Categorical Independence in Nonmonotonic and Probabilistic Reasoning

Roy Ferguson[1,*], Jonas Haldimann[1], Jesse Heynick[2] and Thomas Meyer[1]

[1]University of Cape Town and CAIR, 18 University Avenue, Rondebosch 7701, Cape Town, South Africa
[2]Open Universiteit, the Netherlands

Abstract

Independence in nonmonotonic and probabilistic reasoning, in various forms, has long been an active area of study. At the same time, significant progress has been made in consolidating the study of independence in category theory. In particular, the notion of an *independent product* characterises independence or separability in a remarkable variety of contexts. This paper takes some first steps toward applying this theory to nonmonotonic and probabilistic reasoning. We detail several categories $\mathbb{C}$ of preference structures and introduce the notion of a *preference functor* $P\colon \mathbf{Sig} \to \mathbb{C}$, which assigns signatures to preference structures on possible worlds. We define a notion of semantic independence with respect to a preference functor P in terms of independent products. For ordinal conditional functions, this corresponds to known notions of independence in nonmonotonic reasoning and belief revision. On the other hand, we show that the category of total preorders does not possess an independent product that corresponds to such established notions of independence. In doing so, we demonstrate how established definitions of independence for OCFs are in fact the same phenomenon as usual independence of random variables in probability theory, while also offering a reason why TPOs are not as well-behaved as OCFs in terms of independence. Our proposed framework allows for reasoning about independence in nonmonotonic and probabilistic reasoning at a new level of generality. Furthermore, this opens the door to apply theories of independence from diverse contexts to our setting.

Keywords

Independent products, Preference structures, Preference functors, Nonmonotonic reasoning

1. Introduction

Ideas of independence have proved a significant concern in much of the development of knowledge representation and reasoning (KRR). In the probabilistic setting, they go back to [1], which lists *system independence* as a desirable property satisfied by those information operators based on minimum cross-entropy. Parikh introduced the concept of splitting languages in [2] as well as axiom P for belief revision operators, which expresses the fact that given a splitting across sublanguages, revision by a formula in one sublanguage should only depend on the revision operator of that sublanguage. [3] expresses and refines this axiom in terms of epistemic state revision represented in terms of total pre-orders (TPOs), while [4] introduces a semantic notion for syntax splitting of epistemic states represented in terms of TPOs as well as ordinal conditional functions (OCFs). Finally, in the context of nonmonotonic reasoning (NMR) [5] introduces an analogous condition, **(SynSplit)**, for inductive inference operators as conjunction of two postulates **(Ind)** and **(Rel)**. It is the condition **(Ind)** and its semantic counterparts that will be of focus in this paper. Many of the underlying constructions on which these ideas are built depend on some kind of structure indicating preference (probability distributions, TPOs, OCFs) among worlds or propositions.

Category theory has proved to be a powerful language with which to reason about structures. By

24th International Workshop on Nonmonotonic Reasoning, July 17–19, 2026, Lisbon, Portugal
*Corresponding author.
roy.ferguson@uct.ac.za (R. Ferguson); jonas@haldimann.de (J. Haldimann); jesse.heyninck@ou.nl (J. Heynick); tmeyer@cair.org.za (T. Meyer)
https://jonas.haldimann.de/ (J. Haldimann); https://sites.google.com/view/jesseheyninck (J. Heynick); https://tommiemeyer.org.za/ (T. Meyer)
0000-0001-8822-6642 (R. Ferguson); 0000-0002-2618-8721 (J. Haldimann); 0000-0002-3825-4052 (J. Heynick); 0000-0003-2204-6969 (T. Meyer)

formulating properties in terms of the relationships between structures of a fixed type, this language allows us to abstract beyond the structures themselves, and build a unifying framework to reason at this new level of abstraction. These properties are often formulated in terms of some "meta-structure" that a given class of structures (called a category) possesses. The theory of independence has somewhat recently been developed at the level of categories, in terms of product-like meta-structures [6], called *independent products.* What is quite remarkable is the variety of contexts in which independent products arise: they capture independence of random variables in probability theory, information independence in database theory, and separability of name-binding in formal languages. It is thus a natural question to ask whether independent products are also able to model independence in NMR and more generally in KRR. As a first step to this end, we develop an appropriate language in which to reason over structures of preference on sets of possible worlds. We then define a notion of semantic independence in terms of independent products, which corresponds to existing notions of semantic independence in a number of settings. While the existence of such independent products on the category of probability spaces is known, we show that they also exist for the category of OCFs, while they do not for the category of TPOs.

Our main contributions in this paper are the following conceptual developments, which are reflected in the two theorems of Section 7.

- We demonstrate how the established notion of independence for OCFs described in [4] and [5] can be seen as the same phenomenon as usual probabilistic independence of random variables.
- We offer, at the same level of conceptualisation, a reason for the shortcomings of the established notion of independence for TPOs that is expressed in [4].

These are both made possible by our abstract definition of semantic independence for an arbitrary class of preference structures.

An outline of this paper is as follows. Section 2 sets out the basic concepts to speak about nonmonotonic inference relations based on preference structures, while Section 3 recalls known definitions of independence related to inductive inference operators. Section 4 recalls some prerequisite concepts from category theory. Section 5 recalls the definition of independent products, and details relevant examples. Section 6 introduces our framework of preference structures. In it, we define various categories of preference structures and explore the relationships between them. Section 7 introduces the notion of semantic independence relative to independent products on a category of preference structures.

2. Inference relations based on preference structures

There is a wide range of approaches to reason about statements of the form, "if ϕ then usually ψ will hold" which can be formally captured by *conditionals* $\phi \rightsquigarrow \psi$. [7] provides a comprehensive overview of a number of such approaches, and uses the framework of logical institutions [8] to establish links between them. Many of the ideas in this paper are inspired by this framework. The approaches to inductive reasoning we consider are concerned with propositional statements ϕ, ψ over some finite set $\Sigma = \{x_1, \dots, x_n\}$ of atomic boolean variables, referred to as a *signature*, and an *inference relation* over Σ, which is defined to be any relation $\mid\!\sim\, \subseteq \mathrm{Sen}(\Sigma) \times \mathrm{Sen}(\Sigma)$ on the set $\mathrm{Sen}(\Sigma)$ of all propositional formulas over Σ. This is in contrast to an *inductive inference operator* [5], which associates to each *conditional knowledge base* $\Delta = \{\phi_1 \rightsquigarrow \psi_1, \dots, \phi_s \rightsquigarrow \psi_s\}$ with an inference relation $\mid\!\sim_\Delta$ such that:

- (Direct Inference) $\phi_i \mid\!\sim_\Delta \psi_i$ for each defeasible conditional $\phi_i \rightsquigarrow \psi_i \in \Delta$.
- (Trivial vacuity) If $\Delta = \emptyset$ then $\phi \mid\!\sim_\Delta \psi$ if and only if $\phi \models \psi$.

We will adopt the notation where we suppress conjunction symbol and indicate negation with a bar over the formula. These inference relations are not necessarily monotonic, meaning that $\phi \mid\!\sim \psi$ being true does not necessarily imply that $\phi\xi \mid\!\sim \psi$ is true for a formula ξ. Nonmonotonic reasoning is the study of such inference relations. The proof calculus of a given $\mid\!\sim$ may be governed by a given list of syntactic postulates, such as the System P of [9], which is regarded as core to any nonmonotonic reasoning system. For classical propositional logic, the semantics is determined by the set $\Omega(\Sigma)$ of all complete conjunctions ω, which we call *worlds*, over the atoms $x_1, \ldots, x_n$. The semantics for the nonmonotonic inference relations we consider are based on equipping $\Omega(\Sigma)$ with some structure indicating preference or likelihood of worlds. This structure is then extended to include all formulas in $\mathrm{Sen}(\Sigma)$. We will consider three examples of such preference structures.

