De Finetti Theorems and Models of Linear Logic

Raphaëlle Crubillé*
CNRS, Aix-Marseille Université

One of the benefits of a fully abstract denotational model is that it gives information about how the addition of new primitives to the language will affect contextual equivalence: indeed, as soon as we are able to give in this model an interpretation for a new primitive that extends the interpretation of the original language, we know that this new construct can be added while preserving contextual equivalence. In this work, we are concerned with *randomised* extensions of PCF, meaning that we want to extend the language with probabilistic choices. The simplest way to do so is to add a primitive random(p), where $p \in \mathbb{Q}$, that has Boolean type and that returns true with probability p, and false with probability $(1 - p)$: following [EPT18], we will note $\mathrm{PCF_p}$ for this version of randomised PCF. A fully abstract model is known for $\mathrm{PCF_p}$: the model of *probabilistic coherence spaces* (PCSs), first introduced by Girard [Gir04], and shown to be a fully abstract model of $\mathrm{PCF_p}$ by Ehrhard et al. [EPT18].

As a motivating example of more complex probabilistic primitives that could be added to $\mathrm{PCF_p}$, let's look at arbitrary Boolean *samplers*, i.e. agents that behave as boolean oracles with possibly an internal state, and as a consequence may behave differently the n-th time they are called than the first one: it means that while each call returns a Boolean, there is *no guarantee* that this is done by independent samplings along an identical probability distribution[1]. We call *Bernoulli samplers* the uniform samplers, i.e. the ones where there exists a (sub)-probability distribution μ on $\{0, 1\}$ such that each call consists merely in sampling independently from the distribution μ. Let's now consider a program M of type (Bool $\to \mathbb{N}$) in $\mathrm{PCF_p}$: since the evaluation strategy is call-by-name, during the evaluation of the program $(M)\,N$, the Boolean program N will be evaluated as many times as M uses its argument. It can been shown either by operational or by denotational means that for every $\mathrm{PCF_p}$ program N of type Bool, N models a Bernoulli sampler. The motivating question for this work is to know which – if any – other kinds of samplers can be interpreted denotationally in the model of probabilistic coherent spaces, in a way which is compatible with the existing interpretation for $\mathrm{PCF_p}$.

Since the category of probabilistic coherence spaces (**PCoh**) is a model of Linear Logic, a $\mathrm{PCF_p}$-program $\vdash M : \mathrm{Bool} \to \sigma$ is interpreted as a morphism $[\![M]\!] \in \mathbf{PCoh}(!\,[\![\mathrm{Bool}]\!]\,, [\![\sigma]\!])$, where ! is the exponential modality of Linear Logic. In the following, let us note simply Bool for the PCS $[\![\mathrm{Bool}]\!]$. We would like to interpret a boolean sampler S as a morphism $1 \to !\mathrm{Bool}$, thus S could be passed (in the denotational world) as argument to any program $\vdash M : \mathrm{Bool} \to \sigma$. One simple way of building morphisms $1 \to !\mathrm{Bool}$ – or more generally $!\,[\![\Gamma]\!] \to !\mathrm{Bool}$ is to use the *promotion* rule of linear logic, that transforms any element $x : 1 \to \mathrm{Bool}$ into an element $x^! : 1 \to !\mathrm{Bool}$. In the model of probabilistic coherence spaces, a probability distribution x on $\{0, 1\}$ gives rise to an element $1 \to \mathrm{Bool}$, whose promotion models exactly the associated *Bernoulli sampler*. So our objective can be rephrased as understanding what other kind of elements live in $\mathbf{PCoh}(1, !\mathrm{Bool})$, and which kind

*raphaelle.crubille@lis-lab.fr

[1] We consider Boolean samplers in this work, but our line of reasoning could be extended to natural number samplers, or more generally samplers on ground data types.

of samplers they models.

Our starting point is the fact that–as proved in [CEPT17]–the exponential comonad in **PCoh** is characterised by the layered generic construction of the free exponential by [MTT18]. The core of our work consists in building a close connection between this free exponential construction and the currently very active line of work [JS20, MP22, FGP21] looking at categorical representations of De Finetti's theorem. Via this connection, we show that the samplers that can be added to $\mathrm{PCF}_{\mathrm{p}}$ without altering contextual equivalence – or the denotational interpretation – are the *exchangeable* ones: those for which it is impossible to detect a (finitely supported) permutation in the order of the calls.

