

Article 25fa pilot End User Agreement

This publication is distributed under the terms of Article 25fa of the Dutch Copyright Act (Auteurswet) with explicit consent by the author. Dutch law entitles the maker of a short scientific work funded either wholly or partially by Dutch public funds to make that work publicly available for no consideration following a reasonable period of time after the work was first published, provided that clear reference is made to the source of the first publication of the work.

This publication is distributed under The Association of Universities in the Netherlands (VSNU) 'Article 25fa implementation' pilot project. In this pilot research outputs of researchers employed by Dutch Universities that comply with the legal requirements of Article 25fa of the Dutch Copyright Act are distributed online and free of cost or other barriers in institutional repositories. Research outputs are distributed six months after their first online publication in the original published version and with proper attribution to the source of the original publication.

You are permitted to download and use the publication for personal purposes. All rights remain with the author(s) and/or copyrights owner(s) of this work. Any use of the publication other than authorised under this licence or copyright law is prohibited.

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please contact the Library through email: copyright@ubn.ru.nl, or send a letter to:

University Library
Radboud University
Copyright Information Point
PO Box 9100
6500 HA Nijmegen

You will be contacted as soon as possible.



De Finetti’s Construction as a Categorical Limit

Bart Jacobs¹(✉) and Sam Staton²

¹ Institute for Computing and Information Sciences,
Radboud University, Nijmegen, NL, The Netherlands
bart@cs.ru.nl

² Department of Computer Science, University of Oxford, Oxford, UK
sam.staton@cs.ox.ac.uk

Abstract. This paper reformulates a classical result in probability theory from the 1930s in modern categorical terms: de Finetti’s representation theorem is redescribed as limit statement for a chain of finite spaces in the Kleisli category of the Giry monad. This new limit is used to identify among exchangeable coalgebras the final one.

1 Introduction

An indifferent belief about the bias of a coin can be expressed in various ways. One way is to use elementary sentences like

$$\begin{aligned} & \text{“When we toss the coin 7 times, the probability of 5 heads is } \frac{1}{8} \text{.”} \\ & \text{“When we toss the coin } k \text{ times, the probability of } n \text{ heads is } \frac{1}{k+1} \text{.”} \end{aligned} \quad (1)$$

Another way is to give a probability measure on the space $[0, 1]$ of biases:

$$\text{“The bias of the coin is uniformly distributed across the interval } [0, 1] \text{.”} \quad (2)$$

This sentence (2) is more conceptually challenging than (1), but we may deduce (1) from (2) using Lebesgue integration:

$$\int_0^1 \binom{7}{5} r^5 (1-r)^2 dr = \frac{1}{8}, \quad \int_0^1 \binom{k}{n} r^n (1-r)^{k-n} dr = \frac{1}{k+1}$$

where $\binom{k}{n} = \frac{k!}{n!(k-n)!}$ is the binomial coefficient. Here we have considered indifferent beliefs as a first illustration, but more generally de Finetti’s theorem gives a passage from sentences about beliefs like (1) to sentences like (2).

Theorem 1 (De Finetti, Informal). *If we assign a probability to all finite experiment outcomes involving a coin (such as (1)), and these probabilities are all suitably consistent, then by doing so we have given a probability measure on $[0, 1]$ (such as the uniform distribution, (2)).*

In this paper we give two categorical/coalgebraic formulations of this theorem.

- In Theorem 8, we state the result as: *the unit interval $[0, 1]$ is an inverse limit over a chain of finite spaces in a category of probability channels*. In particular, a cone over this diagram is an assignment of finite probabilities such as (1) that is suitably consistent.
- In Theorem 12, we state the result as: $[0, 1]$ carries a final object in a category of (exchangeable) coalgebras for the functor $2 \times (-)$ in the category of probability channels. A coalgebra can be thought of as a process that coinductively generates probabilities of the form (1). A coalgebra is a little more data than strictly necessary, but this extra data corresponds to the concept of sufficient statistics which is useful in Bayesian inference.

Multisets and Probability Distributions. A multiset is a ‘set’ in which elements may occur multiple times. We can write such a multiset over a set $\{R, G, B\}$, representing red, green and blue balls, as:

$$3|R\rangle + 5|G\rangle + 2|B\rangle.$$

This multiset can be seen as an urn containing three red, five green and two blue balls.

A probability distribution is a convex combination of elements, where the frequencies (or probabilities) with which elements occur are numbers in the unit interval $[0, 1]$ that add up to one, as in:

$$\frac{3}{10}|R\rangle + \frac{1}{2}|G\rangle + \frac{1}{5}|B\rangle.$$

Obviously, distributions, whether discrete (as above) or continuous, play a central role in probability theory. Multisets also play a crucial role, for instance, an urn with coloured balls is very naturally represented as a multiset. Similarly, multiple data items in learning scenarios are naturally captured by multisets (see *e.g.* [12]). The two concepts are related, since any multiset gives rise to a probability distribution, informally by thinking of it as an urn and then drawing from it (sampling with replacement).

This paper develops another example of how the interplay between multisets and distributions involves interesting structure, in an emerging area of categorical probability theory.

Beliefs as Cones. In brief, following an experiment involving K tosses of a coin, we can set up an urn that approximately simulates the coin according to the frequencies that we have seen. So if we see 5 heads out of 7 tosses, we put 5 black balls into an urn together with 2 white balls, so that the chance of black (=heads) is $\frac{5}{7}$. We can describe a belief about a coin by describing the probabilities of ending up with certain urns (multisets) after certain experiments. Clearly, if we draw and discard a ball from an urn, then this is the same as stopping the experiment one toss sooner. So a consistent belief about our coin determines a cone over the chain

$$\mathcal{M}[1](2) \leftarrow \circ \text{---} \mathcal{M}2 \leftarrow \circ \text{---} \mathcal{M}[3](2) \leftarrow \circ \text{---} \dots \quad (3)$$

in the category of finite sets and channels (aka probability kernels). Here $\mathcal{M}[K](2)$ is the set of multisets over 2 of size K , *i.e.* urns containing K balls coloured black and white, and the leftward arrows describe randomly drawing and discarding a ball. We use the notation \multimap for channels, a.k.a. Kleisli map, see below for details. Our formulation of de Finetti’s theorem says that the limit of this chain is $[0, 1]$. So to give a belief, *i.e.* a cone over (3), is to give a probability measure on $[0, 1]$. We set up the theory of multisets and cones in Sect. 3, proving the theorem in Sect. 6.

Exchangeable Coalgebras. In many situations, it is helpful to describe a belief about a coin in terms of a stateful generative process, *i.e.* a channel (aka probability kernel) $h: X \multimap 2 \times X$, for some space X , which gives coinductively a probability for a sequence of outcomes. This is a coalgebra for the functor $2 \times (-)$. An important example is Pólya’s urn, which is an urn-based process, but one where the urn changes over time. Another important example puts the carrier set $X = [0, 1]$ as the set of possible biases, and the process merely tosses a coin with the given bias. We show that $[0, 1]$ carries the final object among the ‘exchangeable’ coalgebras. Thus there is a canonical channel $X \multimap [0, 1]$, taking the belief described by a generative process such as Pólya’s urn to a probability measure on $[0, 1]$ such as the uniform distribution. We set up the theory of exchangeable coalgebras in Sect. 4, proving the final coalgebra theorem in Sect. 7 by using the limit reformulation of de Finetti’s theorem.

2 Background on Multisets and Distributions

Multisets. A *multiset* (or *bag*) is a ‘set’ in which elements may occur multiple times. We shall write $\mathcal{M}(X)$ for the set of (finite) multisets with elements from a set X . Elements $\varphi \in \mathcal{M}(X)$ may be described as functions $\varphi: X \rightarrow \mathbb{N}$ with finite support, that is, with $\text{supp}(\varphi) := \{x \in X \mid \varphi(x) \neq 0\}$ is finite. Alternatively, multisets can be described as formal sums $\sum_i n_i |x_i\rangle$, where $n_i \in \mathbb{N}$ describes the multiplicity of $x_i \in X$, that is, the number of times that the element x_i occurs in the multiset. The mapping $X \mapsto \mathcal{M}(X)$ is a monad on the category **Sets** of sets and functions. In fact, $\mathcal{M}(X)$ with pointwise addition and empty multiset $\mathbf{0}$, is the free commutative monoid on X .