The first example of this kind of inference relation is based on assigning probabilities to worlds. As one would imagine, worlds with higher probabilities here are more preferred or likely. An inference relation is determined by a distribution p on $\Omega(\Sigma)$. Given such a p we set

$$p(\phi) = \sum_{\substack{\omega \in \Omega(\Sigma),\\ \omega \models \phi}} p(\omega)$$

where $\models$ has the classical propositional interpretation. Then we define $\phi \mathrel{|\!\sim}_p \psi$ to mean that $p(\phi\overline{\psi}) < p(\phi\psi)$ or $p(\phi) = 0$. We remark that $\mathrel{|\!\sim}_p$ is rather poorly behaved for many distributions p. Any $\mathrel{|\!\sim}_p$ does however satisfy the System O postulates of [10]. The System P postulates [9] are satisfied if p is a big-stepped probability distribution [11], where the probability of any world is larger the sum of probabilities of all less likely worlds.

A qualitative semantics we will consider is based on the notion of an *ordinal conditional function* (OCF) or *ranking function* [12], which is a function $\kappa\colon \Omega(\Sigma) \to \mathbb{N}_0 \cup \{\infty\}$ where $\kappa^{-1}(0) \neq \emptyset$. Here, worlds that are ranked lower are considered more typical. We extend κ to $\mathrm{Sen}(\Sigma)$ by setting

$$\kappa(\phi) = \min_{\substack{\omega \in \Omega(\Sigma),\\ \omega \models \phi}} \kappa(\omega)$$

and define $\phi \mathrel{|\!\sim}_\kappa \psi$ to mean that $\kappa(\phi\overline{\psi}) > \kappa(\phi\psi)$ or $\kappa(\phi) = \infty$. Inference relations based on ranking functions are generally better-behaved from the point of view of nonmonotonic reasoning: any $\mathrel{|\!\sim}_\kappa$ satisfies the System P postulates.

The final example we consider, which is another qualitative approach, is based on *total preorders* (TPOs). Here, $\Omega(\Sigma)$ is assigned some total preorder structure $\preceq$, and as with OCFs, worlds lower down in the order are considered more typical. $\preceq$ is extended to $\mathrm{Sen}(\Sigma)$ by setting

$$\phi \preceq \phi' \iff \exists_{\omega \in \Omega(\Sigma), \omega \models \phi} \forall_{\omega' \in \Omega(\Sigma), \omega' \models \phi'} \left[\omega \preceq \omega'\right] \tag{1}$$

and we define $\phi \mathrel{|\!\sim}_\preceq \psi$ to mean that $\phi\psi \preceq \phi\overline{\psi}$. The right-hand side of (1) can be replaced with

$$\min_{\prec}\{\omega \in \Omega(\Sigma) \mid \omega \models \phi\} \preceq \min_{\prec}\{\omega \in \Omega(\Sigma) \mid \omega \models \phi'\} \tag{2}$$

Note that since we are not working with partial orders, the expressions on either side of $\preceq$ in (2) are sets, and so is interpreted as everything in the left-hand set is below everything in the right-hand set.

3. Independence for nonmonotonic inference relations

A concept pertinent to independence of the inference relations we consider is that of *marginalisation*. Given a subset $\Theta \subseteq \Sigma$ of a signature, we are able to define a preference structure on $\Omega(\Theta)$ induced by $\Omega(\Sigma)$. For probability distributions, this corresponds to the usual notion of a marginal distribution. More generally, our extension of preference structure to $\mathrm{Sen}(\Sigma)$ is the recipe to follow for computing the marginal. In particular, given an ordinal conditional function κ on $\Omega(\Sigma)$ the *marginal of κ on Θ* is defined on $\omega^\Theta \in \Omega(\Theta)$ by $\kappa|_\Theta(\omega^\Theta) = \kappa(\omega^\Theta)$. Note that in the right-hand side of the equality, ω^Θ is sentence in $\mathrm{Sen}(\Sigma)$ that is not a world in $\Omega(\Sigma)$. Likewise, given a TPO $\preceq$ on $\Omega(\Sigma)$, the marginal of $\preceq$ on Θ is defined by $\omega_1^\Theta \preceq|_\Theta \omega_2^\Theta \iff \omega_1^\Theta \preceq \omega_2^\Theta$.

Inspired by the postulate **(Ind)** for inductive inference operators developed in [5], we propose the following notion of independence on inference relations.

Definition 1. Let $\mid\!\sim$ be an inference relation over a signature Σ. We will say that a partition $\Sigma = \Sigma_1 \sqcup \Sigma_2$ is *independent with respect to* $\mid\!\sim$ if for any $i, j \in \{1, 2\}$ with $i \neq j$ and any formulas $\phi, \psi \in \text{Sen}(\Sigma_i)$ and consistent $\xi \in \text{Sen}(\Sigma_j)$ we have

$$\phi \mid\!\sim \psi \iff \phi\xi \mid\!\sim \psi$$

So, independence here expresses a kind of cross-monotonicity, where inferences in Σ_i are monotonic w.r.t. additional premises from Σ_j. These definitions can be extended to m-ary partitions. Contrary to the postulate **(Ind)** [5], that expresses a property of inductive inference operators, independence in Definition 1 is concerned only with inference relations. The two are, however, closely related: **(Ind)** can be understood as the requirement that if a belief base Δ has a syntax splitting $\Sigma = \Sigma_1 \sqcup \Sigma_2$, then $\Sigma = \Sigma_1 \sqcup \Sigma_2$ must be independent with respect to the inference relation $\mid\!\sim_\Delta$ associated with Δ.

[5] also defines the semantic notion of independence with respect to a TPO $\preceq$: a partition $\Sigma = \Sigma_1 \sqcup \Sigma_2$ is called $\preceq$-*independent* if for any distinct $i, j \in \{1, 2\}$ and any worlds $\omega_i, \upsilon_i \in \Omega(\Sigma_i), \omega_j \in \Omega(\Sigma_j)$, we have that

$$\omega_i \preceq \upsilon_i \iff \omega_i\omega_j \preceq \upsilon_i\omega_j$$

An n-ary version of this definition appeared earlier in [4], but in the context of belief revision, where a TPO on the set of all possible worlds represents an epistemic state, as in the spirit of [13]. Both [4] and [5] state an analogue for ordinal conditional functions κ. In the binary case, κ-*independence* is defined as for each pair of worlds $\omega_1 \in \Omega(\Sigma_1)$ and $\omega_2 \in \Omega(\Sigma_2)$ we have

$$\kappa(\omega_1\omega_2) = \kappa_1(\omega_1) + \kappa_2(\omega_2)$$

where κ_1, κ_2 are the marginalisations of κ to Σ_1 and Σ_2. In [4] the n-ary case is referred to as an *ocf-splitting*.

4. Category theory preliminaries

Category theory is a language that allows us to reason about structures. For a more comprehensive introduction, please see [14] (for the computer scientist) and [15] (for the mathematician). The starting construction here is that of a *category* which we can think of as a collection of structures of the same type. Each such structure we refer to as an *object*, and usually denote these as $X, Y, Z, \ldots$. We usually denote an arbitrary category in blackboard font like $\mathbb{C}$. We typically study objects in the way that they relate to other objects in their category. A given such relationship f from an object X to an object Y we call a *morphism* and denote this as $f\colon X \to Y$ or as an arrow

$$X \xrightarrow{f} Y.$$

In the above display, we refer to X as the *domain* of the morphism f and Y as the *codomain* of f. The class of all objects of $\mathbb{C}$ we denote by $\text{Obj}(\mathbb{C})$ while the class of morphisms of $\mathbb{C}$ we denote by $\text{Mor}(\mathbb{C})$. To make this data a category, we endow the class of morphisms with a partially-defined operation of composition. We say two morphisms g, f are *composable* if the domain of g equals the codomain of f, and we denote their composite as $g \circ f$ or simply as gf. This operation of composition is required to satisfy the following laws.