1 Categorical De Finetti theorems

De Finetti's theorem is a foundational result from probability theory, that characterises the *exchangeable* infinite sequences of random variables on some measurable space X. Here, exchangeable roughly means that for any finite prefix of the list, it is not possible to distinguish whether the variables have been permuted inside this prefix. An infinite list of independent, identically distributed random variables (noted i.i.d.-sequence in the following) is obviously exchangeable, but the converse isn't true. De Finetti's theorem says that, under some conditions on X, every infinite exchangeable list of X-valued random variables s can be written as a *mixture* of i.i.d.-sequences. Let's look for instance at $X = \{0, 1\}$: this result means that the probability measures on $\{0, 1\}^\omega$ that verify the exchangeability requirement are in bijection with the probability measures over the set of Bernoulli samplers, identified to $[0, 1]$. [2]

More generally, let us note as usual $\mathcal{G}$ is the Giry monad, thus $\mathcal{G}X$ is the set of probability distribution over X equipped with a structure of measurable space. De Finetti's theorem says that there exists a probability measure μ on $\mathcal{G}X$ such that the *law* of s–which is a probability measure on X^ω, the infinite product of copies of X–coincides with the probability measure obtained by first sampling from μ an element $x \in \mathcal{G}X$, and then taking an i.i.d.-sequence of law x.

Jacobs and Staton presented in [JS20] a categorical version of De Finetti's theorem when $X = \{0, 1\}$. Their setting is the symmetric monoidal category Stoch, where the objects are measurable spaces, and the morphisms are the stochastic kernels. Jacobs and Staton's formalisation is based on the construction a chain in Stoch, that they call the *draw and delete chain*. For any measurable space X, the objects this chain are the $\mathcal{M}_n(X)$ – the multisets on X of size exactly n – that represent *urns* containing n elements. The morphisms are the $\mathrm{d}_n : \mathcal{M}_n(X) \to \mathcal{M}_{n-1}(X)$ obtained by drawing at random an element in the urn, and then removing it from the urn:

$$\mathbf{1} \xleftarrow{\quad \mathrm{d}_0 \quad} \mathcal{M}_1(X) \xleftarrow{\quad \mathrm{d}_1 \quad} \mathcal{M}_2(X) \xleftarrow{\quad \mathrm{d}_2 \quad} \cdots \tag{1}$$

From a categorical point of view, this chain can also be built from the more primitive *discard* chain, whose objects are the $X^{\otimes n}$, and where the discard morphism $X^{\otimes n+1} \to X^{\otimes n}$ arises from the fact that $\mathbf{1}$ is also the terminal object. We can then obtain the draw-and-delete chain by observing that the $\mathcal{M}_n(X)$ are the equalisers of the symmetries on $X^{\otimes n}$, and lifting the discard chain to the level

[2] De Finetti didn't prove the modern form of De Finetti theorem, since he didn't accept the Kolmogorov axiomatisation of probability theory–see for instance [BR+08, Fis86]. But it's named after him because he proposed taking exchangeable sequences of boolean random variables–that talk only about discrete probability theory, where there was no controversy on the formalism–as foundations for probability theory on the continuous interval $[0, 1]$.

of equalisers.

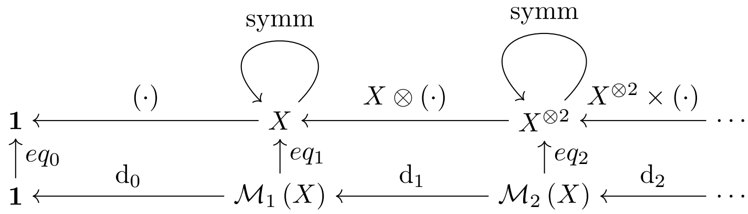

(2)

The law of an exchangeable infinite sequence of random variables on X can then be seen as a *cone* from $\mathbf{1}$ to the draw-and-delete chain: the morphism $1 \to \mathcal{M}_n(X)$ is the law of the n first random variables of the sequence. The fact that the multisets are the equalisers of the symmetries corresponds to the exchangeability requirement. The categorical formalisation of De Finetti's theorem proposed by [JS20] is to state that $\mathcal{G}X$ should be the *limit* of the draw-and-delete chain. In particular, the universal property applies to cones from $\mathbf{1}$, thus those arising from exchangeable sequences of X-valued random variables, which corresponds to the usual De Finetti's theorem.