Frequently we need to have a grip on the total number of elements occurring in a multiset. Therefore we write, for $K \in \mathbb{N}$,

$$\begin{aligned} \mathcal{M}[K](X) &:= \{\varphi \in \mathcal{M}(X) \mid \varphi \text{ has precisely } K \text{ elements}\} \\ &= \{\varphi \in \mathcal{M}(X) \mid \sum_{x \in X} \varphi(x) = K\}. \end{aligned}$$

Clearly, $\mathcal{M}[0](X)$ is a singleton, containing only the empty multiset $\mathbf{0}$. Each $\mathcal{M}[K](X)$ is a quotient of the set of sequences X^K via an accumulator map:

$$X^K \xrightarrow{\text{acc}_X} \mathcal{M}[K](X) \quad \text{with} \quad \text{acc}_X(x_1, \dots, x_K) := \sum_{x \in X} (\sum_i \mathbf{1}_{x_i}(x)) |x\rangle \quad (4)$$

where $\mathbf{1}_{x_i}$ is the indicator function sending $x \in X$ to 1 if $x = x_i$ and to 0 otherwise. In particular, for $X = 2$, $acc_2(b_1, \dots, b_K) = (\sum_i b_i |1\rangle + (K - \sum_i b_i) |0\rangle$. The map acc_X is permutation-invariant, in the sense that:

$$acc_X(x_1, \dots, x_K) = acc_X(x_{\sigma(1)}, \dots, x_{\sigma(K)}) \quad \text{for each permutation } \sigma \in S_K.$$

Such permutation-invariance, or exchangeability, plays an important role in de Finetti's work [5], see the end of Sect. 7.

Discrete Probability. A *discrete probability distribution* on a set X is a finite formal sum $\sum_i r_i |x_i\rangle$ of elements $x_i \in X$ with multiplicity $r_i \in [0, 1]$ such that $\sum_i r_i = 1$. It can also be described as a function $X \rightarrow [0, 1]$ with finite support. We write $\mathcal{D}(X)$ for the set of distributions on X . This gives a monad \mathcal{D} on **Sets**; it satisfies $\mathcal{D}(1) \cong 1$ and $\mathcal{D}(2) \cong [0, 1]$, where $1 = \{0\}$ and $2 = \{0, 1\}$.

Maps in the associated Kleisli category $\mathcal{Kl}(\mathcal{D})$ will be called *channels* and are written with a special arrow \dashrightarrow . Thus $f: X \dashrightarrow Y$ denotes a function $f: X \rightarrow \mathcal{D}(Y)$. For a distribution $\omega \in \mathcal{D}(X)$ and a channel $g: Y \dashrightarrow Z$ we define a new distribution $f \ll \omega \in \mathcal{D}(Y)$ and a new channel $g \circ f: X \dashrightarrow Z$ by Kleisli composition:

$$f \ll \omega := \sum_{y \in Y} \left(\sum_x \omega(x) \cdot f(x)(y) \right) |y\rangle \quad (g \circ f)(x) := g \ll f(x).$$

Discrete Probability over Multisets. For each number $K \in \mathbb{N}$ and probability $r \in [0, 1]$ there is the familiar *binomial* distribution $binom[K](r) \in \mathcal{D}(\{0, 1, \dots, K\})$. It captures probabilities for iterated flips of a known coin, and is given by the convex sum:

$$binom[K](r) := \sum_{0 \leq k \leq K} \binom{K}{k} \cdot r^k \cdot (1 - r)^{K-k} |k\rangle.$$

The multiplicity probability before $|k\rangle$ in this expression is the chance of getting k heads of out K coin flips, where each flip has fixed bias $r \in [0, 1]$. In this way we obtain a channel $binom[K]: [0, 1] \dashrightarrow \{0, 1, \dots, K\}$.

More generally, one can define a multinomial channel:

$$\mathcal{D}(X) \xrightarrow{\text{mulnom}[K]} \mathcal{M}[K](X).$$

It is defined as:

$$mulnom[K](\omega) := \sum_{\varphi \in \mathcal{M}[K](X)} \frac{K!}{\prod_x \varphi(x)!} \cdot \prod_x \omega(x)^{\varphi(x)} |\varphi\rangle.$$

The binomial channel is a special case, since $[0, 1] \cong \mathcal{D}(2)$ and $\{0, 1, \dots, K\} \cong \mathcal{M}[K](2)$, via $k \mapsto k|0\rangle + (K - k)|1\rangle$. The binary case will play an important role in the sequel.

3 Drawing from an Urn

Recall our motivation. We have an unknown coin, and we are willing to give probabilities to the outcomes of different finite experiments with it. From this we ultimately want to say something about the coin. The only experiments we consider comprise tossing a coin a fixed number K times and recording the number of heads. The coin has no memory, and so the order of heads and tails doesn't matter. So we can imagine an urn containing black balls and white balls, and in an experiment we start with an empty urn and put a black ball in the urn for each head and a white ball for each tail. The urn forgets the order in which the heads and tails appear. We can then describe our belief about the coin in terms of our beliefs about the probability of certain urns occurring in the different experiments. These probabilities should be consistent in the following informal sense, that connects the different experiments.

- If we run the experiment with $K + 1$ tosses, and then discard a random ball from the urn, it should be the same as running the experiment with just K tosses.

In more detail: because the coin tosses are exchangeable, that is, the coin is memoryless, to run the experiment with $K + 1$ tosses, is to run the experiment with $K + 1$ tosses and then report some chosen permutation of the results. This justifies representing the outcome of the experiment with an unordered urn. It follows that the statistics of the experiment are the same if we choose a permutation of the results at random. To stop the experiment after K tosses, i.e. to discard the ball corresponding to the last toss, is thus to discard a ball uniformly at random.

As already mentioned: an urn filled with certain objects of the kind x_1, \dots, x_n is very naturally described as a multiset $\sum_i n_i |x_i\rangle$, where n_i is the number of objects x_i in the urn. Thus an urn containing objects from a set X is a multiset in $\mathcal{M}(X)$. For instance, an urn with three black and two white balls can be described as a multiset $3|B\rangle + 2|W\rangle \in \mathcal{M}(2)$, where B stands for black and W for white. We can thus formalise a belief about the K -th experiment, where we toss our coin K times, as a distribution $\omega_K \in \mathcal{D}(\mathcal{M}[K](2))$, on urns with K black and white balls.

For the consistency property, we need to talk about drawing and discarding from an urn. We define a channel that draws a single item from an urn:

$$\mathcal{M}[K + 1](X) \xrightarrow{D} \mathcal{D}(X \times \mathcal{M}[K](X)) \quad (5)$$

It is defined as:

$$D(\varphi) := \sum_{x \in \text{supp}(\varphi)} \frac{\varphi(x)}{\sum_y \varphi(y)} |x, \varphi - 1|x\rangle.$$

Concretely, in the above example:

$$D(3|B\rangle + 2|W\rangle) = \frac{3}{5}|B, 2|B\rangle + 2|W\rangle\rangle + \frac{2}{5}|W, 3|B\rangle + 1|W\rangle\rangle.$$

Notice that on the right-hand-side we use brackets $| - \rangle$ inside brackets $| - \rangle$. The outer $| - \rangle$ are for \mathcal{D} and the inner ones for \mathcal{M} .

We can now describe the 'draw-and-delete' channel, by post-composing D with the second marginal:

$$\mathcal{M}[K+1](X) \xrightarrow{DD := \mathcal{D}(\pi_2) \circ D} \mathcal{D}(\mathcal{M}[K](X))$$

Explicitly, $DD(\varphi) = \sum_x \frac{\varphi(x)}{\sum_y \varphi(y)} |\varphi - 1|x\rangle$.

Now the consistency property requires that the beliefs about the coin comprise a sequence of distributions $\omega_K \in \mathcal{D}(\mathcal{M}[K](2))$ such that $(DD \ll \omega_{K+1}) = \omega_K$. Thus they should form a cone over the sequence

$$\mathcal{M}[0](2) \xleftarrow{DD} \mathcal{M}[1](2) \xleftarrow{DD} \mathcal{M}2 \xleftarrow{DD} \dots \xleftarrow{DD} \mathcal{M}[K](2) \xleftarrow{DD} \dots \quad (6)$$

This means that for all K ,

$$\omega_K(\varphi) = \frac{\varphi(B)+1}{K+1} \cdot \omega_{K+1}(\varphi + |B\rangle) + \frac{\varphi(W)+1}{K+1} \cdot \omega_{K+1}(\varphi + |W\rangle). \quad (7)$$

Known Bias is a Consistent Belief. If we know for sure that the bias of a coin is $r \in [0, 1]$, then we would put $\omega_K = \text{binom}[K](r)$. This is a consistent belief, which is generalised in the following result to the multinomial case. This cone forms the basis for this paper.

Proposition 2. *The multinomial channels form a cone for the chain of draw-and-delete channels, as in:*

$$\begin{array}{ccccccc} \mathcal{M}[0](X) & \xleftarrow{DD} & \dots & \xleftarrow{DD} & \mathcal{M}[K](X) & \xleftarrow{DD} & \mathcal{M}[K+1](X) & \xleftarrow{DD} & \dots \\ & & & & \uparrow & & \uparrow & & \\ & & & & \text{mulnom}[K] & & \text{mulnom}[K+1] & & \dots \\ & \swarrow & & \searrow & \mathcal{D}(X) & \swarrow & \searrow & & \\ & \text{mulnom}[0] & & & & & & & \end{array} \quad (8)$$

Our reformulation of the de Finetti theorem explains the sense in which above cone is a limit when $X = 2$, see Theorem 8.