- Every object X of $\mathbb{C}$ has as associated *identity morphism* $1_X\colon X \to X$ with the property that $f1_X = f$ and $1_Xe = e$ for any morphisms f, e composable with 1_X.
- Composition is associative, meaning that for each composable triple g, f, e we have that $g(fe) = (gf)e$.

We say a morphism $f\colon X \to X'$ is an *isomorphism* if it possesses an inverse $f^{-1}\colon X' \to X$ such that $f^{-1}f = 1_X$ and $ff^{-1} = 1_{X'}$. In this case we say that X and X' are *isomorphic* and denote this as $X \cong X'$.

We list some examples demonstrating these ideas. The prototypical example is the category of sets.

Example 2. The category of sets, denoted **Set**, has sets for objects and functions for morphisms. Here domain, codomain, and composition have their usual interpretations. Identity morphisms are identity functions $1_X : x \mapsto x$. The isomorphisms are precisely the bijections.

Most classical examples of categories have objects sets equipped with some structure and morphisms functions that preserve that structure. We include an example from categorical probability theory, which will serve as a starting point for constructing categories of preference structures.

Example 3. We can see a finite discrete probability distribution as a function $p\colon X \to [0,1]$ where X is a finite set and $\sum_{x\in X} p(x) = 1$. Given two distributions $(X,p),(Y,q)$ we say that a function $f\colon X \to Y$ preserves the distribution if for each $y \in Y$ we have that $q(y) = \sum_{x\in f^{-1}(y)} p(x)$. The category $\mathbf{FinProb}$ has objects finite discrete probability distributions and morphisms probability-preserving surjective maps. A morphism $f\colon (X,p) \to (Y,q)$ in $\mathbf{FinProb}$ can be seen as a random variable over the sample space X with events or attributes in Y. The only morphisms in $\mathbf{FinProb}$ that have underlying function identity are the identity morphisms. We illustrate the above ideas with some small sub-examples. We could take $X = \{\text{dove}, \text{tern}, \text{chicken}, \text{duck}, \text{ostrich}\}$ to be a set of types of birds we survey in a given area and equip it with probability distribution p where

$$\begin{array}{lllllllll} p\colon \text{dove} & \mapsto & 0.28 & p\colon \text{tern} & \mapsto & 0.12 & p\colon \text{chicken} & \mapsto & 0.4 \\ p\colon \text{duck} & \mapsto & 0.18 & p\colon \text{ostrich} & \mapsto & 0.02 & & & \end{array}$$

Then (X,p) is an object of $\mathbf{FinProb}$, which we can consider a sample space. Now let $Y = \{\text{small}, \text{not\_small}\} = \{s, \overline{s}\}$ be a set of events governing the size of a random type of bird. If we restrict ourselves to just the sample space (X,p), then the distribution over Y could be q given by

$$q\colon s \;\mapsto\; 0.4 \qquad q\colon \overline{s} \;\mapsto\; 0.6$$

The probability-preserving function $f\colon X \to Y$ given by

$$\begin{array}{lllllllll} f\colon \text{dove} & \mapsto & s & f\colon \text{tern} & \mapsto & s & f\colon \text{chicken} & \mapsto & \overline{s} \\ f\colon \text{duck} & \mapsto & \overline{s} & f\colon \text{ostrich} & \mapsto & \overline{s} & & & \end{array}$$

expresses precisely the data of the random variable `is_small` over the sample space (X,p).

A *functor* is a structure-preserving map between categories. That is, a functor $F : \mathbb{C} \to \mathbb{D}$ consists of a function $F_0 : \mathrm{Obj}(\mathbb{C}) \to \mathrm{Obj}(\mathbb{D})$ between classes of objects and a function $F_1 : \mathrm{Mor}(\mathbb{C}) \to \mathrm{Mor}(\mathbb{D})$ between classes of morphisms (we typically suppress the subscripts) satisfying

- F preserves domain and codomain: for any $f : X \to Y$ in $\mathbb{C}$ we have $F(f)\colon F(X) \to F(Y)$.
- F preserves identity: for every object X in $\mathbb{C}$ we have $F(1_X) = 1_{F(X)}$.
- F preserves composition: for any composable g, f in $\mathbb{C}$ we have $F(gf) = F(g)F(f)$.

The class of all categories (with appropriate size restriction applied) form the objects of the category $\mathbf{Cat}$, which has morphisms functors. Functors of interest are those that preserve some property or construction in $\mathbb{C}$ or reflect some property or construction in $\mathbb{D}$. Categories $\mathbb{C}$ with objects sets equipped with additional structure are those that possess a functor U to **Set** that forgets the additional structure. We call such a pair $(\mathbb{C}, U)$ a *concrete category.*

Example 4. The forgetful functor $U\colon \mathbf{FinProb} \to \mathbf{Set}$ maps each object (X,p) to the underlying set X and each morphism $f\colon (X,p) \to (Y,q)$ to the underlying function $f\colon X \to Y$, where we just forget the fact that it preserves the distribution p.

In category theory, we typically reason in terms of *diagrams.* A diagram is a collection of objects and arrows, usually denoted by a directed graph. We say that a diagram commutes if composing the morphisms in each directed path with the same starting and ending objects results in the same morphism. An example of a construction that is expressed diagrammatically is that of a *product* (it is a kind of meta-structure described in the introduction). We will be interested in generalisations of products, which, as we will see, are used to reason about independence. The simplest case is a binary product. We say a category $\mathbb{C}$ has *binary products* if for each pair of objects X_1, X_2 in $\mathbb{C}$ there exists a diagram

$$X_1 \xleftarrow[\pi_1]{} X_1 \times X_2 \xrightarrow[\pi_2]{} X_2$$

with the property that for any pair of morphisms f_1, f_2 as displayed below, there exists a unique morphism f making the diagram

$$\begin{tikzcd} & W \arrow[dl, "f_1"'] \arrow[d, dotted, "f"] \arrow[dr, "f_2"] & \\ X_1 & X_1 \times X_2 \arrow[l, "\pi_1"] \arrow[r, "\pi_2"'] & X_2 \end{tikzcd} \tag{3}$$

commute. The morphisms π_1, π_2 are referred to as the projections of the product $X_1 \times X_2$. A key feature of product projections is that they form a *jointly monic pair*, which expresses the fact that the f making Diagram (3) commute is unique when it exists for the pair (f_1, f_2). A nullary version of the product is called a *terminal object*, typically denoted 1, and has the property that every object X has a unique morphism to it. In **Set**, products are given by usual Cartesian product of sets. Similarly, products in $\mathbf{Vec}_{\mathbb{R}}$ are given by usual products of vector spaces. The product $\mathbb{C}_1 \times \mathbb{C}_2$ in $\mathbf{Cat}$ has objects all $(X_1, X_2) \in \mathrm{Obj}(\mathbb{C}_1) \times \mathrm{Obj}(\mathbb{C}_2)$ and morphisms all $(f_1, f_2) \in \mathrm{Mor}(\mathbb{C}_1) \times \mathrm{Mor}(\mathbb{C}_2)$ with identity and composition taken pointwise. $\mathbf{FinProb}$ is an example of a category that does not have all binary products.