A less stratified formalisation of De Finetti theorem [MP22, FGP21] has also been proposed by Perrone et al. This version requires explicitly the existence of the infinite product X^ω, axiomatised categorically as the limit of the discard chain. Then instead of looking at the equalisers of symmetries at each layers, Perrone et al.'s De Finetti formalisation is that $\mathcal{G}X$ should be the equaliser of all *finitely supported* symmetries on X^ω:

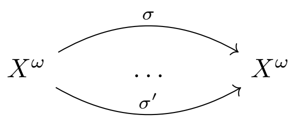

$\forall \sigma, \sigma'$ permutations of $\mathbb{N}$ with finite support.

While X^ω can always be defined as the infinite Cartesian product in Meas, its limit characterisation in Stoch is a consequence of Kolmogorov's extension theorem that does not hold in general for measurable spaces–but hold for all standard Borel spaces, i.e. those measurable spaces whose σ-algebra has been built as the Borel σ-algebra of a Polish space. This formalisation of De Finetti theorem has been shown in the case where X is any *standard Borel space*, and also in a more general case where Stoch is replaced by any *Markov category*.

2 The free exponential layered construction

Starting from a symmetric monoidal closed category $\mathcal{C}$, there are several possible axiomatisations for $\mathcal{C}$ to be able to model the exponential modality of Linear Logic, thus to be a model of (Intuitionistic) Linear Logic [Mel09]. The strongest one, proposed by Lafont, consists in requiring the existence of the free commutative comonoid generated by A–noted $!A$–for every object A of the category. When it is the case, ! can be extended in a comonad that verifies all requirements for a categorical model of Linear Logic. A layered construction is given in [MTT18] for the free exponential modality, that applies to many–but not all– known Lafont models of linear logic. First, it asks for the existence, for every A, of a cartesian product $A_\bullet = A \,\&\, \mathbf{1}$. Secondly, it requires the existence of the equalisers of the symmetries on ${A_\bullet}^{\otimes n}$ for every n. This leads to the *approximants chain*:

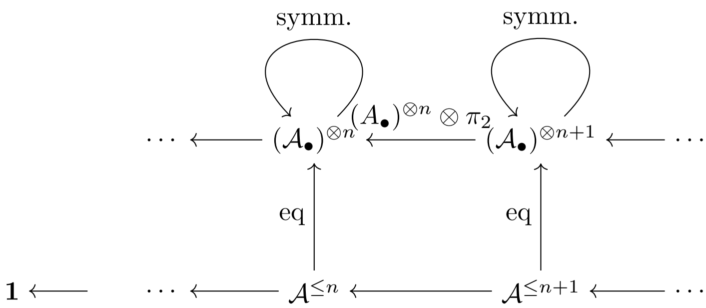

(3)

As proved in [MTT18], the limit of this chain–when it exists, and when moreover the limit commutes with the tensor product–is the free commutative comonoid generated by A, and thus gives a Lafont model of Linear Logic. It has been proved in [CEPT17] by this method that the exponential modality of probabilistic coherent spaces is indeed the free one.

3 Our results

From a formal point of view, there is a striking similarity between the draw-and-delete chain from Section 1 and the approximants chain of Section 2, that reflects the fact that the question of *exchangeability* – and how to manage symmetries between different copies of a resource – is fundamental in Linear Logic semantics. Technically, the only difference between the two chain constructions is that $\mathbf{1}$ is not required to be the terminal object in [MTT18], and accordingly $X^{\otimes n}$ in the draw-and-delete chain is replaced by ${X_\bullet}^{\otimes n}$ in the construction of the approximant chain. The natural place to explore this connection is the category of *integrable cones* (**ICones**) introduced by [EG25] because both the category of stochastic kernels and the category of probabilistic coherent spaces can be faithfully embedded there.

We now sum up the main contributions of our work:

- First, we showed that both chains can indeed be built in the category of integrable cones, and that they indeed have for limit $\mathcal{G}X$ and $!X$ respectively: for the draw-and-delete chain, we proved this for any **ICones**-object X which is the image in **ICones** of a Standard Borel Space; for the approximants chain, we proved this only for Bool, but we hope to be able to show it in the future for at least all the X that are the image of standard Borel spaces. It's interesting to mention that in order to prove the limit of the draw-and-delete chain in **ICones**, we've needed to show that the faithful functor from Stoch to **ICones** preserve both tensor product and *connected* limits, which is a relevant result in itself, and also the more technically involved part of our work.