Proof. For $\omega \in \mathcal{D}(X)$ and $\varphi \in \mathcal{M}[K](X)$ we have:

$$\begin{aligned} & (DD \circ \text{mulnom}[K+1])(\omega)(\varphi) \\ &= \sum_{\psi} \text{mulnom}[K+1](\omega)(\psi) \cdot DD[K](\psi)(\varphi) \\ &= \sum_x \text{mulnom}[K+1](\omega)(\varphi + 1|x) \cdot \frac{\varphi(x)+1}{K+1} \\ &= \sum_x \frac{(K+1)!}{\prod_y (\varphi+1|x)(y)!} \cdot \prod_y \omega(y)^{(\varphi+1|x)(y)} \cdot \frac{\varphi(x)+1}{K+1} \\ &= \sum_x \frac{K!}{\prod_y \varphi(y)!} \cdot \prod_y \omega(y)^{\varphi(y)} \cdot \omega(x) \\ &= \frac{K!}{\prod_y \varphi(y)!} \cdot \prod_y \omega(y)^{\varphi(y)} \cdot (\sum_x \omega(x)) \\ &= \text{mulnom}[K](\omega)(\varphi). \end{aligned} \quad \square$$

Unknown Bias (Uniform) is a Consistent Belief. If we do not know the bias of our coin, it may be reasonable to say that all the $(K + 1)$ outcomes in $\mathcal{M}[K](2)$ are equiprobable. This determines the discrete uniform distribution $v_K \in \mathcal{D}(\mathcal{M}[K](2))$,

$$v_K := \sum_{k=0}^K \frac{1}{K+1} |k|B\rangle + (K - k)|W\rangle.$$

Proposition 3. *The discrete uniform distributions v_K form a cone for the chain of draw-and-delete channels, as in:*

$$\begin{array}{ccccccc}
 \mathcal{M}[0](X) & \xleftarrow{DD} & \dots & \xleftarrow{DD} & \mathcal{M}[K](X) & \xleftarrow{DD} & \mathcal{M}[K+1](X) & \xleftarrow{DD} & \dots \\
 & & & & \uparrow & & \uparrow & & \\
 & & & & v_K & & & & \\
 & & & & | & & & & \\
 & & & & 1 & & & & \\
 & & & & \downarrow & & \downarrow & & \\
 & & & & v_0 & & v_{K+1} & & \dots
 \end{array} \tag{9}$$

Proof. For $\varphi \in \mathcal{M}[K](2)$ we have:

$$\begin{aligned}
 (DD \ll v_{K+1})(\varphi) &= \sum_{\psi} v_{K+1}(\psi) \cdot DD[K](\psi)(\varphi) \\
 &= v_{K+1}(\varphi + |W\rangle) \cdot DD[K](\varphi + |W\rangle)(\varphi) + \\
 &\quad v_{K+1}(\varphi + |B\rangle) \cdot DD[K](\varphi + |B\rangle)(\varphi) \\
 &= \frac{1}{K+2} \cdot \frac{\varphi(W)+1}{K+1} + \frac{1}{K+2} \cdot \frac{\varphi(B)+1}{K+1} \\
 &= \frac{\varphi(W)+\varphi(B)+2}{(K+2)(K+1)} \\
 &= \frac{1}{K+1} \quad \text{since } \varphi(W) + \varphi(B) = K. \quad \square
 \end{aligned}$$

4 Coalgebras and Pólya’s Urn

In the previous section we have seen two examples of cones over the chain (6): one for a coin with known bias (Proposition 2), and one where the bias of the coin is uniformly distributed (Proposition 3).

In practice, it is helpful to describe our belief about the coin in terms of a data generating process. This is formalised as a channel $h: X \multimap 2 \times X$, *i.e.* as a coalgebra for the functor $2 \times (-)$ on the category of channels. The idea is that our belief about the coin is captured by some parameters, which are the states of the coalgebra $x \in X$. The distribution $h(x)$ describes our belief about the outcome of the next coin toss, but also provides a new belief state according to the outcome. For example, if we have uniform belief about the bias of the coin, we predict the outcome to be heads or tails equally, but afterward the toss our belief is refined: if we saw heads, we believe the coin to be more biased towards heads.

Pólya’s urn is an important example of such a coalgebra. We describe it as a multiset over $2 = \{0, 1\}$ with a certain extraction and re-fill operation.

The number 0 corresponds to black balls, and 1 to white ones, so that a multiset $\varphi \in \mathcal{M}(2)$ captures an urn with $\varphi(0)$ black and $\varphi(1)$ white ones. When a ball of color $x \in 2$ is extracted, it is put back together with an extra ball of the same color. We choose to describe it as a coalgebra on the set $\mathcal{M}_*(2)$ of non-empty multisets over 2. Explicitly, $pol: \mathcal{M}_*(2) \rightarrow 2 \times \mathcal{M}_*(2)$ is given by:

$$pol(\varphi) := \frac{\varphi(0)}{\varphi(0) + \varphi(1)} |0, \varphi + 1|0\rangle\rangle + \frac{\varphi(1)}{\varphi(0) + \varphi(1)} |1, \varphi + 1|1\rangle\rangle.$$

So the urn keeps track of the outcomes of the coin tosses so far, and our belief about the bias of the coin is updated according to successive experiments.

4.1 Iterating Draws from Coalgebras

By iterating a coalgebra $h: X \rightarrow 2 \times X$, we can simulate the experiment with K coin tosses $h_K: X \rightarrow \mathcal{M}[K](2)$, by repeatedly applying h . Formally, we build a sequence of channels $h^K: X \rightarrow \mathcal{M}[K](2) \times X$ by $h^0(x) := (\mathbf{0}, x)$ and:

$$h^{K+1} := X \xrightarrow{h^K} \mathcal{M}[K](2) \times X \xrightarrow{\text{id} \times h} \mathcal{M}[K](2) \times 2 \times X \xrightarrow{\text{add} \times \text{id}} \mathcal{M}[K+1](2) \times X$$

where the rightmost map $add: \mathcal{M}[K](2) \times 2 \rightarrow \mathcal{M}[K+1](2)$ is $add(\varphi, b) = \varphi + 1|b\rangle$.

(Another way to see this is to regard a coalgebra $h: X \rightarrow 2 \times X$ in particular as coalgebra $X \rightarrow \mathcal{M}(2) \times X$. Now $\mathcal{M}(2) \times -$ is the writer monad for the commutative monoid $\mathcal{M}(2)$, i.e. the monad for $\mathcal{M}(2)$ -actions. Iterated Kleisli composition for this monad gives a sequence of maps $h: X \rightarrow \mathcal{M}(2) \times X$.)

For example, starting with Pólya's urn $pol: \mathcal{M}_*(2) \rightarrow 2 \times \mathcal{M}_*(2)$, we look at a few iterations:

$$\begin{aligned} (pol)^0(b|0\rangle + w|1\rangle) &= 1|\mathbf{0}, b|0\rangle + w|1\rangle\rangle \\ (pol)^1(b|0\rangle + w|1\rangle) &= pol(b|0\rangle + w|1\rangle) \\ &= \frac{b}{b+w} |1|0\rangle, \varphi + |0\rangle\rangle + \frac{w}{b+w} |1|1\rangle, \varphi + |1\rangle\rangle \\ (pol)^2(b|0\rangle + w|1\rangle) &= \frac{b}{b+w} \cdot \frac{b+1}{b+w+1} |2|0\rangle, (b+2)|0\rangle + w|1\rangle\rangle \\ &\quad + \frac{b}{b+w} \cdot \frac{w}{b+w+1} |1|0\rangle + 1|1\rangle, (b+1)|0\rangle + (w+1)|1\rangle\rangle \\ &\quad + \frac{w}{b+w} \cdot \frac{b}{b+w+1} |1|0\rangle + 1|1\rangle, (b+1)|0\rangle + (w+1)|1\rangle\rangle \\ &\quad + \frac{w}{b+w} \cdot \frac{w+1}{b+w+1} |2|1\rangle, b|0\rangle + (w+2)|1\rangle\rangle \\ &= \frac{b(b+1)}{(b+w)(b+w+1)} |2|0\rangle, (b+2)|0\rangle + w|1\rangle\rangle \\ &\quad + \frac{2bw}{(b+w)(b+w+1)} |1|0\rangle + 1|1\rangle, (b+1)|0\rangle + (w+1)|1\rangle\rangle \\ &\quad + \frac{w(w+1)}{(b+w)(b+w+1)} |2|1\rangle, b|0\rangle + (w+2)|1\rangle\rangle. \end{aligned} \tag{10}$$

It is not hard to see that we get as general formula, for $K \in \mathbb{N}$:

$$(pol)^K(b|0\rangle + w|1\rangle) = \sum_{0 \leq k \leq K} \binom{K}{k} \cdot \frac{\prod_{i < k} (b+i) \cdot \prod_{i < K-k} (w+i)}{\prod_{i < K} (b+w+i)} |k|0\rangle + (K-k)|1\rangle, (b+k)|0\rangle + (w+(K-k))|1\rangle\rangle.$$

We are especially interested in the first component of $h^K(x)$, which is in $\mathcal{M}[K](2)$. Thus we define the channels:

$$h_K := \left(X \xrightarrow{h^K} \mathcal{M}[K](2) \times X \xrightarrow{\pi_1} \mathcal{M}[K](2) \right)$$