Finally, we make note on duality, which features when we move from signatures to preference structures. For any construction based on given diagram, the *dual* or *opposite* construction is exactly the same but with the direction of all arrows reversed. Given any category $\mathbb{C}$, the opposite category $\mathbb{C}^{\mathrm{op}}$ has the same objects as $\mathbb{C}$, but a morphism $f\colon Y \to X$ is a morphism in $\mathbb{C}^{\mathrm{op}}$ if and only if $f\colon X \to Y$ is a morphism in $\mathbb{C}$. The dual construction to terminal object 1 is called an *initial object*, denoted 0.

5. Independent products

There are a number of approaches to capturing notions of independence and conditional independence categorically. One such is the independent product structure of [6], which can be formally defined as a semicartesian monoidal structure $(\mathbb{C}, \otimes, 1)$ with jointly monic projections (see [16] Sections 2.1, 3.1, and 3.2 where these definitions are recalled and explained). This can be seen as a binary operation $\otimes\colon (X_1, X_2) \mapsto X_1 \otimes X_2$ on the objects (and morphisms) of $\mathbb{C}$ which is associative in the sense that $(X_1 \otimes X_2) \otimes X_3 \cong X_1 \otimes (X_2 \otimes X_3)$ and which the terminal object as unit: $X \otimes 1 \cong X \cong 1 \otimes X$. Such a (meta-)structure implies the existence of projections

$$X_1 \xleftarrow[\pi_1]{} X_1 \otimes X_2 \xrightarrow[\pi_2]{} X_2 \tag{4}$$

and these are required to be jointly-monic. We call (4) an $\otimes$-diagram. From the binary operation $\otimes$ we can construct an n-ary operation for each $n \in \mathbb{N}$ by repeated application of $\otimes$. A functor $F\colon (\mathbb{C}, \otimes, e) \to (\mathbb{B}, \boxtimes, d)$ between independent product structures may preserve the structure in the sense that $F(X_1 \otimes X_2) = F(X_1) \boxtimes F(X_2)$ and $F(e) = d$. Such a functor is called *strict monoidal*.

Each independent product structure comes equipped with a binary relation $\perp\!\!\!\perp \subseteq \mathrm{Mor}(\mathbb{C}) \times \mathrm{Mor}(\mathbb{C})$ where $f_1 \perp\!\!\!\perp f_2$ whenever f_1 and f_2 have the same domain and there exists a dotted morphism making

$$\begin{tikzcd} & X \arrow[dl, "f_1"'] \arrow[d, dotted] \arrow[dr, "f_2"] & \\ Y_1 & Y_1 \otimes Y_2 \arrow[l, "\pi_1"] \arrow[r, "\pi_2"'] & Y_2 \end{tikzcd}$$

commute (which is unique, given the joint-monicity of π_1 and π_2). So, when $\perp\!\!\!\perp = \mathrm{Mor}(\mathbb{C}) \times \mathrm{Mor}(\mathbb{C})$, then our monoidal structure in question is just that of product. In general, $\perp\!\!\!\perp$ gives us an indication of how far off $\otimes$ is from being the actual product [1]. $\perp\!\!\!\perp$ extends to a family of relations of each arity $n \in \mathbb{N}$, where $\perp\!\!\!\perp(f_1, \ldots, f_n)$ if and only if all $f_i \colon X \to Y_i$ have the same domain and there exists a morphism $f \colon X \to Y_1 \otimes \ldots \otimes Y_n$ such that each $f_i = \pi_i f$. The extended $\perp\!\!\!\perp$ is an example of an *independence structure*. We will be interested in two properties of independence structures $\perp\!\!\!\perp$:

(IS1) Suppose we have that $\perp\!\!\!\perp(f_i \colon X \to Y_i)_{i \in I}$ and for each $i \in I$ that $\perp\!\!\!\perp(g_{ij} \colon Y_i \to Z_{ij})_{j \in J_i}$. Then $\perp\!\!\!\perp(g_{ij} f_i \colon X \to Z_{ij})_{i \in I, j \in J_i}$.

(IS2) $\perp\!\!\!\perp(f)$ for any singleton family (f).

One of the primary motivating examples of this construction is in probability theory.

Example 5. Note that **FinProb** has as terminal object any one-element set equipped with the uniform distribution. Now given objects $(X_1, p_1), (X_2, p_2)$ in **FinProb**, let $(X_1, p_1) \otimes (X_2, p_2)$ have underlying set $X_1 \times X_2$ and distribution p where $p(x_1, x_2) = p_1(x_1)p_2(x_2)$. Then $(\mathbf{FinProb}, \otimes, 1)$ is an independent product structure, where $f_1 \perp\!\!\!\perp f_2$ expresses mutual independence of f_1 and f_2 as random variables. We continue our example from before. Let $Y_2 = \{u, \overline{u}\}$ be another set of events governing a random selection of birds, with one event for "does feed underwater" and another for "does not feed underwater". If we equip Y_2 with the distribution q_2 given by

$$q_2 \colon u \mapsto 0.3 \qquad q_2 \colon \overline{u} \mapsto 0.7$$

then the function $f_2 \colon X \to Y_2$ given by

$$\begin{array}{lclclc} f_2 \colon \text{dove} \mapsto \overline{u} & & f_2 \colon \text{tern} \mapsto u & & f \colon \text{chicken} \mapsto \overline{u} \\ f_2 \colon \text{duck} \mapsto u & & f_2 \colon \text{ostrich} \mapsto \overline{u} & & \end{array}$$

is a morphism in **FinProb** corresponding to the random variable `feeds_underwater` over the sample space (X, p). Writing yz for the pair (y, z) and renaming $(Y, q), f$ from Example 3 to $(Y_1, q_1), f_1$, the independent product $(Y_1, q_1) \otimes (Y_2, q_2)$ has distribution over $Y_1 \times Y_2$ given by

$$su \mapsto 0.12 \qquad s\overline{u} \mapsto 0.28 \qquad \overline{s}u \mapsto 0.18 \qquad \overline{s}\,\overline{u} \mapsto 0.42$$

Now the functions f_1 and f_2 induce a function $f \colon X \to Y_1 \times Y_2$ given by $f \colon x \mapsto f_1(x)f_2(x)$. The mutual independence of the corresponding random variables is precisely the fact that this f is probability-preserving, that is, a morphism in **FinProb**.

A similar kind of construction exists for **Prob**, the category of probability measures, where again morphisms are considered as random variables, independent products are given by the product probability measure, and $\perp\!\!\!\perp$ expresses independence of random variables. For an example expressing a different kind of independence, we have the following.

Example 6. **FinSur** is the category of finite non-empty sets and surjections. The usual Cartesian product here is an independent product which is not the categorical product. Here $f_1 \perp\!\!\!\perp f_2$ means $f_1^{-1}(y_1) \cap f_2^{-1}(y_2)$ is non-empty for every pair (y_1, y_2) in the product of the codomains of f_1 and f_2. This is referred to as *variation independence* in [17] and is closely related to embedded multivalue dependency in database theory [18].

Another example arises in a theory of name-binding called nominal sets [19]. Here, the independent product is referred to as a *separated product*.

The dual construction to an independent product is that of a *monoidal sum structure* $(\mathbb{C}, \oplus, 0)$. These were independently discovered in [20]. A simple example of interest is **Inj**, where objects are sets, but morphisms are only the injective functions. The disjoint union here makes a sum structure $(\mathbf{Inj}, \sqcup, \emptyset)$ where $f \perp\!\!\!\perp g$ expresses the fact that the images of f and g are disjoint.

[1] To be more precise, $\mathrm{Mor}(C) \times \mathrm{Mor}(C) \setminus \perp\!\!\!\perp$ gives us such an indication.