- In a second part, we built a chain morphism that by the universal property of the limit gives rise to a canonical morphism $\mathcal{G}X \rightarrow !X$, expressing the fact that any exchangeable sampler can be interpreted in $!X$:

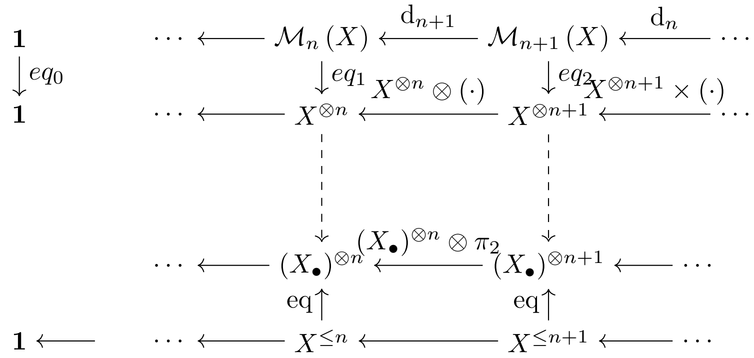

(4)

It's not possible to build an inverse to this chain morphism, and indeed there are some elements in !Bool that aren't the image of elements from $\mathcal{G}X$.

- Finally, we showed that in the case $X = \text{Bool}$, all *total* elements – in the sense of the theory of probabilistic coherence spaces with totality, see [EFP25] – in !Bool are in the image of this morphism $\mathcal{G}\text{Bool} \cong [0,1] \to !\text{Bool}$, thus that the total elements in !Bool are *exactly* the exchangeable samplers.

References

[BR+08] Federico Bassetti, Eugenio Regazzini, et al. The unsung de finetti's first paper about exchangeability. *Rendiconti di Matematica e delle sue applicazioni*, 28:1–17, 2008.

[CEPT17] Raphaëlle Crubillé, Thomas Ehrhard, Michele Pagani, and Christine Tasson. The free exponential modality of probabilistic coherence spaces. In *Foundations of Software Science and Computation Structures: 20th International Conference, FOSSACS 2017, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2017, Uppsala, Sweden, April 22-29, 2017, Proceedings*, pages 20–35. Springer, 2017.

[EFP25] Thomas Ehrhard, Claudia Faggian, and Michele Pagani. Variable elimination as rewriting in a linear lambda calculus. *arXiv preprint arXiv:2501.15439*, 2025.

[EG25] Thomas Ehrhard and Guillaume Geoffroy. Integration in cones. *Logical Methods in Computer Science*, 21, 2025.

[EPT18] Thomas Ehrhard, Michele Pagani, and Christine Tasson. Full abstraction for probabilistic PCF. *J. ACM*, 65(4):23:1–23:44, 2018.

[FGP21] Tobias Fritz, Tomáš Gonda, and Paolo Perrone. De finetti's theorem in categorical probability. *Journal of Stochastic Analysis*, 2(4):6, 2021.

[Fis86] Peter C Fishburn. The axioms of subjective probability. *Statistical Science*, 1(3):335–345, 1986.

[Geo24] Guillaume Geoffroy. personal communication, 2024.

[Gir04] Jean-Yves Girard. Between logic and quantic: a tract. *Linear logic in computer science*, 316:346, 2004.

[JS20] Bart Jacobs and Sam Staton. De finetti's construction as a categorical limit. In *Coalgebraic Methods in Computer Science: 15th IFIP WG 1.3 International Workshop, CMCS 2020, Colocated with ETAPS 2020, Dublin, Ireland, April 25–26, 2020, Proceedings 15*, pages 90–111. Springer, 2020.

[Mel09] Paul-André Mellies. Categorical semantics of linear logic. *Panoramas et syntheses*, 27:15–215, 2009.

[MP22] Sean Moss and Paolo Perrone. Probability monads with submonads of deterministic states. In *Proceedings of the 37th Annual ACM/IEEE Symposium on Logic in Computer Science*, pages 1–13, 2022.

[MTT18] Paul-André Melliès, Nicolas Tabareau, and Christine Tasson. An explicit formula for the free exponential modality of linear logic. *Math. Struct. Comput. Sci.*, 28(7):1253–1286, 2018.

[SYA+17] Sam Staton, Hongseok Yang, Nathanael Ackerman, Cameron Freer, and Daniel M Roy. Exchangeable random processes and data abstraction. In *Workshop on Probabilistic Programming Semantics*, page 2017, 2017.