For the Pólya urn pol , the really remarkable fact about $pol_K: \mathcal{M}_*(2) \rightarrow \mathcal{M}[K](2)$ is that if we start with the urn with one white ball and one black ball, these distributions on $\mathcal{M}[K](2)$ that come from iteratively drawing are uniform distributions:

$$pol_K(1|0) + 1|1) = v_K.$$

More generally:

Proposition 4

$$\begin{aligned} & pol_K(b|0) + w|1)(k|0) + (K - k)|1) \\ &= \binom{K}{k} \cdot \frac{(b + k - 1)!}{(b - 1)!} \cdot \frac{(w + K - k - 1)!}{(w - 1)!} \cdot \frac{(b + w - 1)!}{(b + w + K - 1)!}. \end{aligned} \tag{11}$$

Proof. We start from the K -th iteration formula for pol , given just before Proposition 4, and derive:

$$\begin{aligned} & \frac{\prod_{i < k} (b + i) \cdot \prod_{i < K - k} (w + i)}{\prod_{i < K} (b + w + i)} \\ &= \frac{(b + k - 1)!}{(b - 1)!} \cdot \frac{(w + K - k - 1)!}{(w - 1)!} \cdot \frac{(b + w - 1)!}{(b + w + K - 1)!}. \end{aligned} \quad \square$$

From this characterization, we can deduce that the channels $pol_K: \mathcal{M}_*(2) \rightarrow \mathcal{M}[K](2)$ form a cone over the draw-and-delete chain (6), *i.e.* that $DD \circ pol_{K+1} = pol_K$, by a routine calculation.

In what follows we provide two other ways to see that pol_K forms a cone. First we deduce it using the fact that the coalgebra pol is exchangeable (Lemma 6). Second we deduce it by noticing that $pol_K(b|0) + w|1)(k|0) + (K - k)|1)$ is distributed as $betabin[K](b, w)$, the beta binomial distribution, and hence factors through the binomial cone (Proposition 2) via the beta channel (Lemma 7).

4.2 Exchangeable Coalgebras and Cones

We would like to focus on those coalgebras $h: X \rightarrow 2 \times X$ that describe consistent beliefs, *i.e.* for which $h_K: X \rightarrow \mathcal{M}[K](2)$ is a cone over the chain (6):

$$DD \circ h_{K+1} = h_K. \tag{12}$$

For example, Pólya’s urn forms a cone. But the deterministic alternating coalgebra

$$alt: 2 \rightarrow 2 \times 2 \quad alt(b) := 1|b, \neg b) \tag{13}$$

does not form a cone, since for instance $alt_1(0) = 1|0\rangle$, while $DD \circ alt_2(0) = \frac{1}{2}|0\rangle + \frac{1}{2}|1\rangle$.

The following condition is inspired by de Finetti's notion of exchangeability, and gives a simple coinductive criterion for a coalgebra to yield a cone (12). It is also related to ideas from probabilistic programming (e.g. [1]).

Definition 5. A coalgebra $h: X \multimap 2 \times X$ will be called exchangeable if its outputs can be reordered without changing the statistics, as expressed by commutation of the following diagram.

$$\begin{array}{ccc}
 X & \xrightarrow{h} & 2 \times X \xrightarrow{id \otimes h} 2 \times 2 \times X \\
 & & \cong \downarrow \text{swap} \times id_X \\
 X & \xrightarrow{h} & 2 \times X \xrightarrow{id \otimes h} 2 \times 2 \times X
 \end{array} \quad (14)$$

Returning to Pólya's urn coalgebra $pol: \mathcal{M}_*(2) \multimap 2 \times \mathcal{M}_*(2)$, we see that it is exchangeable, for, by a similar argument to (10),

$$\begin{aligned}
 (id \times pol)(pol(b|0) + w|1)) &= \frac{b}{b+w} \cdot \frac{b+1}{b+w+1} |(0, 0), (b+2)|0\rangle + w|1\rangle\rangle \\
 &+ \frac{b}{b+w} \cdot \frac{w}{b+w+1} |(0, 1), (b+1)|0\rangle + (w+1)|1\rangle\rangle \\
 &+ \frac{w}{b+w} \cdot \frac{b}{b+w+1} |(1, 0), (b+1)|0\rangle + (w+1)|1\rangle\rangle \\
 &+ \frac{w}{b+w} \cdot \frac{w+1}{b+w+1} |(1, 1), b|0\rangle + (w+2)|1\rangle\rangle
 \end{aligned}$$

Reordering the results does not change the statistics, because $b \cdot w = w \cdot b$. On the other hand, the deterministic alternating coalgebra $alt: 2 \multimap 2 \times 2$ from (13) is not exchangeable, since:

$$\begin{aligned}
 (id \times alt)(alt(0)) &= |((0, 1), 0)\rangle \\
 (swap \times 2)((id \times alt)(alt(0))) &= |((1, 0), 0)\rangle.
 \end{aligned}$$

Before we state and prove our theorem about exchangeable coalgebras, we also consider a different sequence, which keeps track of the order of draws. We define a sequence $h^{\sharp K}: X \multimap 2^K \times X$, with $h^{\sharp 0}(x) = ((), x)$ and $h^{\sharp(K+1)}(x) = (2^K \otimes h) \ll h^{\sharp K}(x)$. Then we define channels $h^{\sharp}_K: X \multimap 2^K$ by

$$h^{\sharp}_K(x) := \mathcal{D}(\pi_1)(h^{\sharp K}(x)) \quad (15)$$

These constructions are related similar to the earlier constructions h^K and h_K by

$$\begin{aligned}
 h^K &= \left(X \xrightarrow{h^{\sharp K}} 2^K \times X \xrightarrow{acc \times X} \mathcal{M}[K](2) \times X \right) \\
 h_K &= \left(X \xrightarrow{h^{\sharp}_K} 2^K \xrightarrow{acc} \mathcal{M}[K](2) \right),
 \end{aligned}$$

where $acc: 2^K \rightarrow \mathcal{M}[K](2)$ is the accumulator map from (4). The difference is that the h^{\sharp} variants keep track of the order of draws, which will be useful in our

arguments about exchangeability, since h is exchangeable (Definition 5) if and only if $h^{\sharp 2} = \mathcal{D}(\text{swap} \times X) \circ h^{\sharp 2} : X \multimap 2 \times 2 \times X$.

Lemma 6. *Let $h : X \multimap 2 \times X$ be an exchangeable coalgebra. Then*

1. *the map h_K can be expressed as:*

$$h_K(x) = \sum_{k=0}^K \binom{K}{k} \cdot h_K^\sharp(x) \left(\underbrace{1, \dots, 1}_{k \text{ times}}, \underbrace{0, \dots, 0}_{K-k \text{ times}} \right) |k|1\rangle + (K-k)|0\rangle. \quad (16)$$

2. *the collection of maps (h_K) forms a cone for the chain of draw-and-delete maps DD (6).*

Proof. 1. Among $b_1, \dots, b_K \in 2$ there are $\binom{K}{k}$ -many possibilities of having k -many ones. By exchangeability of h there is a normal form with all these ones upfront

$$h_K^\sharp(x)(b_1, \dots, b_K) = h_K^\sharp(x) \left(\underbrace{1, \dots, 1}_{k \text{ times}}, \underbrace{0, \dots, 0}_{K-k \text{ times}} \right)$$

2. Using (7) we get, for any $\varphi = (k|1\rangle + (K-k)|0\rangle) \in \mathcal{M}[K](2)$,

$$\begin{aligned} & (DD \circ h_{K+1})(x)(\varphi) \\ &= \frac{k+1}{K+1} \cdot h_{K+1}(x)(\varphi + |1\rangle) + \frac{K+1-k}{K+1} \cdot h_{K+1}(x)(\varphi + |0\rangle) \\ &\stackrel{(16)}{=} \frac{k+1}{K+1} \cdot \binom{K+1}{k+1} \cdot h_{K+1}^\sharp(x) \left(\underbrace{1, \dots, 1}_{k+1 \text{ times}}, 0, \dots, 0 \right) + \\ & \quad \frac{K+1-k}{K+1} \cdot \binom{K+1}{k} \cdot h_{K+1}^\sharp(x) \left(\underbrace{1, \dots, 1}_{k \text{ times}}, 0, \dots, 0 \right) \\ &= \binom{K}{k} \cdot \left(h_{K+1}^\sharp(x) \left(\underbrace{1, \dots, 1}_{k \text{ times}}, 0, \dots, 0, 1 \right) + h_{K+1}^\sharp(x) \left(\underbrace{1, \dots, 1}_{k \text{ times}}, 0, \dots, 0 \right) \right) \\ &= \binom{K}{k} \cdot \sum_y h^K(x) \left(\underbrace{1, \dots, 1}_{k \text{ times}}, 0, \dots, 0, y \right) \cdot \left(h_1^\sharp(y)(1) + h_1^\sharp(y)(0) \right) \\ &= \binom{K}{k} \cdot \sum_y h^K(x) \left(\underbrace{1, \dots, 1}_{k \text{ times}}, 0, \dots, 0, y \right) \cdot 1 \\ &= \binom{K}{k} \cdot h_K^\sharp(x) \left(\underbrace{1, \dots, 1}_{k \text{ times}}, 0, \dots, 0 \right) \\ &\stackrel{(16)}{=} h_K(x)(\varphi). \quad \square \end{aligned}$$

Since the Pólya’s urn coalgebra pol is exchangeable, the lemma provides another way to deduce that the channels $pol_K : \mathcal{M}_*(2) \multimap \mathcal{M}[K](2)$ form a cone (12).