6. Categories of preference structures

We introduce a framework to study preference structures categorically. Similarly to what is done for all examples in [7], we define the category $\mathbf{Sig}$ to have as objects finite signatures (sets of atomic variables) and to have as morphisms injections. Now, $\mathbf{Sig}$ is a full subcategory of **Inj**, and it keeps the sum structure there.

Lemma 7. $(\mathbf{Sig}, \sqcup, 0)$ *is a sum structure where* $\iota_1 \perp\!\!\!\perp \iota_2$ *if and only if the images of* ι_1 *and* ι_2 *are disjoint.*

Next we must fix some category $\mathbb{C}$ of structures that will give us an indication of preference or typicality among worlds. We want to construct an assignment of each signature to such a structure. Now, at the level of sets, this assignment will be $\Sigma \mapsto \Omega(\Sigma)$. In fact, Ω defines a functor $\mathbf{Sig}^{\mathrm{op}} \to \mathbf{Set}$. Moreover, given a sum-structure diagram

$$\Sigma_1 \longrightarrow \Sigma_1 \sqcup \Sigma_2 \longleftarrow \Sigma_2$$

in $\mathbf{Sig}$, we have that

$$\Omega(\Sigma_1) \longleftarrow \Omega(\Sigma_1 \sqcup \Sigma_2) \longrightarrow \Omega(\Sigma_2)$$

is a product diagram in **Set**. So Ω is a strict monoidal functor.

Definition 8. Let $(\mathbb{C}, U \colon \mathbb{C} \to \mathbf{Set})$ be any concrete category, intended to be a category of preference structures (for example, sets equipped with probability distributions ($\mathbf{FinProb}$), sets equipped with OCFs, or sets equipped with TPOs). We call a functor $P\colon \mathbf{Sig}^{\mathrm{op}} \to \mathbb{C}$ a *preference functor* if $PU = \Omega$. So, on objects, a preference functor P assigns each signature Σ to the set of worlds $\Omega(\Sigma)$ equipped with the type of preference structure given by $\mathbb{C}$.

We think of an inclusion $\Sigma_1 \hookrightarrow \Sigma$ in $\mathbf{Sig}$ being sent to some kind of projection $P(\Sigma) \to P(\Sigma_1)$ in $\mathbb{C}$. We next consider the appropriate $\mathbb{C}$ for each example inference relation described in Section 2. For our probabilistic inference relation, the appropriate category of preference structures is $\mathbf{FinProb}$. In this category, the morphisms are exactly marginalisations. We use this idea to construct an analogous category for OCFs.

Definition 9. Let Obj($\mathbf{OCF}$) be the class of all (X, κ) where X is a set and κ is an ordinal conditional function. Let Mor($\mathbf{OCF}$) contain all arrows $f\colon (X,\kappa) \to (Y,\lambda)$ where $f\colon X \to Y$ is a surjection such that

$$\lambda(y) = \min(\kappa(x) | x \in f^{-1}(y))$$

for each $y \in Y$.

Lemma 10. $\mathbf{OCF}$ *equipped with the forgetful functor to* ***Set*** *is a concrete category. The only morphisms in* $\mathbf{OCF}$ *with underlying function identity are the identity morphisms.* $\mathbf{OCF}$ *has terminal object given by a one-element set where the element has rank 0.*

Example 11. Suppose we have the sets X, Y as in Example 3, but now equip X with ranking function κ given by

$$\begin{array}{lll} \kappa\colon \text{dove} \mapsto 1 & \kappa\colon \text{tern} \mapsto 1 & \kappa\colon \text{chicken} \mapsto 0 \\ \kappa\colon \text{duck} \mapsto 1 & \kappa\colon \text{ostrich} \mapsto 2 & \end{array}$$

and Y with ranking function λ given by

$$\lambda\colon s \mapsto 1 \qquad \lambda\colon \overline{s} \mapsto 0$$

Then $(X,\kappa), (Y,\lambda)$ are objects of $\mathbf{OCF}$, while the function f defined in Example 3 defines a morphism $f\colon (X,\kappa) \to (Y,\lambda)$ in $\mathbf{OCF}$. A preference functor $P\colon \mathbf{Sig} \to \mathbf{OCF}$ may assign the signature $\Sigma = \{s\}$ to $(Y,\lambda) = (\Omega(\Sigma),\lambda)$, whereas (X,κ) will not be in the image of P since X is not a set of complete conjunctions over some signature.

It turns out that **OCF** behaves much like **FinProb**. So, one would expect that there would be a functor between the two categories that preserves much of the structure. Ideally such a functor F would be *concrete*, meaning that any X and $F(X)$ have the same underlying set and f and $F(f)$ have the same underlying function. Furthermore, one might require such a functor to *reverse relative order among objects* in the sense that if $\kappa(x) < \kappa(x')$ then $F(\kappa)(x') < F(\kappa)(x)$. To see that this is not the case, one can consider the OCF τ on two elements that gives a rank of 1 to one element. Now, we can construct a τ' with arbitrarily many elements ranked 0 and one element ranked 1 and will have a morphism $f\colon \tau' \to \tau$ in **OCF** where $\tau' = \tau f$. Considering $F(f)$ gives us the following.

Proposition 12. *There are no concrete functors from* **OCF** *to* **FinProb** *that reverse relative order among all objects.*

A candidate for a functor $F\colon \mathbf{FinProb} \to \mathbf{OCF}$ is one where $F(X,p)$ has underlying set X equipped with OCF κ given by the opposite of the order on X induced by p. That is $\kappa(x) = 0$ iff $p(x)$ is the highest probability in the distribution. Then $\kappa(x_1) = \kappa(x_2)$ whenever $p(x_1) = p(x_2)$, otherwise $\kappa(x_1) = \kappa(x_2) + 1$ whenever $p(x_2) > p(x_1)$ and there is no x such that $p(x_2) > p(x) > p(x_1)$. Unfortunately, F defined this way is not a concrete functor. To see this, one can easily construct a counter-example where the sum of smaller probabilities overtake a larger probability, and the corresponding morphism in **FinProb** is not sent to a rank-preserving function by F. One might then wonder whether the situation is redeemed when we consider big-stepped probability distributions. This is not the case, as illustrated by another counter-example.

Example 13. Consider the distribution p on $X = \{x_1, x_1', x_1'', x_2\}$ and a morphism $f\colon (X,p) \to (\{x_1, x_2\}, q)$ in **FinProb** given by

$$
\begin{array}{ccccccccccc}
p & : & x_1 & \mapsto & 0.65 & & f & & q & : & x_1 & \mapsto & 0.9 \\
p & : & x_1' & \mapsto & 0.2 & & \to & & q & : & x_2 & \mapsto & 0.1 \\
p & : & x_2 & \mapsto & 0.1 & & & & & & & & \\
p & : & x_1'' & \mapsto & 0.05 & & & & & & & &
\end{array}
$$

The image under F is

$$
\begin{array}{ccccccccccc}
F(p) & : & x_1 & \mapsto & 0 & & F(q) & : & x_1 & \mapsto & 0 \\
F(p) & : & x_1' & \mapsto & 1 & & F(q) & : & x_2 & \mapsto & 1 \\
F(p) & : & x_2 & \mapsto & 2 & & & & & & \\
F(p) & : & x_1'' & \mapsto & 3 & & & & & &
\end{array}
$$

Here the rank of x_2 is not preserved, and so f is not a morphism in **OCF**.

Letting **BigStep** be the full subcategory of **FinProb** with big-stepped distributions, Example 13 shows us that F is not a concrete functor from **BigStep** to **OCF**, but note that the relative order is preserved by F. This takes us to TPOs, where we do the same construction: we consider objects to be sets equipped with a TPO and morphisms to be marginalisations of TPOs.

