unknown mechanism m_t and random noise U_t . To be clear, this means that \mathcal{F} is fully flexible, while \mathcal{P}_z^t is parametrized by an initial condition P_1 , the distributions for U_t , and the mechanisms m_t . In what follows, we will assume a fixed P_1 and $U_t \sim U[0, 1]$ as in (Ahuja et al., 2022a), and leave \mathcal{P}_z^t parametrized purely by the mechanisms m_t .

Denote the marginal distribution of Z_t as P_t . In (Ahuja et al., 2022a), identifiability of the generator f is established up to pre-composition of some transformation A such that $A \circ m_a(z, U) \stackrel{d}{=} m_b(A(z), U)$ for $U \sim U[0, 1]$, for all possible values of z , and $m_a, m_b \in \mathbf{M}$. Using our framework, we are able to show the following stronger identifiability result, using only observations from two time points $t = 1, 2$.

Proposition D.10. *The model described by (67) is identifiable up to $A_{a,b} \in \mathcal{A}(\mathcal{M})$ satisfying*

$$A_{a,b}(m_a(Z, U)) \stackrel{d}{=} m_b(A_{a,b}(Z), U), \quad (68)$$

for any $m_a, m_b \in \mathbf{M}$, $U \sim U[0, 1]$ and any random variable Z independent of U .

Proof. We can analyze this model in our framework using just two time-points, $t = 1, 2$. We work on an augmented latent space $\tilde{\mathbf{Z}} = \mathbf{Z} \times [0, 1]$ and treat the random variables U_t as additional latent variables (i.e., as a “noiseless” case under our framework). For a generator $f : \mathbf{Z} \rightarrow \mathbf{X}$, we extend $\tilde{f} : \tilde{\mathbf{Z}} \rightarrow \mathbf{X} \times [0, 1]$, $\tilde{f}(z, u) = (f(z), u)$. The identity extension ensures that \tilde{f} is still injective, and is unique to f . Now suppose f_a and f_b are such that the distribution of X_1 and $X_2|X_1$ match. Note the marginal and conditional uniquely determine the joint, and hence we simply assume that the joint and hence marginal distributions of X_1 and X_2 match.

Let the joint distribution of Z_1 and U be denoted $\pi_{Z_1, U}$. Since they are independent, we have that $\pi_{Z_1, U} = P_1 \otimes U[0, 1]$.⁵ We also extend the mechanism m as $\tilde{m}(z, u) = (m(z, u), u)$, implying that $\tilde{m}^{-1}(B_z \times B_u) = m^{-1}(B_z) \times B_u$. Since $Z_2 = m(Z_1, U_1)$, this then implies that $P_2 = \pi_{Z_1, U} \circ \tilde{m}^{-1} = (P_1 \circ m^{-1}) \otimes (U[0, 1])$ (note the standard m in the right-hand-side). The same applies to an extended indeterminacy transformation, i.e., $P_2 \circ \tilde{A}_{a,b}^{-1} = (P_1 \circ A_{a,b}^{-1}) \otimes (U[0, 1])$.

We now apply Lemma 2.1 to $t = 1$, where Z_1 has fixed distribution P_1 (i.e., it is a singleton), and to $t = 2$, where the latent distribution may vary with the mechanism m_a or m_b , denoted $P_{2,a}, P_{2,b}$. As a result, we obtain

$$P_1 = A_{a,b\#}P_1, \quad P_{2,b} = \tilde{A}_{a,b\#}P_{2,a}. \quad (69)$$

Applying these identities simultaneously to $P_{2,b}$ gives

$$P_{2,b} = (m_{b\#}P_1) \otimes (U[0, 1]) = (m_{b\#}A_{a,b\#}P_1) \otimes (U[0, 1]) \quad (70)$$

$$P_{2,b} = \tilde{A}_{a,b\#}P_{2,a} = (A_{a,b\#}m_{a\#}P_1) \otimes (U[0, 1]), \quad (71)$$

⁵This means that for a Borel product $B_z \times B_u$, where B_z, B_u are Borel sets in their respective domains, we have $\pi_{X_1, U}(B_z \times B_u) = (P_1(B_z))(U[0, 1](B_u))$.