5 Background on Continuous Probability

De Finetti's theorem is about the passage from discrete probability to continuous probability. Consider the uniform distribution μ on $[0, 1]$. The probability of any particular point is 0. So rather than give a probability to individual points, we give probabilities to sets; for example, the probability of a point in the half interval $[0, \frac{1}{2}]$ is $\frac{1}{2}$.

In general, a σ -algebra on a set X is a collection of sets Σ that is closed under countable unions and complements. A measurable space (X, Σ) is a set together with a σ -algebra; the sets in Σ are called measurable. In this paper we mainly consider two kinds of σ -algebra, and so they will often be left implicit:

- on a countable set (e.g. 2 , $\mathcal{M}(2)$, $\mathcal{M}[K](2)$) we consider the σ -algebra where every set is measurable.
- on the unit interval $[0, 1]$, the Borel σ -algebra, which is the least σ -algebra containing the intervals.

On a measurable space (X, Σ) , a *probability measure* is a function $\omega: \Sigma \rightarrow [0, 1]$ with $\omega(X) = 1$ and with $\omega(\bigcup_i U_i) = \sum_i \omega(U_i)$ when the measurable subsets $U_i \in \Sigma$ are pairwise disjoint. On a finite set X the probability measures coincide with the discrete distributions.

The morphisms of measurable spaces are the functions $f: X \rightarrow Y$ for which the inverse image of a measurable set is measurable. Crucially, if $f: X \rightarrow [0, 1]$ is measurable and $\omega: \Sigma \rightarrow [0, 1]$ is a probability measure, then we can find the Lebesgue integral $\int f(x)\omega(dx)$, which is the expected value of f . When X is finite, the Lebesgue integral is a sum: $\int f(x)\omega(dx) = \sum_X f(x)\omega(\{x\})$. When $X = [0, 1]$ and ω is the uniform distribution, then the Lebesgue integral is the familiar one.

We write $\mathcal{G}(X, \Sigma)$ for the set of all probability measures on (X, Σ) , where we often omit the σ -algebra Σ when it is clear from the context.: $\mathcal{D}(X) \cong \mathcal{G}(X)$, when all subsets are seen as measurable. The set $\mathcal{G}(X, \Sigma)$ itself has a σ -algebra, which is the least σ -algebra making the sets $\{\omega \mid \omega(U) < r\}$ measurable for all fixed U and r .

Like \mathcal{D} , the construction \mathcal{G} is a monad, commonly called the Giriy monad, but \mathcal{G} is a monad on the category of measurable spaces. Maps in the associated Kleisli category $\mathcal{Kl}(\mathcal{G})$ will also be called *channels*. Thus $f: X \multimap Y$ between measurable spaces can be described as a function $f: X \times \Sigma_Y \rightarrow [0, 1]$ that is a measure in Σ_Y and measurable in X . For a measure $\omega \in \mathcal{G}(X)$ and a channel $g: Y \multimap Z$ we define a new distribution $f \ll \omega \in \mathcal{G}(Y)$ and a new channel $g \circ f: X \multimap Z$ by Kleisli composition:

$$(f \ll \omega)(U) := \int_X f(x)(U) \omega(dx) \quad (g \circ f)(x) := g \ll f(x).$$

Notice that a discrete channel is in particular a continuous channel, according to our terminology, and so we have an inclusion of Kleisli categories $\mathcal{Kl}(\mathcal{D}) \hookrightarrow \mathcal{Kl}(\mathcal{G})$.

Illustrations from the Beta-Bernoulli Distributions. At the end of Sect. 4.1, we briefly mentioned the beta-bernoulli distribution as it arises from the coalgebra for Pólya’s urn. In general, this beta-binomial channel is defined for arbitrary $\alpha, \beta \in \mathbb{R}_{>0}$ as:

$$\text{betabin}[K](\alpha, \beta) := \sum_{0 \leq k \leq K} \binom{K}{k} \cdot \frac{B(\alpha + k, \beta + K - k)}{B(\alpha, \beta)} |k|0\rangle + (K - k)|1\rangle, \quad (17)$$

This description involves the function $B: \mathbb{R}_{>0} \times \mathbb{R}_{>0} \rightarrow \mathbb{R}$ defined as:

$$B(\alpha, \beta) := \int_0^1 x^{\alpha-1} \cdot (1 - x)^{\beta-1} dx \quad (18)$$

This B is a well-known mathematical function satisfying for instance:

$$B(\alpha + 1, \beta) = \frac{\alpha}{\alpha + \beta} \cdot B(\alpha, \beta) \quad B(n + 1, m + 1) = \frac{n! \cdot m!}{(n + m + 1)!}, \quad (19)$$

for $n, m \in \mathbb{N}$. The latter equation shows that (11) is an instance of (17).

The function B is also used to define the Beta distribution on the unit interval $[0, 1]$. We describe it here as a channel in the Kleisli category of the Giry monad:

$$\mathbb{R}_{>0} \times \mathbb{R}_{>0} \xrightarrow{\text{beta}} \mathcal{G}([0, 1]).$$

It is defined on $\alpha, \beta \in \mathbb{R}_{>0}$ and measurable $M \subseteq [0, 1]$ as:

$$\text{beta}(\alpha, \beta)(M) := \int_M \frac{x^{\alpha-1} \cdot (1 - x)^{\beta-1}}{B(\alpha, \beta)} dx. \quad (20)$$

We prove some basic properties of the beta-binomial channel. These properties are well-known, but are formulated here in slightly non-standard form, using channels.

Lemma 7. *1. A betabinomial channel can itself be decomposed as:*

$$\begin{array}{ccc} & & \mathcal{M}[K](2) \\ & \nearrow \text{betabin}[K] & \uparrow \text{binom}[K] \\ \mathbb{R}_{>0} \times \mathbb{R}_{>0} & \xrightarrow{\text{beta}} & [0, 1] \end{array}$$

2. The beta-binomial channels

$$\mathbb{R}_{>0} \times \mathbb{R}_{>0} \xrightarrow{\text{betabin}[K]} \mathcal{M}[K](2)$$

for $K \in \mathbb{N}$, form a cone for the infinite chain of draw-and-delete channels (6).

Proof. 1. Composition in the Kleisli category $\mathcal{Kl}(\mathcal{G})$ of the Giry monad \mathcal{G} yields:

$$\begin{aligned}
 & (\text{binom}[K] \circ \text{beta})(\alpha, \beta)(k|0) + (K - k|1) \\
 &= \int_0^1 \binom{K}{k} \cdot x^k \cdot (1 - x)^{K-k} \cdot \frac{x^{\alpha-1} \cdot (1-x)^{\beta-1}}{B(\alpha, \beta)} dx \\
 &= \binom{K}{k} \cdot \frac{\int_0^1 x^{\alpha+k-1} \cdot (1-x)^{\beta+(K-k)-1} dx}{B(\alpha, \beta)} \\
 &\stackrel{(18)}{=} \binom{K}{k} \cdot \frac{B(\alpha + k, \beta + K - k)}{B(\alpha, \beta)} \\
 &\stackrel{(17)}{=} \text{betabin}[K](\alpha, \beta)(k|0) + (K - k|1).
 \end{aligned}$$

2. The easiest way to see this is by observing that the channels factorise via the binomial channels $\text{binom}[K]: [0, 1] \rightarrow \mathcal{M}[K](2)$. In Proposition 2 we have seen that these binomial channels commute with draw-and-delete channels DD . \square

6 De Finetti as a Limit Theorem

We first describe our ‘limit’ reformulation of the de Finetti theorem and later on argue why this is a reformulation of the original result.

Theorem 8. *For $X = 2$ the cone (8) is a limit in the Kleisli category $\mathcal{Kl}(\mathcal{G})$ of the Giry monad; this limit is the unit interval $[0, 1] \cong \mathcal{D}(2)$.*

This means that if we have channels $c_K: Y \rightarrow \mathcal{M}[K](2) \cong \{0, 1, \dots, K\}$ with $DD \circ c_{K+1} = c_K$, then there is a unique mediating channel $c: Y \rightarrow [0, 1]$ with $\text{binom}[K] \circ c = c_K$, for each $K \in \mathbb{N}$.

In the previous section we have seen an example with $Y = \mathbb{R}_{>0} \times \mathbb{R}_{>0}$ and $c_K = \text{betabin}[K]: Y \rightarrow \mathcal{M}[K](2)$ and mediating channel $c = \text{beta}: Y \rightarrow [0, 1]$.

Below we give the proof, in the current setting. We first illustrate the central role of the draw-and-delete maps $DD: \mathcal{M}[K+1](X) \rightarrow \mathcal{M}[K](X)$ and how they give rise to hypergeometric and binomial distributions.