Definition 14. Let Obj(**TPO**) be the class of all $(X, \preceq)$ where X is a finite set and $\preceq$ is a total preorder. Let Mor(**TPO**) contain all arrows $f\colon (X, \preceq_X) \to (Y, \preceq_Y)$ where $f\colon X \to Y$ is a surjection such that

$$y \preceq_Y y' \iff \min(f^{-1}(y)) \preceq_X \min(f^{-1}(y'))$$

Lemma 15. **TPO** *equipped with the forgetful functor to* ***Set*** *is a concrete category. The only morphisms in* **TPO** *with underlying function identity are the identity morphisms.* **TPO** *has terminal object given by a one-element set equipped with the reflexive relation.*

The assignment of an OCF to the TPO induced by its image's relative ordering does in fact define a functor, essentially by construction.

Lemma 16. *Let $R\colon$ Obj$(\mathbf{OCF}) \to$ Obj$(\mathbf{TPO})$ be the defined by $(X, \kappa) \mapsto (X, R(\kappa))$ where $xR(\kappa)x'$ if and only if $\kappa(x) \leq \kappa(x')$. Then R defines a concrete functor $\mathbf{OCF} \to \mathbf{TPO}$.*

Going in the opposite direction, we do not have many different kinds of options. Recall that a *chain* is a total partial order. Any concrete functor $G\colon \mathbf{TPO} \to \mathbf{OCF}$ is entirely determined by the image of the 2-element chain. This gives us the following.

Proposition 17. *For each $a \in \mathbb{N}_0 \cup \infty$, let $G_a\colon (X, \preceq) \mapsto (X, \kappa)$ where the lowest equivalence class of $\preceq$ is assigned to 0 by κ and all other equivalence classes of $\preceq$ are assigned to a by κ. Then $G_a\colon \mathbf{TPO} \to \mathbf{OCF}$ is a concrete functor. All concrete functors $\mathbf{TPO} \to \mathbf{OCF}$ are of this form.*

The kind of reasoning that shows that F is not a functor $\mathbf{FinProb} \to \mathbf{OCF}$ also tells us that taking the opposite of the relative order of a probability distribution does not define a concrete functor $\mathbf{FinProb} \to \mathbf{TPO}$. However, Example 13 seems to suggest that big-stepped distributions evade the issue. This is indeed the case!

Proposition 18. *Let B assign any big-stepped probability distribution (X, p) to the TPO $(X, \preceq)$ where $x_1 \preceq x_2$ if and only if $p(x_2) \leq p(x_1)$. Then B defines a concrete functor $\mathbf{BigStep} \to \mathbf{TPO}$.*

In summary, we have the following diagram of concrete functors in $\mathbf{Cat}$:

$$\begin{tikzcd} \mathbf{FinProb} & \mathbf{OCF} \arrow[d, shift right, "R"'] \\ \mathbf{BigStep} \arrow[u, hook] \arrow[r, "B"'] & \mathbf{TPO} \arrow[u, shift right, "G_a"'] \end{tikzcd}$$

7. Independence among preference structures

We now have all the requisite ingredients to define a notion of semantic independence of preference structures. This turns out to often be a stronger notion of independence than Definition 1 requires for the associated inference relation.

Definition 19. Let $(\mathbb{C}, U\colon \mathbb{C} \to \mathbf{Set})$ be a concrete category where $\mathbb{C}$ is equipped with an independent product $\otimes$. Let $P\colon \mathbf{Sig} \to \mathbb{C}$ be a preference functor. We will say that a partition $\Sigma = \Sigma_1 \sqcup \Sigma_2$ is *semantically independent with respect to P and $\otimes$* if the diagram

$$\Sigma_1 \longrightarrow \Sigma_1 \sqcup \Sigma_2 \longleftarrow \Sigma_2$$

in $\mathbf{Sig}$ maps to an $\otimes$-diagram

$$P(\Sigma_1) \longleftarrow P(\Sigma_1 \sqcup \Sigma_2) \longrightarrow P(\Sigma_2)$$

in $\mathbb{C}$. We will say that P is *semantically independent w.r.t. $\otimes$* if it is strict monoidal, which is equivalent to requiring that each partition in $\mathbf{Sig}$ is semantically independent with respect to P and $\otimes$.

Since we have iterated products, we can extend the above straightforwardly to a definition of semantic independence for an n-fold partition $\Sigma = \Sigma_1 \sqcup \ldots \sqcup \Sigma_n$. We now study this notion for our examples. For the probability-based preference structure, semantic independence is simply the mutual independence of the random variables $P(\Sigma) \to P(\Sigma_1)$ and $P(\Sigma) \to P(\Sigma_2)$.

Proposition 20. *Suppose a partition $\Sigma = \Sigma_1 \sqcup \Sigma_2$ is semantically independent with respect to a preference functor $P\colon \mathbf{Sig} \to \mathbf{FinProb}$ and the usual independent product on $\mathbf{FinProb}$. Then $\Sigma_1 \sqcup \Sigma_2$ is independent with respect to $\mid\!\sim_{P(\Sigma)}$.*

Being semantically independent here is a strictly stronger notion than that of Definition 1.

To do the same for ordinal conditional functions, we must define an independent product on $\mathbf{OCF}$. The needed description is given by the notion of κ-independence.

Proposition 21. *Consider the assignment of OCFs*

$$\otimes\colon ((X_1,\kappa_1),(X_2,\kappa_2)) \mapsto (X_1 \times X_2, \kappa_1+\kappa_2)$$

where $(\kappa_1+\kappa_2)(x_1,x_2) = \kappa_1(x_1)+\kappa_2(x_2)$. *Then* $\otimes$ *defines a functor* $\mathbf{OCF} \times \mathbf{OCF} \to \mathbf{OCF}$ *which is an independent product.*

Example 22. Renaming Σ, λ of Example 11 to Σ_1, λ_1 and, for example, setting $\Sigma_2 = \{u\}$ and λ_2 on $\Omega(\Sigma_2)$ to

$$\lambda_2\colon u \;\mapsto\; 1 \qquad\qquad \lambda_2\colon \overline{u} \;\mapsto\; 0$$

the independent product $(\Omega(\Sigma_1),\lambda_1) \otimes (\Omega(\Sigma_2),\lambda_2)$ has OCF λ over $\Omega(\Sigma_1)\times\Omega(\Sigma_2)$ given by

$$\lambda\colon su \;\mapsto\; 2 \qquad \lambda\colon s\overline{u} \;\mapsto\; 1 \qquad \lambda\colon \overline{s}u \;\mapsto\; 1 \qquad \lambda\colon \overline{s}\,\overline{u} \;\mapsto\; 0$$

If a preference functor $P\colon \mathbf{Sig} \to \mathbf{OCF}$ sends $\Sigma_1 \sqcup \Sigma_2$ to $(\Omega(\Sigma_1 \sqcup \Sigma_2),\lambda)$ then the splitting $\Sigma_1 \sqcup \Sigma_2$ is semantically independent w.r.t. P and $\otimes$.

The binary part of the independence structure $\perp\!\!\!\perp$ associated to the independent product on $\mathbf{OCF}$ corresponds to κ-independence. Likewise, $\perp\!\!\!\perp$ corresponds to ocf-splitting. These are the case by definition. As indicated by Proposition 7 of [5], semantic independence implies independence of the associated inference relation. These observations are made precise below.