which by the properties of a product measure, means that

$$m_{b\#}A_{a,b\#}P_1 = A_{a,b\#}m_{a\#}P_1. \quad (72)$$

Writing the above in terms of their random variables, we have $m_b(A_{a,b}(Z), U) \stackrel{d}{=} A_{a,b}(m_a(Z, U))$ for Z with any fixed distribution P_1 independent of U . \square

Compared to the original proof, we are able to strengthen the result while weakening the assumptions due to our measure-theoretic framework as follows:

- Letting P_1 be any point mass recovers the original identifiability result in (Ahuja et al., 2022a).
- Our proof structure, which can be found in the Appendix, follows the intuition originally laid out in (Ahuja et al., 2022a), but we do not assume a diffeomorphic generator.

E TRIANGULAR TRANSPORT MAPS

This section describes the triangular maps in Section 5 in more detail, and proves the results therein. First, we define $f : \mathbb{R}^d \rightarrow \mathbb{R}^d$ to be a monotone increasing triangular map. This means that:

$$f(x) = \begin{bmatrix} f_1(x_1) \\ f_2(x_1, x_2) \\ \vdots \\ f_d(x_1, \dots, x_d) \end{bmatrix},$$

where each $x_d \rightarrow f_d(x_{1:d-1}, x_d)$ is monotone increasing (hence invertible) for any $x_{1:d-1}$. The inverse of f is as follows:

$$f^{-1}(x) = \begin{bmatrix} f_1^{-1}(x_1) \\ f_2^{-1}(f_1^{-1}(x_1), x_2) \\ \vdots \\ f_d^{-1}(f_1^{-1}(x_1), f_2^{-1}(f_1^{-1}(x_1), x_2), \dots, x_d) \end{bmatrix}.$$

This is also a monotone increasing triangular map—the inverses of monotone increasing maps are also monotone increasing. Note the map described above is lower-triangular—upper-triangular maps are analogously defined. For the purposes of this section, a triangular map refers to a lower-triangular map. As long as all maps considered are either lower, or upper triangular, the same closure properties apply.

E.1 Proofs of Theorem 5.1 and Proposition 5.2, and Knöthe–Rosenblatt transports

It is well-known that if $\mu = \nu \circ f^{-1}$, where μ, ν have strictly positive density and f is a monotone increasing triangular map, then f is equivalent to the Knöthe–Rosenblatt (KR) transport almost everywhere (see Jaini et al. 2019, Theorem 1, for example). The KR transport is described recursively as follows. Let $F_\mu(x_m | x_{1:m-1})$ be the conditional CDF of the m -th component of μ on the preceding components. Because μ has strictly positive density, F_μ is monotone increasing. Then, the m -th component of the KR transport is as follows:

$$K_m(x_{1:m-1}, x_m) = F_\nu^{-1}\{F_\mu(x_m | x_{1:m-1}) | K_1(x_1), \dots, K_{m-1}(x_{1:m-1})\}.$$

That is, K_m sends x_m through the conditional CDF of μ on $x_{1:m-1}$, and back through the inverse conditional CDF of ν on $y_{1:m-1} = (K_1(x_1), \dots, K_{m-1}(x_{1:m-1}))$. This CDF transform is the unique (almost everywhere) monotone increasing transport map between the 1-dimensional unique (almost everywhere) regular conditional probabilities.

It is clear that the map K defined by its components K_m is monotone increasing triangular. Since it is the unique such map transporting μ to ν , and triangular monotone increasing maps are closed under inverses and compositions, it must be that:

- For K the KR transport from μ to ν , K^{-1} is the KR transport from ν to μ .

- For measures μ, ν, π , if K_1 is the KR transport from μ to ν , K_2 is the KR transport from ν to π , then $K_1 \circ K_2$ is the KR transport from μ to π .

Now, it is clear that if f_a, f_b are KR transports, their generator transformation $\vec{A}_{a,b}$, and hence all indeterminacy transformations, are also KR transports. This drives our results from the main paper—the proofs of our results are trivial given the observations above.