Via the basic isomorphism $\mathcal{M}[K](2) \cong \{0, 1, \dots, K\}$ the draw-and-delete map (7) becomes a map $DD: \{0, \dots, K+1\} \rightarrow \{0, \dots, K\}$, given by:

$$DD(\ell) = \frac{\ell}{K+1}|\ell - 1\rangle + \frac{K+1-\ell}{K+1}|\ell\rangle. \tag{21}$$

Notice that the border cases $\ell = 0$ and $\ell = K + 1$ are handled implicitly, since the corresponding probabilities are then zero.

We can iterate DD , say N times, giving $DD^N: \{0, \dots, K+N\} \rightarrow \{0, \dots, K\}$. If we draw-and-delete N times from an urn with $K + N$ balls, the result is the same as sampling without replacement K times from the same urn. Sampling without replacement is modelled by the *hypergeometric distribution*. So we can calculate DD^N explicitly, using the hypergeometric distribution formula:

$$DD^N(\ell) = \sum_{0 \leq k \leq K} \binom{\ell}{k} \cdot \frac{\binom{K+N-\ell}{K-k}}{\binom{K+N}{K}} |k\rangle.$$

Implicitly the formal sum is over $k \leq \ell$, since otherwise the corresponding probabilities are zero.

As is well known, for large N , the hypergeometric distribution can be approximated by the binomial distribution. To be precise, for fixed K , if $(\ell/K+N) \rightarrow p$ as $N \rightarrow \infty$, then the hypergeometric distribution $DD^N(\ell)$ approaches the binomial distribution $\text{binom}[K](p)$ as N goes to infinity. This works as follows. Using the notation $(n)_m = \prod_{i < m} (n - i) = n(n - 1)(n - 2) \cdots (n - m + 1)$ we can write $\binom{n}{m} = (n)_m/n!$ and thus:

$$\begin{aligned} \frac{\binom{\ell}{k} \cdot \binom{K+N-\ell}{K-k}}{\binom{K+N}{K}} &= \frac{(\ell)_k}{k!} \cdot \frac{(K+N-\ell)_{K-k}}{(K-k)!} \cdot \frac{K!}{(K+N)_K} \\ &= \binom{K}{k} \cdot \frac{(\ell)_k \cdot (K+N-\ell)_{K-k}}{(K+N)_K} \\ &= \binom{K}{k} \cdot \frac{\prod_{i < k} (\ell - i) \cdot \prod_{i < K-k} (K+N-\ell - i)}{\prod_{i < K} (K+N - i)} \\ &= \binom{K}{k} \cdot \frac{\prod_{i < k} (\ell/K+N - i/K+N) \cdot \prod_{i < K-k} (1 - \ell/K+N - i/K+N)}{\prod_{i < K} (1 - i/K+N)} \\ &\rightarrow \binom{K}{k} \cdot p^k \cdot (1-p)^{K-k} \quad \text{as } N \rightarrow \infty \\ &= \text{binom}[K](p)(k). \end{aligned}$$

The proof of Theorem 8 makes use of a classical result of Hausdorff [10], giving a criterion for the existence of probability measures. It is repeated here without proof (but see *e.g.* [6] for details).

Recall that the *moments* of a distribution $\mu \in \mathcal{G}([0, 1])$ form a sequence of numbers $\mathbf{m}_1, \dots, \mathbf{m}_K, \dots \in [0, 1]$ given by $\mathbf{m}_K = \int_0^1 r^K \mu(dr)$: the expected values of r^K . In general, moments correspond to statistical concepts such as mean, variance, skew and so on. But in this specific case, the K -th moment of μ is the probability ($\text{binom}[K] \ll \mu$)(K) that K balls drawn will *all* be black.

Recall that a sequence $a = (a_n)$ is *completely monotone* if:

$$(-1)^k \cdot (\Delta^k a)_n \geq 0 \quad \text{for all } n, k.$$

This formulation involves the difference operator Δ on sequences, given by $(\Delta a)_n = a_{n+1} - a_n$. This operator is iterated, in the standard way as $\Delta^0 = \text{id}$ and $\Delta^{m+1} = \Delta \circ \Delta^m$.

Theorem 9 (Hausdorff). *A sequence $a = (a_n)_{n \in \mathbb{N}}$ of non-negative real numbers a_n is completely monotone if and only if it arises as sequence of moments: there is a unique measure $\mu \in \mathcal{G}([0, 1])$ with, for each n ,*

$$a_n = \int_0^1 x^n \mu(dx). \quad \square$$

Proof (of Theorem 8). As a first step, we shall reason pointwise. Let, for each $K \in \mathbb{N}$, a distribution $\omega_K \in \mathcal{D}(\{0, 1, \dots, K\})$ be given, with $DD \ll \omega_{K+1} = \omega_K$.

We need to produce a unique (continuous) distribution $\omega \in \mathcal{G}([0, 1])$ such that $\omega_K = \text{binom}[K] \ll \omega$, for each K .

The problem of finding ω is essentially Hausdorff's moment problem, see Theorem 9. The connection between Hausdorff's moments problem and de Finetti's theorem is well known [6, §VII.4], but the connection is particularly apparent in our channel-based formulation.

We exploit Hausdorff's criterion via two claims, which we elaborate in some detail. The first one is standard.

Claim 1. If $\mu \in \mathcal{G}([0, 1])$ has moments \mathbf{m}_K , then:

$$\int_0^1 x^n \cdot (1-x)^k \mu(dx) = (-1)^k \cdot (\Delta^k \mathbf{m})_n.$$

This equation can be proven by induction on k . It holds by definition of \mathbf{m} for $k = 0$. Next:

$$\begin{aligned} & \int_0^1 x^n \cdot (1-x)^{k+1} \mu(dx) \\ &= \int_0^1 x^n \cdot (1-x) \cdot (1-x)^k \mu(dx) \\ &= \int_0^1 x^n \cdot (1-x)^k \mu(dx) - \int_0^1 x^{n+1} \cdot (1-x)^k \mu(dx) \\ &\stackrel{\text{(IH)}}{=} (-1)^k \cdot (\Delta^k \mathbf{m})_n - (-1)^k \cdot (\Delta^k \mathbf{m})_{n+1} \\ &= (-1)^{k+1} \cdot \left((\Delta^k \mathbf{m})_{n+1} - (\Delta^k \mathbf{m})_n \right) \\ &= (-1)^{k+1} \cdot (\Delta^{k+1} \mathbf{m})_n. \end{aligned}$$

The next claim involves a crucial observation about our setting: the whole sequence of distributions $\omega_K \in \mathcal{D}(\{0, 1, \dots, K\})$ is entirely determined by the probabilities $\omega_K(K)$ of drawing K balls which are all black.

Claim 2. The numbers $\mathbf{b}_K := \omega_K(K) \in [0, 1]$, for $K \in \mathbb{N}$, determine all the distributions ω_K via the relationship:

$$\omega_{k+n}(n) = \binom{k+n}{n} \cdot (-1)^k \cdot (\Delta^k \mathbf{b})_n.$$

In particular, the sequence \mathbf{b}_K is completely monotone.

The proof works again by induction on k . The case $k = 0$ holds by definition. For the induction step we use the cone property $DD \ll \omega_{k+1} = \omega_k$. Via (21) it gives:

$$\omega_k(i) = \frac{i+1}{k+1} \cdot \omega_{k+1}(i+1) + \frac{k+1-i}{k+1} \cdot \omega_{k+1}(i),$$

so that:

$$\omega_{k+1}(i) = \frac{k+1}{k+1-i} \cdot \omega_k(i) - \frac{i+1}{k+1-i} \cdot \omega_{k+1}(i+1).$$

We use this equation in the first step below.

$$\begin{aligned}
 & \omega_{(k+1)+n}(n) \\
 &= \frac{k+1+n}{k+1} \cdot \omega_{k+n}(n) - \frac{n+1}{k+1} \cdot \omega_{k+1+n}(n+1) \\
 &\stackrel{\text{(IH)}}{=} \frac{k+1+n}{k+1} \cdot \binom{k+n}{n} \cdot (-1)^k \cdot (\Delta^k \mathbf{b})_n - \frac{n+1}{k+1} \cdot \binom{k+1+n}{n+1} \cdot (-1)^k \cdot (\Delta^k \mathbf{b})_{n+1} \\
 &= \binom{k+1+n}{n} \cdot (-1)^{k+1} \cdot \left((\Delta^k \mathbf{b})_{n+1} - (\Delta^k \mathbf{b})_n \right) \\
 &= \binom{k+1+n}{n} \cdot (-1)^{k+1} \cdot (\Delta^{k+1} \mathbf{b})_n.
 \end{aligned}$$