Theorem 23. *Suppose we have a partition* $\Sigma = \Sigma_1 \sqcup \Sigma_2$, *a preference functor* $P\colon \mathbf{Sig} \to \mathbf{OCF}$ *and the usual independent product* $\otimes$ *on* $\mathbf{OCF}$. *Then the following hold.*

- $\Sigma_1 \sqcup \Sigma_2$ *is semantically independent w.r.t* P *and* $\otimes$ *if and only if it is* $P(\Sigma)$*-independent, which in turn is the case if and only if* $P(\Sigma)$ *ocf-splits over* (Σ_1,Σ_2).
- *If* $\Sigma_1 \sqcup \Sigma_2$ *is semantically independent w.r.t* P *and* $\otimes$ *then it is independent with respect to* $\mid\!\sim_{P(\Sigma)}$.

Again, semantic independence here is a stronger form of independence than that of Definition 1. A natural question to ask whether for specific OCFs κ of interest, it is the case that semantic independence for κ is equivalent to independence of $\mid\!\sim_\kappa$. There exist examples of c-representations [21] where this is not the case.

We now work towards showing that the same cannot be done for TPOs. That is, we cannot construct some independent product on $\mathbf{TPO}$ such that the associated independence structure corresponds to $\preceq$-independence. Two morphisms $f_1\colon (X,\preceq_X) \to (Y_1,\preceq_{Y_1})$ and $f_2\colon (X,\preceq_X) \to (Y_2,\preceq_{Y_2})$ in $\mathbf{TPO}$ here would need to be independent if and only if for any distinct $i,j \in \{1,2\}$ and $x_i, x_i' \in X$ such that $f_j(x_i) = f_j(x_i')$ we have

$$f_i(x_i) \preceq f_i(x_i') \quad\iff\quad x_i \preceq x_i'$$

Let us denote the corresponding family of relations on $\mathrm{Mor}(\mathbf{TPO})$ as $\perp\!\!\!\perp^{tpo}$. Now, it is easy to find morphisms f_1, f_2, e in $\mathbf{TPO}$ such that $f_1 \perp\!\!\!\perp^{tpo} f_2$, but not $f_1 e \perp\!\!\!\perp^{tpo} f_2 e$. This means that $\perp\!\!\!\perp^{tpo}$ cannot satisfy both (IS1) and (IS2). Since any independent product must have associated independence structure satisfying (IS1) and (IS2), we get the following.

Theorem 24. *There is no independent product on* $\mathbf{TPO}$ *such that the associated independence structure is given by* $\perp\!\!\!\perp^{tpo}$.

This supports the view expressed in [4] that a shortcoming of epistemic state Ψ represented by TPOs is that there is no unique and systematic way to recover Ψ from marginalised states Ψ_1, Ψ_2. As suggested by Proposition 7 of [5], we have that our relative order functor $R\colon \mathbf{OCF} \to \mathbf{TPO}$ preserves independence.

Proposition 25. *Suppose* $P\colon \Sigma \to \mathbf{OCF}$ *is some preference functor and that a partition*

$$\Sigma_1 \xrightarrow[\iota_1]{} \Sigma \xleftarrow[\iota_2]{} \Sigma_2$$

is semantically independent w.r.t. P. *Then* $RP(\iota_1) \perp\!\!\!\perp^{tpo} RP(\iota_2)$.

One might wonder whether R reflects independence in the sense that if any $R(f_1) \perp\!\!\!\perp^{tpo} R(f_2)$ in $\mathbf{TPO}$ then $f_1 \perp\!\!\!\perp f_2$ in $\mathbf{OCF}$. This is not the case, as demonstrated by the following.

Example 26. Let $X = \{x_1, x_2, x_3\}$ and $Y = \{y_1, y_2, y_3, y_4\}$. Consider the TPO $\preceq$ on $X \times Y$ given by

x_1y_1	x_2y_1 x_1y_2	x_3y_1 x_2y_2	x_1y_3	x_2y_3 x_1y_4	x_3y_2	x_3y_3	x_2y_4	x_3y_4

with projections p_1 and p_2 to $\boxed{x_1 \mid x_2 \mid x_3}$ and $\boxed{y_1 \mid y_2 \mid y_3 \mid y_4}$. In the above depictions, the smallest equivalence class is on the left, and we increase in order toward the right. Then $p_1 \perp\!\!\!\perp^{tpo} p_2$. Now suppose there are some OCFs κ on X and λ on Y such that $R((X, \kappa) \otimes (Y, \lambda)) = (X \times Y, \preceq)$. Then from the 3rd and 4th levels of $\preceq$ we get that $\kappa(x_2) - \kappa(x_1) < \lambda(y_3) - \lambda(y_2)$. On the other hand, inspecting the 5th to 8th levels of $\preceq$ we get $\kappa(x_2) + \lambda(y_4) - (\kappa(x_1) + \lambda(y_4)) = \kappa(x_2) - \kappa(x_1)$ is larger than $\kappa(x_3) + \lambda(y_3) - (\kappa(x_3) + \lambda(y_2)) = \lambda(y_3) - \lambda(y_2)$ contradicting the inequality we got from the 3rd and 4th levels.

To close off, we remark that independent products allow for a generalisation of the notion of semantic independence of a partition. Via the independence structure $\perp\!\!\!\perp$, we can talk about independence of disjoint $\Sigma_1, \Sigma_2 \subseteq \Sigma$ even if they don't cover Σ. We can then define that Σ_1 and Σ_2 are *semantically independent with respect to P and $\otimes$* if $P(\iota_1) \perp\!\!\!\perp P(\iota_2)$. The results discussed in this section apply to semantically independent subsets just as they do for semantically independent partitions.

8. Further research

Having expressed κ-independence in terms of independent products, it will be of great interest to survey the other examples of independent products to see which ideas and results carry across to OCFs. In the other direction, it will also be of interest to establish which other kinds of preference structures, for example strict partial orders, also possess independent products. Of interest is also extending our results to infinite signatures.

Conditional syntax splittings is an active area of research [22] [23]. We wish to develop conditional variants of what we have introduced in this paper and compare these to existing constructions. A wide variety of approaches to capturing conditional independence categorically have been recently proposed, for example in [6], [24], and [25]. Here, instead of independent pairs of morphisms that factor through "partial products", conditional independence structures are based on the idea of partial *fibred products* and are expressed in terms of systems of commuting squares. Of interest would be which of these capture semantic conditional independence best in the framework of preference functors.

A different avenue would be to extend our framework of preference functors and consolidate it with the approach based on institutions. For a given category of preference structures, $\mathbb{C}$, the analogue of the Mod functor would be the subcategory $\mathbf{Pref}_{\mathbb{C}}$ of the functor category $\mathbb{C}^{\mathbf{Sig}^{\mathrm{op}}}$ consisting of all preference functors. In fact, for those examples of inference relations in [7] that we consider here, Mod and $\mathbf{Pref}_{\mathbb{C}}$ express exactly the same data. How best to construct an inference relation from this data is an open question.

Yet another avenue is to explore uniqueness of independent product decompositions in this context. For products, this is closely related to the notion of coextensivity [26] via the strict refinement property [27]. We can formulate notions of coextensivity for independent products, and those considered here exhibit appropriate partial forms of coextensivity, similar to what is studied in [28]. This allows us to explore irreducibility and coconnectedness for our independent products. These ideas can be used to prove general results for finest syntax splittings [29].

9. Conclusion

We have proposed an abstract definition of semantic independence for a good class of preference structures. Here, good means that the corresponding category with marginalisations as morphisms possesses independent products. Probability-based preference structures are known to be good in this sense, while we have proved that OCFs are too. In doing so, we have demonstrated how the established notion of κ-independence or ocf-splitting is the same phenomenon as usual independence of random variables in probability theory. We have also shown that TPOs are not good in this sense, justifying the ideas expressed in [4].