Proof of Theorem 5.1. $\mathcal{A}(\mathcal{F})$ is the set of functions equal almost everywhere to KR maps transporting between measures in \mathcal{P}_z . Since $\mathcal{P}_z = \{P_z\}$, $\mathcal{A}(\mathcal{F})$ is the set of functions equal almost everywhere to the KR transport from P_z to itself—the identity map. By Theorem 2.2, $\mathcal{A}(\mathcal{M}) = \tilde{\text{id}}_z$. \square

Proof of Proposition 5.2. $\mathcal{A}(\mathbf{X})$ is the set of functions equal almost everywhere to KR transports between measures in \mathcal{P}_z . Let $P_{z,a}$ and $P_{z,b}$ be two such measures, which by assumption have independent components. By its construction, it is clear that K_m depends only on x_m in any KR transport K between $P_{z,a}$ and $P_{z,b}$. Such a map is monotone increasing and diagonal—hence by Theorem 2.2, $\mathcal{A}(\mathcal{M})$ consists of invertible, component-wise transformation. \square

F LINEAR EXAMPLES

We conclude the appendix by analyzing the identifiability of the linear examples of Section 1.1 in detail. The factor analysis example provides intuition on why multiple environments reduces the indeterminacy set, and the ICA example provides intuition on how the two intersecting sets in Theorem 2.2 can be manipulated in synergy.

F.1 Example: Linear-Gaussian

We present a simple example of using multiple environments, and a basis of priors, to obtain identifiability, using only linear algebra concepts. This example also provides intuition for the minimality of the number of environments. That is, for this 2-d latent space, three environments is enough to obtain strong identifiability, while two environments is insufficient.

Suppose two competing linear generative models for a random vector $x \in \mathbb{R}^{10}$ with latent vector $z \in \mathbb{R}^2$, for data arising from three environments indexed by $e = 1, 2, 3$:

$$\begin{aligned} z^{(e)} &\sim N(\mu, I_{2 \times 2}) & z^{(e)} &\sim N(\mu_e, I_{2 \times 2}) \\ \epsilon &\sim N(\mu, I_{10 \times 10}) & \epsilon &\sim N(\mu, I_{10 \times 10}) \\ y^{(e)} &= \alpha + Fz^{(e)} + \epsilon & x^{(e)} &= \alpha + Fz^{(e)} + \epsilon. \end{aligned} \quad (73) \qquad (74)$$

The left model is a single environment model, while in the right model, two of the μ_e are linearly independent, i.e., a multiple environment model. Note that the generator function here is

$$g(z) = \alpha + Fz, \quad (75)$$

where F is a full rank 10×2 matrix, and α is an offset vector in data space, fixed for all environments. For each environment we have the marginal distribution under the multiple environment model:

$$x^{(e)} \sim N(\alpha + F\mu_e, FF^\top + I_{10 \times 10}) \quad (76)$$

Recall the Gaussian distribution is characterized entirely by its mean and covariance—that is, for marginal distributions parametrized by $\theta_1 = (\alpha_1, F_1), \theta_2 = (\alpha_2, F_2)$:

$$P_{\theta_1, e} = P_{\theta_2, e} \iff \alpha_1 + F_1\mu_e = \alpha_2 + F_2\mu_e, \quad F_1F_1^\top = F_2F_2^\top. \quad (77)$$

To say that this model is strongly identifiable means that the right-hand-side equalities for each e imply $\alpha_1 = \alpha_2$ and $F_1 = F_2$.

In the single environments model, there are the following constraints:

$$\alpha_1 + F_1\mu = \alpha_2 + F_2\mu \quad (78)$$

$$F_1F_1^\top = F_2F_2^\top \quad (79)$$

The single environments model is not identifiable. For example, let R be an orthogonal (rotation) matrix, then, let $F_2 = F_1R$ and $\alpha_2 = \alpha_1 - F_1R\mu + F_1\mu$. We have

$$\alpha_2 + F_2\mu = \alpha_1 - F_1R\mu + F_1\mu + F_1R\mu = \alpha_1 + F_1\mu \quad (80)$$

$$F_2F_2^\top = F_1RR^\top F_1^\top = F_1F_1^\top, \quad (81)$$

where the last equality is due to R being an orthogonal matrix. This is a classical case of exploiting the rotational invariance of the Gaussian to construct a non-identifiable example.