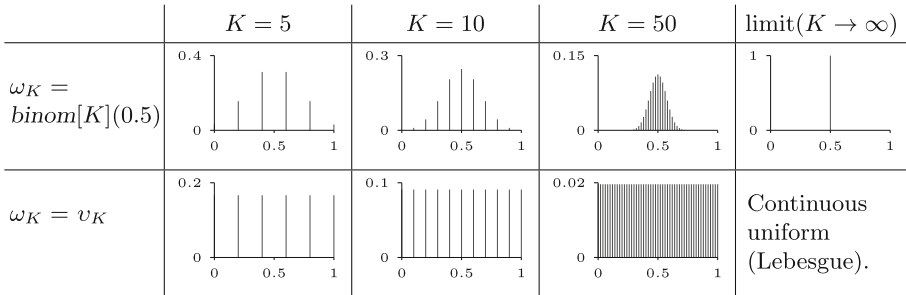
This second claim implies that there is a unique distribution $\omega \in \mathcal{G}([0, 1])$ whose K -th moments equals $\mathbf{b}_K = \omega_K(K)$. But then:

$$\begin{aligned}
 \omega_K(k) &= \omega_{(K-k)+k}(k) = \binom{K}{k} \cdot (-1)^{K-k} \cdot (\Delta^{K-k} \mathbf{b})_k && \text{by Claim 2} \\
 &= \binom{K}{k} \cdot \int_0^1 x^k \cdot (1-x)^{K-k} \, d\omega(x) && \text{by Claim 1} \\
 &= (\text{binom}[K] \ll \omega)(k).
 \end{aligned}$$

As an aside, it is perhaps instructive to look at how the distribution $\omega \in \mathcal{G}([0, 1])$ is built in the proof of Hausdorff’s criterion. One way is as a limiting measure of the sequence of finite discrete distributions on $[0, 1]$,

$$\bar{\omega}_K := \sum_{k=0}^K \omega_K(k) \Big| \frac{k}{K}. \tag{22}$$

This is illustrated in the following table, for the examples of the binomial cone and the uniform cone. The graphs show $\bar{\omega}_K$ (plotted as densities with respect to the counting measure).



We conclude by explaining why this construction gives a measurable mediating map. Suppose we have a measurable space (Y, Σ) and a sequence of channels $c_K : Y \rightarrow \{0, \dots, K\}$ forming a cone over (8), i.e. $c_K = DD \circ c_{K+1}$ for all $K \in \mathbb{N}$. We define a mediating channel $c : Y \rightarrow [0, 1]$ by putting the distribution $c(y)$ as the unique one with moments $c_K(y)(K)$, as above. It remains to show that c , regarded as a function $Y \rightarrow \mathcal{G}([0, 1])$, is a measurable function. It is immediate that the sequence-of-moments function $\mathbf{m} : Y \rightarrow [0, 1]^\infty$ given by:

$$\mathbf{m}(y)_K := \int_0^1 r^K c(x)(dr) = c_K(y)(K).$$

is measurable, because the c_K 's are assumed to be channels: let $V \subseteq [0, 1]$ be measurable; for associated basic measurable subsets of $[0, 1]^\infty$ one has:

$$\mathfrak{m}^{-1}([0, 1]^K \times V \times [0, 1]^\infty) = \{y \in Y \mid c_K(y)(K) \in V\} = c_K^{-1}([0, 1]^K \times V).$$

Further, $c: Y \rightarrow \mathcal{G}([0, 1])$ factors through \mathfrak{m} ; indeed $\mathcal{G}([0, 1])$ can be thought of as a subset of $[0, 1]^\infty$, when each distribution is regarded as its sequence of moments. We conclude by the less well-known fact that $\mathcal{G}([0, 1])$ is a *subspace* of $[0, 1]^\infty$. \square

Theorem 8 is a reformulation of a classical result of de Finetti [5]. It introduced the concept of *exchangeability* of random variables which has disappeared completely from our reformulation, but is implicit in our use of multisets. This requires some explanation.

A (discrete) random variable is a pair $\rho \in \mathcal{D}(A), U: A \rightarrow \mathbb{R}$ of a distribution and an observable. Usually the distribution ρ is implicit. Instead, the notation $P(U)$ is used for the image distribution $\mathcal{D}(U)(\rho)$ on \mathbb{R} , also sometimes written as $P(U = r)$, for $r \in \mathbb{R}$.

Let $(U_1, \rho_1), \dots, (U_n, \rho_n)$ be a finite sequence of random variables. It induces a joint distribution on \mathbb{R}^n ,

$$P(U_1, \dots, U_n) := \mathcal{D}(U_1 \times \dots \times U_n)(\rho_1 \otimes \dots \otimes \rho_n),$$

where $U_1 \times \dots \times U_n: A_1 \times \dots \times A_n \rightarrow \mathbb{R}^n$ and $\rho_1 \otimes \dots \otimes \rho_n \in \mathcal{D}(A_1 \times \dots \times A_n)$.

Such a sequence (U_i, ρ_i) is called *exchangeable* if for each permutation π :

$$P(U_1, \dots, U_n) = P(U_{\pi(1)}, \dots, U_{\pi(n)}).$$

An infinite sequence is called exchangeable if each finite subsequence is exchangeable.

The original form of de Finetti's theorem involves binary (0–1) random variables with $U_i: A_i \rightarrow 2 = \{0, 1\} \hookrightarrow \mathbb{R}$. The distribution $P(U_i)$ is then on a 'Bernoulli' distribution on 2. When random variables are exchangeable their order doesn't matter but only how many of them are zero or one. This is captured by their sum $\sum_i U_i$. But this sum is a *multiset* over 2. Indeed, in a multiset, the order of elements does not matter, only their multiplicities (numbers of occurrences). This is a crucial observation underlying our reformulation, that already applied to the accumulator map acc in (4). But let's first look at the usual formulation of de Finetti's (binary) representation theorem, see [5] and the textbooks [6, 16].

Theorem 10. *Let U_1, U_2, \dots be an infinite sequence of exchangeable binary random variables. Then there is a probability measure ω on $[0, 1]$ such that for all n ,*

$$P(U_1, \dots, U_n) = \int_0^1 x^{\sum_i U_i} (1-x)^{(n-\sum_i U_i)} \omega(dx).$$

It is clear that there is a considerable gap between this original formulation and the inverse limit formulation that has been developed in this paper. There is room for narrowing this gap.

The essence that is used in de Finetti’s theorem about n exchangeable binary random variables are their outcomes via the accumulator map acc from (4), giving a map $\mathcal{D}(acc): \mathcal{D}(2^n) \rightarrow \mathcal{D}(\mathcal{M}[n](2))$.

In our approach we have reformulated the requirements about random variables and about exchangeability and worked directly with distributions on multisets, see Proposition 2 and Theorem 8. De Finetti’s (binary) representation Theorem 10 becomes a limit statement in a Kleisli category. We think that this approach brings forward the mathematical essentials.

7 A Final Exchangeable Coalgebra

We return to the setting of coalgebras, begun in Sect. 4. Any coalgebra $h: X \multimap 2 \times X$ for the functor $2 \times (-)$ on $\mathcal{Kl}(\mathcal{D})$ or on $\mathcal{Kl}(\mathcal{G})$ is a probabilistic process that produces a random infinite sequence of boolean values. Following [14], we can understand this induced sequence as arising from the fact that the pair of head and tail function

$$2^\infty \xrightarrow{\cong} 2 \times 2^\infty$$

is a final coalgebra in the category $\mathcal{Kl}(\mathcal{G})$, and so there is a unique channel $h^\sharp: X \multimap 2^\infty$ assigning to each element of X the random sequence that it generates. Here 2^∞ is the measurable space of infinite binary streams, regarded with the product σ -algebra.

It is helpful to spell this construction out a little. The sequence of first projections

$$2^0 \longleftarrow 2^1 \longleftarrow 2^2 \longleftarrow 2^3 \longleftarrow \dots \tag{23}$$

has as a limit the space 2^∞ , with the limiting cone being the obvious projections. And this limit is also a limit in $\mathcal{Kl}(\mathcal{G})$. (This is a categorical formulation of Kolmogorov’s extension theorem, see [4, Thm. 2.5].)

Now, as usual in coalgebra theory, we find $h^\sharp: X \multimap 2^\infty$ for any coalgebra $h: X \multimap 2 \times X$ by noticing that the channels $h_K^\sharp: X \multimap 2^K$ (15) define a cone over the chain (23). The limiting channel $h^\sharp: X \multimap 2^\infty$ for the cone is the unique coalgebra homomorphism to the final coalgebra.

Example 11. We define the *bernoulli coalgebra* to be the channel $bern: [0, 1] \multimap 2 \times [0, 1]$ given by $bern(r) := r \cdot |1, r\rangle + (1 - r) \cdot |0, r\rangle$. In words, $bern(r)$ makes a coin toss with known bias r and returns the outcome together with the same bias r unchanged.

The final channel $iid := bern^\sharp: [0, 1] \multimap 2^\infty$ takes a probability r to a random infinite sequence of independent and identically distributed binary values, all distributed as $r|1\rangle + (1 - r)|0\rangle$. More explicitly, for basic measurable subsets, at each stage n ,

$$iid(r)(2^n \times \{1\} \times 2^\infty) = r \quad \text{and} \quad iid(r)(2^n \times \{0\} \times 2^\infty) = 1 - r.$$

So, any coalgebra defines a cone over (23). We have shown in Lemma 6 that any *exchangeable* coalgebra defines a cone over the draw-delete chain (6). We use this to show that the bernoulli coalgebra $bern: [0, 1] \dashrightarrow 2 \times [0, 1]$ is the final exchangeable coalgebra.