We have also defined and explored the new categories of preference structures $\mathbf{OCF}$ and $\mathbf{TPO}$. In doing so, we have clarified the relationships between these and the known categories $\mathbf{FinProb}$ and $\mathbf{BigStep}$ in terms of the existence and non-existence of concrete functors. To summarise, the relative order defines a functor $\mathbf{OCF} \to \mathbf{TPO}$ and opposite relative order defines a functor $\mathbf{BigStep} \to \mathbf{TPO}$. The opposite relative ranking is not a functor $\mathbf{BigStep} \to \mathbf{OCF}$ and we have a limited range of concrete functors $\mathbf{TPO} \to \mathbf{OCF}$. We also have no concrete functors $\mathbf{OCF} \to \mathbf{FinProb}$ that reverse order on objects. In line with [7], this affirms the idea that, while probabilities and ranking functions behave similarly, they are quite different in terms of marginalisations. We have also clarified that TPO-independence is something strictly larger than the image of OCF-independence under the relative ordering functor.

The constructions we have defined here open the door to a categorical treatment of independence in KRR, with much further research to be conducted, including conditional variants, consolidation with the institutions framework, and applications to finest splittings. Another investigation of deep interest is also the extent to which we can transfer ideas and concepts from other known contexts where independent products exist.

Acknowledgments

This work is based on the research supported in part by the National Research Foundation of South Africa (REFERENCE NO: SAI240823262612).

Declaration on Generative AI

The authors have not employed any Generative AI tools.

References

[1] J. Shore, R. Johnson, Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy, IEEE Trans. Inf. Theor. 26 (2006) 26–37. doi:10.1109/TIT.1980.1056144.

[2] R. Parikh, Beliefs, belief revision, and splitting languages, Center for the Study of Language and Information, USA, 1999, p. 266–278.

[3] P. Peppas, M.-A. Williams, S. Chopra, N. Foo, Relevance in belief revision, Artificial Intelligence 229 (2015) 126–138. doi:10.1016/j.artint.2015.08.007.

[4] G. Kern-Isberner, G. Brewka, Strong syntax splitting for iterated belief revision, in: Proceedings of IJCAI-17, 2017, pp. 1131–1137. doi:10.24963/ijcai.2017/157.

[5] G. Kern-Isberner, G. Beierle, G. Brewka, Syntax splitting = relevance + independence: New postulates for nonmonotonic reasoning from conditional belief bases, in: Proceedings of KR 2020, 2020, pp. 560–571. doi:10.24963/kr.2020/56.

[6] A. Simpson, Category-theoretic structure for independence and conditional independence, volume 336, 2018, pp. 281–297. The Thirty-third Conference on the Mathematical Foundations of Programming Semantics (MFPS XXXIII).

[7] C. Beierle, G. Kern-Isberner, Semantical investigations into nonmonotonic and probabilistic logics, Annals of Mathematics and Artificial Intelligence 65 (2012) 123–158. doi:10.1007/s10472-012-9310-1.

[8] J. Goguen, R. Burstall, Institutions: Abstract model theory for specification and programming, Journal of the ACM 39 (1992) 95–146.
[9] S. Kraus, D. Lehmann, M. Magidor, Nonmonotonic reasoning, preferential models and cumulative logics, Artificial Intelligence 44 (1990) 167–207. doi:10.1016/0004-3702(90)90101-5.
[10] J. Hawthorne, On the logic of nonmonotonic conditionals and conditional probabilities, Journal of Philosophical Logic 25 (1996) 185–218. doi:10.1007/BF00247003.
[11] S. Benferhat, D. Dubois, H. Prade, Possibilistic and standard probabilistic semantics of conditional knowledge bases, Journal of Logic and Computation 9 (1999) 873–895. doi:10.1093/logcom/9.6.873.
[12] W. Spohn, Ordinal Conditional Functions: A Dynamic Theory of Epistemic States, Springer Netherlands, Dordrecht, 1988, pp. 105–134. doi:10.1007/978-94-009-2865-7_6.
[13] A. Darwiche, J. Pearl, On the logic of iterated belief revision, Artificial Intelligence 89 (1997) 1–29. doi:10.1016/S0004-3702(96)00038-0.
[14] M. Barr, C. Wells, Category Theory for Computing Science, number v. 1 in Prentice-Hall international series in computer science, Prentice Hall, 1995. URL: https://www.math.mcgill.ca/triples/Barr-Wells-ctcs.pdf.
[15] J. Adamek, H. Herrlich, G. Strecker, Abstract and Concrete Categories: The Joy of Cats, Pure and Applied Mathematics: A Wiley Series of Texts, Monographs and Tracts, Wiley, 1990. URL: http://www.tac.mta.ca/tac/reprints/articles/17/tr17.pdf.
[16] R. Ferguson, Sums and covers in categories, Ph.D. thesis, Stellenbosch University, 2026. URL: https://scholar.sun.ac.za/server/api/core/bitstreams/29b72861-2a64-469b-94f9-8c2c817f02a8/content.
[17] A. Dawid, Separoids: A mathematical framework for conditional independence and irrelevance, Annals of Mathematics and Artificial Intelligence 32 (2001) 335–372.
[18] Y. Sagiv, S. F. Walecka, Subset dependencies and a completeness result for a subclass of embedded multivalued dependencies, Journal of the ACM 29 (1982) 103–117.
[19] A. M. Pitts, Nominal Sets: Names and Symmetry in Computer Science, Cambridge Tracts in Theoretical Computer Science, Cambridge University Press, 2013.
[20] Z. Janelidze, Cover relations on categories, Applied Categorical Structures 17 (2009) 351–371. doi:10.1007/s10485-008-9134-7.
[21] G. Kern-Isberner, Conditionals in Nonmonotonic Reasoning and Belief Revision: Considering Conditionals as Agents, Lecture Notes in Artificial Intelligence, Springer, 2001.
[22] J. Heyninck, G. Kern-Isberner, T. Meyer, J. Haldimann, C. Beierle, Conditional syntax splitting for non-monotonic inference operators, Proceedings of the AAAI Conference on Artificial Intelligence 37 (2023) 6416–6424. doi:10.1609/aaai.v37i5.25789.
[23] L.-P. Spiegel, J. Haldimann, J. Heyninck, G. Kern-Isberner, C. Beierle, Generalized safe conditional syntax splitting of belief bases, in: J. Kwok (Ed.), Proceedings of IJCAI-25, 2025, pp. 4678–4686. doi:10.24963/ijcai.2025/521.
[24] A. Simpson, Equivalence and conditional independence in atomic sheaf logic, J. ACM (2026). doi:10.1145/3809163.
[25] M. D. Meglio, C. Heunen, J.-S. P. Lemay, P. Perrone, D. Stein, Dagger categories of relations: the equivalence of dilatory dagger categories and epi-regular independence categories, 2025. arXiv:2508.01146.
[26] A. Carboni, S. Lack, R. F. C. Walters, Introduction to extensive and distributive categories, Journal of Pure and Applied Algebra 84 (1993) 145–158. doi:10.1016/0022-4049(93)90035-R.
[27] M. Hoefnagel, M-coextensive objects and the strict refinement property, Journal of Pure and Applied Algebra 224 (2020) 106381.
[28] M. Hoefnagel, E. Theart, On extensivity and coextensivity of morphisms, Theory and Applications of Categories 44 (2025) 617–642.
[29] J. P. Haldimann, C. Beierle, Finest Syntax Splittings of Ranking Functions and Total Preorders on Worlds, in: Proceedings of KR 2023, 2023, pp. 747–751. doi:10.24963/kr.2023/75.