Now, we analyze the multiple environment model. To be explicit, the three environments impose the following constraints in the multiple environments model:

$$\alpha_1 + F_1\mu_1 = \alpha_2 + F_2\mu_1 \quad (82)$$

$$\alpha_1 + F_1\mu_2 = \alpha_2 + F_2\mu_2 \quad (83)$$

$$\alpha_1 + F_1\mu_3 = \alpha_2 + F_2\mu_3 \quad (84)$$

$$F_1F_1^\top = F_2F_2^\top, \quad (85)$$

We can show directly that these constraints imply that $\alpha_1 + F_1z = \alpha_2 + F_2z$. First, assume that μ_1 and μ_2 are the linearly independent pair. Then, taking differences,

$$F_1(\mu_1 - \mu_3) = F_2(\mu_1 - \mu_3) \quad (86)$$

$$F_1(\mu_2 - \mu_3) = F_2(\mu_2 - \mu_3) \quad (87)$$

$$F_1F_1^\top = F_2F_2^\top. \quad (88)$$

Written in matrix form, the first two constraints read

$$F_1M = F_2M \implies F_1 = F_2, \quad (89)$$

since $\mu_1 - \mu_3$ and $\mu_2 - \mu_3$ remain linearly independent, and hence M is invertible. It immediately follows from the original constraints that $\alpha_1 = \alpha_2$ also.

The above analysis showed that, for identifiability, a single environment was insufficient, while three environments was adequate. This begs the question, what about two environments? In other words, is the three environment constraint minimal?

Consider a model with two environments with means μ_1, μ_2 . By the arguments above, it imposes the following constraints:

$$\alpha_1 + F_1\mu_1 = \alpha_2 + F_2\mu_1 \quad (90)$$

$$\alpha_1 + F_1\mu_2 = \alpha_2 + F_2\mu_2 \quad (91)$$

$$F_1F_1^\top = F_2F_2^\top. \quad (92)$$

Can we construct a non-identifiable example? Let $F_2 = F_1R$, $\alpha_2 = \alpha_1 - F_1R\mu_1 + F_1\mu_1$ as in the single-environment case. Clearly, these satisfy the first and third constraint for any orthogonal matrix R . We aim to find a specific rotation matrix that also satisfies the second constraint. Observe that:

$$\alpha_2 + F_2\mu_2 = \alpha_1 - F_1R\mu_1 + F_1\mu_1 + F_1R\mu_2 \quad (93)$$

$$= \alpha_1 + F_1\mu_1 + F_1R(\mu_2 - \mu_1). \quad (94)$$

Let x be a vector orthogonal to $\mu_2 - \mu_1$, standardized such that $\|x\|^2 = \|\mu_2 - \mu_1\|^2$. Consider

$$R = \frac{1}{\|\mu_2 - \mu_1\|^2} \begin{bmatrix} | & | & | \\ \mu_2 - \mu_1 & x & | \\ | & | & | \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \begin{bmatrix} | & | & | \\ \mu_2 - \mu_1 & x & | \\ | & | & | \end{bmatrix}^\top. \quad (95)$$

This is the eigendecomposition of an orthogonal matrix (it is the product of orthogonal matrices) with eigenvalues 1 and -1 , and corresponding eigenvectors $\mu_2 - \mu_1$ and x . Since it is an eigenvector, we have $R(\mu_2 - \mu_1) = \mu_2 - \mu_1$.⁶ Then, we have

$$\alpha_2 + F_2\mu_2 = \alpha_1 + F_1\mu_1 + F_1\mu_2 - F_1\mu_1 = \alpha_1 + F_1\mu_2, \quad (96)$$

⁶R is essentially a rotation matrix with axis $(\mu_2 - \mu_1)$.

which satisfies the second constraint as desired. This shows that three environments are required, and hence minimal for strong identifiability of this model.

Note that such a construction will not work for the three-environment model. For three environments, the rotation has to satisfy both

$$R(\mu_3 - \mu_1) = \mu_3 - \mu_1 \quad (97)$$

$$R(\mu_2 - \mu_1) = \mu_2 - \mu_1, \quad (98)$$

that is, the eigenspace of R associated to the eigenvalue 1 spans \mathbb{R}^2 , i.e., it is the identity.