The final coalgebra $2^\infty \dashrightarrow 2 \times 2^\infty$ is not exchangeable. To see this, consider the infinite alternating sequence $01010101 \dots \in 2^\infty$, so that the left hand side of (14) at the sequence gives $(0, 1, 010101 \dots)$, whereas the right hand side of (14) at the sequence gives $(1, 0, 010101 \dots)$, which is different.

The bernoulli coalgebra $bern: [0, 1] \dashrightarrow 2 \times [0, 1]$ is exchangeable, because

$$\begin{aligned} (\text{id} \otimes \text{bern}) \ll \text{bern}(r) &= r^2 |1, 1, r\rangle \\ &\quad + (1-r)r |0, 1, r\rangle + r(1-r) |1, 0, r\rangle \\ &\quad + (1-r)^2 |0, 0, r\rangle. \end{aligned}$$

Notice that reordering the outputs does not change the statistics.

Theorem 12. *The bernoulli coalgebra $bern: [0, 1] \dashrightarrow 2 \times [0, 1]$ is the final exchangeable coalgebra.*

Proof. Let $h: X \dashrightarrow 2 \times X$ be an exchangeable coalgebra. By Lemma 6 the associated maps $h_K: X \dashrightarrow \mathcal{M}[K](2)$ form a cone. By de Finetti's theorem in limit form—that is, by Theorem 8—we get a unique mediating channel $f: X \dashrightarrow [0, 1]$ with $\text{binom}[K] \circ f = h_K$ for each K . Our aim is to show that f is a unique coalgebra map: $bern \circ f = (\text{id} \otimes f) \circ h$.

We do know that the limiting map $X \xrightarrow{f} [0, 1] \xrightarrow{iid} 2^\infty$ is a unique coalgebra homomorphism. So to show that f is a unique coalgebra homomorphism, it suffices to show that iid is a monomorphism in $\mathcal{Kl}(\mathcal{G})$.

To see this, first notice that the K -th moment of some distribution $\omega \in \mathcal{G}([0, 1])$ is the probability $(iid \ll \omega)(\{1\}^K \times 2^\infty)$, and so if two random sequences are equal, $iid \ll \omega_1 = iid \ll \omega_2$, then in particular the moments of ω_1 and ω_2 are equal and so $\omega_1 = \omega_2$. Moreover the endofunctor $2 \times -$ preserves monomorphisms. So if the final channel $X \dashrightarrow 2^\infty$ factors through $iid: [0, 1] \dashrightarrow 2^\infty$ as a channel then it also factors through as a coalgebra homomorphism, necessarily uniquely, by the definition of monomorphism. \square

Example 13. We return to the example of Pólya's urn coalgebra $pol: \mathcal{M}_*(2) \dashrightarrow 2 \times \mathcal{M}_*(2)$. We have shown in Proposition 4 that the cone $pol_K: \mathcal{M}_*(2) \rightarrow \mathcal{M}[K](2)$ defined by iterated draws is given by the beta-bernoulli distributions: $pol_K(\varphi) = \text{betabern}[K](\varphi(0), \varphi(1))$, and we have shown in Proposition 7 that the beta-bernoulli cone over (6) has mediating map the beta distribution. So the final channel $beta: \mathcal{M}_*(2) \dashrightarrow [0, 1]$ is given by putting $beta(\varphi)$ as the beta distribution with parameters $(\varphi(0), \varphi(1))$, and this is a coalgebra homomorphism. The commuting diagram

$$\begin{array}{ccc} \mathcal{M}_*(2) & \xrightarrow{pol} & 2 \times \mathcal{M}_*(2) \\ \text{beta} \circ \downarrow & & \downarrow \circ 2 \times \text{beta} \\ [0, 1] & \xrightarrow{bern} & 2 \times [0, 1] \end{array}$$

expresses the ‘conjugacy’ between the bernoulli and beta distributions, see also [13].

8 Conclusions and Further Work

This paper has translated the classical representation theorem of de Finetti into the language of multisets and distributions and turned it into a limit theorem. This limit is subsequently used for a finality result for exchangeable coalgebras.

We have concentrated, as is usual for the de Finetti theorem, on the binary case, giving a limit statement for the special case $X = 2 = \{0, 1\}$ of the cone in Proposition 2. There is an obvious question whether this formulation can be extended to $X = \{0, 1, \dots, n - 1\}$ for $n > 2$, or to more general measurable spaces. We see two possible avenues:

- extension to finite sets via a multivariate version of Hausdorff’s moments theorem, see [15];
- extension to Polish spaces, following [2] and its categorical justification for extending de Finetti’s theorem to such spaces (esp. Theorems 35 and 40) in terms of exchangeability.

Another point for future work is to investigate whether the categorical formulation of de Finetti’s theorem could be taken as an axiom for a ‘synthetic probability theory’ (e.g. following [7, 8, 17]), and indeed whether it holds in some proposed categorical models which do support de-Finetti-like arguments [11, 18].

Finally, we note that chains of channels between multisets also arise in recent exponential modalities for linear logic [3, 9]; the connections remain to be worked out.

Acknowledgements. Staton is supported by a Royal Society Fellowship, and has enjoyed discussions about formulations of de Finetti’s theorem with many people including coauthors on [18].

References

1. Ackerman, N., Freer, C., Roy, D.: Exchangeable random primitives. In: Proceedings of PPS 2016 (2016)
2. Dahlqvist, F., Danos, V., Garnier, I.: Robustly parameterised higher-order probabilistic models. In: Desharnais, J., Jagadeesan, R. (eds.) CONCUR 2016 – Concurrency Theory. LIPIcs, vol. 59, pp. 23:1–23:15. Schloss Dagstuhl (2016)
3. Dahlqvist, F., Kozen, D.: Semantics of higher-order probabilistic programs with conditioning. In: Principles of Programming Languages, pp. 57:1–57:29. ACM Press (2020). <https://doi.org/10.1145/3371125>
4. Danos, V., Garnier, I.: Dirichlet is natural. In: Ghica, D. (ed.) Mathematical Foundations of Programming Semantics, number 319 in Electronic Notes in Theoretical Computer Science, pp. 137–164. Elsevier, Amsterdam (2015)

5. de Finetti, B.: Funzione caratteristica di un fenomeno aleatorio. *Memorie della R. Accademia Nazionale dei Lincei*, IV, fasc. 5, 86–113 (1930). www.brunodefinetti.it/Opere/funzioneCaratteristica.pdf
6. Feller, W.: *An Introduction to Probability Theory and Its applications*, vol. II. Wiley, Hoboken (1970)
7. Fritz, T.: A synthetic approach to Markov kernels, conditional independence, and theorems on sufficient statistics. arxiv.org/abs/1908.07021 (2019)
8. Fritz, T., Rischel, E.: The zero-one laws of Kolmogorov and Hewitt-Savage in categorical probability. arxiv.org/abs/1912.02769 (2019)
9. Hamano, M.: A linear exponential comonad in s-finite transition kernels and probabilistic coherent spaces. [arxiv:1909.07589](https://arxiv.org/abs/1909.07589), September 2019
10. Hausdorff, F.: Summationsmethoden und Momentfolgen I. *Math. Zeitschr.* **9**, 74–109 (1921)
11. Heunen, C., Kammar, O., Staton, S., Yang, H.: A convenient category for higher-order probability theory. In: *Logic in Computer Science*, pp. 1–12. IEEE Computer Society (2017)
12. Jacobs, B.: *Structured probabilistic reasoning* (2019, forthcoming book). <http://www.cs.ru.nl/B.Jacobs/PAPERS/ProbabilisticReasoning.pdf>
13. Jacobs, B.: A channel-based perspective on conjugate priors. *Math. Struct. Comput. Sci.* **30**(1), 44–61 (2020)
14. Kerstan, H., König, B.: Coalgebraic trace semantics for probabilistic transition systems based on measure theory. In: Koutny, M., Ulidowski, I. (eds.) *CONCUR 2012*. LNCS, vol. 7454, pp. 410–424. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-32940-1_29
15. Kleiber, C., Stoyanov, J.: Multivariate distributions and the moment problem. *J. Multivar. Anal.* **113**, 7–18 (2013)
16. Klenke, A.: *Probability Theory*. Springer, London (2013). <https://doi.org/10.1007/978-1-4471-5361-0>
17. Ścibior, A., et al.: Denotational validation of higher-order Bayesian inference. In: *Principles of Programming Languages*, pp. 60:1–60:29. ACM Press (2018)
18. Staton, S., Stein, D., Yang, H., Ackerman, N., Freer, C., Roy, D.: The Beta-Bernoulli process and algebraic effects. In: Chatzigiannakis, I., Kaklamanis, C., Marx, D., Sannella, D. (eds.) *International Colloquium on Automata, Languages and Programming*. LIPIcs, vol. 107, pp. 141:1–141:15. Schloss Dagstuhl (2018)