F.2 Linear, non-Gaussian ICA

Consider a generative model (Equation (1)) with $\mathbf{Z} = \mathbb{R}^{d_z}$ and $\mathbf{X} = \mathbb{R}^{d_x}$. Assume $d_x \geq d_z$. Let the generator parameter space be $\mathcal{F} = \{A \in \mathbb{R}^{d_x \times d_z}; \text{rank}(A) = d_z\}$. That is, the generators are full-rank linear transformations, and hence injective. Let the prior parameter space be

$$\mathcal{P}_z = \{p(z) = \prod_{i=1}^{d_z} p_i(z); p_i \text{ are non-Gaussian, and not a point mass}\}, \quad (99)$$

i.e., probability distributions on \mathbb{R}^{d_z} with a density, and the density factorizes as independent, non-Gaussian components.

The identifiability of this problem was first studied in (Comon, 1994).⁷ In their analysis, identifiability is established up to pre-multiplication of a diagonal matrix and a permutation. That is, for generators $F_a, F_b \in \mathcal{F}$ with $P_{z,a}, P_{z,b} \in \mathcal{P}_z$, if the marginal distributions on \mathbf{X} match, then $F_a = F_b \Lambda P$, where Λ is an invertible diagonal matrix and P is a permutation matrix.

Under our framework, i.e., Lemma 2.1, we must have that $P_{z,b} = P_{z,a} \circ \vec{A}_{a,b}^{-1}$, where $F_a = F_b \circ \vec{A}_{a,b}$. Using our framework, we now show that $\vec{A}_{a,b} = \Lambda P$ as above. The identifiability result obtained in (Comon, 1994) rests on the following result (restated and re-proved to match our notation):

Theorem F.1 (Theorem 10, (Comon, 1994)). *Let z be a random vector with factorized density. Let $x = Cz$, such that x also has factorized density. Then, z_j is non-Gaussian if the j -th column has at most one non-zero entry.*

Proof. We require Theorem 19 from (Comon, 1994).

Lemma F.2 ((Comon, 1994), Darmois' Theorem). *Define two random variables Z_1 and Z_2 as*

$$Z_1 = \sum_i a_i z_i \quad Z_2 = \sum_i b_i z_i, \quad (100)$$

where z_i are independent random variables, i.e., their joint distribution factorizes. Then, if Z_1 and Z_2 are independent, all variables z_j for which $a_j b_j \neq 0$ are Gaussian.

Now, let z be a random vector with factorized density and $x = Cz$, where x has factorized density also. Note that this implies any x_i, x_k are independent for $i \neq k$. We have that

$$x_i = \sum_j C_{ij} z_j \quad x_k = \sum_j C_{kj} z_j, \quad (101)$$

and hence by Lemma F.2, if z_j is non-Gaussian, it must be that $C_{ij} C_{kj} = 0$. This holds for each $i \neq k$, and hence, the j -th column has at most one non-zero entry. \square

Recall the definition of \mathcal{P}_z is such that any prior must factorize and be non-Gaussian. Then, Theorem F.1 implies that

$$\mathcal{A}(\mathcal{P}_z) \cap \mathbb{R}^{d_z \times d_z} = \{A \in \mathbb{R}^{d_z \times d_z} \mid A \text{ has no column with more than one nonzero element}\}. \quad (102)$$

⁷In the original analysis, the model is fit according to a criteria maximizing the independence between components, and also one component of the prior is allowed to be Gaussian. For simplicity, we will simply study the implications of matching observational marginal distributions (i.e., maximum likelihood) and where all prior components are non-Gaussian.

That is, any linear isomorphisms between two priors must have no column with more than one nonzero element. Now, note that for any $\vec{A}_{a,b} \in \mathcal{A}(\mathcal{F})$, we have

$$\vec{A}_{a,b} = f_b^{-1} \circ F_a, \tag{103}$$

where f_b^{-1} is the restriction of the linear map represented by the pseudoinverse F_b^\dagger to the range of \mathcal{F} . By Lemma D.1, $\vec{A}_{a,b}$ is an invertible linear map and hence full rank. Finally, we conclude that for any $\vec{A}_{a,b} \in \mathcal{A}(\mathcal{F}) \cap \mathcal{A}(\mathcal{P})$, $\vec{A}_{a,b}$ must have exactly one nonzero element in each column. We can then apply a permutation P such that $P\vec{A}_{a,b} = \Lambda$, where Λ is diagonal. Finally, we obtain $\vec{A}_{a,b} = P^\top \Lambda$, where P^\top is a permutation matrix and Λ is diagonal and invertible.